



RESEARCH REPORT

Artificial Intelligence Use Cases

215 Use Case Descriptions, Examples, and Market Sizing and Forecasts Across Enterprise, Consumer, and Government Markets

Published 3Q 2017

JESSICA GROOPMAN
Principal Analyst

ADITYA KAUL
Research Director

SECTION 1

EXECUTIVE SUMMARY

1.1

INTRODUCTION

Defining artificial intelligence (AI) is a lot like defining intelligence; it is rarely agreed upon and manifests differently in different contexts. Tractica defines AI as an information system that is inspired by a biological system designed to give computers the human-like abilities of hearing, seeing, reasoning, and learning. These capabilities are powered by a range of technologies, such as machine learning (ML), deep learning (DL), computer vision (CV), natural language processing (NLP), machine reasoning (MR), and strong AI, all of which fall under the AI umbrella.

Vast amounts of data, faster processing power, and increasingly smarter algorithms are powering AI applications and use cases across consumer, enterprise, and government markets around the world. Based on our research and forecasting, Tractica believes the opportunity for AI spans a wide range of industries and geographies and is particularly disruptive in highly domain-specific markets with high-volume data needs and ontologies, as well as those with growing applications for machine perception. From autonomous robotics to algorithmic news stories, from product recommendations to processing patient data, and from virtual assistants to voice recognition, AI is widely considered one of, if not *the next big technological shift*, on par with past shifts like the industrial revolution, the computer age, and the smartphone revolution.

Across 29 industries, Tractica's research into AI has identified more than 200 use case categories, each of which is explored in this report. The report defines, contextualizes, and offers real-world examples and revenue forecasts for each use case organized by industry. It serves as a referential compendium to Tractica's ongoing market forecasting of the AI space, offering an overview and analysis for each use case included in the model.

1.2

ARTIFICIAL INTELLIGENCE EXPANDS ACROSS INDUSTRIES

Although AI has been around for decades, it is the convergence of three independent trends that has brought about an explosion in the market. More data, faster hardware, and better algorithms are accelerating research, development, and commercial investment in AI applications at lightning speeds. Those sectors already leading in the digital space are accelerating in AI adoption, as the question of how to better use and monetize data persists. Tractica's quantitative market assessment forecasts that annual revenue generated from the direct and indirect application of AI software will increase from \$1.38 billion in 2016 to \$59.75 billion by 2025.

As the scope and velocity of the AI market expands, it can be challenging for suppliers and adopters alike to keep up. The dynamics or developments in one sector or technology can influence another; opportunities for multi-disciplinary collaboration or risk mitigation are coalescing; and the very definition of digital transformation is evolving. In the age of colossal data and rapidly shifting customer expectations, companies must navigate the hype, adopt new capabilities, and adapt their strategies, all while proving efficiencies and new revenue.

Tractica's in-depth analysis of more than 200 use cases highlights the emergence of a number of overarching themes, illustrating critical dynamics to watch across the broader AI market. A summary of these trends includes:

- **All AI falls into three macro categories:** Big Data, vision, and language. Although most think AI is driven by Big Data analytics, the larger growth area has to do with vision and language perception capabilities, which will feed longer-term growth and strong AI.
- **AI applications mark the next evolutionary step in digital transformation:** Computing, sensing, networking, and data generation were only the beginning. The ability to process data more quickly and intelligently across systems, leveraging hardware, sensors, and cameras, and to digitize language itself marks the next era of organizational transformation.
- **AI is shorthand for a combination of technologies:** Use cases most often consist of multiple types of AI applied or configured in conjunction with one another and other technologies. For example, ML, CV and sensors; or DL and NLP.
- **AI can be overt and visible or implicit and invisible:** For end users, AI interactions like robotics or autonomously moving machines are obvious, even tangible; but AI can also support Big Data analysis, real-time responses, systems management, and many other invisible means of processing data.
- **AI-driven personalization and operations automation will become interconnected:** Advanced AI deployments will be marked by the ability to infuse both user-facing services and interactions with back-end or enterprise process and supply chain optimization, such as in retail, financial services, energy, and healthcare.
- **AI maturity is highly fragmented:** Maturity and the metric for success vary widely from application to application. Relatively low-stakes applications, such as movie recommendations, are widely accepted and optimized, while others like credit scoring or medical treatment recommendations remain regulatory grey areas and face significant barriers to widespread adoption.
- **AI's ability to pass the Turing Test is also fragmented:** When it comes to machines' abilities to seamlessly interact as a human would, the jury is still out. While social media bots have effectively passed for millions of Twitter or Facebook users, neither robots nor chatbots are very close to disguising their code-based composition.
- **AI's manifestation will shift alongside other technology macrotrends:** AI is not the only show in town; numerous other technologies (e.g., the Internet of Things (IoT), augmented reality (AR), virtual reality (VR), cameras, blockchain, renewable energy, genomics, three-dimensional (3D) printing, etc.) will both leverage and influence AI's development, adoption, and regulation
- **AI is an extension of brand interactions:** As more companies deploy AI, specifically virtual agents to power consumer-facing functions, services, products, and touchpoints, brands must balance unprecedented opportunities for personalization with significant risk of failure, faux pas, or backlash.

- **AI is alluring, particularly in hyper-competitive markets:** It is not just greater automation and operational efficiencies that AI suppliers promise adopters, it is the ability to illuminate hidden patterns and big “dark” unstructured data sets, to simulate scenarios for decision-making, and enable altogether new products. Beware the many ways AI is oversold.
- **AI promises both diverse benefits and diverse challenges.** Across use cases, profound opportunities lie in forecasting, empirical decision-making, operations automation, product optimization, new business models, greater access to services, targeted services, enhanced user experiences, and even improved environmental and public health. Simultaneously, it poses urgent challenges: data integrity, reskilling workforces, diverse ethical uncertainties, privacy concerns, unchartered legal and regulatory questions or standards, and the explainability and accountability of deep neural networks, among others.
- **AI will have a complex relationship with humans that will change over time:** While certain jobs will become automated, AI is more often poised to augment human labor and decision-making. Longer-term, many applications will be designed to empower humans with non-human capabilities, memory, experiences, and knowledge. Many ethical, philosophical, cultural, societal, and business norms will be forced into re-assessment.

1.3

MARKET FORECAST

Tractica’s market forecast is focused on identifying the software, hardware, and services revenue opportunity for AI. Using a bottom-up, use case-based model that classifies and estimates the revenue potential for each use case, rolled up by industry, technology, and world region, Tractica estimates overall AI market revenue from 2016 to 2025.

The revenue for each use case described in this report represents software revenue, which is accounted for as direct or indirect revenue. Direct revenue represents the income derived from the sales of an AI-led solution, where AI is the key value being sold and marketed. For example, emotion analysis, legal contract analysis, or cybersecurity threat estimation are services where AI is being sold as the key value proposition. Indirect revenue is counted in cases where AI is not necessarily the key value proposition, but AI is a layer or plugin that enhances an existing application or service. In other words, for indirect revenue, the use cases are AI-enabled rather than AI-led. For example, Google search, Amazon product recommendations, and Facebook news feeds are existing services where AI does not define the end product, but is a way of enhancing it.

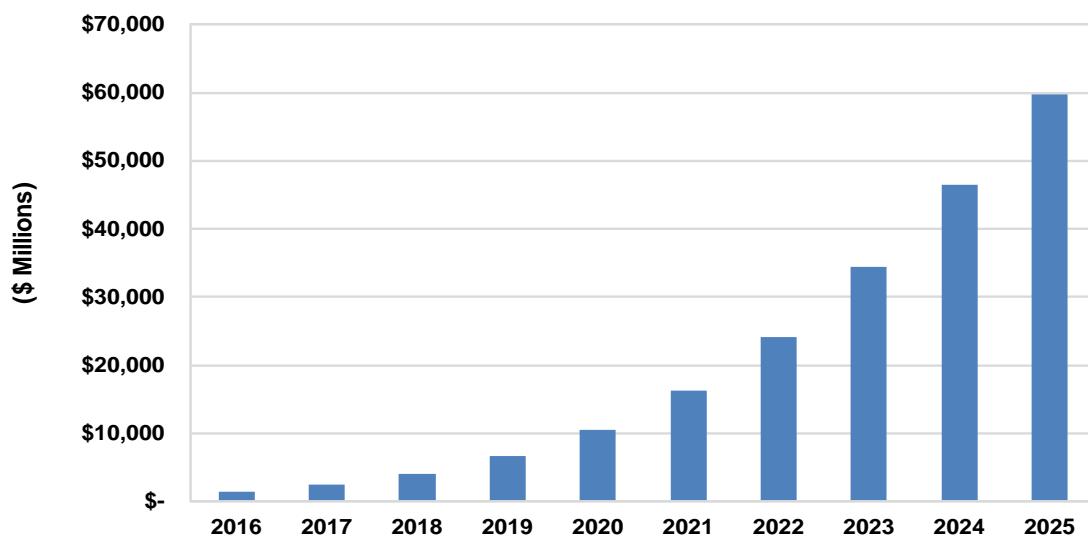
The forecasts throughout this report are snapshots from Tractica’s 3Q 2017 edition of the [Artificial Intelligence Market Forecasts](#) report.

1.3.1

TOTAL REVENUE FOR ARTIFICIAL INTELLIGENCE

Tractica forecasts that the revenue generated from the direct and indirect application of AI software will grow from \$1.38 billion in 2016 to \$59.75 billion by 2025. This represents a significant growth curve for the forecast period with a compound annual growth rate (CAGR) of 52%.

Chart 1.1 Artificial Intelligence Software Revenue, World Markets: 2016-2025



(Source: Tractica)

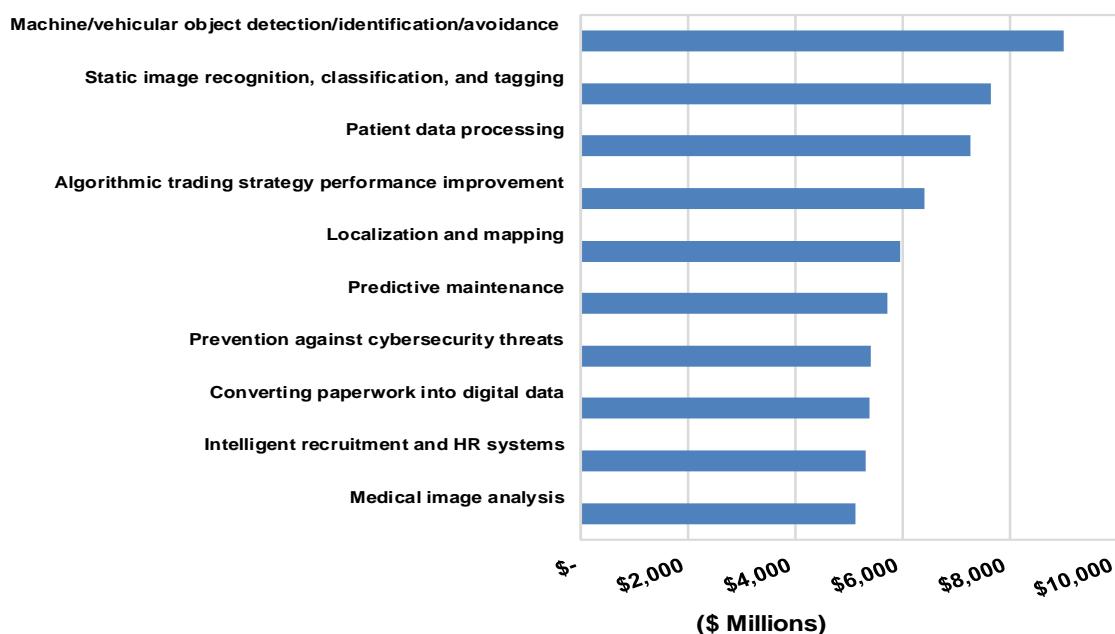
1.3.2

TOP 10 USE CASES FOR ARTIFICIAL INTELLIGENCE

This report provides a qualitative assessment of the market opportunity for AI across more than 200 distinct use cases in 29 industries. To view 2016 to 2025 revenue for each use case, reference the table at the end of each respective use case description.

Across all applications for AI, Tractica has also ranked the top 10 use cases, ranked by cumulative revenue accrued during the period from 2016 to 2025.

Chart 1.2 Cumulative Artificial Intelligence Software Revenue, Top 10 Use Cases, World Markets: 2016-2025



(Source: Tractica)

SECTION 2

ARTIFICIAL INTELLIGENCE USE CASES

2.1 OVERVIEW

This report provides a qualitative overview for each use case Tractica has identified in the AI market. Each use case assessment includes a description, industry context, and considerations for AI application, an example, and global revenue forecasts for 2016 to 2025. Use cases are organized and revenue is calculated according to industry. Many use cases appear across multiple industries. Tractica's assessment of each use case involved both primary and secondary research and was conducted between March and August of 2017.

2.2 ADVERTISING

2.2.1 AD INSERTIONS INTO IMAGES AND VIDEO

Advertisers and brands have been working to optimize ad placement on the internet for years, but the industry remains rife with challenges around tracking, visibility, data management, and attribution.

Companies are increasingly using AI and DL to detect patterns and infer opportunities for ad insertion into images and video consumed by customers and prospects. Using image recognition, classification, and tagging helps companies automate what ads to place where, when, and for whom, and to drive intended actions. Facebook uses AI to look for text on an image used for advertising, and labels it as “high text,” “medium text,” or “low text,” helping advertisers achieve a higher success rate with ads that have low text.

CV specialist GumGum uses AI to embed ads or links into photos where it finds relevance and helps brands target and expand their advertising. It has used the technique to post ads about an upcoming TV series on targeted photos that featured the star of the show. Kaltura is using AI to power similar real-time ad placements for live video streams like games or concerts, without requiring media companies to pre-schedule placements. Providers like these and others are working across web-based channels today, while developing similar techniques in connected TVs and VR headsets.

This use case represents significant market opportunity given its appeal in both automating and increasing control in programmatic advertising workflows. Companies do run the risk of oversaturating consumers with advertisements, particularly as data inputs mature to make advertising more personalized, and new channels (e.g., VR) emerge where gratuitous advertising could slow adoption.

Table 2.1 Ad Insertions into Images and Video in Advertising, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.02	8.34	22.29	47.83	91.83	161.07	256.68	368.91	480.11	576.01	102.3%

(Source: Tractica)

2.2.2

HUMAN EMOTION ANALYSIS

It is no secret that humans are emotional creatures, often motivated more by emotion than pragmatism when making purchase decisions. Economists and advertisers have understood this for years. But as advertising has become increasingly digitized and as companies of all sizes seek advertising at scale, staying emotionally in-tune with consumer segments has grown more difficult (and has not been without many gaffs going viral).

Although computers are far better at calculating statistical probabilities than anything resembling emotion, developers are working to train models to recognize, categorize, and tag human emotions so that algorithms can make decisions based on such categorizations. Techniques could involve CV, DL, or NLP, or even robotics depending on the use case. While this is an emerging and controversial area of AI, early studies show computers are very adept at identifying human emotions. As a result, more and more companies are turning to AI to aid in the quest to better understand, predict, advertise, and display ads based on human emotions.

RealEyes uses AI to tell how people feel when they see static or video content or hear audio content. Using webcams, the company sources an audience of 300 in a targeted geography and use algorithms to process and analyze facial expressions. The company then delivers reports with insights and content, distribution, and targeted recommendations around creative testing and media planning. In a recent partnership with Heineken and its media partner, AOL, testing the emotional resonance of video content to inform ad spending, the insights generated in RealEyes' report led the companies to reduce spending on short trailers and invest in longer-form episodic content instead, which drove 2X click-thru-rates (CTRs), 3X more social action conversion, and 6X actions taken on the content itself. At the time of this report's writing, the episodes had more than 35 million views over 208 countries.

The broader use case of emotion analysis, whether in advertising, investment, healthcare, or otherwise, is a controversial one, rife with human doubt and privacy concerns. Can we trust algorithms to accurately identify how we truly feel? Given the great diversity of cultural and social nuance, how can we, and advertisers, know with certainty that algorithms are accurate at scale? While advertising for goods and services based on emotion may be relatively low stakes, how could the same technology be applied in contexts that limit access: insurance, education, employment eligibility, etc.? These remain critical questions for all constituencies to explore when using AI to ascertain our most qualitative states: feelings.

Tractica forecasts that the annual revenue for human emotion analysis in advertising will increase from \$18.07 million worldwide in 2016 to \$378.52 million in 2025.

Table 2.2 Human Emotion Analysis in Advertising, World Markets: 2016-2025

Units											CAGR (2016- 2025)
	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	
(\$ Millions)	18.07	23.26	32.56	48.99	76.80	120.13	179.69	249.48	318.68	378.52	40.2%

(Source: Tractica)

2.2.3

INTERACTIVE WINDOW DISPLAYS

Perhaps one of the oldest forms of advertising, window displays have been marketing goods and services to passersby for centuries. The problem with this historically analog approach is limitations in dynamism: content must be manually replaced and the message stays the same for all.

With the rise of more intelligent TVs, equipped with sensors, cameras, and software integrations, the digital signage industry has grown significantly in the last 5 years. Retailers, restaurants, hotels, and even municipal environments like parks and train stations are all adopting digital signage solutions that facilitate fresh content, real-time messaging, and dynamic displays. AI enhances this through applications in computer vision, sensor data, voice recognition, and other techniques enabling user interactions.

Figure 2.1 Nike Window Display Gamifies Shoppers' Interactions Using Motion Detection



(Source: Nike & CoDesign)

Nike recently created a series of interactive window displays to snag the attention of the thousands of people who walk by its brick and mortar (B&M) location in central London. One such display, developed with staat, a Dutch creative agency, invited passersby to stand on a blue dot and jump as high as possible, then through Kinect-powered motion tracking, users can save their scores, see them on the screen, compare their ranking to others, essentially connecting offline and online experiences.

The opportunity for interactive window displays is one of top-of-funnel brand engagement and recognition. Using AI to enhance consumer experiences in the physical world, while fostering learning and digital insights on the back end, is one way brands are beginning to integrate brick with click. While installation costs can be high, driving awareness of and foot traffic into stores pays dividends.

Tractica forecasts that the annual revenue for interactive window displays in advertising will increase from \$0.12 million worldwide in 2016 to \$6.95 million in 2025.

Table 2.3 Interactive Window Displays in Advertising, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.12	0.26	0.50	0.92	1.57	2.53	3.75	5.04	6.16	6.95	56.5%

(Source: Tractica)

2.2.4

PERFORMANCE REPORTING AND ANALYTICS OF AD CAMPAIGNS

Advertisers and publishers have been automating the planning, purchase, and optimization of ad placement for years, indeed more than half of display ads purchased in the United States today are done so automatically already. Yet, programmatic advertising and the ad tech industry have struggled with analytics and reports that can adequately convey campaign performance, particularly across multiple screens and devices. One reason is that the industry is, by and large, not saturated with adequate talent in data science.

Anyone in advertising understands that the massive amounts of data generated in advertising are far beyond human capability for analysis, yet well suited for AI. AI is being used for reporting and analytics for cross-screen targeted advertising. Specifically, ML and DL help process more diverse data more efficiently for better targeted advertising across multiple screens, to analyze, recommend, and automatically optimize customer and prospect profiles; targeting, channeling tactics, and increasing conversion toward intended actions.

Appier is using AI to resolve the complexity and difficulty of effectively advertising across multiple screens through enhanced analytics and reporting. One of its clients, Estee Lauder, recently leveraged the product to gain brand awareness and increase customer insights and mailing list conversion, while keeping cost-per-click (CPC) and cost-per-lead (CPL) costs at a minimum. Appier used AI to identify all devices owned by individual users and run advanced re-marketing based on behaviors, while also using a “look-alike” simulation feature to source prospects with similar attributes and profiles. All of these interactions were analyzed and optimized over time, yielding the following results:

- Reduced CPC by 43% and CPL by 63% compared to targets
- Increased # of clicks by 74%; # of leads by 167% compared to targets
- Cross screen conversion across three devices were:
 - 11X higher than conversions on PCs
 - 4X higher than conversions on tablets
 - 3X higher than conversions on smartphones

Applying AI to advertising analytics is rife with opportunity, given the massive amounts of data and investment already flowing through the programmatic ad space. As with other advertising and consumer-facing use cases, risks around creep, over-personalization, and explainability of algorithms remain critical for advertisers and publishers to address.

Tractica forecasts that the annual revenue for performance reporting and analytics of ad campaigns in advertising will increase from \$4.88 million worldwide in 2016 to \$353.25 million in 2025.

Table 2.4 Performance Reporting and Analytics of Ad Campaigns in Advertising, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	4.88	9.46	18.05	33.62	60.34	102.27	160.11	227.96	295.21	353.25	60.9%

(Source: Tractica)

2.2.5

QUERYING IMAGE CONTENT

For many years, images and video content were largely muted when it came to the ability of advertisers and publishers to extract analytics and insights about their composition, or to use that to inform search results. Instead, text analytics and structured data analysis (e.g., conversions, quantitative assessments, etc.) dominated how advertisers and search engines determined what content to serve to users.

The text query of images is specifically related to search-based and social media advertising with AI. But rather than the general use case of tagging and classifying images, and using a search term to offer pre-classified and tagged images as end results for a search-based or social media ad, text query of images is related to understanding what the image contains. For example, someone wants to know the car featured in a particular image. Typing a text query like “What brand is the red car in this image?” should return the brand result and a targeted ad featuring that brand. In this case, AI goes one step further than using tagged and classified images, performing an analysis of the image itself in real time, and using the results to offer ads. There could be other ads targeted, such as someone asking “Where is this beach located?” or possibly go a step further by asking “Can you find me tickets in December to the location featured in this photo?”, where AI is trying to understand the meaning of the sentence, which is finding good airfare deals, while trying to parse through the location featured in the image.

This technology is still in its infancy today, but is expected to play a big role in the coming years as computers begin to have a deeper understanding of images and what is contained within them. Advertising giants like Google and Facebook are leading research in this area of context-based understanding of images and text.

Tractica forecasts that the annual revenue for querying image content in advertising will increase from \$0.61 million worldwide in 2016 to \$964.3 million in 2025.

2.2.6

STATIC IMAGE RECOGNITION, CLASSIFICATION, AND TAGGING

Recognizing, classifying, and tagging images has been a job left to humans for years. Humans, after all, are able to exercise discretion and simple look and tag tasks were cheap relative to the value they enabled at scale. But it is precisely this sort of tedious job that AI now threatens.

Thanks to various combinations of ML, DL, NLP, and CV, computers are now powering many types of image recognition at scale. In advertising, many algorithms are used to improve advertising by tagging and classifying images, or suggesting improvements to calculate the optimal ad to show to the current user at the present time. Typically, these algorithms focus on variations of the ad, optimizing different properties, such as background color, image size, or a set of images. Companies like Google, Facebook, and Yahoo are actively using image recognition and classification algorithms to improve advertising by tagging and classifying images, or suggesting improvements.

CV specialist GumGum uses AI to embed ads or links into photos where it finds relevance and helps brands target and expand their advertising. It has used the technique to post ads about an upcoming TV series on targeted photos that featured the star of the show. Facebook uses AI to look for text on an image used for advertising, and labels it as “high text,” “medium text,” or “low text,” helping advertisers achieve a higher success rate with ads that have low text.

Another company, Ditto Labs, uses DL to identify company brand logos in photos posted on social media. The software then evaluates the environments and related sentiments in which the brands appear, and then offers companies the ability to target advertising campaigns accordingly and compare brand performance against competitors.

While models occasionally categorize images incorrectly, even absurdly, these are typically distant outliers compared to the number of images processed in a given application. In advertising, the risks for AI-driven image classification are fairly low, but new ethical concerns arise in terms of how categorizations involving users could be used in unforeseen ways. Tractica's analysis of cumulative revenue across world markets finds image recognition, classification, and tagging as the top grossing use case over the next decade.

Tractica forecasts that the annual revenue for static image recognition, classification, and tagging in advertising will increase from \$30.53 million worldwide in 2016 to \$1.15 billion in 2025.

Table 2.6 Static Image Recognition, Classification, and Tagging in Advertising, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	30.53	45.73	73.69	123.92	209.65	343.84	528.71	745.50	960.39	1,145.98	49.6%

(Source: Tractica)

2.2.7 TARGETED ADVERTISING USING MULTI-DOMAIN CUSTOMER DATA (SOCIAL, WEB, CONTEXT)

Pulling together as many diverse and disparate data sets to ascertain customer behavior has long been a chief objective for advertisers and marketers. After all, leveraging all of these data to target the right person with the right offer at just the right time remains the proverbial “holy grail” of data-driven advertising. But given the vast amounts and varied channels of data generation, and the inadequate training in data science, advertisers have struggled to effectively wield data across multiple domains.

AI offers new solutions to this challenge in the form of language processing and Big Data analysis, both of which help to more rapidly process unstructured data. AI is also being applied to network selection, sometimes called ad mediation, where bots can now automate the (once manual) process of sorting ads according to website and, ultimately, the consumer. When it comes to AI-enhanced programmatic ad targeting, ML is designed to increase the likelihood a user will click. This is accomplished through optimizing content displayed when retargeting and determining what copy is most effective when, where, and for whom. Algorithms are also designed to optimize bids for advertisers in order to achieve the best cost-per-acquisition (CPA) from the available inventory. Only the most relevant ads deployed, instructed by keywords associated with contextual data like a website, past history, geography, and timing, will be pulled in from multiple sources. Between more accurate ad positioning and insertions (outlined in section 2.21), and more efficient network

selection, ads are optimized based on actual data and human error is reduced.

Dole Packaged Foods Asia recently ran an AI-driven campaign with enterprise marketing platform, Albert. The campaign was designed to drive awareness and sales of its “Seasons” fruit cocktail brand and gain a stronger market share hold among competitors. The campaign was simultaneously deployed across social, display, and search, and included in the design was the ability for the Albert AI to manage the campaigns ad spending budget, allowing it to automatically bid, buy, place, and optimize all of the creative materials. “It behooves a company to create lots of raw creative materials for the AI to play with. The more creative options that are thrown at it, the better it is able to operate, because it’s constantly optimizing between different creative choices based on how users are interacting with the material,” offers Ashvin Subramanyam, Dole Packaged Food Asia’s Vice President (VP) of Marketing and Innovation in an interview with PYMNTS.com. The AI could execute both efforts and decision-making autonomously across channels, and monitor competitors’ bidding, results, and competing campaigns. Within the first 8 weeks of the campaign, in-store business grew 87% and reached some 60 million impressions.

Although advertisers again risk poor, redundant, or creepy user experiences, the push to match the right users with the right offers is only growing. Currently, this is a very small proportion of the overall market, but this is likely to become the predominant proposition for advertising, with data from multiple domains collated together to provide microadvertising campaigns. Tractica expects advertisers, brands, and service providers of all sorts to increase adoption of AI to better wield customer data across numerous contexts in order to provide offers that feel less like pushy ads and more like useful, well-tailored, and contextually relevant content.

Tractica forecasts that the annual revenue for targeted advertising using multi-domain customer data in advertising will increase from \$10.12 million worldwide in 2016 to \$626.48 million in 2025.

Table 2.7 Targeted Advertising Using Multi-Domain Customer Data in Advertising, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	10.12	18.28	33.52	61.11	108.40	182.57	284.87	404.88	523.83	626.48	58.2%

(Source: Tractica)

2.2.8 VIDEO CONTENT ANALYSIS

From an analytics perspective, video content has been largely invisible to advertisers and publishers since the earliest days of online streaming. This left a massive void in the ability of companies to understand the composition of videos, their impacts, engagement, resonance, and return on investment (ROI).

Just as AI can now query and analyze images, it can do so for videos as well. AI, particularly ML, DL, and CV, are being used to analyze the influence of video performance, resonance, and monetization across channels, users, and spending.

Valossa offers a video analysis application programming interface (API) that uses AI to understand the contents of video, generating descriptive tags, categories, and overview descriptions automatically. These metadata are extracted to summarize content, serve targeted ads, and search within videos themselves. They also enable clients to train their

own recognition models based on specific market or company needs.

Video content analysis using AI is already in deployment and will continue to grow rapidly as content creators and publishers catch up to the analytics they have enjoyed for text and websites. Tractica expects this capability will emerge as an adjacent feature to many existing digital advertising analytics platforms as these tools embrace AI across their product suites.

Tractica forecasts that the annual revenue for video content analysis in advertising will increase from \$1.22 million worldwide in 2016 to \$194.41 million in 2025.

Table 2.8 Video Content Analysis in Advertising, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.22	3.71	8.42	17.02	31.82	55.08	87.18	124.86	162.20	194.41	75.8%

(Source: Tractica)

2.2.9 VOICE/SPEECH RECOGNITION

Until recently, voice and speech recognition were hardly a viable mode of interaction with computers, or even meaningful dialog to which advertising campaigns would be attached. Advertisements and marketing campaigns were delivered based on general demographic data and text-based inquiry. With the advent of voice and speech recognition, AI offers advertisers new capabilities in new (hands-free) environments, and potentially new market share. Advertisements can now be voice interactive, wherein users could respond to an ad and drive conversion through voice prompts. Advertisements could be delivered based on voice-delivered queries, where a user may search using a voice service, such as Apple's Siri or Amazon's Alexa, while they are cooking, walking, or driving for instance. This capability also opens up searchability and interactions to user segments previously limited in their abilities to use computers, such as disabled, blind, or elderly folks.

IBM Watson recently announced the launch of voice interactive ads, leveraging its acquisition of The Weather Company, in which users can ask questions and receive real-time responses highly tailored to their contexts. A user might ask what to make for dinner, at which point Watson would sort through recipes and deliver recommendations based on time of day, weather, location, and ingredients surfaced via dynamic ads. IBM Watson Ads kicked off on Weather.com and its associated mobile app, but plans to expand to Unilever, GSK, Campbell's Soup, and beyond.

As with other use cases involving voice and speech recognition, opportunities exist in the ease of use and intuitive interface of voice over text input. This allows for audio advertising that may feel more contextually appropriate (i.e., while driving) than purely visual push advertising. Still, advertisers and manufacturers must be mindful of privacy concerns with devices listening by default or advertising in inappropriate contexts.

Tractica forecasts that the annual revenue for voice/speech recognition in advertising will increase from \$0.26 million worldwide in 2017 to \$20.78 million in 2025.

Table 2.9 Voice/Speech Recognition in Advertising, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.26	0.77	1.69	3.28	5.78	9.24	13.29	17.31	20.78	N/A

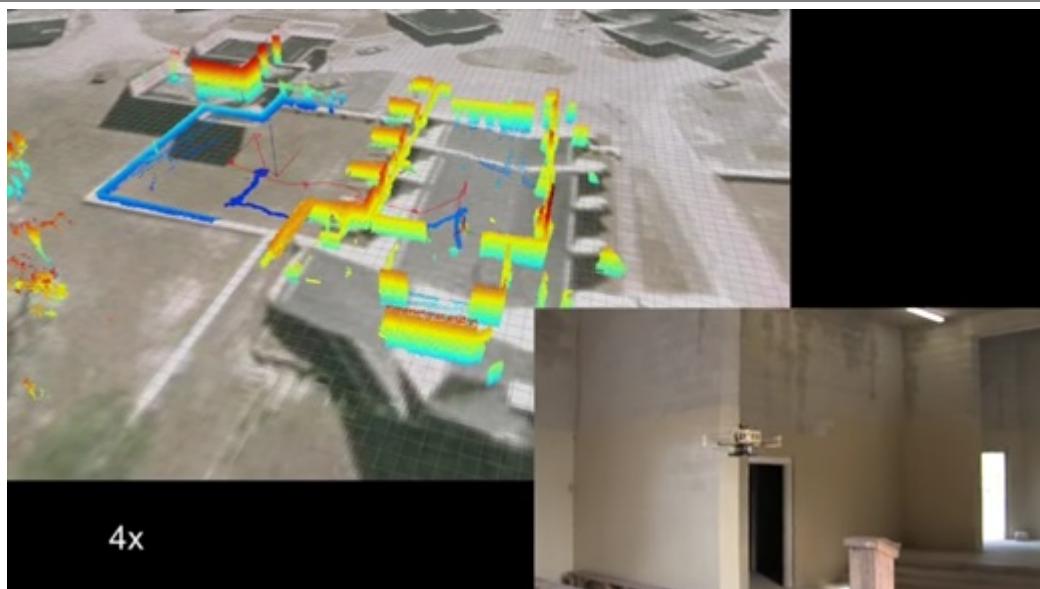
(Source: Tractica)

2.3 AEROSPACE

2.3.1 LOCALIZATION AND MAPPING (AIRCRAFT AND DRONES)

Localization and mapping concerns the need and computational ability to simultaneously construct maps of the immediate environment while updating both the agent's position on that map and movement therein. In the context of aerospace, localization and mapping is a core technique for autonomous movement of airplanes, drones, or any other unmanned aerial vehicle (UAV).

While machine navigation has historically relied on human sight and perception, certain aircraft are now almost entirely operated autonomously using simultaneous localization and mapping (SLAM). In airplanes, SLAM techniques must be fail-proof, accounting for weather, nearby objects, physical changes in environment, and high-precision depth perception. A variety of available algorithms and statistical techniques support SLAM that vary by different types of maps, image data, sensing/sensors, kinematics, 3D modeling, etc.

Figure 2.2 Drones Use Mapping and Localization to Fly Indoors


In the image, a drone is launched outside where it begins rendering the map. It is then flown indoors, where the map and its positioning continue to update in real time. The drone flies through rooms within the indoor environment.

(Source: Intel)

This technology is somewhat less mature in the drone space, as drones are significantly more constrained in size and power supply and often must navigate in very tight or unpredictable spaces. As fly spaces become larger, uncertainties and computational costs increase. The computational intensity of SLAM in 3D environments is due to the use of complex real-time particle filters, sub-mapping strategies, or the hierarchical combination of metric topological representations. The level of computational power and degree of certainty required for reliable SLAM makes this one of the most fundamental challenges in robotics and autonomous navigation.

Tractica forecasts that the annual revenue for localization and mapping in aerospace will increase from \$15.52 worldwide in 2017 to \$663.54 million in 2025.

Table 2.10 Localization and Mapping in Aerospace, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	15.52	38.41	71.81	119.71	186.47	275.75	388.56	521.04	663.54	N/A

(Source: Tractica)

2.3.2 MACHINE/VEHICULAR OBJECT DETECTION/IDENTIFICATION/AVOIDANCE (AIRCRAFT, DRONES, SATELLITES)

In more than 100 years of history, safety and accident avoidance has improved drastically in the aerospace market. Due to the high variability of data, as well as the criticality of aircraft being able to reliably and accurately detect objects, many existing techniques fell short. But as airplanes and other aerial vehicles grow more sophisticated in autonomous operations, object detection, identification, and avoidance is paramount to the success of the technology.

With the introduction of CV and DL, new commercial services are emerging that allow organizations to detect, measure, and monitor objects, improving resolution, thermodynamics, and accurate assessment of contents, as well as modeling and predicting patterns. For aircraft, drones, and satellites alike, DL and CV are becoming critical tools to enable or optimize sight and detection.

Neurala, a Boston-based company specializing in AI, is tackling the problem of drone collisions with the help of DL technology. The company trained its software by feeding it video images of potential collisions from the Microsoft Flight Simulator. Neurala's software notifies drone users and operators whenever it recognizes similar, real-time images from a single camera mounted on the drone.

This technology is being applied to support a variety of use cases, including image detection, segmentation, and classification, as well as to support navigation, search, change monitoring, and research, and to identify, find, and track specific types of objects.

Tractica forecasts that the annual revenue for machine/vehicular object detection/identification/avoidance in aerospace will increase from \$28.55 million worldwide in 2016 to \$677.77 million in 2025.

Table 2.11 Machine/Vehicular Object Detection/Identification/Avoidance in Aerospace, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	28.55	44.90	68.81	103.36	152.29	219.53	308.08	418.05	544.69	677.77	42.2%

(Source: Tractica)

2.3.3 PREDICTIVE MAINTENANCE (AIRCRAFT, DRONES, SATELLITES)

Maintaining fleets of aircraft, drones, satellites, or most any other vehicle is a costly endeavor, and one historically reactive in nature; an issue occurs, service providers respond. Predictive maintenance uses data inputs from disparate streams to predict failures in machinery. Unlike preventive maintenance or condition-based maintenance, which is triggered by the occurrence of one or more indicators, predictive maintenance helps to predict failures beforehand. Both predictive and condition-based maintenance use real-time data as feeds. While condition-based maintenance is much more widely used today, predictive maintenance is gaining popularity, especially for mission-critical assets.

AI is being applied using numerous techniques to support this use case, in aerospace among many other industries. ML algorithms are used to identify failure patterns and detect anomalies, often triggering automated maintenance actions, such as service upgrades, scheduling service engineers, or managing spare parts in inventory chains. DL is particularly useful in its ability to automatically extract features from raw data that are most suitable. This has historically been a manual, non-scalable, bias-prone process (requiring significant physical and mechanical expertise) of constructing the right features from the data set for detection, as well as derivative features for solving tasks.

Airbus is working with EasyJet to provide predictive maintenance capabilities for its fleet of more than 200 aircraft. Airbus is using EasyJet fleet data in conjunction with data from other carriers to improve prognostic tools and predict when parts need to be replaced, ultimately helping carriers like EasyJet improve fleet performance and reduce maintenance costs. Microsoft recently announced a product suite designed to monitor aircraft and predict the remaining useful life of aircraft engine components, based on analyzing large public data sets from past aircraft engine life performances.

The ability to predict failures before they happen and systematically address them helps increase safety and reduce mishaps, delays, and costs associated with broader downtime in the event issues are not preemptively identified and addressed.

Tractica forecasts that the annual revenue for predictive maintenance in aerospace will increase from \$15.86 million worldwide in 2016 to \$513.11 million in 2025.

Table 2.12 Predictive Maintenance in Aerospace, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	15.86	27.00	43.60	68.04	103.33	152.84	219.47	304.21	404.42	513.11	47.2%

(Source: Tractica)

2.3.4

SENSOR DATA FUSION IN MACHINERY (AIRCRAFT, DRONES, SATELLITES)

Sensor data fusion is the process of combining data from multiple sensors in order to improve machine performance, situational, or environmental awareness. Sensor data fusion is fundamentally about creating a “whole” picture of an environment that is greater than the sum of its data streams, making more intelligent decisions about how the machine functions. This is a critical underpinning for helping various types of heavy, mission-critical machinery adjust and adapt to their environments in real time, as well as broader product development and improvement. Sensor data fusion using traditional methods use fixed or hard-wired algorithms to combine data from multiple sensors, and then provide a real-time assessment of the environment, to make adjustments that go beyond object avoidance and navigation.

AI-based sensor fusion exploits statistical interdependencies between disparate data sources, using Bayesian networks and probabilistic graphical models. DL, in particular, is being used to merge samples from diverse sensor types (e.g., accelerometer, gyroscope, magnetometer, barometer, satellite receiver, etc.) and account for high dynamism.

Numerous research efforts characterize development in this space, including notable work from ONERA’s French Aerospace Lab, which is using DL to perform optical and laser sensor data fusion, assess remote sensing images, introduce multi-kernel convolutional neural networks (CNNs) for fast aggregation, and prediction of scene labeling and segmentation for urban areas.

Tractica has identified sensor data fusion in machinery as one of the most significant use cases in terms of revenue potential, in aerospace, automotive, and numerous other industries. Tractica forecasts that the annual revenue for sensor data fusion in machinery in aerospace will increase from \$15.53 million worldwide in 2017 to \$673.78 million in 2025.

Table 2.13 Sensor Data Fusion in Aerospace, Machinery World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	15.53	38.48	72.08	120.38	187.86	278.34	392.97	528.01	673.78	N/A

(Source: Tractica)

2.3.5

SWARMING DRONES

Like ants or bees, the notion of a swarm emerges naturally in biology as a means for a group to coordinate a task that no single member could execute alone. A single ant cannot do much on its own, but an entire colony can render profound impacts and solve complex problems. “Swarm intelligence” is the concept in which decentralized, self-organized systems, either biological or artificial, engage in collective behavior. Broadly speaking, “robotic swarms” is the application of this concept on groups of autonomous devices.

As AI powers onboard processing for UAV and drone technology, it does so in the context of coordinating entire fleets called “swarms” in unison. Swarming drones involve a variety of AI technologies, from CV, SLAM, and object detection, to DL for data analysis and predictive behavior. Swarms of drones could also coordinate collectively to achieve tasks, such as lifting objects, building 3D models, gathering geospatial intelligence, surveying sites, etc. The Texas Advanced Computing Center is developing a drone swarm designed specifically to solve 3D modeling wherein a group of drones coordinates together to create a high-definition (HD) 3D model of structures and geographic features.

While swarming drones are primarily under development in the defense sector, a number other use cases are beginning to emerge. One of the primary areas is emergency response (e.g., floods, fires, earthquakes), wherein a swarm of drones could be deployed and communicate among each other for search and rescue, to monitor or respond to environmental disasters, or to cover a large area of land quickly, thoroughly, and with relatively greater efficiency than helicopters, airplanes, blimps, or other aerial vehicles. Other applications may include aiding lifeguards, animal herding in fields or on farms, police hunts, games, or stage entertainment.

While swarming drones may seem somewhat frightening and uncontrollable to the average consumer, Tractica expects adoption across numerous industries over the longer-term. This is due to the significant potential benefits in rescue, prevention of the loss of life, time and cost efficiencies, and safety (both of professionals that have historically been employed for the same tasks, and individuals whose lives could be endangered if not rescued).

Tractica forecasts that the annual revenue for swarming drones in aerospace will increase from \$0.04 million worldwide in 2016 to \$36.61 million in 2025.

Table 2.14 Swarming Drones in Aerospace, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
0.04	1.02	2.42	4.42	7.23	11.05	16.05	22.22	29.27	36.61	111.2%	(Source: Tractica)

2.3.6

VEHICLE NETWORK AND DATA SECURITY (AIRCRAFT, DRONES, SATELLITES)

The airline industry has been grappling with the nightmarish threat of cyber-hacking or terrorism of its planes since such systems came online. Even today, many systems within planes are separated so as to avoid penetration scenarios, where malicious actors enter through one system and attack another. There are two broad areas of vulnerability: network security, including command and control systems, databases, communications (which all rely on network security); and platform security, including operational systems, combat systems, and engineering plants. Then there remains the constant internal threat, in the event an employee knowingly or unknowingly uploads malware into a critical system. There are also threats along the ecosystem: air traffic control, pilots' mobile devices, in-cabin Wi-Fi, third-party vendors, etc. As manufacturers and operators gain increasing visibility into fleets of machines, sensors, data, and networks simultaneously open up new vulnerabilities and new security methods. For example, cybersecurity experts at Airbus cite the threat of drones sending radio signals to confuse an aircraft's flight or landing.

AI can be applied in an IoT security context, in which various techniques, such as ML, MR, sensor data fusion, DL, and CV, can be used to enhance machine, network, and device security by monitoring sensor and environmental data, analyzing systems and anomalous events, and acting accordingly. AI could pull in data from aircraft in flight, detect a new threat, and automatically issue the appropriate updates to every aircrafts' software for real-time defense intelligence. The AI could also update maps of where threats were and automatically reroute both manned and unmanned aircraft around them.

In a scenario in which multiple vehicles within a network communicate with one another, such techniques are being explored to simulate human intelligence in situation awareness by powering security schemes in which beacons and signatures are validated based on specific contexts. When discrepancies arise, systems alert security analysts or execute tasks

to mitigate threats.

Raytheon is developing a project aimed at helping aviators counter potential cyberattacks that could arise mid-flight. The software is designed to detect anomalies in MIL-STD 1553 networks, which are standard for most military and commercial aircrafts. When the system detects anomalies, it analyzes them for signatures and profiles of cyberattack. From there, the system involves operators to dialog in order to gain deeper understanding for the level of threat and what the system needs to do to assist. The project remains in development as Raytheon works to optimize interface (and trust) between pilot and system.

Tractica forecasts that the annual revenue for vehicle network and data security in aerospace will increase from \$11.49 million worldwide in 2017 to \$495.62 million in 2025.

Table 2.15 Vehicle Network and Data Security in Aerospace, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	11.49	28.45	53.25	88.87	138.59	205.19	289.49	388.68	495.62	N/A

(Source: Tractica)

2.3.7 WEATHER FORECASTING

Weather monitoring and analysis is inherent to successful flights, as well as on-the-ground operations. Local inclement weather contributes in a direct and measurable way to congestion at airports, flight performance, route, time, fuel, and an array of safety considerations.

AI and sensor data from hundreds of thousands of sources collected and monitored in real time (and over many years) is transforming the level of understanding and ability to forecast conditions. With accurate insights into local weather, airlines can better predict congestion, turbulence, wind, etc. to make more precise decisions about exactly how much fuel to put on any given plane. Planes themselves, equipped with sensors and software, are also being used for weather forecasting.

Microsoft powers a wind prediction service called Windflow, which is used by airplane carriers to precisely predict and optimize flight times. The service is, in part, powered by a network of thousands of planes flying every day, providing real-time data and sufficient Big Data for predictive analytics about atmospheric conditions, optimal routes, turbulence, large-scale weather processes, storm tracking, etc. The tool offers wind conditions at altitudes as low as 6,000 feet, and as high as 39,000 feet.

Tractica forecasts that the annual revenue for weather forecasting in aerospace will increase from \$0.57 million worldwide in 2016 to \$34.54 million in 2025.

Table 2.16 Weather Forecasting in Aerospace, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.57	1.11	1.97	3.30	5.35	8.38	12.70	18.55	25.92	34.54	57.7%

(Source: Tractica)

2.4 AGRICULTURE

2.4.1 FOOD SAFETY

Ensuring the safety and sanitation of food is inherent to all agricultural production, but challenging given the fact that threats to food safety can emerge across various points of the supply chain. Moreover, there is an increasing demand for transparency and more fresh food in the developed world.

AI is now being used to analyze both crops and food at the molecular level, which offers benefits to food safety. At the crop level, the ability to identify via CV or image scanning, and protect crops from various diseases can help prevent bad batches from entering the market. Meanwhile, food producers are being held accountable for outbreaks of food-related diseases. Using ML and DL, companies or government agencies like the Food and Drug Administration (FDA) could more accurately and efficiently conduct inspections and on-the-spot testing, as well as monitor data over time.

In 2015, food at some Chipotle Mexican Grill restaurants was the cause for two separate outbreaks of E. coli food poisoning. In the first outbreak, 55 people in 11 states were infected by the foodborne illness, of which 21 were hospitalized. The second, smaller outbreak infected five people from three states, of which one was hospitalized. The company was forced to close restaurants, change safety procedures, and work to try to win back public confidence. As a result, Chipotle implemented high-resolution Deoxyribonucleic acid (DNA)-based testing and bacterial recognition of many ingredients in its food.

Nuritas uses AI and DNA analysis to identify within food peptides with antimicrobial capabilities that can be used as natural food preservatives to enhance food safety and extend shelf life.

Tractica forecasts that the annual revenue for food safety in agriculture will increase from \$5.48 million worldwide in 2016 to \$596.47 million in 2025.

Table 2.17 Food Safety in Agriculture, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	5.48	13.18	25.58	45.27	76.02	123.04	192.84	292.33	426.74	596.47	68.4%

(Source: Tractica)

2.4.2 LIVESTOCK MANAGEMENT

Managing livestock goes back to the dawn of agriculture and animal domestication. It underscores our relationships with land and other animals like dogs and horses, and how we think about animal and meat production at scale. Today, animal-based proteins represent some 20% of the global caloric intake, and has been increasing steadily for years. Animal producers are looking for better tools to measure and manage their livestock in order to yield more with less.

AI is now being deployed to aid livestock producers with more efficient management. AI-enhanced livestock management may also apply CV techniques, monitoring animals with sensors capturing 3D images to recognize indicators of animal conditions like illness, fertility, feeding regime, muscle, fat deposits, etc. Companies like Vance.io are using AI for data mining and analysis to power laborless rotations, rotating livestock via a mobile app and

without the need for all-terrain vehicles (ATVs), horses, dogs, or other additional labor. They are also using AI to monitor, analyze, and make recommendations around the health and fertility of animals. Another benefit is the relatively healthier lives that livestock can lead, with fewer antibiotics and less stress from human interaction.

BovControl, a 5-year-old startup, aims to create the “internet of cows.” Farmers enter cow data (e.g., weight, birthdate, medication, vaccinations) and connect the app to the monitoring device they use to track the animals (e.g., smart collars, ear tags, etc.). Then the app uses AI to analyze and make predictions about each cow, predicting due dates for pregnant cows, milk production or anomalies in production, medication and vaccination needs, etc. The company is also expanding features in meat sourcing and provenance, compliance adherence, export, inventory, and integration with other farm management systems.

Tractica forecasts that the annual revenue for livestock management in agriculture will increase from \$1.90 in 2026 to \$190.84 in 2025.

Table 2.18 Livestock Management in Agriculture, Annual Revenue, 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.90	4.37	8.34	14.63	4.47	39.50	61.82	93.63	136.59	190.84	66.9%

(Source: Tractica)

2.4.3

MACHINE/VEHICULAR OBJECT DETECTION/IDENTIFICATION/AVOIDANCE

The ability to see and detect specific traits and anomalies in crops, animals, and land is essential for agricultural productivity, and a core value cherished by farmers for generations. As many farming tasks have always required human intuition and perception, agriculture has remained a relatively conservative industry in new technology adoption. Meanwhile, population, economic, and competitive shifts are forcing agricultural producers to adopt technologies to keep up.

One of the central use cases for AI in the agricultural space is in object detection, identification, and avoidance. This use case shows up in a range of applications, from self-driving farming equipment to image recognition for identifying and killing weeds, to harvesting tomatoes based on physical attributes, to the detection of defects in poultry eggs. As more equipment, machines, and devices are developed with CV and DL techniques, agricultural producers can leverage their ability to see attributes or things that were previously too big, too small, or too obscure to see.

Startup Prospera uses DL to “see” threats that farmers and drones cannot. It uses a device equipped with CV and proximal red, green, blue (RGB) cameras to assess water and nutrients, detect pestilence and disease, and monitor current yields. It uses DL to process all of this information and predict output, recommend nutrient optimization, conserve resources, and analyze plant development approaches.

The use of AI technology is projected to provide large cost savings, as well as reductions in pesticide and fertilizer use. In certain cases, such as robotics that “analyze” every strawberry before picking, rich data can be collected and used to optimize the entire cultivation process. The challenge for farmers today is the relatively high investments required to leverage these advanced technologies, particularly given the need to maintain them and stay competitive against the larger suppliers.

Tractica forecasts that the annual revenue for machine/vehicular object detection/identification/avoidance in agriculture will increase from \$6.71 million worldwide in 2016 to \$13.57 million in 2025.

Table 2.19 Machine/Vehicle Object Detection/Identification/Avoidance in Agriculture, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	6.71	6.71	6.80	7.01	7.39	8.00	8.90	10.14	11.72	13.57	8.1%

(Source: Tractica)

2.4.4 SATELLITE IMAGERY FOR GEO-ANALYTICS

Satellite imagery has long been a closed domain with high-resolution image databases only available to a select few companies and organizations, such as weather centers, government agencies, the military, and oil & gas companies. Being able to track changes on the ground from space has been vital for these industries, but required human analysis for years. Rapid increases in the availability and improvement in the level of detail of satellite imagery, and advancements in AI, CV, and DL have created new ways of identifying features, tracking changes, and extracting value from satellite imagery.

Apart from providing a way for humans to track the planet on a daily basis, this also means that image processing will have to be automated, in order to take advantage of this quick refresh rate and trove of imagery data. New commercial AI-driven methods offer updates to this information once every day or two with county-level accuracy. Using DL and CV, satellite imagery is captured and analyzed, also in conjunction with other data sets, such as weather or historical data. Still, some basic challenges remain when it comes to weather, viewpoint, lighting, and atmospheric unpredictability.

Farmers and agricultural suppliers have traditionally relied on periodic (monthly, end of season, or less frequent) releases of forecast data. With constant image refreshes, satellite images can be used to assess crop health, to aid/validate in precision agriculture, identify areas of new resources, and estimate deforestation and investment. It is also useful in terms of monitoring land, predicting seasonal performance, and analyzing geographic influences. More generally, satellite imagery can help track a bounded area, providing alerts and updates when something changes in that specific area, or for historical changes over said area. These are not just new applications, but new business models that provide country-wide, or object-specific analysis of satellite imagery to vertical markets.

Startup Descartes Labs is using 4 petabytes of satellite imaging data to assess crop health from space. The company uses spectral information (not visible to the human eye) to measure chlorophyll levels and inform models for crop yield. Spaceknow and Orbital Insight are two other companies that are using satellite imagery data and applying AI techniques to provide analytics around economic or environmental indicators to aid in forecasting.

Tractica forecasts that the annual revenue for satellite imagery for geo-analytics in agriculture will increase from \$0.38 million worldwide in 2016 to \$524.12 million in 2025.

Table 2.20 Satellite Imagery for Geo-Analytics in Agriculture, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.38	7.04	17.88	35.20	62.35	103.96	165.82	254.09	373.40	524.12	123.6%

(Source: Tractica)

2.4.5 SENSOR DATA ANALYTICS

Pervasive sensor application and related networked services, often termed the IoT, has been transforming the agriculture sector for the last decade. As sensors are now being used to monitor everything from soil, rain, air quality, plants, livestock, and fertilizers, to tractors, forklifts, containers, and beyond, agricultural suppliers' visibility into their operations is changing the way goods are produced. As in other industries, the next phase of sensor application is measuring and analyzing it at scale.

Sensor data analytics in AI concerns the analysis of multiple sensor data sources together for "ecosystem-level" intelligence. Whereas sensor data fusion pulls together sensor data for the performance of a machine or fleet of machines, sensor data analytics applies learnings to a broader context, such as a farm or smart city.

AI, ML, and DL are enhancing sensor technology by analyzing and producing insights. AI has been used for irrigation scheduling by using rainfall and drip irrigation sensors, a phenotype measuring system for greenhouse climate control, or for predicting the fermentation process of cattle, which can be used to determine their nutritional feed. Crop health monitoring, both in outdoor fields and within greenhouses, is another area where AI tools gather data from multiple sensors like temperature sensors, soil sensors, pressure sensors, light sensors, water sensors, and wind sensors. This data is then combined and analyzed to predict the health of crops, identify pests that could damage yields, and provide suggestions for improving crop yield. These analyses, combined with aerial images from satellites and drones, build detailed models regarding a wide range of environments: machinery performance, soil viability, weather models, etc.

IBM's The Weather Company recently announced Deep Thunder, a hyper-local weather forecaster that harnesses diverse data inputs, including historical weather reports, to predict and model future conditions. Deep Thunder is tuned for forecasts at a 0.2 to 1.2 mile resolution. Not only does this enhance the depth and dimension of weather data, it allows The Weather Company (and IBM) to offer highly personalized and hyper-local farm-management-as-a-service, as well as other business applications.

Sensors play a fundamental role in agricultural monitoring, workflows, and production. They will continue to be integrated as critical inputs across broader supply chain automation and transparency efforts. Thus, using AI to enhance the analysis of this data, and in conjunction with other data, helps justify ROI and gives users a competitive edge as data is used for ongoing optimization.

Tractica forecasts that the annual revenue for sensor data analytics in agriculture will increase from \$29.61 million worldwide in 2016 to \$1.313 billion in 2025.

Table 2.21 Sensor Data Analytics in Agriculture, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	29.61	46.95	74.46	117.74	184.94	287.31	438.89	654.64	945.80	1,313.29	52.4%

(Source: Tractica)

2.4.6 SENSOR DATA FUSION IN MACHINERY

Sensor data fusion is the technique used to aggregate or “fuse together” multiple sensor data feeds and other data feeds in order to ascertain a more complete or multi-dimensioned picture of operations. The resulting multi-dimensional data offers less uncertainty than if the data feeds were viewed individually. Sensor data fusion in agriculture is about extracting data from multiple data sources to facilitate optimal positioning or function of autonomous vehicles or devices. Unlike sensor data analysis, which assesses a broader context (e.g., farm, smart city), sensor data fusion is geared toward the understanding and performance optimization of the machine itself.

Sensor data fusion in agriculture might monitor temperature, vibrations, speed, wear and tear, weather, crop interactions, or fuel efficiency in order to optimize the machine itself. As onboard processing increases in throughput, DL will be used to more accurately detect, classify, model, and “learn” from environmental context and impacts.

Particularly as agricultural machine manufacturers consider new business models involving leasing or time-based access, sensor data fusion to support machine uptime and predictive maintenance will be key. Like in aerospace or energy, these applications are often mission-critical and precision, speed, and reliability are paramount to adoption.

Tractica forecasts that the annual revenue for sensor data fusion in machinery in agriculture will increase from \$0.24 million worldwide in 2017 to \$13.34 million in 2025.

Table 2.22 Sensor Data Fusion in Machinery in Agriculture, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.24	0.60	1.15	1.95	3.13	4.78	7.03	9.91	13.34	N/A

(Source: Tractica)

2.4.7 LOCALIZATION AND MAPPING

Localization and mapping concerns the need and computational ability to simultaneously construct maps of the immediate environment while updating both the agent’s position on that map and movement therein. In the context of agriculture, localization and mapping is a core technique for autonomous movement of any UAV.

While machine navigation has historically relied on human sight and perception, agricultural machinery are increasingly growing more autonomous using SLAM. Using drones and/or as applied to tractors, forklifts, or any other farming vehicle, the technology could aid agricultural producers in planning, planting, cultivation, pesticide application, harvesting, transshipping, and beyond.

ASI, in partnership with New Holland, CNH Industrial, and Case IH, is developing self-driving farm vehicles that do not just operate autonomously, but that can deploy tandem field

coverage. This means multiple machines can operate in a field to accomplish tasks and coordinate between each other for navigation, tilling, and planting. Advanced control systems allow farmers to see, manage, and synchronize in-field working between different machines.

Approximately 52% of farmers use some form of auto steer with projections of 64% by 2018. Auto-steering aside, like consumers, many farm operators remain skeptical about altogether abandoning the cab and leaving precious crops to the precision of autonomous machinery. However, such vehicles offer advantages of addressing labor shortages, 24/7 overnight performance, path optimization, and safety.

Tractica forecasts that the annual revenue for localization and mapping in agriculture will increase from \$0.55 million worldwide in 2017 to \$26.12 million in 2025.

Table 2.23 Localization and Mapping in Agriculture, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.55	1.35	2.52	4.21	6.61	9.92	14.30	19.77	26.12	N/A

(Source: Tractica)

2.4.8 WEATHER FORECASTING

By 2050, the world will need to feed some 10 billion people—a 70% increase in food production compared to today. While there are numerous seen and unseen forces affecting agricultural food production and demand, weather has, is, and will always be a critical influence. From almanacs and meteorology to sensors, methods for tracking and predicting weather have evolved alongside technology since the beginning of agriculture.

AI is being used for weather forecasting in agriculture to aid farms and organizations with more accurate forecasting, and apply reinforcement learning on past predictions and actual outcomes. As farmers aim for greater precision in all phases of cultivation, applying specific pesticides and fertilizers and very specific points in crop lifecycles to maximize yield, for instance, the ability to accurately forecast and pinpoint environmental conditions is key. By comparing predictions with accuracies, the model is able to learn and improve simulation capabilities, as well as forecast much further into the future.

AI can be used to perform weather pattern detection, such as cyclonic activity or other extreme weather events. The U.S. National Energy Research Computing Center (NERSC) has used CNNs to classify threatening climate events like cyclones. This work was performed on a central processing unit (CPU)-only Cray XC30 supercomputer, where both the training and inference was run on the same platform, although there was some effort involved in adapting the CNN algorithm to the climate data. The main goal for NERSC was to have a model learn the characteristics of a cyclone and classify it, an area where human decision-making variance is an issue. With the algorithm having between 80% and 90% accuracy in identifying extreme weather events, this is only the start and shows that AI techniques can be used for classification and identification of more complex weather systems and events.

There is a profound implication in our ability to better forecast weather events, not just for agricultural producers' ability to foresee and plan, but to benefit from ongoing weather data and what it can tell us about impacts on crops, soil, livestock, water sources, air quality, and many other variables, even commodities and market forces.

Tractica forecasts that the annual revenue for weather forecasting in agriculture will increase from \$0.01 million worldwide in 2016 to \$6.3 million in 2025.

Table 2.24 Weather Forecasting in Agriculture, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.01	0.09	0.22	0.43	0.76	1.26	2.00	3.06	4.49	6.30	100.9%

(Source: Tractica)

2.4.9 WEED IDENTIFICATION

Weeds compete with productive crops, pastures, nutrients, water, and light, and ultimately can impact yield. They can be poisonous, distasteful, create shelter for pests, cause damage, or interfere with the use and management of desirable plants by contaminating harvests or interfering with livestock. Worldwide, about 3% of the entire plant species population is weeds. Because they are so common, it is difficult to entirely quantify the extent of their impact, particularly considering the billions of dollars spent every year on herbicides and other methods of weed control.

AI, particularly CV, can be applied to help mitigate risks and losses incurred by weeds. Other techniques are collecting huge volumes of images and training neural networks to recognize images from other plant species and crop seedlings. By identifying weeds early on, and gathering data on weed-related patterns over time, agricultural producers can better forecast and fight against their emergence.

Hummingbird Technologies specializes in imagery analytics, captured through drones, for precision agriculture. In addition to general field and crop analytics, including early detection of crop diseases, it is also using CV to map weed patterns within fields. This helps farmers target weeds with leaf-level precision, and optimize nutrients and planting as a result. Success is a function of increased productivity and profits, as well as decreased costs of operations.

Tractica forecasts that the annual revenue for weed identification in agriculture will increase from \$4.11 million worldwide in 2016 to \$286.66 million in 2025.

Table 2.25 Weed Identification in Agriculture, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	4.11	7.85	13.82	23.28	38.01	60.52	93.88	141.42	205.61	286.66	60.3%

(Source: Tractica)

2.5 AUTOMOTIVE

2.5.1 AUTOMATED ON-ROAD CUSTOMER SERVICE

One of the key benefits of the great expansion in sensor technology is the ability to automatically monitor and alert the driver when a problem occurs with the vehicle, as well as any other party that is connected to the vehicle via wireless communication.

This can include providing automated alerts when the vehicle's operating status falls out of the norm, automatically contacting a towing service to retrieve the vehicle, and automatically

alerting and scheduling an appointment with a dealer or repair shop, sending along information about the parts that need to be repaired or replaced. Other data, such as the vehicle's position could be fed into traffic systems, sending out an alert that a lane is blocked, or that an emergency medical response is required. Ultimately, the ability to automatically analyze, classify, and send out an appropriate response or action is driven by AI systems.

From Tesla to Ford to Toyota, just about every auto manufacturer is working on developing channel strategies for delivering customer service to drivers and passengers. While driving or riding in connected cars, will drivers welcome customer support from manufacturers, hyper-local marketing from brands, personalized alerts via context-aware virtual assistants? Automated on-road customer service will also likely tie in with personalized services available in cars, another use case outlined in Section 2.5.6. The question is who will support or co-develop which business models. Will the answer be the manufacturer, dealership, network service provider (NSPs), insurance provider, advertiser, technology giant, or city? The market is too nascent to determine a clear winner, although manufacturers and NSPs are collaborating closely. The market is still too early to determine which model will be most successful.

Tractica forecasts that the annual revenue for automated on-road customer service in automotive will increase from \$1.88 million worldwide in 2017 to \$113.72 million in 2025.

Table 2.26 On-Road Customer Service in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	1.88	4.96	9.93	17.77	29.58	46.21	67.38	91.04	113.72	N/A

(Source: Tractica)

2.5.2

BUILDING GENERATIVE MODELS OF THE REAL WORLD

The concept of strong AI is the idea that AI is able to exhibit behavior and act as skillfully and flexibly as humans can. Today, this concept remains largely fiction, as it entails a vast interconnected understanding of the physical laws, taxonomies, consequences, and even social constructs that govern our world, which is a far cry from any AI application to date. Building generative models of the real world is a small but important step in this direction.

At a high level, AI is being used to help generate models and maps of the real world. By using a combination of sensing technology, including HD cameras, ultrasonic sensors, radar, light detection and ranging (LIDAR), and global positioning system (GPS) mapping technology, highly accurate maps can be generated, with accuracy within a few centimeters. This high degree of accuracy is especially important in enabling autonomous vehicles, which may use this data to establish position while on the road, in conjunction with onboard sensors. This is an essential step toward enabling vision-based systems in things like cars and robots, so they can start to understand the physics of the world.

In addition to two-dimensional (2D) data, 3D modeling allows features like curbs, bumps, and other features to be accurately captured, which can be integral to the safe operation of an autonomous vehicle. Indeed, without this information, a vehicle may not realize that there are speed bumps ahead, or an automated plow may not know where road plates or other anomalies in the road may impede its progress.

Today, the mode of operation for autonomous vehicles' "vision" is a function of object detection and generally lacks information beyond category. For example, a ball rolls into the

street. If the car has not seen the object before, it makes a guess and likely stops as a safety precaution. By contrast, with a generative model, the AI powering vision-based systems would understand the movement, what it is doing, in what direction it is moving, how fast, etc.). Generative models of the real world help develop the context with which to make a decision about how to maneuver the vehicle.

Tractica forecasts that the annual revenue for building generative models of the real-world in automotive will increase from \$11.16 million worldwide in 2017 to \$622.09 million in 2025.

Table 2.27 Building Generative Models of the Real-World in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	11.16	29.12	57.76	102.30	168.61	260.69	376.24	503.14	622.09	N/A

(Source: Tractica)

2.5.3

DRIVER FACE ANALYTICS AND EMOTION RECOGNITION

Human faces are able to convey various states, such as happiness, sadness, stress, and other emotions, but designing technology to accurately read and interpret these emotions and states is exceptionally challenging. Sophisticated algorithms that measure changes in a face's features and position, the presence of other cues (such as laughing, crying, or shouting), can be used to help determine one's state, which has been shown to have an impact on decision making, particularly when conducting a complex task like driving.

AI systems that can accurately and reliably ascertain a driver's emotional or physical state can be extremely valuable, in terms of both safety and convenience. They can be used to monitor the driver's condition to make sure they are alert and focusing on the task of driving, and to trigger other types of actions, such as suggesting a rest stop or turning on the entertainment sound system to a particular artist or genre to match the person's mood. By incorporating AI algorithms with driver-facing cameras and sensors to measure driver inputs (such as acceleration or steering inputs), it will be easier to ascertain if and when a driver begins to be fatigued, issuing an alert to the driver to snap to attention. Furthermore, if the driver continues to exhibit these signs, the vehicle could begin to take some control, such as limiting the ability of the truck to accelerate, forcing the driver to pull over to the side of the road. This alone could help prevent accidents due to human miscalculation of how tired or fatigued they may be.

Eyeris is focused on this area.

Tractica forecasts that the annual revenue for driver face analytics and emotion recognition in automotive will increase from \$0.01 million worldwide in 2016 to \$3.93 million in 2025.

Table 2.28 Driver Face Analytics and Emotion Recognition in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.01	0.08	0.19	0.36	0.64	1.04	1.62	2.34	3.15	3.93	88.2%

(Source: Tractica)

2.5.4

GESTURE RECOGNITION

Another key area in which AI is required to bring full functionality to a technology is with gesture recognition. The ability to accurately track and recognize gestures made by humans, which, by their nature, are not capable of repeating a gesture repeatedly using the exact same speed, position, or trajectory, requires algorithms that can account for these variances, as well as understand context. This is critical if gesture recognition is deployed as a tool for drivers to use, when their primary focus should be on driving. It is more likely that gesture recognition tools will be deployed primarily on passengers or in driverless vehicles.

Figure 2.3 BMW's 2016 7-Series Incorporates Gesture Recognition for Six Commands



(Source: BMW)

BMW has introduced this feature into its 2016 7-Series, which uses 3D sensors and gesture recognition to respond to cues from drivers, including:

- Increasing or decreasing volume by circling finger (clockwise for up; counter-clockwise for down)
- Accepting a phone call by pointing toward the dashboard touch-screen
- Rejecting a call by swiping hand to the right
- Changing the camera angle of the multi-camera view by making a circle with thumb and finger
- Custom command (e.g., “navigate home”) using two-finger point toward touchscreen

Tractica forecasts that the annual revenue for gesture recognition in automotive will increase from \$0.02 million worldwide in 2017 to \$1.17 million in 2025.

Table 2.29 Gesture Recognition in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.02	0.05	0.10	0.18	0.30	0.48	0.69	0.94	1.17	N/A

(Source: Tractica)

2.5.5 MACHINE/VEHICULAR OBJECT DETECTION/IDENTIFICATION/AVOIDANCE

Perhaps the most valuable use of AI in vehicles is the use of object detection and classification, which takes sensor data, often from cameras, and then uses complex algorithms to classify these objects so that the AI system can then “learn” their characteristics, and recognize them in real time.

The challenge is not in capturing images, as today’s HD cameras can present images in stunningly clear detail. However, in a moving environment, objects can appear to change size as a vehicle or camera approaches. The angle at which an object is viewed can also skew its appearance, and the presence of other factors (rain, bright sunlight, low lighting, glare, dirt, snow, or any other number of obstructions) can alter the appearance of an object, making it hard to accurately and consistently identify the object.

This is an area where machine vision and ML can provide invaluable support. By capturing a wide range of images of objects from a variety of vantage points, angles, and in different conditions, a repository of images that can be definitively classified as that object can be created, and used to “train” a ML system to identify and classify objects that resemble objects in the repository. By then assigning various other attributes to each object, such as whether the object is informational like a traffic sign, whether or not it is permanent or temporary like a road barrier, or whether or not it has the capability of motion and how it typically moves, the system can begin to develop logical rules on handling each object and the rules for dealing with them. Of course, all of this takes massive amounts of processing power. This is why most of this initial training is done at processing centers, rather than onboard the vehicle in real time. Every auto manufacturer developing autonomous or semi-autonomous vehicles is working on this as it is a vital building block for successful deployment.

Tractica forecasts that the annual revenue for machine/vehicle object detection/identification/avoidance in automotive will increase from \$75.96 million worldwide in 2016 to \$561.52 million in 2025.

Table 2.30 Machine/Vehicle Object Detection/Identification/Avoidance in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	75.96	82.59	94.62	115.12	148.33	199.07	270.87	362.37	464.33	561.52	24.9%

(Source: Tractica)

2.5.6 PERSONALIZED SERVICES IN CARS

Driving has always been a somewhat personal experience, but as data collection and generation begin to infuse the driver’s experience, the ability to personalize driving will become more automated.

AI is playing a major role in the development of vehicle personalization, including learning the preferences of drivers. This could encompass a variety of types of personalization, from basic infotainment preferences, such as favorite radio stations, automatically selecting preferred routes to destinations, or suggesting favorite services along a route.

Self-learning capabilities may also extend to driving tasks like automatically adjusting how the engine responds to an individual's driving style, such as automatically engaging a "sport" mode, based on how individuals accelerate, or more fully engaging an advanced driver assistance system (ADAS), such as lane-keeping assist, for drivers who tend to have difficulty keeping their vehicles in the center of the lane while driving.

Personalized services available in cars will also run in conjunction with automated on-road customer service, another use case outlined in Section 2.5.1.

Tractica forecasts that the annual revenue for personalized services in cars in automotive will increase from \$11.21 worldwide in 2017 to \$624.85 million in 2025.

Table 2.31 Personalized Services in Cars in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	11.21	29.25	58.02	102.75	169.36	261.85	377.91	505.38	624.85	N/A

(Source: Tractica)

2.5.7 TRUCK PLATOONING

Although fully autonomous trucks are years away, highway-based autonomous driving can provide significant benefits to the trucking industry, taking control of the vehicle during long stretches of highway driving, when it is sometimes monotonous and difficult for a human driver to maintain absolute attention. While a driver would still be able to respond to an emergency situation, having an autopilot system could help avoid accidents that may occur due to a driver drifting off to sleep, or being hypnotized by the relatively unchanging scenery or conditions.

In particular, the concept of truck platooning, powered by autonomous driving technologies like object detection and localization and mapping, is when multiple trucks are deployed to drive in formation in a convoy of five to six, or "in one platoon." Altogether, the use of multiple trucks in a platoon can reduce wind resistance, time-to-break, and introduce significant fuel savings and CO₂ emissions, as aerodynamics encountered by the first truck decrease friction in the subsequent second, third, and so on down the line. Platooning trucks are still mostly under experimentation, but growing quickly.

In California, Volvo, Caltrans, and Partners for Advanced Transportation Technology (PATH) at U.C. Berkeley recently conducted an experiment involving a three-truck semi-autonomous platoon where trucks were 50 feet apart from one another and robots were controlling the pedals in two of the three vehicles. In traffic and planning simulations supporting the experiment, the companies found that platoons could help facilitate up to 50% more trucks using the same lane.

Tractica forecasts that the annual revenue for truck platooning in automotive will increase from \$0.62 million worldwide in 2016 to \$77.85 million in 2025.

Table 2.32 *Truck Platooning in Automotive, World Markets: 2016-2025*

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.62	1.86	3.89	7.19	12.41	20.33	31.53	45.90	62.10	77.85	71.0%

(Source: Tractica)

2.5.8 PREDICTING DEMAND FOR ON-DEMAND TAXIS

Taxi demand, particularly in congested city areas, is largely driven by external factors, such as major events, overall traffic patterns, the availability of other transportation options like public transit systems, weather (extreme heat or cold, or precipitation), etc. Managing such variables can be complex for humans to handle, given that each of these inputs are changing constantly.

On-demand taxi services, whether privately operated like Uber or Lyft or traditional taxi services, are increasingly embracing technology designed to make hailing a car simpler and easier. By bringing AI into the system, it will be easier to forecast and manage demand for taxis by matching real-time conditions with historical data patterns. AI systems can capture and crunch data more quickly than humans, and can often recognize data patterns earlier than humans, thereby helping to realign or redeploy taxis to meet demand. As the algorithms are used over time, they can also learn from their past experiences, and refine themselves to become more accurate.

Companies like Uber are using ML and DL, taking driver and demand monitoring to new heights, tracking drivers (and passengers) wherever they go while using the service, and also using DL to detect patterns of potential driver misbehavior, bad driving, or fraud. Uber and Lyft are also increasingly gamifying drivers' experience to incentivize them to supply probable demand, by rewarding drivers with in-app badges, setting earnings goals, and alerting drivers of the next offer before the current ride has ended.

Similarly, when demand is able to be accurately quantified, surge pricing can be deployed with the confidence that it actually matches demand, as outlined in Section 2.5.12.

Tractica forecasts that the annual revenue for predicting demand for on-demand taxis in automotive will increase from \$3.89 million worldwide in 2016 to \$215.07 million in 2025.

Table 2.33 *Predicting Demand in On-Demand Taxis in Automotive, World Markets: 2016-2025*

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	3.89	6.07	9.91	16.57	27.86	46.23	74.27	113.37	161.94	215.07	56.2%

(Source: Tractica)

2.5.9 PREDICTIVE MAINTENANCE

As vehicles become more and more digitized in operations, composition, and supply chain interactions, the need to monitor and preemptively address potential failures or downtime becomes critical. Common and costly failures occur across the supply chain, from manufacturing critical parts for cars to car operations to applications within the car.

Just about every auto manufacturer has been using predictive maintenance, but increasingly, these capabilities are leveraging more sophisticated techniques, such as DL

and CV. Techniques like sequence analysis can be used to understand failure patterns and follow-on failures, while ML and DL can be used to perform predictive models or recurrent event models. Such models tend to leverage (in training and inference) the following basic parameters: failure history, maintenance and repair history, machine performance, conditions, telemetry data, and operating environment, among a range of other external inputs. AI systems that can take into account additional data, such as the style of driving, as well as more detailed assessments of wear, can create more accurate models for scheduling maintenance on a particular vehicle's systems. They can also identify patterns of use that may be impacting a specific component or system, and alert the owner (such as illustrating how excessive braking may be accelerating wear on the brake pads).

Figure 2.4 Predictive Maintenance Dashboard for Connected Cars



(Source: Data RPM)

Hewlett Packard Enterprise (HPE) works with auto manufacturers on predictive maintenance programs using ML to source and mine existing and new data sources that provide relevant information—not just related to the car's componentry and systems, but across dealership data, manufacturing line defects, warranty data, and even search and social data. For example, an uptick in Google searches for “fan belt” associated with a specific car model may be an early indicator of a bigger issue. All data is integrated into models that are then integrated into production processes for unique to specific models and fleets. One customer forecasted a 4% decrease in warranty costs based on the early detection of defects alone.

Tractica forecasts that the annual revenue for predictive maintenance in automotive will increase from \$6.27 worldwide in 2017 to \$398.23 million in 2025.

Table 2.34 Predictive Maintenance Taxis in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	6.27	16.61	33.46	60.20	100.84	158.51	232.65	316.50	398.23	N/A

(Source: Tractica)

2.5.10

SENSOR DATA FUSION IN MACHINERY

Sensor data fusion is the technique used to aggregate, or “fuse together” multiple sensor data feeds and other data feeds in order to ascertain a more complete or multi-dimensioned picture of operations. The resulting multi-dimensional data offers less uncertainty than if the data feeds were viewed individually. Sensor data fusion in automotive contexts is about extracting data from multiple data sources to facilitate optimal positioning or function of autonomous vehicles or devices. Unlike sensor data analysis, which assesses a broader context (e.g., farm, smart city), sensor data fusion is geared toward the understanding and performance optimization of the machine itself.

One of the key challenges with today’s automobiles and the cars of the future is the vast number of disparate sensors used to control or augment the vehicle. AI systems can alleviate many of the challenges, allowing the systems to process information based on specific algorithms that may prioritize or de-emphasize information, depending on the system using a specific sensor or group of sensors. This will result in greater efficiency and ensure that the systems remain responsive, with little or no latency. As onboard processing increases in throughput, DL will be used to more accurately detect, classify, model, and “learn” from environmental context and impacts.

Particularly as automotive manufacturers consider new business models involving leasing or time-based access, sensor data fusion to support machine uptime and predictive maintenance will be key. Like in aerospace or energy, these applications are often mission-critical and precision, speed, and reliability are paramount to adoption.

Tractica forecasts that the annual revenue for sensor data fusion in machinery in automotive will increase from \$1.07 million worldwide in 2016 to \$210.66 million in 2025.

Table 2.35 Sensor Data Fusion in Machinery in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.07	4.58	10.31	19.56	34.11	56.01	86.74	125.75	169.20	210.66	79.8%

(Source: Tractica)

2.5.11

SIMULATING WORLDS FOR ARTIFICIAL INTELLIGENCE TRAINING

In order to simulate the functionality and form of an object, manufacturers and designers used to have to undergo significant, costly, and sometimes dangerous product prototyping periods. Extensive testing, evaluation, re-configuring, re-testing, and repeat, often in conjunction with manually processed data sources (e.g., road data, safety compliance, etc.) was the status quo in order to advance features, functions, and designs to a point of reliability, security, and safety. Even in web-based environments, annotating real-world data for training is difficult to scale.

AI is influential in this area, particularly as it can power very precise and highly programmable environments that can be used to simulate worlds for AI training. The benefits to testing in simulated worlds are manifold. One, costs are often lower, as no hardware prototyping is required, and where safety risks are involved, there are none in simulations. Second, such environments can be programmed with many (one day, infinite) variables and parameters, so that a wide range of scenarios can be incorporated, learned, and tested over and over. Examples include: lighting and climate; the physical dynamics of certain surfaces like brick or forces like wind; differences in acceleration for different kinds of roads or turns; etc. Using games like Pac-Man, chess, or other board games as a way to test and train AI systems

(through reinforcement learning) has helped accelerate algorithms and model development for years, but only recently has the focus turned to training AI for real-world applications. The technique is driven by a reward function, only instead of points as rewards in a game, reward functions in the physical world might be a vehicle stopping for a dog or a robot successfully picking up a cup.

Researchers can develop (or even leverage, as in the case of Grand Theft Auto (GTA)) thousands of dynamics to help train AI. GTA has become a favored choice among some autonomous car manufacturers as its settings are rich virtual environments containing all manner of municipal contexts. The settings of Los Santos and San Andreas, for instance, feature hundreds of different vehicles, various traffic signs, multiple road types, bridges, tunnels, and thousands of characters roaming around.

OpenAI, an AI research foundation, recently unveiled Universe, an open-source “digital playground” where developers can virtually test and train AI using games, apps, and websites. Universe contains thousands of environments with an expanding catalog of everything from space to biological science apps. The software also enables “transfer learning,” in which an agent takes what it has learned in one application and applies it to another, enabling what OpenAI calls “general-purpose” knowledge about the world. This is a small but significant step toward more generalized AI.

Improbable, micropsi, and Prowler.io support game-like simulated environments for AI training specific to autonomous vehicle development. While game or software-based training is unlikely to ever fully replace the unpredictable chaos of the physical world, it poses an interesting supplement to the learning process.

Tractica forecasts that the annual revenue for simulating worlds for AI training in automotive will increase from \$5.99 million worldwide in 2016 to \$127.82 million in 2025.

Table 2.36 Simulating Worlds for AI Training in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	5.99	8.10	10.98	15.78	23.54	35.53	52.83	75.49	101.65	127.82	40.5%

(Source: Tractica)

2.5.12

SURGE PRICING FOR ON-DEMAND TAXIS

Taxi demand, particularly in congested city areas, is largely driven by external factors, such as major events, overall traffic patterns, the availability of other transportation options like public transit systems, weather (extreme heat or cold, or precipitation), etc. Managing such variables can be complex for humans to handle, given that each of these inputs changes constantly.

As on-demand taxi services are increasingly embracing technology designed to make hailing a car simpler and easier, they are working to unite supply and demand to reduce wait times, while increasing market share. By bringing AI into the system, it will be easier to forecast and manage demand for taxis by matching real-time conditions with historical data patterns. Surge pricing is a concept in which pricing temporarily increases given high demand; AI now determines when and the extent to which prices surge in a given area or time period.

Tractica separates surge pricing for on-demand taxis as this is typically a consumer-facing interface for delivering on demand. Whereas the use case of “predicting demand” outlined

in Section 2.5.8 typically involves back-end systems and more objectives than meeting passenger demand, surge pricing is one way consumers are interacting in real time with AI as it determines whether supply is short enough to increase prices temporarily.

Tractica forecasts that the annual revenue for surge pricing for on-demand taxis in automotive will increase from \$6.15 million worldwide in 2016 to \$41.39 million in 2025.

Table 2.37 Surge Pricing for On-Demand Taxis in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	6.15	6.76	7.64	8.96	11.00	14.14	18.76	25.09	32.89	41.39	23.6%

(Source: Tractica)

2.5.13 LOCALIZATION AND MAPPING

Despite the popularity of GPS navigation, consumer-grade GPS accuracy is limited to about 2 to 3 meters, and GPS is rendered completely useless if a car enters a tunnel or the line of sight to satellites is otherwise compromised. This is unsuitable for autonomous vehicles that must know their location at all times, and within a much more accurate level. AI is also being used to further the development of vehicle localization and mapping elements.

Localization and mapping concerns the need and computational ability to simultaneously construct maps of the immediate environment while updating both the agent's position on that map and movement therein. In the context of automotive, localization and mapping is a core function for the autonomous movement of cars, trucks, or any other autonomous machine that moves.

AI systems can provide that visibility via a model through two variables: an unknown variable, which is the location of the car, and observations about the car's location based on the sensor inputs at that given time. The AI component takes these two variables and, based on a randomized algorithm that repeatedly samples possible scenarios, returns a best estimate for where the vehicle currently is situated. These models can be refined over time by also incorporating HD, 3D maps, which provide more accuracy than typical 2D maps provided by Google and others.

Tractica forecasts that the annual revenue for localization and mapping in automotive will increase from \$3.36 million worldwide in 2016 to \$378.51 million in 2025.

Table 2.38 Localization and Mapping in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	3.36	9.64	19.91	36.46	62.48	101.61	156.54	226.31	304.09	378.51	69.0%

(Source: Tractica)

2.5.14 VEHICLE NETWORK AND DATA SECURITY

As the automotive and transportation industries develop more connected and autonomous vehicles, they grapple with the nightmarish threat of cyber-hacking or terrorism of its fleets. Today's vehicles have more control units, computing power, lines of code, and wireless connections with the outside world than ever before, which is why vehicles of the future are

cause for great security concerns. A recent study by Munich Re, the world's second-largest reinsurer, found that 55% of corporate risk managers surveyed named cybersecurity as their top concern for autonomous vehicles.

Even today, many systems within cars are separated so as to avoid penetration scenarios, where malicious actors enter through one system and attack another. There are two broad areas of vulnerability: network security, including command and control systems, databases, and communications (which all rely on network security); and platform security, including operational systems, engineering plants, and applications. Then there remains the constant internal threat, in the event an employee knowingly or unknowingly uploads malware into a critical system. Data security of drivers and their devices also cannot be ignored. There are also threats along the ecosystem: traffic controls, mobile devices, in-vehicle Wi-Fi, third-party vendors, etc. As manufacturers and operators gain increasing visibility into fleets of machines, sensors, data, and networks simultaneously open up new vulnerabilities and new security methods.

AI can be applied in an IoT security context, in which various techniques, such as ML, sensor data fusion, DL, CV, and MR, can be used to enhance machine and device security by monitoring sensor and environmental data, analyzing systems and anomalous events, and acting accordingly. AI could pull in data from vehicles in transport, detect a new threat, and automatically issue the appropriate updates to every other vehicle's software for real-time defense intelligence. The AI could also update maps of where threats were and automatically reroute both manned and unmanned vehicles around them.

Karamba Security is another cybersecurity company focused on connected vehicles. Its Autonomous Security technology works to secure electronic control units (ECUs) by allowing any car's ECU to protect itself from any potential threat by automatically locking it to the ECU's factory settings. This blocks any operations that are not part of basic performance and safety, preventing hackers from accessing critical systems through adjacent systems like infotainment or dongles. Its in-memory protection blocks memory-based attacks like buffer overrun or return-oriented programming (ROP). Default factory instructions are "good" by design, so the system does not have to guess about a command it has not encountered, avoiding the risk of false alarms or false positives.

Volvo recently announced it would be acquiring a 40% stake in CYMOTIVE Technologies, an Israeli cybersecurity platform that specializes in automobiles. CYMOTIVE's approach involves a multi-layered car security architecture, incorporating security solutions for in-vehicle, backend, mobile services and other connected functions, and uses simulation of attack vectors to "reverse engineer" potential hacks and penetrations.

Tractica forecasts that the annual revenue for vehicle network and data security in automotive will increase from \$0.99 million worldwide in 2017 to \$61.45 million in 2025.

Table 2.39 Vehicle Network and Data Security in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.99	2.62	5.26	9.43	15.76	24.70	36.14	49.01	61.45	N/A

(Source: Tractica)

2.5.15

VIRTUAL TESTING AND SIMULATION FOR RACING CARS

Simulating race car driving scenarios was largely an in-car, on-road, hands-on experience until recently. AI is influential in the area of testing and simulation, particularly as it can power very precise and highly programmable environments that can be used to simulate worlds for AI training.

AI can also be used to help test and simulate the performance of racing cars, without requiring human drivers to risk injury or death while testing new technology. AI can take previous performance data and combine it with modeled data, using known information, such as a racecar's speed, traction, braking performance, and other attributes. Then, the intelligence engine can run hundreds or thousands of simulations, while "learning" what happens when a variable is changed. The result is a scenario that allows a virtual car to test the limits of certain systems without placing a human driver at risk.

Groups like Roborace are testing driverless racecars that use AI to pilot a vehicle around a track. While the cars are equipped with streamlined versions of LIDAR sensors, HD cameras, optical speed sensors, and ultrasonic sensors, they still need refinement with respect to interacting with other vehicles on the track.

Tractica forecasts that the annual revenue for virtual testing and simulation for racing cars in automotive will increase from \$.04 million worldwide in 2016 to \$6.4 million in 2025.

Table 2.40 Virtual Testing and Simulation for Racing Cars in Automotive, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.04	0.10	0.21	0.41	0.75	1.30	2.15	3.33	4.79	6.40	77.1%

(Source: Tractica)

2.6

BUILDING AUTOMATION

2.6.1

BUILDING AUTOMATION AND ENERGY MANAGEMENT

A building controlled by a building automation system (BAS) is often referred to as an intelligent building, smart building, or (if a residence) a smart home. Building automation is enabled through devices that control a building's heating, ventilation, and air conditioning, lighting, and other systems.

AI enables smart control systems to learn about human habits and facility environments without being programmed. ML and DL are being applied in building automation, leveraging sensors like motion detectors, photocells, temperature, air quality, smoke detection, cameras, and vibration. Companies are using these inputs to closely identify and track environmental dynamics and threats, and recommend spatial optimization in and around buildings.

PointGrab is a company that provides sensing hardware and software that use DL and CV embedded into IoT devices for edge processing. Specifically, the company uses object tracking algorithms for background modeling, novelty detection, motion estimation, and non-rigid object detection, coupled with proprietary ML classifiers and training pipelines to support learning and modeling of office/work space management, staff planning, retail analytics, and occupant safety to track movement of building occupants, and to drive energy savings, smarter allocation, and cost savings for commercial environments.

AI is a natural extension to the building automation market, given the large and diverse data collected and the constant need to increase efficiency and decrease costs.

Tractica forecasts that the annual revenue for building automation and energy management in the building automation sector will increase from \$3.35 million worldwide in 2017 to \$255.22 million in 2025.

Table 2.41 Building Automation and Energy Management, World Markets: 2016-2025

Units (\$ Millions)											CAGR (2016- 2025)
	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	
-	3.35	8.93	17.95	32.18	53.89	85.63	129.63	186.67	255.22	N/A	

(Source: Tractica)

2.7 BUSINESS SERVICES

2.7.1 AGENT-BASED SIMULATIONS FOR DECISION-MAKING

Complex organizations like businesses and governments have long understood the importance of long-term strategy. Entire industries of strategic consulting firms, financial and industry analysts, and competitive intelligence brokers exist to help such organizations plan for future scenarios. In business contexts, strategic decision-making might include areas like competitive, economic, consumer, and technological forces, security planning, disaster response, market expansion, regulatory or policy impacts, employee turnover, distribution requirements, inventory, etc. In any of these contexts, businesses are faced with the challenge of understanding highly complex systems and designing sophisticated financial, service, and technical schema, governance frameworks, and feasible outcomes while balancing costs and “what-if” scenarios.

In perhaps one of the most alluring applications for AI, agent-based simulation for decision-making is useful in simulating and predicting the behavior of complex systems, where millions of individual entities or agents (humans, economies, transactions, cars, viruses etc.) can have multiple dynamic characteristics. Each of the entities interacts with each other and behavior can be simulated using AI techniques like reinforcement learning. To understand and plan for complex systems benefits from simulation, developers and planners in the past were limited by compute power, and ability to scale or introduce new elements in real time. Graphics processing units (GPUs) and high-speed processors are helping make virtual simulation possible.

One example of a company using AI to help businesses with new software deployment is ScenGen (short for Scenario Generation) by Scorpion Computer Services. ScenGen is designed to generate all possible scenarios for a given situation, and then simulate the execution of all user actions, messaging, data problems, and tests for new software releases. Financial services firms, aerospace and defense, and utilities companies use ScenGen to reduce issues, damages, downtime, misinformation, bugs, memory exceptions, crashes, and failed installations. In financial services applications, clients use the software to test quality assurance for equity trading systems, e-commerce decision engines for credit underwriting, and model insurance problems. Using AI to accelerate scenario planning and testing is a powerful way to mitigate risks.

Tractica forecasts that the annual revenue for agent-based simulation for decision-making in business will increase from \$0.02 million worldwide in 2018 to \$0.66 million in 2025.

Table 2.42 Agent-Based Simulations for Decision-Making in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.00	0.02	0.04	0.08	0.13	0.21	0.31	0.42	0.55	0.66	117.1%

(Source: Tractica)

2.7.2 AUDIO AND VIDEO MINING

With the rise of digital media, high volumes of content are now a business imperative for marketing, sales, customer support, and engagement. In business, audio and video are useful media for storytelling, brand awareness, and education, but until recently, such efforts, whether owned, paid, or earned, have been difficult, if not impossible to mine for insights. In sales and support, there is not a lot of data being recorded in terms of the conversations employees are having.

Now, organizations can begin to leverage these insights at scale. As an extension of image recognition and analysis, AI is now also being used by organizations to aid in audio and video mining. In a marketing or market analysis context, speech and voice recognition can be mined for specific moments, such as a user posting a video about a product. In a call center context, AI can be used to transcribe, identify keywords, and mine phone calls, video footage, or online media. DL can also be applied here for auto-generated speech-to-text transcription.

Chorus uses NLP to analyze sales calls with the intent to improve sales outcomes for internal sales teams. The model identifies moments that impact selling outcomes and that can be used for real-time sales coaching, for collaboration, and ongoing learning and improvement. “Studies show that win rates increase by 33% with a proper coaching program in place, yet most managers don’t have the time to sit in on calls, and no one has the capacity to learn from the thousands of meetings that take place each quarter,” said Roy Raanani, Chief Executive Officer (CEO) and co-founder of Chorus, in a statement. According to Raanani, there are few, if any objections from sales teams. “It’s less about the salesperson and more about what the customer is saying,” said Raanani, “The meeting is recorded, and now you have notes that are time stamped. That’s helpful for the salesperson because they don’t have to remember everything that was said.” The system will also issue simple reminders, such as checking halfway through the meeting to ask the salesperson if they are getting what they need. The platform integrates with Salesforce and automatically captures meetings in WebEx, GoToMeeting, Zoom, Join.Me, UberConference, BlueJeans, and ClearSlide. Customers have used Chorus to analyze more than 500,000 sales conversations over the past year, according to a press release.

Another company leading in this space is Deepgram. Businesses are using their platform for discovery, wherein a call center agent can quickly search through troves of old audio datasets to surface relevant calls and solutions. Call centers are also using the tool to track keywords, phrases, and mine past call center calls for quality assurance and compliance. Other customers are using the Deepgram for keyword discovery for marketing, sales, and other areas internally to rapidly source relevant content or context. The company also builds custom models for clients to automatically analyze and classify their audio or video streams. Deepgram also provides an API that allows users to apply audio and video mining to calls, meetings, podcasts, video clips, and lectures, and then rapidly search them. This is a very promising use case for AI, particularly given the relative darkness of audio and video content compared to text analytics to date.

Tractica forecasts that the annual revenue for audio and video mining in business will increase from \$.04 million worldwide in 2016 to \$1.81 million in 2025.

Table 2.43 Audio and Video Mining in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.04	0.08	0.14	0.24	0.39	0.59	0.85	1.16	1.49	1.81	54.5%

(Source: Tractica)

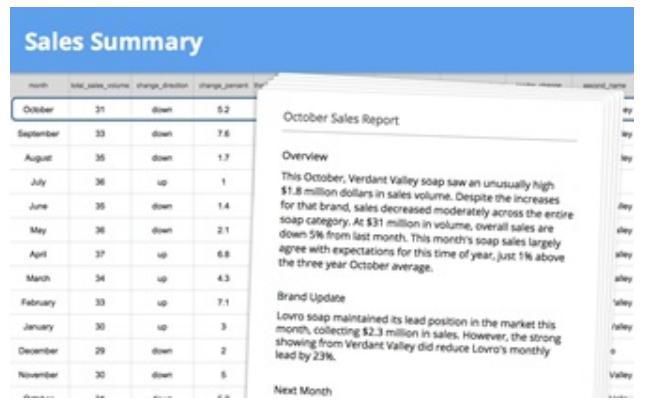
2.7.3 AUTOMATED REPORT GENERATION

Companies generate reports for internal stakeholders, as parts of client programs, or even as formal products. They do so on everything from advertising performance to sales to employee satisfaction, with every level of frequency. As the amount of data flowing into and across organizations grows more and more massive, the problem is not just one of content distribution, but of the time it takes to comprehensively identify and organize insights that are useful and consumable.

AI is now a tool well suited for report generation. Using NLP, ML, and DL in some cases, companies are using AI to collate reports far more rapidly than humans. AI-generated reports can surface relevant metrics, tables and charts, and generate multiple paragraphs of narrative. Automated report generation tools generally support the following tasks:

- **Data Sourcing:** Identify and extract data from relevant internal and external sources, including industry news and reports, social media listening, and competitor intelligence
- **Data Interpretation:** Upon consolidating data in standardized formats, the solution aligns the data in templates, codes and prepares it for analysis using ML
- **Data Analytics:** Defines business rules and correlation/causality at scale; with predictive modelling and data enrichment, solutions can run hundreds of “what if” scenarios and perform trend analysis
- **Narrative and Semantic Commentary:** Using NLP and generation, solutions can sometimes automate variance analysis and commentary writing in a systematic and structured way

Figure 2.5 Sample Sales Summary Populated by Artificial Intelligence



The figure shows a screenshot of a Sales Summary report. The top section is titled "Sales Summary" and contains a table with columns: month, total_sales_volume, change_direction, and change_percent. The table data is as follows:

month	total_sales_volume	change_direction	change_percent
October	31	down	-5.2
September	33	down	-7.6
August	35	down	-1.7
July	36	up	1
June	35	down	-1.4
May	36	down	-2.1
April	37	up	6.8
March	34	up	4.3
February	33	up	7.1
January	30	up	9
December	29	down	2
November	30	down	5
October

The bottom section is titled "October Sales Report" and includes a "Overview" paragraph and a "Brand Update" section. The "Overview" paragraph discusses the overall sales volume and trends for October. The "Brand Update" section highlights L'Oréal's market position and its lead over Verdant Valley.

(Source: Econsultancy)

Automated Insights has a product called Wordsmith, which specializes in auto-generating client reports for marketers and agencies. Clients can customize the fields they want and the frequency at which reports are run, and the AI runs the rest. Reports analyze millions of data points and are delivered in standardized format. The company estimates some 4 to 6 hours of labor saved for each report generated. It also works with media organizations like the Associated Press to deliver custom reports on finance, sports, politics, and beyond. The company also provides an API for developers to take data and convert it to consumable reports with narrative.

A number of other companies are emerging in this space (e.g., Arria, Genpact, Narrative, Narrative Science) and a variety of other vertical specialists as the need is almost universal.

Tractica forecasts that the annual revenue for automated report generation in business will increase from \$0.5 million worldwide in 2016 to \$0.89 million in 2025.

Table 2.44 Automated Report Generation in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.50	0.51	0.52	0.54	0.58	0.62	0.68	0.75	0.82	0.89	6.6%

(Source: Tractica)

2.7.4

AUTOMATED WORKFORCE SCHEDULING

Whether scheduling internally or as part of service programs, scheduling takes time and is almost entirely reactive in nature. Entire roles are created to expedite the scheduling process. But with the rise of data-driven service programs, scheduling field services becomes part of a company's differentiation and can influence major cost factors, such as productivity, time to resolve a specific problem or machine, or downtime.

Automated workforce scheduling is now a task handled by AI. As systems collect more data on machine/product/service performance, malfunction patterns, and employee or field service whereabouts, and data integration grows more sophisticated, ML is used to facilitate faster and more optimized scheduling, as well as more preemptively.

ServiceMax, a platform recently acquired by GE, is an automated dispatch, scheduling, and workforce optimization tool that uses AI to constantly improve upon field service scheduling. The model consumes huge amounts of data and schedules specific field technicians, parts, and inventory needs based on location, skill sets, the type of job, duration of job, etc. It also uses service data to "learn from" past interactions, schedule parts more efficiently, optimize routes, and decrease time to resolution. The workforce optimization service is part of the broader service automation platform. The sheer amount of time tools like this save renders this a very promising use case. The ability to monitor and learn from fix rates, technician needs, and performance paves the way for better customer experiences and effectiveness across the service delivery chain.

Tractica forecasts that the annual revenue for automated workforce scheduling in business will increase from \$0.05 million worldwide in 2016 to \$4.39 million in 2025.

Table 2.45 Automated Workforce Scheduling in Business, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
	0.05	0.15	0.31	0.55	0.91	1.40	2.05	2.82	3.63	4.39	64.4%

(Source: Tractica)

2.7.5

CHATBOT-BASED BRAND/SERVICE INTERACTIONS

The advent of the internet and social media brought brands into dialog with consumers. What was once a one-way “broadcast to many” model evolved into a model in which brands were able to communicate directly with customers; customers could communicate directly with brands in real time and on social platforms and communities. The struggle for brands has been how to effectively scale this kind of brand-customer interaction, while maintaining personalized and “authentic” interactions and accounting for numerous micro and macro contexts.

Enter chatbots, loosely defined as AI-enhanced computer programs able to hold audio or text-based conversations that simulate convincingly how a human would interact. Chatbots are now being used by brands for service interactions, including simple outreach, education, feedback and survey collection, questions and answers (Q&A), tips and advice, etc. Many brands are using chatbots to extend the brand as a “friend,” easing pressures to buy by developing such bots with personality and the ability to engage far beyond the scope of sales or customer support.

Figure 2.6

Mark Zuckerberg on Messenger Business at F8 Conference in 2016

“I don’t know anyone who likes calling a business. And no one wants to have to install a new app for every business or service that they interact with. We think you should be able to message a business, in the same way you would message a friend.”—Mark Zuckerberg at F8 in 2016.

(Source: Facebook)

Perhaps one of the most notable examples of a brand-developed chatbot capable of interacting on topics far beyond the brand’s product and services is Amazon Alexa and the Echo product. Primarily a voice-interactive interface, users can ask Alexa everything from “tell me a joke” to “sing me a song” to “how many milileters are in a cup” to “turn on the lights” to “when is the next baseball game.” While the development of the Alexa is ongoing, and no small feat, Amazon’s in-home, device-based chatbot is feeding Amazon incredibly rich and valuable insights about their customers.

Beauty brand Sephora recently launched a chatbot that works on messaging app Kik in which users can message the brand for personalized beauty tips. By taking a short quiz, the bot offers up specific recommendations for make-up, hair care, and self-care techniques, and surfaces relevant reviews of relevant products. According to the company, the same quiz offered on Kik generated a 40% higher completion rate than similar campaigns run on other platforms.

Companies can also use chatbots as personalized triage mechanisms, wherein bots answer the “easier” questions asked many times before, thereby profiling and recommending users to the right person. Healthcare startup HealthTap offers an interesting example wherein, via

Facebook Messenger, users can ask health and wellness-related questions, and the bot determines whether the question may be addressed with existing content developed from similar interactions, or merits interaction with a human doctor, one from a network of more than 100,000 doctors on its platform.

Highly regulated industries, such as financial and insurance services, are deploying AI for customer service interactions, even beyond chatbots. Wells Fargo, for instance, is using a software called Mattersight, which analyzes callers' tone, tempo, keywords, and grammar to triage calls based on specific parameters and words. The company claims to reduce call times by 23%.

While chatbot-based brand interactions that do not involve sales or customer issue resolution may not yield the most immediate returns and revenue streams, delightful brand experiences, accessible anytime, anywhere, on the device and apps of choice carry tremendous value. As consumers engage with chatbots on Facebook Messenger, Kik, or WhatsApp, for example, brands are gaining real estate for far less cost than developing and hosting their own mobile applications. Of course, brands must navigate wisely in these contexts, as they are indeed immediate extensions of the brand itself, subject to mishap, liability, and public relations (PR) crises just as humans would be.

Tractica forecasts that the annual revenue for chatbot-based brand/service interactions in business will increase from \$8.91 million worldwide in 2016 to \$716.65 million in 2025.

Table 2.46 Chatbot-Based Brand/Service Interactions in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	8.91	25.37	51.13	90.49	148.50	229.60	335.01	459.73	591.92	716.65	62.8%

(Source: Tractica)

2.7.6 CHATBOT-BASED E-COMMERCE AND SALES

As brands constantly work to both scale and personalize customer interactions, chatbots are also penetrating e-commerce and service interactions. (See Section 2.7.5 for an overview of non-sales-oriented brand utilizations of chatbots.)

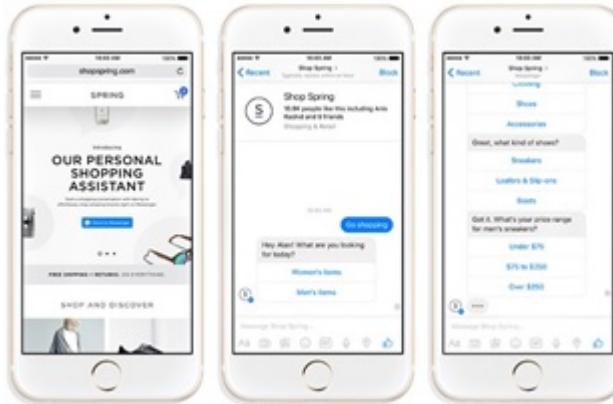
NLP, ML, and DL are powering an explosion of chatbot development and activity from companies across all industries and all sizes. Chatbots are handling a range of e-commerce related tasks, including but not limited to:

- Product search and discovery
- Product customization
- Product/service account alerts
- Product selection, purchase
- Appointment reservation, booking
- Location search and discovery
- Customer service and support
- Customer triage

- Integration with customer account or loyalty programs

Facebook, Kik, and WeChat are platforms that have opened up messenger APIs to allow brands to use their platforms for chatbot-based marketing, support, e-commerce, and sales interactions. Brands like 1-800-Flowers provide Messenger chatbots to quickly search and order flowers; users can just as easily order an Uber, or search and buy plane tickets from KLM or AirMéxico right from within Facebook.

Figure 2.7 Retailers are Integrating with Facebook Messenger App to Tie E-Commerce Directly to Facebook Experience



In the image above, users can search retailer Spring's catalog directly from the Facebook Messenger app. The chatbot serves up a series of questions to quickly tailor recommendations based on user inputs.

(Source: Facebook)

Other brands, such as Whole Foods, Pizza Hut, Disney, and The North Face, have developed their own chatbots, available on their own mobile apps, websites, short messaging service (SMS), and messenger apps alike. The North Face, for instance, built an expert personal shopper (XPS) bot that helps match specific customer needs with specific products. The company recently reported that customer engagements with the bot averaged about 2 minutes in length and the platform had a 60% CTR for product recommendations.

These sorts of conversational interfaces have taken off in recent years, thanks in part to AI advancements robust enough to support them, but also due to the scale, personalization, and significant saved hassle and time they promise. No more hold times, annoying phone tree loops, or repetitive conversations with different call center agents; instead, more rapid search and discovery, and potential for faster sales and service conversion.

Tractica forecasts that the annual revenue for chatbot-based e-commerce and sales in business will increase from \$19.6 million worldwide in 2016 to \$794.66 million in 2025.

Table 2.47 Chatbot-Based e-Commerce and Sales in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	19.60	37.63	65.84	108.95	172.47	261.29	376.72	513.31	658.07	794.66	50.9%

(Source: Tractica)

2.7.7

CROWDSOURCED MARKET RESEARCH

Since the dawn of business, market research has enabled insight into groups of people at scale. As technologies have evolved, so too have market research techniques, from focus groups to online surveys and panels to mobile apps and far beyond.

AI is now permeating the market research space by using ML, DL, and CV to capture human insights at scale more rapidly than ever possible. Using authenticated mobile devices, both machine and human intelligence can yield highly nuanced, while still data-driven insights around very specific problems, demographics, or market questions. AI is also being used to drive secondary market research at scale, in media, finance, and government, among other industries, as outlined in sections 2.17.9 and 2.19.3.

Premise specializes in using a panel of mobile information gatherers on the ground to help companies collect data and insights about specific problems. CV enables new insights when, for example, companies can conduct market research on products into which they have low visibility. A global consumer packaged goods (CPG) company with market presence worldwide lacked metrics and data into where products were traded and sold after distribution centers. The company used Premise for a study in Vietnam to field human information gatherers, equipped with mobile cameras, to verify product availability, stock keeping unit (SKU) pricing, shelf placement and share, and quantities. As a result, the brand was able to optimize local brand and customer strategies.

Tractica forecasts that the annual revenue for crowdsourced market research in business will increase from \$.03 million worldwide in 2016 to \$2.91 million in 2025.

Table 2.48 *Crowdsourced Market Research in Business, World Markets: 2016-2025*

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
	0.03	0.10	0.21	0.37	0.60	0.93	1.36	1.86	2.40	2.91	63.5%

(Source: Tractica)

2.7.8

ENTERPRISE CHATBOTS FOR PRODUCTIVITY AND COLLABORATION

Enterprises and organizations large and small struggle to maintain and evolve their internal strategies and tactics to drive and improve workforce productivity and collaboration. Many organizations invest in cultural analysis, training, and tools to enhance group productivity and enable personal productivity as well. Enterprise social networks and mobile app ecosystems are two important examples of how software has already transformed the way employees communicate and work horizontally across each other, and vertically within hierarchies.

Chatbots are now being applied to workforce productivity and collaboration tasks. Given the wide range of tasks employees now do and document online, AI is being used to collect and mine this information across teams, then trigger specific messaging, actions, and reports.

Slack, the enterprise messaging platform, which now has upward of 2.3 million visitors a day, is developing “manager bots.” These bot-enabled digital assistants can automate managerial tasks, such as communicating with team members, sending reminders, due dates, and collecting and sending status updates to others in the organization.

Chatbots are a notable AI application when it comes to workforce productivity, and may be particularly effective in their ability to tailor tone, messaging, and timing to better suit

personalities and workstyles. In conjunction with business and productivity-related AI applications, such as report generation, real-time news analysis, or predictive sales, employees at all levels may increasingly leverage bots for more efficient communications and workflows.

Tractica forecasts that the annual revenue for enterprise chatbots for productivity and collaboration in business will increase from \$26.76 million worldwide in 2016 to \$44.79 million in 2025.

Table 2.49 Enterprise Chatbots for Productivity and Collaboration in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	26.76	26.84	27.24	28.12	29.63	31.88	34.84	38.30	41.78	44.79	5.9%

(Source: Tractica)

2.7.9

INTELLIGENT CUSTOMER RELATIONSHIP MANAGEMENT SYSTEMS

Customer relationship management (CRM) systems have been helping organizations track and make sense of customer sales and interactions for years. What was born primarily as a sales tracking tool has expanded, with the advent of digital and social media, into robust platforms with tools designed to unify insights around broader customer interactions and transactions, beyond just sales. Functionality tends to support at least four areas: contact management, customer acquisition, sales, and customer service. The goal of these systems is to facilitate “a single 360° view” of any individual customer, although this has been easier said than done given the complexity of integrating online and offline customer profiles and behaviors. As the internet has forced businesses to prioritize customer experiences, CRM has become the critical tool for the job.

AI is now infusing all aspects of CRM systems, and CRM more broadly. When it comes to **contact management**, companies are using ML and DL to mine large data sets for cleanliness and data integrity, purging bad data, helping process incomplete contacts, suggesting those to de-duplicate, etc.. AI can be used to suggest potential contacts worth outreach as well. This is a particularly useful tool for sales enablement and **customer acquisition**. When it comes to sourcing, analyzing, prioritizing, and predicting prospective customers, AI is being applied for predictive lead scoring, suggested prioritization for sales outreach, and optimizing related sales workflows. ML and DL, in conjunction with NLP, are being applied for content curation and strategic outreach, wherein models process large data sets and then recommend specific content, offers, and outreach that may resonate with particular kinds of prospects or customers.

AI-enabled CRMs are also helping companies assess which customers could be the most profitable and likely to respond to sales outreach. AI is also being used for **sales enablement, even predictive sales**. Similar to predictive or proactive customer service, AI can help scale sales agents’ ability to read, triage, and respond to inbound prospects; to analyze and predict the most appropriate action to take based on behavior and conversion trends; and even to filter, score, and prioritize similar leads. Not only do AI models take into account customer trends, but some companies, such as AgilOne, fuse CRM data with external data from news, social media, weather, etc. to come up with sales leads and predictive pitches.

Finally, the post-purchase phase of the customer lifecycle is being enhanced by AI-enabled CRM systems as well. **Customer service**-related use cases enhance efficiencies on both enterprise and consumer sides. For consumers, the benefit should be more pain-free support experiences, void of redundant conversations and repetitive troubleshooting, and even delight through preemptive service actions that prevent downtime or failure. When tools like chatbots are effective, they can save customers time and energy. On the enterprise side, call centers and service agents are using AI to automate simple Q&A through chatbots; to automate triage and service escalation, activity capture, case classification, recommended responses, etc. AI is also increasingly used by service organizations to more efficiently allocate resources.

USAA is working on initiatives with Intel's Saffron platform, which analyzes thousands of factors in order to match broad patterns of customer interactions and behaviors to model when, how, and why a customer might reach out for support needs. Using this information, USAA allocates call center agents accordingly (e.g., how many people are needed for chat support, phone support, etc.) It is also able to use the same data to inform more personalized communications. At the time of this report's publications, this initiative has helped USAA improve its guess rate for how, when, why, and for which product users will next contact support, from 50% to 88%.

The sum of these capabilities works toward marketers' goals for deeper customer insights and stronger relationships. Salesforce.com, SugarCRM, Infor, NetSuite, and many other CRM software providers are developing AI capabilities across their product suites.

Tractica forecasts that the annual revenue for intelligent CRM systems in business will increase from \$12.21 million worldwide in 2016 to \$242.35 million in 2025.

Table 2.50 Intelligent CRM Systems in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	12.21	17.56	25.94	38.74	57.60	83.97	118.25	158.80	201.79	242.35	39.4%

(Source: Tractica)

2.7.10 INTELLIGENT RECRUITING AND HUMAN RESOURCES SYSTEMS

Companies spend many millions a year optimizing their recruiting and engagement efforts in order to make for happy workers and workplaces, and avoid costly turnover and re-hiring processes. With the advent of the internet and social media, human resources (HR) departments realized huge opportunities lie in the online channels and individual data available for sourcing the right talent. An entire market of recruiting and HR software helps many businesses with their recruitment and employee engagement efforts; typically, functionality on these platforms supports candidate and talent sourcing, recruitment, and employee engagement.

AI is now being applied to save time, energy, and money during the **talent sourcing and recruitment processes**. Models mine large data sets, third-party job sites, and social media to source candidates with higher likelihood of interest and hiring potential. These models help sort multiple resumes, mine text for specific needs, prioritize, and surface candidates.

Textio uses AI to develop content. Recruiters have been using the ML platform to provide real-time suggestions for job postings designed to be gender neutral and appeal to broader pools of candidates. According to Textio, clients who maintain a score of 90 or higher attract

applicant pools that are, on average, 24% more qualified and 12% more diverse.

Belong Technologies is working on **predictive talent**, which proactively source candidates most likely to move forward with personalized interactions. Predictive talent-finding also includes sourcing candidates that align with business requirements like business performance, growth, and resource allocation. 12grapes has candidates undergo questionnaire and video screening, then uses AI to analyze facial and emotional cues, develop a profile, and recommend job prospects. A variety of companies are developing AI-enabled software to support this, including Belong, Connectifier, Wade & Wendy, and others.

Similar to AI applications for CRM and workforce collaboration, some systems are also using ML to drive **employee engagement**. BetterWorks, for example, focuses on using AI to ease the employee-manager feedback loop. It does this by building work profiles, which it calls “Work Graphs” based on data integrations across Google Apps, Office 365, Salesforce, JIRA, email, and Slack, then track employees’ goal progress, alignment, comments, cheers, budgets, cross-functional collaboration, etc. to inform employee engagement strategies. Specifically, they use ML to prompt contextually appropriate feedback, recognition, council, questions, and learn from employees’ preferred channels, time of day, etc.

SkillSurvey predicts individuals’ turnover and performance based on words used by the people listed as references. References are presented with an online series of behavioral-science-based questions tailored to the specific job and inputs are graded and averaged. The results can be compared with a large database of candidates for the same position. HealthSouth, which employs 24,000 people, reported a 17% decrease in employee terminations, a 10% drop in people quitting, and 92% less time spent checking references after one year of using SkillSurvey. The tool is also used by other large brands like Adidas, Keurig, and Reebok.

This technology has great potential given the strong desire on both the part of the employer and the candidates to find the right fit. It is critical that such a system take into account risks concerning racial, gender, age, or any other inherent bias, disenfranchisement, or other unfair advantages to which code could be blind.

Tractica forecasts that the annual revenue for intelligent recruiting and HR systems in business will increase from \$8.31 million worldwide in 2016 to \$1.44 billion in 2025.

Table 2.51 Intelligent Recruiting and Human Resources Systems in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	8.31	41.72	94.00	173.88	291.60	456.18	670.09	923.20	1,191.47	1,444.58	77.4%

(Source: Tractica)

2.7.11

PREVENTION AGAINST CYBERSECURITY THREATS

Maybe the single greatest threat to any business today is cybersecurity. While this is nothing new, the proliferation of systems, cloud technologies, apps, devices, and distributed endpoints has only exacerbated cybersecurity threats. With global cyber spending expected to reach \$170 billion by 2020, eyes are on the cybersecurity industry to innovate better and better solutions.

Companies are now turning to AI to aid in security protection for their business assets. Many techniques developed in defense and military programs may now be applied to business problems and processes. Specifically, companies are using ML, DL, and MR to review massive amounts of data (billions of log files a day, for instance) to detect suspicious behavior.

DarkTrace is a startup in this space that aspires to mimic the human immune system in its response to security threats. Its Enterprise Immune System technology has the ability to detect previously unidentified anomalies and potential threats in real time, which other legacy approaches either fail to see or take longer to eradicate. By applying its unsupervised ML system, DarkTrace claims it has identified 30,000 previously unknown threats in over 2,400 networks, including zero-days, corporate espionage, IoT hacks, criminal campaigns, insider threats, and more stealth attacks.

A number of other startups, such as DeepInstinct, BlueVector, Cylance, Jask, Harvest.ai, PatternEx, and others, are developing AI tools for cybersecurity. Enterprises are also involved. As part of a year-long research project, IBM's Watson for CyberSecurity partnered with numerous universities and institutions to train the model on security language by learning the nuances of security research findings and discovering patterns and evidence of hidden cyberattacks and threats that might otherwise go unseen. This use case has the potential to become an essential enterprise tool to thwart cyberthreats. It is critical that these tools themselves do not become vulnerable to attack or proliferating attacks.

Tractica forecasts that the annual revenue for prevention against cybersecurity threats in business will increase from \$1.39 million worldwide in 2016 to \$38.73 million in 2025.

Table 2.52 Prevention Against Cybersecurity Threats in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	\$1.39	\$2.25	\$3.61	\$5.69	\$8.75	\$13.03	\$18.59	\$25.17	\$32.15	\$38.73	44.8%

(Source: Tractica)

2.7.12

PROCUREMENT MANAGEMENT

Procurement is the act of finding, acquiring, and buying goods, services, or works from an external source, often via a tendering or competitive bidding process. For years, the process was manual, later becoming somewhat more digitized through procurement software systems.

As the entire supply chain management challenge grows evermore digitized and automated, procurement tools are becoming more AI-enabled. This is in part due to the vast amount of data, mostly unstructured, now critical to supply chain visibility—images, voice, sensor data, video, etc. Automating the supply chain is not only focusing on automating supplies, but insights into demand as well.

Companies like SMART by GEP and Coupa's Spend360 are two examples of AI-powered spend management platforms, which ingest vast amounts of spend data (e.g., invoices, engagements, travel and expense data, etc.), learn from these data, and serve up recommendations for optimization. These platforms focus on optimization in areas like spend analysis, consolidation, sourcing, compliance adherence, assigning probabilities for classification, supplier management, etc.

SAP's Ariba software recently introduced a procurement bot (named Procurement), which introduces a conversational interface to communicate with buyers and sellers about orders. The bot is designed to learn about users' specific preferences and companies' policies and procedures to enable faster processing, fewer errors, and easier compliance. SAP also released a Slack bot, which communicates with Concur and SuccessFactors, to support broader employee expense management.

Broadly speaking, procurement management enabled by AI carries a host of considerations around the efficiencies versus risks enabled through automated procurement. Supply chain visibility and agility is becoming an essential means of efficiency gains for businesses in any industry. But the race to digitize every aspect faces hurdles: for instance, when machines themselves begin processing or negotiating spending, new legal and regulatory questions will emerge.

Tractica forecasts that the annual revenue for procurement management in business will increase from \$.26 million worldwide in 2016 to \$14.21 million in 2025.

Table 2.53 Procurement Management in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.26	0.58	1.09	1.86	3.01	4.61	6.68	9.14	11.75	14.21	56.2%

(Source: Tractica)

2.7.13 PROJECT AND STAKEHOLDER MANAGEMENT

Companies use project and stakeholder management tools to automate all kinds of workflows, from internal communications and planning, to client relationships and project delivery. A large market of project management tools exists today to help organizations keep minute-to-minute tabs on the status and milestones of projects, stakeholder engagement, and on-time delivery. But still, data about projects (both present and past), stakeholders, projects, or market changes that take place remains fragmented, poorly disseminated, and can stifle efficiencies.

AI is now being applied to these needs in a number of areas, often supported by ML, DL, NLP, and CV. AI in project management typically plays the role of assistant, facilitator, or expert. AI is also a vehicle for transmitting large bodies of knowledge and project management to specific users in highly customized modes and dashboards. The list of project management problems to which AI can be applied depends on the industry, and varies widely. Some include, but are not limited to managing scope, time, costs, and operations across the following phases:

- **Planning:** AI could aid in minimizing errors in developing project plans, scoping, benchmarking, etc. AI can be used to pull in historical and/or market data to simulate planning scenarios
- **Resource Allocation:** AI could aid not only in optimizing resource allocation based on historical data, but could potentially be applied to efficiently source, run outreach, or negotiate specialist or freelancer talent.
- **Tracking:** Monitoring, prompting check-ins or feedback, consolidating insights, and tracking interactions are some of the areas in which AI can help unify disparate data sets for a “complete view” into project status.

As an assistant, project management bots, for example, can be used to offer recommendations, provide real-time snapshots or updates, manage expectations, and assemble learnings for future projects. They can also recognize key characteristics of stakeholders and provide recommendations on how to engage with them, maximizing alignment of “hard and soft” client engagement factors. Many AI applications for project management will augment, but not replace humans given sensitive client relationships.

Palisade has a product called @RISK, which supports planning and risk modeling for project management using ML in conjunction with Monte Carlo simulations.

Tractica separates this use case from project management because it involves outside stakeholders across the value chain, rather than internal employees only. Reference Section 2.7.17 for an overview of AI-enabled project management in which clients and external partners are not involved.

Tractica forecasts that the annual revenue for project and stakeholder management in business will increase from \$0.25 million worldwide in 2016 to \$6.7 million in 2025.

Table 2.54 Project and Stakeholder Management in Business, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
0.25	0.43	0.70	1.10	1.67	2.45	3.43	4.55	5.69	6.70	43.9%	

(Source: Tractica)

2.7.14 REAL-TIME NEWS ANALYSIS AND COMPETITIVE INTELLIGENCE

Companies across industries are in a constant race to stay up-to-the-minute on news and competitive intelligence that directly impact their reputations, services, and bottom lines. In the past, businesses have relied on traditional market research methods, forecasting, consultants, analysts, and other fairly ad-hoc techniques to assess their strengths, weaknesses, opportunities, and threats in the context of their respective market landscapes.

New capabilities powered by AI are helping companies monitor and analyze a profoundly greater range of inputs to guide strategies, messaging, and product development. ML and DL are being used to generate reports on competitors or market trends, to monitor market trends to provide real-time and personalized reports for specific users, and to improve forecasting for distribution requirements, inventory, market penetration, etc.

ai-one, in partnership with KDD Analytics provides competitive intelligence analytics as a software-as-a-service (SaaS) tool for large enterprises. This collaboration in science, aerospace, and academia helped lay the groundwork for the commercial tool. It compiles massive amounts of data across SEC filings, financial data, and social data, among others. It then processes and standardizes the data so that it is presented in a visually digestible manner based on user, and made into a repeatable and consistent format for quarterly (or more) reports. Its Financial Analyst Toolbox (FaTbx) is a beta solution for enterprises that includes comparatives for three publicly traded competitors, suppliers, or customers. The service offers 30 presentation-ready Tableau dashboards that can be custom configured to deliver financial categories and growth metrics, disclosures, trends, topic heat maps, etc.

This saves clients vast amounts of time, resources, and stress compared to the analog mode of individual analysts. The tool is used primarily by analysts today to accelerate the process of compiling more information, giving them more time to analyze and enhance findings and recommendations. Tractica also found various consulting and systems integration firms

adopting or building tools to provide AI-powered competitive intelligence to their clients.

Tractica forecasts that the annual revenue for real-time news analysis and competitive intelligence in business will increase from \$0.12 million worldwide in 2016 to \$2.04 million in 2025.

Table 2.55 Real-Time News Analysis and Competitive Intelligence in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.12	0.17	0.24	0.34	0.50	0.72	1.00	1.34	1.70	2.04	36.6%

(Source: Tractica)

2.7.15 SOCIAL MEDIA PUBLISHING AND MANAGEMENT

Since the emergence of social media, a vast array of tools has emerged in order to help brands effectively identify, monitor, engage, and learn from user-generated content related to their company or market. The market for social media management tools has evolved fairly rapidly, and many legacy CRM, content, and email marketing software solutions are now re-positioning themselves as “customer experience” platforms. These “full suite” customer engagement tools do not just handle social media publishing, but include the gamut of listening, flagging, engaging, upselling, and optimizing content, ad spend, products, services, support channels, and positioning based on the voice of the customer.

As customer experience management (CEM) software (including social media) coalesces, it is no surprise that providers are using more and more AI to enhance every part of digital content production and distribution. ML, DL, and NLP are growing rapidly as tools for mining big unstructured data sets (e.g., social media posts, comments, reddit threads, online communities, etc.). Plugging first- and third-party data, such as weather or loyalty data, into clustering algorithms and using results in CRM and customer engagement platforms is a rapidly expanding use case for AI.

Some other ways AI helps augment companies’ abilities to use social media to improve customer experience include:

- Detect disgruntled customers through sentiment analysis
- Triage or engage directly with customers using social media for support using chatbots or support agents equipped with AI-recommendation systems
- Automatically tag, classify photos, logo placements, and brand mentions using image recognition
- Offer alerts, information, product updates, campaign reminders, loyalty incentives, etc.
- Recommend specific products, resources, content, etc.
- Analyze competitive content performance, resonance, and score accordingly
- Communicate in multiple languages

AgilOne helps marketers integrate customer data from across digital, mobile, social, and even physical channels (e.g., in-store, car), and uses ML to predict customer behaviors, and then tailor hyper-personalized and targeted interactions. AgilOne uses customer algorithms

to surface specific behaviors or group company products; propensity modeling to score leads and gauge likelihood to buy, convert, unsubscribe, whether they have money left to spend; recommendation engines for products and services; and marketing spend analysis, which looks at customer attributions to model and match prospects with the most successful acquisition and retention approach.

One of its customers, Peter Glenn, a ski and sports retailer, was able to extract unique trends across buyer segments and channel patterns. This informed advanced segmentation configurations that they could test and optimize promotional, in-moment, and lifecycle campaigns to drive engagement, in-store traffic, and sales during non-peak months. With automated targeting and personalization campaigns for cart abandonment, upsell, next-sell, catalog sends, reactivation, high-value customers, and customer churn, Peter Glenn is able to deliver personalized marketing at scale. The company reports that prior to the partnership, more than 80% of its customer base had lapsed; since then, the average order value (AOV) of campaigns has increased 30%.

Tractica forecasts that the annual revenue for social media publishing and management in business will increase from \$2.62 million worldwide in 2016 to \$44.38 million in 2025.

Table 2.56 Social Media Publishing and Management in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
											(\$ Millions)
	2.62	3.59	5.11	7.43	10.86	15.64	21.86	29.22	37.02	44.38	36.9%

(Source: Tractica)

2.7.16 TRAVEL CONCIERGE AND BOOKING SERVICES

For many years, consumers and businesses alike relied on travel agencies and booking companies to assist in their travel searches, recommendations, and reservations. The rise of the web, aggregation services, such as Travelocity and Kayak, and social media have all but entirely disrupted traditional travel agencies. Instead of relying on a local or corporate broker, finding and booking travel is as simple as using a search engine. AI and machine learning are taking the digitization of travel booking to the next level. Companies are beginning to use the vast amounts of travel data (e.g., seasonal trends, reservation data, pricing data, ratings and review data, social media data, and demographic data, among others) to mine for patterns and correlations across the data in order to serve up more accurate and better-timed recommendations. These applications are almost always designed to drive speedier and higher-value conversion.

WayBlazer is developing chatbots as voice-enabled agents to aid travel agents with super specific results based on inquiry. It uses supervised and unsupervised DL, NLP, and image processing to mine vast amounts of data, running sentiment analysis across text reviews, tagging images and activity content, and other data. From an end-user perspective, a query such as "we want a romantic weekend getaway" is served by inferring properties that have been tagged and force ranked as "romantic," then tailoring recommendations for the individual based on internet protocol (IP) address, user identification (ID), and geographic coordinates. The longer-term vision for WayBlazer is to make the most of sentiment and purchase intention at just the right time, in order to enable proactive booking. Travel systems should know when the kids are on spring break, when you went skiing last year, how it went, etc. The idea is to use this data to proactively offer travel experiences and to offer better conversion by delivering recommendations in advance.

In business contexts, travel is generally less concerned with adventure and serendipity, and more with efficiency, comfort, adequate workspaces, and scheduling. Unlike in consumer markets, where 80% to 85% accuracy is still likely to lead to great vacations, AI-driven booking services for business travelers have somewhat higher thresholds to meet in order to gain the trust many reserve for humans today.

Tractica forecasts that the annual revenue for travel concierge and booking services in business will increase from \$5.79 million worldwide in 2018 to \$248.77 million in 2025.

Table 2.57 Travel Concierge and Booking Services in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	5.79	14.84	28.68	49.07	77.57	114.62	158.47	204.93	248.77	N/A

(Source: Tractica)

2.7.17

WORKFLOW AND PROJECT MANAGEMENT

Workflow and project management can be defined as the process of managing all internal employees, workflows, and collaborations around a specific task or project. Companies rely on all kinds of tools to enhance collaboration on projects, many of which also involve client stakeholders or outside partners, as described in Section 2.7.13.

Supporting internal workflow collaborations and project management is a use case that is ripe for NLP adoption for some redundant tasks given the large amount of complex data that needs to be analyzed and monitored. Particular value can come from the ability of NLP combined with ML to automate monitoring of a broad range of devices and platforms for a real-time view of status and reducing the size of larger project teams. These applications aim to increase productivity, identify appropriate individuals for specific tasks, surface events, content, or messaging to support project management itself.

According to a Harvard Business Review survey, administrative duties of a project, such as determining work schedules and checking on shipments take up 54% of a project manager's time. For teams that work within Slack, there are many chatbots and applications for workflow and project management, including Fireflies.ai, Trello, Asana, Wunderlist, and Pivotal Tracker just to name a few.

Tractica forecasts that the annual revenue for workflow and project management in business will increase from \$3.86 million worldwide in 2016 to \$30.06 million in 2025.

Table 2.58 Workflow and Project Management in Business, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	3.86	4.51	5.54	7.10	9.37	12.50	16.49	21.08	25.80	30.06	25.6%

(Source: Tractica)

2.8 CONSTRUCTION

2.8.1 SATELLITE IMAGERY FOR GEO-ANALYTICS

Satellite imagery has long been a closed domain with high-resolution image databases only available to a select few companies and organizations, such as weather centers, government agencies, the military, and oil & gas companies. Being able to track changes on the ground from space has been vital for these industries, but required human analysis for years. But with rapid increases in the availability and improvement in the level of detail of satellite imagery, and advancements in AI, CV, and DL have created new ways of identifying features, tracking changes, and extracting value from satellite imagery.

Apart from providing a way for humans to track the planet on a daily basis, this also means that image processing will have to be automated, in order to take advantage of this quick refresh rate and trove of imagery data. Collecting information through aerial imaging may be cheaper than a full networked sensor and connectivity implementation, for example. DL is particularly helpful given that it requires low or no feature engineering. Some basic challenges do remain when it comes to weather, viewpoint, lighting, and atmospheric unpredictability.

In construction, satellite imagery is used to assess project feasibility, detect changes within a bounded area, and project progress at a given site. Interestingly, construction sites themselves are taking on new meaning in other industries. Satellite images are being mined for real estate development, conservation efforts estimating deforestation, and forecasting growth by analyzing construction sites, for instance. More generally, satellite imagery can help track a bounded area with alerts and updates provided when something changes in that specific area, or for historical changes over said area. These are not just new applications, but new business models that provide country-wide, or object-specific analysis of satellite imagery to vertical markets.

Companies like Orbital Insight, SpaceKnow, Descartes Labs, and RS Metrics are building solutions that allow anyone to analyze satellite imagery and perform geo-analytics using computer vision and AI.

Tractica forecasts that the annual revenue for satellite imagery for geo-analytics in construction will increase from \$0.45 million worldwide in 2016 to \$6.23 million in 2025.

Table 2.59 Satellite Imagery for Geo-Analytics in Construction, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.45	0.54	0.68	0.91	1.26	1.79	2.55	3.57	4.83	6.23	33.9%

Source: Tractica)

2.9 CONSUMER

2.9.1 AUTOMATED TOUR GUIDE AND ITINERARY SERVICE

Instead of people providing tours, there is potential that social robots could replace humans as tour guides in some instances, such as museums and zoos, where personnel budgets are limited. It is also possible that autonomous cars, vans, buses, trains, or even boats could support tour services, in which the text or content of the tour is delivered by a voice-enabled bot and even uses sensor data to detect real-time context like location, weather, human interactions, etc.

The challenge for social robots is similar to issues that virtual digital assistants (VDAs) like Siri have, which is filtering out ambient noise. In a social robot's case, that would also include filtering and categorizing priority in those speaking to it. When the stakes for social interaction are lower or non-existent, as in the case of building a simple itinerary, NLP, ML, and DL could offer other opportunities. AI could be used to mine relevant data sets, such as past travel, purchase, social, and location data, as well as third-party data on foot traffic, weather, events, etc. to deliver personalized itineraries based on an individual's unique travel contexts.

There is a walking tour for visitors in Helsinki, Finland. The 140-minute tour mixes education, sight-seeing, and adventure with an AI-powered computer game. Participants walk along a route wearing headphones and the AI guides them along streets, shops, on the metro, and beyond.

IBM partnered with Local Motors to develop Olli, an autonomous van that uses IBM Watson IoT for Automotive to provide a “chauffeur” experience. Olli can take passengers to requested destinations, while answering questions about the area, the journey, and providing recommendations for nearby places to see.

Tractica forecasts that the annual revenue for automated tour guide and itinerary services in consumer markets will increase from \$0.03 million worldwide in 2017 to \$0.97 million in 2025.

Table 2.60 Automated Tour Guide and Itinerary Services in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.03	0.07	0.13	0.22	0.34	0.49	0.66	0.82	0.97	N/A

(Source: Tractica)

2.9.2

BUILDING GENERATIVE MODELS OF THE REAL WORLD

The concept of strong AI is the idea that AI is able to exhibit behavior and act as skillfully and flexibly as humans can. Today, this concept remains largely fiction, as what it entails—a vast interconnected understanding of the physical laws, taxonomies, consequences, and even social constructs that govern our world—are a far cry from any AI application to date. Building generative models of the real world is a small but important step in this direction.

At a high level, AI is being used to help generate models and maps of the real world. This is an essential step toward enabling vision-based systems in things like cars and robots, so they can start to understand the physics of the world. By using a combination of sensing technology, including HD cameras, ultrasonic sensors, radar, LIDAR, and GPS mapping technology, highly accurate maps can be generated, with accuracy within a few centimeters. In consumer applications like robotics, a high degree of accuracy is especially important in enabling autonomous devices, which may use this data to establish position while in a home, room, or in-store environment.

Today, the mode of operation for autonomous devices’ “vision” is a function of object detection and generally lacks information beyond category. By contrast, with a generative model, the AI powering vision-based systems would understand the movement (what an object or person is doing). Generative models of the real world help develop the context with which to make a decision about how an autonomous device maneuvers itself.

A German company, Micropsi, is developing robotics for industrial applications, but its techniques are relevant in consumer applications as well. The company trains robots via reinforcement learning and uses models to simulate gravity and movement, and uses gaming reward functions to reward or punish the AI based on task. The company says it has used this technique to power robots capable of painting, using screwdrivers, applying labels to packages, polishing, and other highly-refined tasks. OpenAI, as well as companies like Prowler.io and Improbable, is taking steps to support this use case.

Tractica forecasts that the annual revenue for building generative models of the real world in consumer markets will increase from \$0.33 million worldwide in 2016 to \$35.67 million in 2025.

Table 2.61 Building Generative Models of the Real World in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.33	1.37	2.93	5.20	8.43	12.76	18.16	24.24	30.34	35.67	68.3%

(Source: Tractica)

2.9.3 CALENDAR, MEETING, EVENT SCHEDULING, AND REMINDERS

Scheduling is time-consuming, not just for business engagements, but for consumers as well. AI is introducing numerous possibilities for accelerating the tedious job of scheduling and ensuring successful attendance to events. NLP, in particular, helps analyze text, natural language understanding (NLU) helps process voice for easier spoken interactions to schedule or add reminders. ML can also be applied to help learn from past patterns and recommend times or venues to ensure successful meetings or event attendance.

Startups like x.ai, Zoom, Clara Labs use natural language to understand scheduling requests and to sift through unstructured data, such as email addresses and contacts to automatically propose, correspond, and confirm in person, on line or telephone meetings. Offerings from internet giants like Google and Microsoft can perform some calendar scheduling, but as of now, they are not as seamless and full-service as the specialists mentioned.

Tractica forecasts that the annual revenue for calendar, meeting, event scheduling, and reminders in consumer markets will increase from \$1.91 million worldwide in 2016 to \$11.46 million in 2025.

Table 2.62 Calendar, Meeting, Event Scheduling, and Reminders in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.91	2.14	2.51	3.08	3.92	5.06	6.52	8.20	9.92	11.46	22.0%

(Source: Tractica)

2.9.4 CHILD BEHAVIORAL ANALYTICS

Parents and teachers are instrumental in children's development and such interactions have long been a core focus of child psychologists. With the birth of mobile apps came a wave of child behavior tracking apps, which typically supported classic engagement or monitoring techniques, such as incident tracking, reward systems, suggestions for encouragement, and

even basic money management. An emerging application for AI is to augment such apps and caretakers' abilities to analyze children's digital activity to gain insights. ML, in combination with DL, NLP, facial recognition, gesture recognition, and a variety of technologies, take app-based child behavior analytics to the next level.

An app called Bark uses ML to help parents assess teenagers' online interactions, risks of cyberbullying, sexting, and depression, while protecting their privacy. ML, NLP, and a team of youth advisory specialists can identify language and develop specific tags that may be of concern. For example, codes like "CP9" can mean "parents are nearby" or "53X" for sex. Its software works by monitoring teens' social media accounts for specific behavioral signals without storing or sharing any of the data, which Bark promotes as a welcome trust-builder between parents and teens.

Another example is Light.House, an in-home camera marketed as an in-home assistant. The device uses 3D sensing technology, NLP, and DL to distinguish between adults, kids, pets, objects, and actions, known and unknown. The objective of the device is to aid parents in three areas of insight: what has happened, what is happening, and what is happening that should not be happening. The device supports interactive voice and gesture recognition so parents can communicate with those in the house remotely, custom-design activity alerts, security actions, and historical search of activities or video feed, etc. A few examples commands the Light.House product supports include:

- What did my kids do while I was out today?
- Ping me if the kids are playing outside
- If you haven't seen the kids by 4 p.m., send me an alert.
- What did the babysitter do with the kids?

Tractica forecasts that the annual revenue for child behavioral analytics in consumer markets will increase from \$0.02 million worldwide in 2016 to \$0.52 million in 2025.

Table 2.63 Child Behavioral Analytics in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.02	0.04	0.07	0.12	0.18	0.26	0.35	0.44	0.52	N/A

(Source: Tractica)

2.9.5 COMPUTER-AIDED ART

People have been using computers to generate drawings, images, sounds, music, 3D designs, and a host of other art forms since the dawn of PCs. In these contexts, software programs were developed to aid as a sort of canvas for creation—for drawing, editing images, arranging melodies or rhythms, analyzing the dimensions of objects or other 3D assets for manipulation, manufacturing, visualizing, or other industry-specific needs.

AI is the next evolutionary step to computer-aided art, only instead of simply providing a canvas and tools to automate designs, AI itself contributes to or even fully develops designs. A variety of technologies can support this depending on the application, but today, computer-aided design often leverages ML and NLP to learn from and suggest unique renderings based on training data.

Google recently released Sketch-RNN, a tool that allows users to collaborate with neural networks to suggest different ways to complete your drawing. Start with a shape, and the software then predicts auto-completes for the drawing based on its experience having analyzed millions of user-generated examples. Sketch-RNN follows an earlier Google tool called Quick, Draw! in which it used DL to guess what people were drawing while they were drawing it. Another app called AutoDraw identifies hand-drawn doodles and suggests clip-art replacements.

But it is not just about doodles. Neural network art is becoming a niche genre, as programmers continue to build on each other's work to develop algorithmic-generated images. The AI Painter Artwork tool allows users to upload a photo followed by a photo of a painting, and the app automatically turns the photo into a painting of that style.

Figure 2.8 *AI Painter, a Neural Network that Renders Photos as Paintings*



(Source: Deep Dream Generator)

Other efforts, such as Pix2Pix, DeepWarp, and Google's DeepDream project, all use neural networks to create art based on user inputs.

Tractica forecasts that the annual revenue for computer-aided art in consumer markets will increase from \$0 worldwide in 2017 to \$4.61 million in 2025.

Table 2.64 Computer-Aided Art in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.09	0.24	0.47	0.82	1.33	2.0	2.82	3.73	4.61	N/A

(Source: Tractica)

2.9.6 CONTEXTUAL INTELLIGENCE FOR MOBILE

Pulling together as many diverse and disparate data sets to ascertain user behavior has long been a chief objective for mobile operators and brand marketers. After all, leveraging all of this data to target the right person with the right message or experience at the just the right time remains the proverbial "holy grail" of customer experience. Yet this has been challenging to deliver due to a host of reasons: data integrity, data ownership, privacy/creep concerns, connectivity constraints, and limited demand from end users.

AI presents a number of efficiencies and opportunities to enabling contextually relevant and sensitive services via mobile. Using ML, NLP, and, in some cases, DL to analyze diverse data sets is one efficiency, but the real value comes in training models to identify ways for

users to improve their own efficiencies.

Google's vast hardware and software ecosystems are working toward this. By integrating any single user's native apps, third-party apps, and Android-enabled hardware, they are building smart agents (called Google Assistant) designed to have a sort of aerial view of all data and behaviors.

- **Examples of Native Apps:** Calendar, personal preferences, email, photos, weather, reservations, file share, Google+, Allo messaging bot, and other Google services
- **Examples of Third-Party Apps:** Social media, e-commerce, media, travel, sports, news, etc.
- **Examples of Android-Enabled Hardware:** Mobile phones, watches, Google Home, Android TV, soon will integrate with cars, etc.

Like other virtual assistants, Google Assistant is designed to develop ways to proactively suggest or offer helpful actions in context. For example, a user might be driving, and the assistant might suggest offering an alternative route given sudden traffic, routing the driver by a nearby gas station because gas is running low. Users can summon Google Assistant to quickly offer in-context suggestions, such as in the middle of texting with a friend about going to see a movie, writing "Okay Google, what movies are playing near me tonight?" The more users interact with Google Assistant, the more it learns user preferences and behaviors. Google touts the product as a sort of personal Google, "a Google for your own world."

Google is one of many large (and smaller) companies working on contextual intelligence for mobile. Microsoft's Cortana, Amazon's Alexa, and Apple's Siri are highly competitive plays in contextual intelligence. Meanwhile, many startups and apps are targeting this space with specialized applications; Trevorai is working to manage time to help users align activities and to-do lists toward personal goals and habit builders. Hound is another app that claims superior voice services that deliver results faster. In an adjacent trend, Bragi is developing an earpiece designed to be a self-contained computer for your ear, which uses sensors to offer real-time context via voice.

Tractica forecasts that the annual revenue for contextual intelligence in mobile in consumer markets will increase from \$2.16 million worldwide in 2016 to \$55.31 million in 2025.

Table 2.65 Contextual Intelligence in Mobile in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.16	3.51	5.58	8.71	13.25	19.50	27.50	36.80	46.45	55.31	43.4%

(Source: Tractica)

2.9.7 FACIAL RECOGNITION

Facial recognition is a computer or machine's ability to identify or verify a person based on their facial characteristics. Computer applications use digital images, video frames, and video feeds to recognize people's faces. AI supports facial recognition through various ML and DL techniques, sometimes involving CV. Recognition algorithms are commonly divided into two main approaches:

- **Geometric:** Looks at distinguishing features (face, nose, shape of eyes)
- **Photometric:** Takes a statistical approach by processing an image into values, then eliminates variances by comparing the values with templates

Advancements in processing power and in other adjacent technologies have brought about complementary techniques to enhance facial recognition. Some of these include:

- **3D Facial Recognition:** Using 3D sensors to capture information about shape, depth, lightfall
- **Skin Texture Analysis:** Uses image recognition to turns unique lines, spots into a mathematical space
- **Thermal Analysis:** Uses thermal cameras to detect head shape, while accessories such as glasses or make-up are undetected
- **Eye and Retina Recognition:** Detects unique features of a person's eyes
- **Emotion Recognition:** Facial expressions or physical features are analyzed against databases to determine the subject's disposition

Facial recognition is a verifiable biometric and useful in a variety of security and identity authentication applications. Recent advancements in the technology have also opened up a host of new commercial applications in marketing, service, and customer experience. In consumer markets, facial recognition is being used to unlock software on mobile devices, to organize personal photo collections or tag friends, to search for friends or even lost children, to streamline the e-commerce process of trying on glasses or sampling make-up, or even to authenticate identity on smart home devices or participating in online services, such as educational courses.

An early example is a product called Chui, which is an intelligent doorbell that uses facial recognition to enable keyless, secure, and individual-specific entry. It also scans faces and alerts users of who is at the door. Similar authentication is supplementing verification in other areas like online educational courses, in healthcare check-ins, and in “video-banking” cases.

Disney recently integrated facial recognition into its MyMagic+ systems, to authenticate season pass users, and automatically place photos taken on rides and cruises into personalized albums.

A variety of uses exists for facial recognition in consumer-facing markets about which consumers may be less informed, aware, or accustomed. Some examples include:

- Shopping malls for security and shopper identification, driving personalized in-store engagement from sales associates or mobile apps
- Casinos, hotels, restaurants, stadiums, other high-profile events for security and loyalty member identification, to identify celebrity stalkers, criminals, etc.
- Social media companies for product, service, algorithmic improvement, for targeted advertising
- Advertisers for digital and billboard ad targeting
- Insurance companies to assess risk, health, life expectancy, etc.
- Automotive manufacturers to enable security notifications and individual driver profiles storing preferences for seats, radio, calendar, address book, etc.

- Churches to analyze the frequency with which congregants were attending services

Facial recognition does not have to be individually identifying, and can serve as an effective way to detect useful signals. For instance, South African coffee company Douwe Egberts set up a coffee machine in an airport, and used facial recognition to simply dispense free coffee to those who yawned.

Still, it may not be surprising that consumers often express reticence or even outrage about facial recognition and surveillance without their knowledge or consent. A recent study found that 75% of consumers would not shop in a store that used facial recognition surveillance if the data was used for marketing purposes, according to research firm, First Insight. Unlike other authentication techniques, such as fingerprinting, iris scans, or speech recognition, faces can be recognized without a person's awareness or participation. Furthermore, facial recognition can be used to unearth *additional* personal data about an individual, such as social networking profiles, blog posts, travel patterns, internet behavior, and other areas where individuals' photos may appear. What remains critical for consumer-facing commercial use cases is to inform users of both the use of facial recognition technologies and how data is used thereafter.

Tractica forecasts that the annual revenue for facial recognition in consumer markets will increase from \$1.16 million worldwide in 2016 to \$22.8 million in 2025.

Table 2.66 Facial Recognition in Consumer, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
1.16	1.74	2.64	3.98	5.90	8.51	11.80	15.56	19.38	22.80	39.3%	

(Source: Tractica)

2.9.8 LANGUAGE TRANSLATION SERVICES

Language barriers are not always easy to overcome and can sometimes create barriers for business and consumer relationships. While the kaleidoscope of language and linguistics will remain, AI presents fascinating potential for accelerating the process of translation.

Legacy machine translation has depended on rules-based and statistical models, but accuracy has been an issue, and machine translation has not been accurate enough to replace professional translators. Progress has been made for language translation services through leveraging NLP in combination with ML and DL. Google and Bing use NLP for translation and offer APIs. Open-source software is now available from Harvard, called Open NMT.

Lilt is focused on language translation, a market which Lilt CEO Spence Green estimates is a \$40 billion market opportunity. The solution is used to assist human translation; a human translator looks at a piece they are translating, Lilt looks at the words and makes suggestions. "We estimate Lilt is making translations between 2-7 times faster than unassisted human translation," said Green. The company's largest customer is Canada's Hudson Bay Company, which has 20 translators on staff. Lilt also works with many translation agencies. The solution, which is predictive, marries machine translation and its NLP API.

Tractica forecasts that the annual revenue for language translation services in consumer markets will increase from \$4.29 million worldwide in 2016 to \$127.32 million in 2025.

Table 2.67 Language Translation Services in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	4.29	6.80	10.81	17.06	26.47	39.94	57.88	79.74	103.71	127.32	45.8%

(Source: Tractica)

2.9.9 LOCAL SEARCH AND DISCOVERY

Long ago, search engines realized that part of a successful search was not merely the relevance of the results to the inquiry, but the relevance of the results to the individual searching. Local search and discovery has long been moving toward greater personalization, primarily by taking into account browsing history and any available location data.

What is new is the application of AI and DL into local search, offering far more nuanced personalization than mere browsing history or zip code. Social data, mobile data, IoT data, e-commerce data, demographic data, and so forth are all fed into DL algorithms to deliver hyper-personalized local results based on user preferences. The more search engines know about you, the more relevant the results and user experience (UX) of local search and discovery will become. Consider the difference, for example, between searching:

- Mexican food near me: Resulting in a list of Mexican restaurants within 2 miles
- Mexican food near me: Resulting in top three nearby Mexican restaurants visited within the last 4 weeks, with images of dishes you have ranked on Yelp, notable food allergies, typical restaurant spend, optimal time of day, and coupon based on level of previous engagement

To power such personalization, brands and search engines will use digital assistants, sometimes called virtual agents. These assistants—early examples include Amazon Alexa, Apple's Siri, and Google Assistant—will pervade our homes, cars, office, and smartphones, and possess rich contextual data about users. Voice and speech recognition will only enhance the usability (and our reliability) on these services. Conversational interfaces do not just increase engagement, but they involve voice biometrics, which will drive digital assistants to integrate with multiple devices, pushing consistent UX tailored to individuals' typical behaviors, needs, and interactions. Google Home, for example, can now recognize up to six distinct voices.

In addition to voice, beacons and AR will power local search and discovery as screens become less desirable and hyper-local proximity tracking (e.g., in the bread aisle at the grocery store) become possible. Both AR and beacons will enable users to signal interest and preferences by interacting with the real world, not through touchscreens. This explains why search giants like Google, Facebook, and Microsoft are leading innovations in AR.

Tractica forecasts that the annual revenue for local search and discovery in consumer markets will increase from \$0.82 million worldwide in 2016 to \$33.12 million in 2025.

Table 2.68 Local Search and Discovery in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.82	1.47	2.52	4.16	6.63	10.16	14.88	20.62	26.92	33.12	50.9%

(Source: Tractica)

2.9.10

MOVIE RECOMMENDATIONS

The internet, channels like YouTube, and mobile have forever changed the movie industry from ideation to production to distribution, the mode of movie creation and consumptions are far less centralized. As has happened in print and music media, disruptive new platforms and business models have emerged that offer “all-you-can-watch” movie viewing.

Companies like Netflix, Hulu, YouTube, and Amazon have all but replaced traditional B&M movie rental stores. What these services have in common are advanced uses for data and ML and DL. Netflix, in particular has pioneered DL for movie recommendations supporting users in over 190 countries worldwide. It has developed neural networks to support hyper-personalized suggestions and rankings, search, similarity, and page generation as far back as 2014. Sophisticated models do not just ingest viewing patterns, favored actors, themes, film locations, and languages, but also must account for complex content licensing agreements and term limits, local/cultural variations in taste, global interest communities, device viewing preference, multi-lingual input patterns, and even optimal metrics to measure quality. They also run extensive ML to optimize the look, feel, and organization of personalized home pages.

Figure 2.9 *Netflix Uses Artificial Intelligence for Personal Homepage Optimization and A/B Testing for Page Generation*



(Source: Netflix Tech Blog)

These efforts balance competing agenda: tens of thousands of videos to show; targeting with personalized content that caters to specific interests but is not overly narrow; easy-to-use navigation like search and lists; discovering new content; etc.

Tractica forecasts that the annual revenue for movie recommendations in consumer markets will increase from \$85.68 million worldwide in 2016 to \$509.94 million in 2025.

Table 2.69 *Movie Recommendations in Consumer, World Markets: 2016-2025*

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	85.68	95.49	110.31	132.64	165.55	212.00	273.44	347.94	429.53	509.94	21.9%

(Source: Tractica)

2.9.11

MUSIC RECOMMENDATIONS

The internet and digital media have forever changed the music industry, from production to distribution and even performance and concerts. This is also particularly true for music discovery. It used to be discovering music happened by word of mouth, through magazine subscriptions, or traditional print media. Then arose digital channels like Napster, MySpace, Apple music, and iTunes, and more recently, Pandora, Spotify, and a broad range of “all-you-can-eat” online radio subscription services.

The next evolutionary step in digital music discovery involves AI, in which models are trained on large data sets of listening data, user data, artist data, etc. These models tailor recommendations to individual subscribers or users, for specific songs, artists, playlists, and so forth.

Spotify is a global leader in online music streaming and has pioneered the use of ML and DL for hyper-personalized song and playlist recommendations. The service streams about a billion songs every day and uses that data to optimize its service. It is important to recognize how Spotify’s AI relies on human curation: when users interact with songs (e.g., download, add to playlist, shares, “recent” versus past listens, etc.), Spotify takes careful note of these interactions to tweak algorithms for both music discovery and “hit” prediction. The company is placing big bets on AI to drive its strategy moving ahead as well, having just acquired French startup Niland, an AI company that it bought to continue optimizing music search and recommendations, while focusing on innovative products for *both fans and artists*.

Tractica forecasts that the annual revenue for music recommendations in consumer markets will increase from \$10.15 million worldwide in 2016 to \$49.47 million in 2025.

Table 2.70 Music Recommendations in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	10.15	11.09	12.50	14.59	17.66	21.96	27.64	34.52	42.05	49.47	19.2%

(Source: Tractica)

2.9.12

MACHINE/VEHICULAR OBJECT DETECTION/IDENTIFICATION/AVOIDANCE

The ability for consumer electronics, vehicles, or any other appliance to “see” has been locked in the realm of science fiction until very recently. From driving a car to vacuuming the floor, consumers relied on themselves or other humans to guide machines using their eyesight alone.

With advances in ML and CV, which are becoming DL enabled, the ability to more accurately and precisely detect and identify specific features in physical spaces automates tasks like navigation, obstacle identification, and avoidance. As these techniques continue to grow more reliable, and in some cases, more “edge-based” (i.e., processing occurs locally on the device), more and more consumer devices will be equipped with object detection, avoidance, and navigation.

Tractica’s research finds **cleaning robots** are an early adopter in this space, with a variety of autonomous vacuums, pool cleaners, floor washers, and other similar products now using AI-enabled LIDAR technologies to self-direct, localize, and even charge themselves. One example is Dyson’s Eye 360 robotic vacuum, which plans its routes according to floor type and charging needs. Much of the processing for this device happens locally, with an

architecture not only to ensure more reliable functionality and less latency, but one built with privacy in mind as well. After all, autonomous robots may collect very personal imagery while traveling around the home. Reference Tractica's [Robotics Market Forecasts](#) report for a deeper discussion on robots.

Personal robots are another area in which object (and person) identification and navigation will be both AI-powered and define many of the use cases these devices promise. For autonomous personal robots, movement around home or in-store environments will be predicated on doing so without running into users or damaging furniture or infrastructure. Even for fixed personal robots, i.e., those that do not move around, facial recognition might be intertwined with use cases. An early example is the Clone Robot, an autonomous personal robot that uses object recognition, as well as facial and emotion recognition, to power an in-home "personal assistant." The device can assist with in-home automation and security across other devices, personal photographer or videographer, storyteller for kids, videoconferencing platforms, and autonomously navigates and maps its way through the house. Personal robots also include elderly care robots, educational or toy robots, or other household robots designed to aid in or accomplish specific tasks like interacting with users, contacting others, moving or delivering objects, mowing the lawn, and beyond.

Even beyond robots, many objects that consumers use will shift toward more **autonomous machines**. Connected cars are one obvious example, but other appliances, such as TVs, refrigerators, security systems, lighting fixtures, etc., may become equipped with CV capabilities to detect specific users or objects based on facial, image, or object recognition, or understand the difference between a dog and a child. The Natatmo security camera is able to detect people, cars, and animals, and send alerts accordingly. Baidu's DuLight is an early release of a device designed for blind or visually impaired users that uses CV and image recognition to identify what is in front of the wearer and describe it to them in real-time. Wearables or mobile devices that support AR could also employ this technology, as AR relies on sophisticated recognition of landscapes, including people, in order to accurately overlay holograms.

Tractica forecasts that the annual revenue for machine/vehicle object detection/identification/avoidance in consumer will increase from \$2.18 million worldwide in 2017 to \$75.98 million in 2025.

Table 2.71 Machine/Vehicle Object Detection/Identification/Avoidance in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	2.18	5.44	10.24	17.04	26.24	37.77	50.89	64.16	75.98	N/A

(Source: Tractica)

2.9.13

PERSONALIZED HEALTH, FITNESS, AND WELLNESS IMPROVEMENT

Since the 1990s, pedometers and fitness trackers have been paving the way for people to take their fitness and health matters into their own hands, rather than relying solely on the one-size-fits-all advice of the so-called diet industry. Millions of people worldwide spend billions of dollars in an attempt to improve fitness, health, and wellness. As new technologies emerge, so too do new capabilities for personalizing health, fitness, and wellness programs to individual scenarios.

Thanks to decreasing costs of sensor technology, wearable devices have enjoyed explosive growth over the last 5 years. Innovations in this space have given everyone from kids to seniors greater insight into all manner of wellness, from step-counting and marathon training to diabetes monitoring and location tracking. Advancements in AI are beginning to enhance consumers' abilities in both wearable and non-wearable applications.

In **wearable** applications, AI can be used to analyze movements and biometrics collected from device sensors and recommend specific behaviors, exercises, decisions, etc. In both athletic and health and wellness contexts, data is often mined for highly personalized recommendations around training plans, injury prevention, dietary suggestions, water intake, sleep, etc.

PIQ is a 13-axis sports sensor designed for athletes, which uses multiple motion-capture algorithms to break down any body movement and later associate it accurately with specific sports. The sensor collects thousands of data points during each workout, then compares data to previous performances (individual and at community level) to assess specific areas to focus on in future performances. The app delivers highly personalized recommendations based on individual stats and trends from similar professional athletes.

ML, NLP, and potentially DL can also power new use cases for wearables used in lifestyle contexts. For instance, AI-enabled personal assistance in which a user wearing a smart earpiece—sometimes called a hearable—might receive text alerts or navigation assistance. AI can also power voice recognition, as in the Apple Watch; users can take advantage of it to make calls or request X with Apple's Siri technology. Some suggest the impact of AI on wearables will introduce a “wearables 2.0” era, as AI-powered software will better support new form factors, such as head-mounted displays (HMDs), VR headsets, hearables, glasses, and even shoes and clothing.

In **non-wearable** scenarios, where AI powers mobile or web-based software, AI supports big(ger) data analysis, which often includes wearable data, but draws from various other sources as well.

Sports retailer, Under Armor has partnered with IBM Watson in its app Record, which does not only track and analyze workouts, sleep, and nutrition, but mines other third-party apps and data sources to deliver personal nutrition coaching and training advice. Recommendations tap into Watson's modeling of other similar health/fitness profiles, as well as nutritional databases, psychological, and behavioral data. In the press release for the app, IBM explained the potential for such aggregated wisdom:

A 32-year-old woman who is training for a 5K race could use the app to create a personalized training and meal plan based on her size, goals, lifestyle. The app could map routes near her home/office, taking into account the weather and time of day. It can watch what she eats and offer suggestions on how to improve her diet to improve performance.

Tractica forecasts that the annual revenue for personalized health, fitness, and wellness improvement in consumer markets will increase from \$13.14 million worldwide in 2016 to \$353.5 million in 2025.

Table 2.72 Personalized Health, Fitness, and Wellness Improvement in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	13.14	22.49	36.71	57.85	88.12	129.23	180.99	240.06	300.02	353.50	44.2%

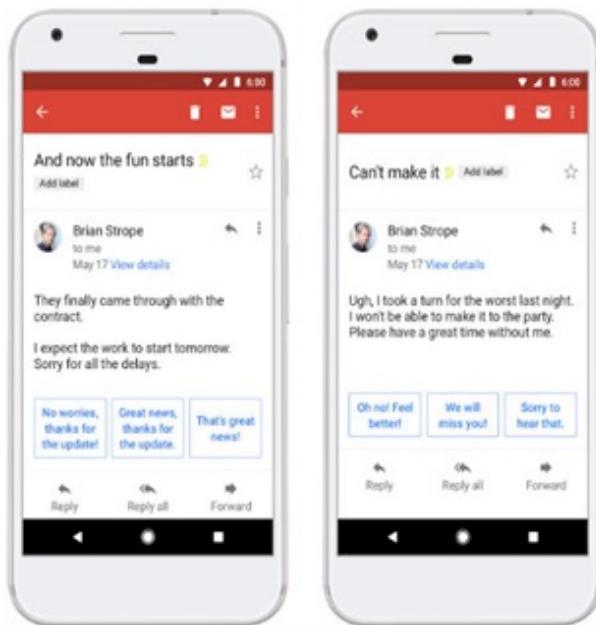
(Source: Tractica)

2.9.14 PREDICTIVE TYPING ASSISTANT

For as long as there has been information input via typing, there have been errors. While keyboard hardware has evolved over the years, the software powering input is evolving, too. Some tools like auto-correct for words, proper names, or phrases have been around for over a decade, but increasingly, AI is starting to support new ways of aiding information input and even response. Predictive typing is an example of ML-based typing assistance wherein the program is able to predict, and even populate a phrase by reading what the first word or two the user inputs. For example, a user might type “Looking” and the machine might suspect, based on the past interactions with the recipient, the remainder of the phrase to be “forward to seeing you.” Sometimes, the AI completes the sentence automatically, but user input overrides the machine if its suggestion is incorrect. Another form of predictive typing assistant was introduced a few years ago with “swype” in which input was not done by typing each letter, but by tracing a line across letters on the keyboard to spell the word.

Google is currently taking the concept of predictive typing to the next level with its new Smart Reply feature. This provides relevant suggestions to quickly respond to incoming messages with the tap of a button, rather than typing at all. For example, if a user receives an email with an invite for dinner at 7:30, the tool might serve up three options: accept, propose a new time or place, and decline.

Figure 2.10 Google's Smart Reply Offers Auto-Generated In-Context Responses



(Source: Google)

While it is unlikely typing will ever go fully extinct, the UX of text-based input remains fairly tedious. Predictive typing assistant functions help consumers communicate more quickly and accurately using keyboard devices across a range of use cases, from messaging and social media to e-commerce and entertainment. As voice interaction, predictive replies, emojis, and a range of other inputs supplement our expressions, we also expect the text input to decrease, or at least grow smarter. Look for predictive word/phrase assistance to infiltrate speech-enabled interfaces over the next few years as well.

Tractica forecasts that the annual revenue for predictive typing assistants in consumer markets will increase from \$9.54 million worldwide in 2016 to \$30.2 million in 2025.

Table 2.73 Predictive Typing Assistants in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	9.54	9.97	10.72	11.92	13.72	16.24	19.45	23.14	26.89	30.20	13.7%

(Source: Tractica)

2.9.15 PRODUCT RECOMMENDATIONS

Companies aiming to recommend products to prospects and customers have relied on ML for years. Amazon pioneered the “You Might Also Like” shopping experiences, which now incorporate everything from past purchasing history to social media connections, environmental data, advertising campaigns, and beyond. With the explosive growth of e-commerce and massive increases in data, more companies are now beginning to apply ML and DL for “right product, right person” high-precision recommendations to incentivize people to buy, using image recognition and NLP and potentially NLU (for voice-generated queries).

Pioneer of online product recommendations, Amazon, decided to open source its DL framework designed for product recommendation engines called DSTNNE in May 2016. Now, any developer can leverage the framework, while also offering clever improvements or data efficiencies Amazon had not thought of itself. Amazons press release stated, “We hope that researchers around the world can collaborate to improve it. But more importantly, we hope that it spurs innovation in many more areas.”

Some companies are using DL to better recommend their own products, while others are applying this use case to open up new business models by selling others’ products on their platforms. Houzz is a home remodeling platform that is using DL to scan photos of its proposed remodels and compare images to furniture and products in its database of some 11 million home photos. It now makes available about 6 million products across 15,000 merchants.

Tractica forecasts that the annual revenue for product recommendations in consumer markets will increase from \$27.8 million worldwide in 2016 to \$393.99 million in 2025.

Table 2.74 Product Recommendations in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	27.80	35.52	47.66	66.43	94.54	134.62	187.94	252.81	323.93	393.99	34.3%

(Source: Tractica)

2.9.16

RELATIONSHIPS AND MATCHMAKING

The internet has drastically altered how people approach dating. Since the early days of the internet, online chatrooms or communities served as forums for striking up [digital] conversation. With the dot com boom, dating websites like Match.com, e-Harmony, OKCupid, and many others not only digitized the “search” for a romantic partner, but created highly niche, even “personalized” platforms based on common interests. Today, millions of people use dating sites and mobile apps to make billions of connections each year. More than a third of marriages between 2005 and 2012 began online, according to research from the University of Chicago.

The massive amount of data required and generated in the online search for a partner renders it an ideal candidate for AI. From Q&A data about highly specific and personal preferences to demographic, location, mobile, and social media data, and beyond. Relationship and matchmaking sites have been an early adopter of ML algorithms to more efficiently wield very Big Data for very personal(ized) recommendations. As AI has evolved over the last few years, sites are experimenting with new techniques using NLP and DL.

Match.com UK recently launched a “dating bot,” a chatbot they call Lara, designed to be a virtual dating assistant that helps users hone their Match profiles (via Facebook integration and interface) to attract potential partners. The bot uses NLP to analyze 50 categories or criteria, from hobbies to astrological sign, to deliver recommendations in the Messenger app, while also tailoring new suggestions by learning from user responses over time.

Tractica forecasts that the annual revenue for relationships and matchmaking in consumer markets will increase from \$0.22 million worldwide in 2016 to \$11.72 million in 2025.

Table 2.75 Relationships and Matchmaking in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.22	0.53	1.01	1.71	2.72	4.10	5.85	7.85	9.89	11.72	55.4%

(Source: Tractica)

2.9.17

SEARCH ENGINE QUERIES

Making information easy to find and access has been the most basic mission for search engines since their inception. But the vast amounts of data and information available (about both the user searching, and their search inquiry), mean augmenting search engines is about optimizing content curation to specific user intentions and profiles. Consumers have become very impatient with search and multiple attempts with keywords has proven to propel users to other engines or to abandon search on a particular topic, reducing advertising opportunities.

ML, DL, NLP, and image recognition, in particular, now power most of the world’s main search engines. NLP technology has improved dramatically over the past few years, and combined with DL, search engines are, according to Search Engine Watch, “able to understand longer, more complex queries, with different components that modify each other and can’t operate independently.”

Most search engines, such as Google, Microsoft Bing, and Baidu, but also internet players like Facebook, have integrated natural language search into their software and are steadily eliminating keyword-based search. Engines have progressed so far as to be able to answer questions in context. Microsoft Bing’s Smart Search claims to be able to answer follow up

questions, which depend on the previous query for context. In the example given, Bing was asked, “who is the president of America” and then in a separate question, “How tall is he?” for which the engine replies with both correct responses.

In 2015, Google rolled out a DL system called RankBrain to power responses to search queries and interpret very large data sets. Additional acquisitions of Deep Mind and api.ai have enabled Google to use DL to continuously optimize its search product and results. Through ongoing refinement of some 12 billion web searches conducted per day, Google then leverages its algorithms to analyze much content across the web to deliver more accurate search results (and advertising). One such tactic is to look at customer reviews and earned media where people use their own real language to articulate sentiments about a product or experience. Andrew Howlett, founding partner of RAIN, explains that “someone might leave a review saying ‘this place has the best chips and salsa anywhere that doesn’t cost a fortune.’ Then that sentence will help now with someone searching for something like ‘I’m on a budget, where is a good restaurant with awesome chips and salsa?’”

Natural language search is also critically important to voice search. Some enterprises like Expedia have already enabled their website search with natural language capabilities.

Tractica forecasts that the annual revenue for search engine queries in consumer markets will increase from \$128.62 million worldwide in 2016 to \$683.88 million in 2025.

Table 2.76 Search Engine Queries in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	128.62	141.72	161.36	190.76	233.94	294.74	375.04	472.34	578.87	683.88	20.4%

(Source: Tractica)

2.9.18

SMART OVEN CONTROL WITH FOOD RECOGNITION

As sensors, connectivity, and networked services have begun to pervade consumer electronics and appliances, these connected devices have begun to penetrate smart home environments. From lightbulbs and thermostats to coffee-makers, ovens, TVs, and security systems, consumers are connecting their home infrastructure to web-based services to improve security, efficiency, and control.

AI is beginning to support these applications across a range of products and use cases, most of which remain in the early stages due to slow adoption, privacy concerns, and device limitations. ML, DL, NLP, CV, and other techniques will power new capabilities and, eventually, new business models in the smart home.

Ovens equipped with food recognition technology are one example of the intersection between AI and connected devices in the smart home. These ovens use image recognition to identify the content inside and adjust themselves accordingly by, for example, decreasing the heat or notifying the user. The June Oven is an example of such a device. Using built-in cameras with CV and image recognition, the oven recognizes the item when placed inside and suggests the best way(s) to cook it. AI powers intelligent recognition that learns with every meal: for meat, the oven offers different options to target raw, medium, etc.; for bagels, it can tell when they are upside down or right-side up. In addition, a built-in digital scale, precision time, and temperature sensors allow users to monitor the cooking process right from their apps.

Figure 2.11 June Oven Uses Image Recognition to Identify, Automate, and Optimize Cooking



(Source: June Life)

While novel, these sorts of AI applications are promising to manufacturers, as they could potentially help enable new business models: integrating with online recipe services, other devices, or grocery retailers, even fitness plans.

Tractica forecasts that the annual revenue for smart oven control with food recognition in consumer markets will increase from \$0.02 million worldwide in 2016 to \$6.13 million in 2025.

Table 2.77 Smart Oven Control with Food Recognition in Consumer, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
	0.02	0.14	0.34	0.65	1.11	1.78	2.68	3.76	4.96	6.13	88.0%

(Source: Tractica)

2.9.19 SOCIAL MEDIA FEED CURATION

The rise of social media and online communities wherein anyone can contribute and share content and ideas has been a revolutionary force around the world. Billions of connections sharing trillions of pieces of content about everything from politics to puppies have transformed the mode and mix of consumers' media consumption diets. Today, Facebook's News Feed is seen by over a billion users, offering over a billion personalized prioritizations for what news we read, whose updates we see, what events we learn about, and all manner of observations about the world around us.

Keeping users on social media platforms (instead of clicking away) is the foremost objective for social media platforms. To do this, social media giants like Facebook, Twitter, and Baidu use real-time interactions and Big Data to constantly analyze, model, and predict what content will incentivize users to stay. Recommendation algorithms analyze every single individual user's interactions (e.g., engagement with content, other users, scrolling, responses, click-thrus, etc.) in order to serve up new content designed to predict what the user will like best. Curation algorithms are the most strategic lever for social media platforms and content publishers. This algorithm is designed to deliver precisely the optimal content for each individual user at precisely the right moment. This creates a feedback mechanism that can then be used to optimize the user's interaction with the site, increasing the user's engagement, while simultaneously maximizing the effectiveness of advertising on the site.

To support social media feed curation, Facebook uses DL, natural language processing, and image/object recognition. Some examples include:

- Analyzing troves of posts and comments to understand semantic language and text
- Translating content across more than 100 languages to erode language barriers in sharing and connections
- Recognizing faces (e.g., friends, family, celebrities, influencers, etc.)
- Identifying objects to supplement context (e.g., typical objects and images associated with wedding, baby shower, funeral, party, or other social gatherings)
- Mining past photos (uploads and interactions)

While enticing and well-targeted social media feed curation is what makes social media so sticky, it raises a number of unprecedented cultural and ethical questions that remain unanswered. In some cases, the inherent “categorizing” effect of curation means that content delivered is prioritized over content not delivered. While this seems obvious enough, mainly raise questions around who or what is designing, controlling, or paying for this prioritization. Do such strategies advertently or inadvertently manipulate users into reading, thinking, feeling, sharing, or purchasing? Researchers point to numerous examples of how content can be used to alter a user’s emotional state, which in turn, algorithms could exploit to drive engagement or purchases with a particular product or cause.

Then there are the cases in which content itself presents major conflicts of interest in advertising business models. A research team from Oxford University cites the rapid emergence of “fake news” during the 2016 U.S. presidential election. During the campaign, companies aiming to increase advertising revenue found that targeted fake news stories were effective as clickbait; through trial and error, they found that for certain users, these faux articles and headlines generated more clicks than verifiable stories. This creates a cycle in which users read (fake) stories, advertisers get more clicks, revenue goes up for advertisers and social media platforms; and as a result, fake news stories and their false narratives are proliferated. This represents a dangerous impact for users, not just in believing falsehoods, but in undermining faith in critical democratic institutions and trust itself.

As DL, in particular, continues to drive social media feed curation, the issue of opacity and poor algorithmic transparency may exacerbate the problem. The question of accountability remains a complicated and poorly understood challenge for technology platforms, advertisers, industry, government, and regulators.

Tractica forecasts that the annual revenue for social media feed curation in consumer markets will increase from \$51.81 million worldwide in 2016 to \$525.06 million in 2025.

Table 2.78 Social Media Feed Curation in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	51.81	62.04	77.96	102.39	138.82	190.62	259.43	343.07	434.75	525.06	29.3%

(Source: Tractica)

2.9.20

STATIC IMAGE RECOGNITION, CLASSIFICATION, AND TAGGING

The primary purpose for consumer-oriented applications for image recognition and classification is to help users automatically segment, tag, and store images for better data mining and retrieval, either stored on-device or in the cloud. Image recognition, often alongside NLP, ML, or DL helps power a range of capabilities for consumer image applications, such as find and search, auto-organize, recommend, social or keyword tag suggestions, and design options.

Photo upload sites like Google Photos, Apple Photos, and Flickr all use AI image recognition and tagging techniques to automate photo classification and tagging. Facebook and Snapchat recognize individual faces. Many of these platforms also power search capabilities as well, wherein a user can search their, others', or public photos by keyword, such as "cats" or "Christmas".

Other novel consumer uses for image tagging include reading photo descriptions aloud for blind people, an approach that was pioneered by Facebook. Google is now able to automatically produce captions for images, predictive search rendering by device type, and even use algorithms to detect spam and prevent redirects.

Tractica forecasts that the annual revenue for static image recognition, classification, and tagging in consumer markets will increase from \$42.09 million worldwide in 2016 to \$901.26 million in 2025.

Table 2.79 Static Image Recognition, Classification, and Tagging in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	42.09	59.82	87.98	131.77	197.57	291.61	416.87	569.35	736.56	901.26	40.6%

(Source: Tractica)

2.9.21

TEXT-BASED AUTOMATED BOTS

Just as the web created fertile ground for websites and mobile devices laid the foundation for mobile apps, consumer-facing software platforms are becoming the foundation for text-based automated bots. Tractica defines these as bots that are designed specifically for messenger or communications platforms, such as Facebook Messenger, Telegram, SMS, Twitter, Viber, WhatsApp, Skype, Slack, etc. (not WeChat/Weixin, which limits chatbots to customer service chatbots for brands).

Text-based automated bots are generally defined more from a consumer perspective and less by specific brands. Often, such bots are more brand-agnostic, such as using WhatsApp to search for a local plumber or book a table at a nearby restaurant. Unlike brand-generated chatbots, these bots are multi-purpose and vary in ability to drive conversions or purchases. Many text-based automated bots are more utilitarian in nature, for example:

- **Do-Not-Pay:** A "lawyer chatbot" that has helped users appeal 250,000 parking tickets and successfully contest more than 160,000 by using AI to assess whether an appeal is possible given the conditions in which the ticket was administered.
- **Better:** A service using both bots and human agents to detect errors and help manage and reduce costs of out-of-network medical bills.

- **Operator:** A “request network” that uses bots to route requests to a network of human “concierges” who can execute any shopping-related request (e.g., order concert tickets, get gift ideas, or furniture recommendations).

Entertainment-focused chatbots on popular platforms like Facebook Messenger, Skype, Viber, Twitter, Telegram, and Kik have gotten off to a rocky start, but Tractica believes entertainment chatbots will eventually find traction with consumers.

Chatbot community/website BotList is a resource for cross-platform chatbot search and discovery. According to VentureBeat, the five most popular chatbots on BotList the week of June 19, 2017 were Beam, a bot that allows Discord to automatically sync Beam subscribers to a role (essentially massively multiplayer online (MMO) game communications); Magic 8, a magic 8 ball Q&A simulator; Nonstop Chuck Norris, which enables users to chat with Nonstop Chuck, create Chuck Norris memes, etc.; AeroBot, another Discord bot that can help expand the functionality of game servers; and Diply, a website to “find the funniest, craftiest, nerdiest, most inspiring content the internet has to offer.”

Tractica forecasts that the annual revenue for text-based automated bots in consumer markets will increase from \$1.49 million worldwide in 2016 to \$505.26 million in 2025.

Table 2.80 Text-Based Automated Bots in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.49	11.52	27.70	53.12	91.56	146.69	220.28	309.97	408.37	505.26	91.0%

(Source: Tractica)

2.9.22 TRAVEL CONCIERGE AND BOOKING SERVICE

For many years, consumers relied on travel agencies and booking companies to assist in their travel searches, recommendations, and reservations. The rise of the web, aggregation services like Travelocity or Kayak, and social media have all but entirely disrupted traditional travel agencies. Instead of relying on a local broker, finding and booking travel is as simple as using a search engine or e-commerce site. The process is still rife with inefficiency; the average travel planner visits some 38 pages before making a booking, according to Expedia.

AI and ML are taking the digitization of travel booking to the next level. Many companies working in this space are developing bots to streamline every part of travel—not just accommodations or activities at the destination, but seat selections for flight, legroom, layovers, post-booking changes, loyalty program points, and countless other parameters. Companies are beginning to use the vast amounts of travel data (e.g., seasonal trends, reservation data, pricing data, ratings and review data, social media data, and demographic data, among others) to mine for patterns and correlations across this data in order to serve up more accurate and better-timed recommendations. Images, recommendations, travel reviews, and even page layouts are personalized depending on customer data and inquiry. These applications are almost always designed to drive speedier and higher-value conversion.

The ex-founder of travel giant Kayak recently developed Lola, an AI-enabled travel agent designed not to replace, but to enhance human travel agents. The app combines chatbot functionality with DL and a team of (human) travel specialists for a “concierge-like, tech-first online travel booking experience.” The idea, says founder Paul English, is “to create super-human travel consultants who are AI-powered and can handle more trips per hour than

regular travel agent can.” The longer-term vision for the company involves selling Lola’s front-end technology to other travel agencies, travel management companies, and global distribution systems.

IBM Watson has emerged as a leading partner for travel apps and is powering AI-based travel agents for travel booking for a variety of AI-enabled travel apps including Wayblazer (outlined in Section 2.7.16), Hilton’s robotic guest assistant Connie, and Baarb app. Baarb is an app that develops individual profiles for each travel shopper, including psychographic, behavioral, and social data from across the web. Online travel review site TripAdvisor is currently using a software called Flyr, which mines pricing data as users search, then allows them to “lock-in” pricing between two and seven days prior to booking. Kayak, Hipmunk, Expedia, and Skyscanner have also recently rolled out chatbots designed to simulate dialog with a human travel agent experience and help customers book more rapidly. Other companies developing similar virtual travel agent tools include Tripfinity, HelloGbye, John Paul, Boxever, and Pana.

Tractica forecasts that the annual revenue for travel concierge and booking services in consumer markets will increase from \$2.47 million worldwide in 2017 to \$124.3 million in 2025.

Table 2.81 Travel Concierge and Booking Services in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	2.47	6.46	12.72	22.21	35.81	53.97	76.10	100.39	124.30	N/A

(Source: Tractica)

2.9.23 VOICE/SPEECH RECOGNITION

Until recently, voice and speech recognition were hardly a viable mode of interaction with computers, not to mention meaningful dialog to which core product or service functionality would be ascribed. In the early days of speech recognition, systems struggled to simply understand “yes” or “no”. Systems then progressed to learn digits, and then to hundreds of words. That took 15 years to happen. Within the last few years, NLU and DL helped boost basic interactive voice recognition (IVR) to become a truly reliable mode of user interaction: conversational user interfaces.

Today, voice control represents a rapidly growing trend, as it vies to become the primary user interface in connected consumer environments (e.g., smart home, connected car, mobile health, etc.) In 2017, some 44% of U.S. broadband households are using voice-controls on internet-connected devices, according to Parks and Associates.

With the advent of voice and speech recognition, AI offers consumers new capabilities in new (hands-free) environments, and potentially new market share (e.g., elderly, kids, blind, disabled, etc.) While speech recognition enables computers to understand what someone says, NLP added to speech recognition enables computers to understand what someone means. Development in DL in both of these areas are pushing models to learn highly nuanced features of speech, such as dialect, slang, native versus non-native, tone, and emotion.

Voice recognition opens up a host of new capabilities that were previously only possible through text or touch input:

- Dialog with devices, machines, or environments
- Search queries
- Identity authentication (security, purchase, medical, etc.)
- Commands to instigate or select services, or other users to act
- Commands to control devices or state changes

Tractica's [Voice and Speech Recognition](#) report offers a deeper analysis of use cases.

Software-powered VDAs, also known as **voice-based personal assistants** in the consumer space include Amazon's Alexa, Apple's Siri, Microsoft Cortana, Google Assistant, and Baidu Doer. These are primary examples of the sophisticated intersection of speech recognition and NLP. Software development kits (SDKs) and APIs for both Alexa and Siri are enabling third-party developers, and in the case of Amazon, third-party devices, to integrate into a wide range of consumer-focused use cases, essentially bringing the conversational user interface into play. This introduces a new kind of "stickiness" to consumer engagement; Amazon Echo owners spend, on average, 10% more dollars with Amazon, and more time engaging, according to a recent Accenture study.

Voice and speech recognition is also rapidly beginning to pervade **connected devices**, particularly in the smart home, such as speakers, TVs, thermostats, locks, garage door openers, and even mirrors. Startup Duo is developing a 27-inch mirror with a full HD touchscreen and a voice-interactive AI named Albert. Albert acts as a sort of personal butler, and through its own app store, will integrate with various products and services.

Robots are another segment in which voice recognition sees growing adoption. **Cleaning robots**, such as LG's Roboking robotic vacuum cleaners, use voice recognition for simple commands. **Social robots** and **kids' robots**, such as Buddy, Rokid, Kuri, and Jibo, use voice recognition to converse with users: responding to requests, reading stories to kids, delivering recipes, etc. Finally, **elderly care robots**, such as Sota or Pillo, use voice recognition to enable simple interaction with seniors or disabled persons for actions like medication reminders, video or audio conferencing, or other tasks.

Tractica believes the most prevalent use cases for the conversational user interface are within the connected home, within the connected car, within the smart city/smart building applications, chatbots for mobile e-commerce and chatbots for customer service and brand interaction, and voice search.

Tractica forecasts that the annual revenue for voice/speech recognition in consumer markets will increase from \$18.49 million worldwide in 2016 to \$871.97 million in 2025.

Table 2.82 Voice/Speech Recognition in Consumer, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
(\$ Millions)	18.49	41.50	76.56	128.83	204.00	306.47	436.03	584.59	736.09	871.97	53.4%

(Source: Tractica)

2.10 DEFENSE

2.10.1 AGENT-BASED SIMULATIONS FOR DECISION-MAKING

When it comes to strategy, warfare has grown exponentially complex over the last 100 years with new technologies, scientific breakthroughs, new capabilities, greater urbanization, the internet, etc. The importance of taking into account massive amounts of data, contexts, and evolving geopolitical forces is a task simply beyond the scope of human bandwidth. In 2009, officials at the Defense Advanced Research Projects Agency (DARPA) discussed the opportunity for DL, as related to image recognition, multi-data set analysis, and decision-making:

Full exploitation of information is a major challenge... Human observation and analysis of [intelligence, surveillance and reconnaissance] assets are essential, but the training of humans is both expensive and time-consuming. Human performance also varies due to individuals' capabilities and training, fatigue, boredom, and human attentional capacity.

As a result, DL is being applied to simulate tactical moves and refine military strategy in real time. Reinforcement learning helps make agents smarter and each agent plays out different strategies and the best strategy could be applied. The U.S. Department of Homeland Security's Synthetic Environment for Analysis and Simulations (SEAS) project is using DL to predict and evaluate future events and courses of action.

The DARPA Visual Media Reasoning (VMR) system aids intelligence analysts in searching, filtering, and exploring visual media through the use of advanced CV and reasoning techniques. "The goal of DARPA's VMR program is to extract mission-relevant information, such as the who, what, where and when, from visual media captured from our adversaries and to turn unstructured, ad hoc photos and video into true visual intelligence," Dr. Jeff Hansberger said.

Agent-based simulation is an emerging, but profound use case with applications in government, business, and beyond. As decision-making shifts to software agents, away from humans, so too might our reliance on such agents. The biggest risk of such reliance is the extent to which such software is vulnerable to hacking or other malicious acts.

Tractica forecasts that the annual revenue for agent-based simulations for decision-making in defense will increase from \$0.02 million worldwide in 2016 to \$586.34 million in 2025.

Table 2.83 Agent-Based Simulations for Decision-Making in Defense, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.02	8.99	23.45	46.26	81.34	133.62	208.32	309.38	437.11	586.34	216.3%

(Source: Tractica)

2.10.2 LOCALIZATION AND MAPPING (AIRCRAFT AND DRONES)

Military and defense programs have been investing in the development of aviation technologies for decades. Many of these investments have been in reliable autonomous aircraft (e.g., fighter jets), as well as drones. The main push behind these efforts has been to lighten or even eliminate pilots' workloads, with the idea being that more autonomous aircraft free up pilots with brainpower to focus elsewhere, thereby offering a cognitive

advantage in warfare.

Localization and mapping concerns the need and computational ability to simultaneously construct maps of the immediate environment, while updating both the agent's position on that map and movement therein. In the context of defense, localization and mapping is a core technique for the autonomous movement of airplanes, drones, or any other UAV.

When it comes to military aircraft, onboard localization and mapping autonomy is essential in the event communication links are disrupted due to cyber-vulnerabilities. Any machine that has to remain in constant communications with an operator or centralized system is simply more hackable than one that is able to perform basic functions regardless of those communications. Onboard localization and mapping (and operations in general) are stealthier, making it much more difficult for adversaries to detect. AI-driven autonomy also allows for more reliable landing on aircraft carriers, as recently demonstrated by Boeing's Advanced F/A-XX Advanced Navy Strike Fighter and the Navy's own X-47B.

Tractica forecasts that the annual revenue for localization and mapping in defense will increase from \$14.02 million worldwide in 2016 to \$155.28 million in 2025.

Table 2.84 Localization and Mapping in Defense, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	14.02	17.15	21.88	28.88	39.02	53.22	72.28	96.40	124.77	155.28	30.6%

(Source: Tractica)

2.10.3 MACHINE/VEHICULAR OBJECT DETECTION/IDENTIFICATION/AVOIDANCE (DEFENSE AIRCRAFT AND DRONES)

Although the military has lots of data, there is a constant need for better, faster intelligence. Consider, for instance, that in 2011, during the height of the Iraq and Afghanistan Wars, the U.S. Air Force was processing 1,500 hours of full-motion video and 1,500 still images taken from aerial drones every day. The ability to [more] rapidly capture, analyze, and predict based on data, especially without relying on cloud processing, is the key enabler for autonomous decision-making.

Militaries are using DL to better enhance target recognition, search-and-rescue missions, and optimize delivery and support during crises. Like in other sectors, models are trained on large amounts of data to detect specific objects with high precision. Such capabilities can be applied in military settings involving object detection and classification for avoidance and navigation in the case of aircraft, drones, robots, tanks, ships, or other vehicles; search-and-rescue to provide humanitarian support; facial recognition; learning from signals or sensor data (e.g., radar, GPS, sound); and more.

A key enabler of these use cases in military contexts is high-performance embedded computing (HPEC) platforms, in which neural networks can run at the chip level instead of in the cloud. Field programmable gate arrays (FPGAs), power-efficient GPUs, and advanced single instruction, multi-data (SIMD) processing units help surpass processing limitations required for real-time compute in data-intensive, connectivity-constrained, and mission-critical contexts. This sort of compute allows, for example, drones to process object and image recognition on board, instead of sending it back to human data analysts, who could be halfway around the world. It would also support critical data processing in disaster zones in order to circumvent malicious threats and compartmentalize sensitive intelligence.

Tractica forecasts that the annual revenue for machine/vehicle object detection/identification/avoidance in defense will increase from \$20.61 million worldwide in 2016 to \$100.41 million in 2025.

Table 2.85 Machine/Vehicle Object Detection/Identification/Avoidance in Defense World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	20.61	22.22	24.81	28.79	34.65	42.88	53.87	67.65	83.60	100.41	19.2%

(Source: Tractica)

2.10.4 PREDICTIVE MAINTENANCE (DEFENSE AIRCRAFT, DRONES, SATELLITES)

Most military programs have stringent requirements around equipment, machinery, vehicle, and systems management, including checks, inspections, services, and technical maintenance performed before, during, and after any movement or event.

As ML and other AI technologies are being applied to help assess and predict maintenance needs, the same is true in the military and defense sector. This is essential not only to ensure that equipment is ready and functional in times of need, but to prevent potential injuries. Techniques like sequence analysis can be used to understand failure patterns and follow-on failures, while ML and DL can be used to perform predictive models or recurrent event models.

Spark Cognition is a company that provides AI-based security and maintenance services to the military and various energy sectors. It uses ML techniques that develop pattern recognition models to monitor mechanical systems within each specific aircraft, and predict failure. The cognitive nature of these algorithms allows insights to adapt to the unique characteristics of that particular plane and develop symptom-based early warnings of impending failures. It also integrates with other systems, such as diagnostic databases, maintenance records, and personnel records, to help classify fault codes, recommend the right personnel, and schedule maintenance in an optimal manner. The system helps prioritize resources, expertise, and scheduling based on critical needs.

Tractica forecasts that the annual revenue for predictive maintenance in defense will increase from \$2.33 million worldwide in 2016 to \$104.76 million in 2025.

Table 2.86 Predictive Maintenance in Defense, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.33	4.65	8.10	13.16	20.46	30.69	44.44	61.89	82.48	104.76	52.6%

(Source: Tractica)

2.10.5 PREVENTION AGAINST CYBERSECURITY THREATS

Cybersecurity represents one of the greatest threats to any society, as government agencies, corporations, and individuals have increasingly become victims of cyberattacks. All computer databases are, to some extent, vulnerable to being hacked. Today's devices, machines, and vehicles (aerial and otherwise) have more control units, computing power, lines of code, and wireless connections with the outside world than ever before. This renders

them more “intelligent” in connectivity, but more vulnerable to hackers.

In these scenarios, ML and DL are used to aid in learning from threats and predicting optimized protection for all types of military activities, assets, and intelligence. Indeed, many techniques developed in defense and military programs may now be applied to business problems and processes. Specifically, companies are using ML and DL, and MR to review massive amounts of data (billions of log files a day, for instance) to detect suspicious behavior. While algorithms have long been used to identify threat types and profiles, AI development is now targeting how to respond to cyberattacks on networks, working to quickly block suspicious communications and analyze malicious behavior and software—tasks still often allocated to humans. When under attack, the system will be able to identify the entry point and stop the attack, as well as patch the vulnerability.

One of the largest sources of funding for AI research came from the DARPA, which is agency of the U.S. Department of Defense responsible for the development of new technologies for use by the military. Military defense contracting firm Lockheed Martin recently invested some \$20 million in startups, including AI-powered cybersecurity company Cybereason, whose ML software detects network attacks as they happen. In the same week, Boeing, one of Lockheed’s largest competitors partnered with Verizon to increase funding (\$32 million) for Spark Cognition. Spark Cognition’s Deep Armor product helps defense and enterprises protect networks from malware attack. It uses ML, NLP, and AI algorithms to analyze files, detect obfuscation, learn from and predict modern attack vectors, and provide signature-free security.

Tractica forecasts that the annual revenue for prevention against cybersecurity threats in defense will increase from \$49.28 million worldwide in 2016 to \$1.05 billion in 2025.

Table 2.87 Prevention Against Cybersecurity Threats in Defense, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	49.28	66.30	92.61	133.01	194.08	284.09	411.81	583.84	800.73	1,053.79	40.5%

(Source: Tractica)

2.10.6 SATELLITE IMAGERY FOR GEO-ANALYTICS

Satellite imagery has long been a closed domain with high-resolution image databases only available to a select few companies and organizations, such as weather centers, government agencies, the military, and oil & gas companies. Being able to track changes on the ground from space has been vital for these industries, but required human analysis for years. Rapid increases in the availability and improvement in the level of detail of satellite imagery, and advancements in AI, CV, and DL have created new ways of identifying features, tracking changes, and extracting value from satellite imagery.

Satellite imaging companies are in the process of launching dozens of new satellites in the next year or so, which will be able to provide a refresh rate of 24 hours for the entire planet. Planet, a startup based in Silicon Valley, recently deployed 88 satellites in a single launch, with plans to have 143 more in orbit soon.

Apart from providing a way for humans to track the planet on a daily basis, this also means that image processing will have to be automated, in order to take advantage of this quick refresh rate and trove of imagery data. Collecting information through aerial imaging may be cheaper than a full networked sensor and connectivity implementation, for example. DL is

particularly helpful, given it requires low or no feature engineering. That said, some basic challenges remain when it comes to weather, viewpoint, lighting, and atmospheric unpredictability. Satellite images are being mined for real estate development, conservation efforts estimating deforestation, and forecasting growth by analyzing construction sites, and a host of other applications outlined throughout this report. These are not just new applications, but new business models that provide country-wide, or object-specific analysis of satellite imagery to vertical markets.

DigitalGlobe is a provider of high-resolution Earth imagery and analytics that processes 4 million square kilometers of satellite imagery every day. The company has been using DL, CV, and ML to more efficiently and accurately identify imagery, objects, and activities. Objects can be fixed, such as infrastructure, buildings, and bridges, or moveable, such as helicopters and airplanes. Activities may be events like wildfires or flooding. In certain cases, such as a recent earthquake in Nepal, the company has had tens of thousands of volunteers pitch in to crowdsource damage assessment over a million tiles of imagery. The company launched an open data initiative called SpaceNet, which provides commercial satellite data, labeled and hosted via the Amazon Web Services (AWS) cloud platform, similar to ImageNet, which is an open database of images for training AI models.

Tractica forecasts that the annual revenue for satellite imagery for geo-analytics in defense will increase from \$1.95 million worldwide in 2016 to \$12.38 million in 2025.

Table 2.88 Satellite Imagery for Geo-Analytics in Defense, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.95	2.17	2.49	2.95	3.61	4.56	5.88	7.63	9.83	12.38	22.8%

(Source: Tractica)

2.10.7 SENSOR DATA FUSION IN MACHINERY (DEFENSE AIRCRAFT, DRONES, SATELLITES)

Sensor data fusion is the technique used to aggregate, or “fuse together” multiple sensor data feeds and other data feeds in order to ascertain a more complete or multi-dimensioned picture of operations. The resulting multi-dimensional data offers less uncertainty than if the data feeds were viewed individually. Sensor data fusion has been deployed in military and defense settings for years, and is increasingly using DL to more accurately detect, classify, model, and “learn” from environmental context and impacts.

In a military setting, sensor data fusion in the traditional sense would fall short if the radar were being jammed in a certain direction, while with an AI closed-loop technique, the radar can adjust the antenna in the jammer’s direction to nullify the effect. In military or marine applications, sometimes a GPS is not very reliable, which makes sensor fusion a useful technique that uses geographic information system (GIS) data to determine a vehicle’s location and predict its future location. Like in aerospace or energy, these applications are often mission-critical and precision, speed, and reliability are paramount to adoption.

Tractica forecasts that the annual revenue for sensor data fusion in machinery in defense will increase from \$13.38 million worldwide in 2016 to \$160.67 million in 2025.

Table 2.89 Sensor Data Fusion in Machinery in Defense, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	13.38	16.50	21.22	28.25	38.50	53.00	72.65	97.82	127.84	160.67	31.8%

(Source: Tractica)

2.10.8 SWARMING DRONES

In the defense sector, one of the core elements of strategy is to create overmatch, in which perceptible, overwhelming power deters adversaries—akin to bringing a gun to a knife fight. As such, swarming drones, including swarms of larger UAVs and are seen as advantageous in that they can systematically operate in groups and carry out specific tasks without risking human life, potentially with lower risk of failure or detection. Drones are not pre-programmed, but rather act in unison, sharing a “distributed brain” and adapting to each other as a collective organism.

In defense contexts, such conditions might include surveillance, investigation, tracking, surrounding, or attacking targets, search and rescue, or targeted assassinations. Similar technology is also being explored in autonomous boats.

Currently, DARPA has partnered with Raytheon to develop drone swarms for mission execution in challenging conditions. The U.S. military recently deployed a test in California in which 103 drones were dropped from 3 F/A-18 Super Hornet fighter jets. Each drone was just 30 centimeters (cm) in length. Figure 2.12 below shows the swarm conducting various formation missions.

Figure 2.12 DARPA's Swarm of Drones Simulates Group Formations over California


(Source: Office of the Secretary of Defense Public Affairs)

See Section 2.3.5 for a broader overview of swarming drone applications.

Tractica forecasts that the annual revenue for swarming drones in defense will increase from \$0.09 million worldwide in 2016 to \$34.42 million in 2025.

Table 2.90 Swarming Drones in Defense, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.09	1.00	2.32	4.20	6.83	10.42	15.12	20.91	27.53	34.42	93.7%

(Source: Tractica)

2.10.9

VEHICLE NETWORK AND DATA SECURITY (DEFENSE AIRCRAFT, DRONES, SATELLITES)

The defense industry has been grappling with the nightmarish threat of cyber-hacking or terrorism of its machinery (land, air, and water) since such systems came online. Even today, many systems within planes, cars, boats, etc. are separated to avoid penetration scenarios, where malicious actors enter through one system and attack another. There are two broad areas of vulnerability: network security, including command and control systems, databases, and communications (which all rely on network security); and platform security, including operational systems, combat systems, and engineering plants. Then there remains the constant internal threat, in the event an employee knowingly or unknowingly uploads malware into a critical system. There are also threats along the ecosystem: ground controls, mobile devices, third-party vendors, etc. As manufacturers and operators gain increasing visibility into fleets of machines, sensors, data, and networks simultaneously open up new vulnerabilities and new security methods. For example, cybersecurity experts at Airbus cite the threat of drones sending radio signals to confuse an aircraft's flight or landing.

AI can be applied in an IoT security context, in which various techniques, such as sensor data fusion, DL, CV, and ML can be used to enhance machine and device security by monitoring sensor and environmental data, analyzing systems and anomalous events, and acting accordingly. AI could pull in data from vehicles in transport, detect a new threat, and automatically issue the appropriate updates to every other vehicle's software for real-time defense intelligence. The AI could also update maps of where threats were and automatically reroute both manned and unmanned vehicles around them.

Raytheon is developing a project aimed to help aviators counter potential cyberattacks that could arise mid-flight. The software is designed to detect anomalies in MIL-STD 1553 networks—standard for most military and commercial aircrafts. When the system detects anomalies, it analyzes them for signatures and profiles of cyberattack. From there, the system involves operators to dialog in order to gain deeper understanding for the level of threat and what the system needs to do to assist. The project remains in development as Raytheon works to optimize interface (and trust) between pilot and system.

Tractica forecasts that the annual revenue for vehicle network and data security in defense will increase from \$16.14 million worldwide in 2016 to \$198.35 million in 2025.

2.11

Vehicle Network and Data Security in Defense, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	16.14	20.13	26.15	35.07	48.00	66.18	90.65	121.77	158.54	198.35	32.1%

(Source: Tractica)

2.11

EDUCATION

2.11.1

PERSONALIZED TUTORING AND ADAPTIVE LEARNING

Everyone learns differently. While a shortage of teacher, technology, and government resources have precluded truly individualized curricula, new tools are shifting the narrative away from highly standardized learning and testing to become far more adaptive and personally tailored.

ML and DL, as well as NLP are now being explored as avenues to help mine vast amount of student and curriculum data to make for more dynamic learning experiences. One of the

greatest advantages to AI-driven adaptive learning systems is that they can gather and analyze large amounts of data, and in a virtuous cycle, dynamically improve models of pedagogy, domain organization, and learner models, while scaling both personalized feedback and techniques.

Alinea, a leading Danish publisher whose math content is used by the majority of students in grades 1 to 7 in Denmark, recently launched CampMat, an adaptive math learning product that tailors recommendations for students in grades 1 to 3 studying numbers and algebra. In partnership with adaptive courseware provider, Knewton, Alinea's content populates the platform, while ML powers a dynamic digital curriculum. Interactions and student data inform individualized instruction based on real-time analysis of what a student knows, how she learns, and her stated learning goals. CampMat also leverages game-based learning strategies to engage and motivate students.

Another company called iTalk2Learn system¹⁶ takes a multi-dimensional approach to helping young students learn about fractions. The system uses a ML -driven "learner model" that includes information about student's math knowledge, cognitive needs, emotional state, and incorporates dialog (feedback loops) from both student and instructor.

Tractica forecasts that the annual revenue for personalized tutoring and adaptive learning in education will increase from \$1.25 million worldwide in 2016 to \$510.44 million in 2025.

Table 2.92 Personalized Tutoring and Adaptive Learning in Education, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
											(\$ Millions)
	1.25	11.91	28.02	52.06	87.17	136.98	204.73	291.83	396.02	510.44	95.1%

(Source: Tractica)

2.11.2

AUTOMATED CLIFFSNOTES, STUDY NOTES, AND QUIZ GENERATORS

Part of the learning process involves distilling information into digestible formats. Teachers use quizzes and abbreviated notes and frameworks to help students make sense of complex subjects. This is an area ripe for AI, not only given the scope of relevant data in any given domain, but the opportunity to incorporate new information, such as research, publications and references, news, etc. AI provides a tool to wield the dynamic nature of information in a way that textbooks and fixed environments cannot. ML, pattern recognition, and NLP are used to mine textbook data and patterns of learning to deliver personalized formats, interfaces, instruction design, and other optimal content to support students' needs.

Cram101 uses AI to mine textbook content and re-configure information into digestible "smart" study guides. The system then generates multiple choice, true-false questions, flashcards, and chapter summaries to suit specific learning types. Other companies, such as Netex Learning, are taking a more platform-driven approach, wherein instructors (in education or business environments) design digital curricula in which multiple media formats, devices, and instruction modes are integrated for adaptive learning. The idea is to scale the content and ways of teaching, while simultaneously personalizing the experience for different learning types. Instructors can update a publisher's content, adding new lessons and activities, while the Netex platform supports gamification, simulations, virtual courses, self-assessments, video conferencing, automated feedback, content discussion, recommendations, and real-time analytics for each student and for class engagement.

Tractica forecasts that the annual revenue for automated CliffsNotes, study notes, and quiz generators in education will increase from \$1.4 million worldwide in 2016 to \$602.89 million in 2025.

Table 2.93 Automated Cliffs Notes, Study Notes, Quiz Generators in Education, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.40	10.65	25.53	48.96	84.98	138.62	215.25	318.89	449.87	602.89	96.2%

(Source: Tractica)

2.11.3 AUTOMATED GRADING OF TESTS

Instead of requiring teachers to read tens or hundreds of responses over and over, more and more educators are now turning to computers and software for assistance. Of course, standardized testing was the first step toward machine-based grading (now widespread), but that limited test formats to multiple choice questions.

Advancements in ML and DL, as well as NLP, will now help educators grade tests more rapidly. This increases the scope of what can be graded by a machine, even including open-ended question formats. While the richness of language and the importance of personalized feedback for qualitative responses may always benefit from human graders, AI helps support more rapid grading. Research has shown that timely, targeted feedback accelerates learning. AI's ability to aid with automated grading of tests, particularly in adaptive learning environments helps educators scale in both feedback and turnaround time.

Research conducted by Pearson and the University College London Knowledge Lab note that, increasingly, the design of model-based adaptive systems is more transparent. Educators can see into paths of learning, identify where confusion arises, and understand how systems select next steps based on student inputs.

EdX, a non-profit founded by Harvard and MIT, introduced the education market to software that used AI to read and grade essays. The tool requires humans to first grade 100 essays, then uses ML to train itself to grade essays automatically. While the technology has been met with much skepticism from educators, underscoring the importance of human assessment of reasoning, adequacy of evidence, good sense, ethical stance, clarity, etc., its introduction in 2013 was only the beginning. Now, very large online education platforms like Udacity and Coursera use advanced AI techniques for automatic and real-time assessment given what they cite as the importance of instant feedback. Netex Learning, outlined in Section 2.11.2, provides real-time grading of quizzes, self-assessments, even offering video conference with instructors for discussion on open-ended questions.

Tractica forecasts that the annual revenue for automated grading of tests in education will increase from \$0.1 million worldwide in 2016 to \$57.55 million in 2025.

Table 2.94 Automated Grading of Tests in Education, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.10	0.98	2.40	4.63	8.07	13.20	20.52	30.42	42.93	57.55	103.6%

(Source: Tractica)

2.11.4

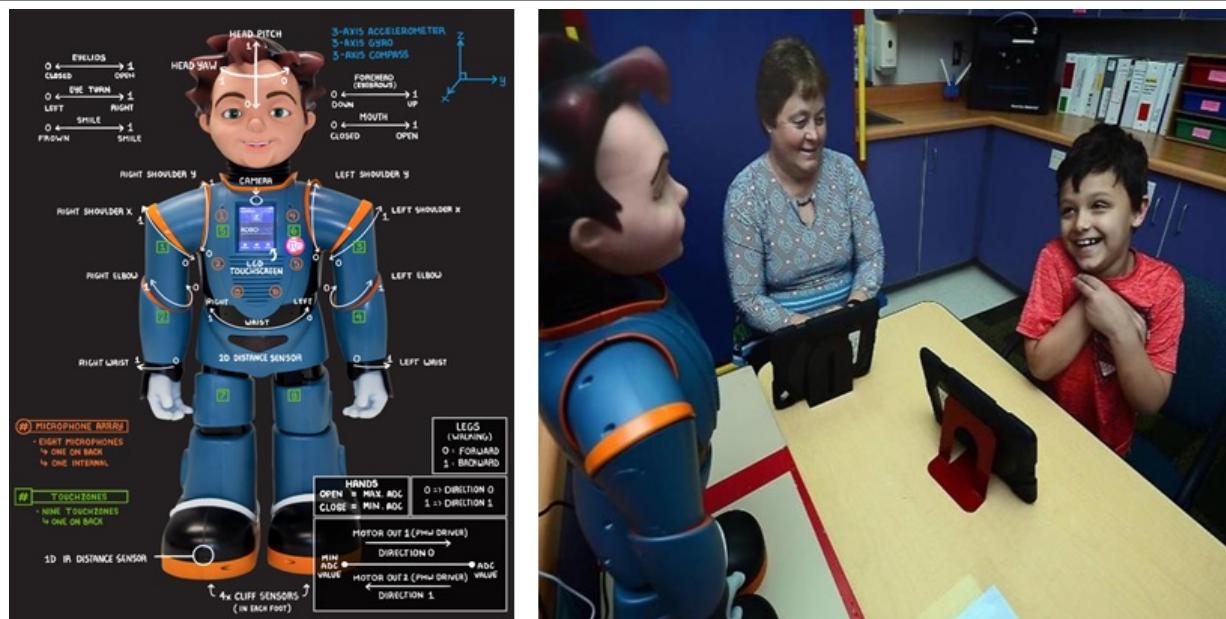
EDUCATION FOR AUTISTIC AND SPEECH DEFICIENT CHILDREN

It is not uncommon for children and people with Autism Spectrum Disorder (ASD) to display high levels of comfort with computers, toys, and robots. Such objects and their programs are logical, generally predictable, and can offer outlets for specialized interests, such as music, shapes, math, or other specialized activities. Furthermore, they never insult, pass judgement, or make fun of users, a benefit these technologies offer for people with speech deficiencies as well.

While AI can power more personalized learning and instant assessment, as outlined in Section 2.11.1 above, it also shows promise for ASD users in the area of social robots. The idea is that social robots, powered by various CV, DL, object recognition, and voice recognition techniques can help users learn speech and conversational skills, while also incorporating techniques to build emotional intelligence. Using social robots to develop skills early on helps with social integration for human interactions. In addition, robots can offer teachers and parents video feeds, analytics, the capture of an individual student's progress for specific tasks, and personalized lessons over time.

Research in this area has been underway for over a decade, and results thus far show progress. In 2013, a primary school in Britain introduced NAO, a 2-foot-tall humanoid robot able to converse fluidly with students and imitate human patterns of speech. Interactions with NAO showed numerous social breakthroughs, such as eye contact, hand-holding, hugs, and other tender gestures, according to video feeds. Over the course of NAO's testing, some children have even transitioned out of Autism programs to regular classes.

Figure 2.13 Milo, a Humanoid Robot, Helps ASD Students Identify Human Emotions



Used in over 50 American schools, Milo can walk, dance, speak, and simulate human facial expressions and social skills. Throughout interactions, symbols are displayed on his chest and tablets reinforce lessons with 4-5 second video clips.

(Source: Robots4Autism.com)

Tractica forecasts that the annual revenue for education for ASD children in education will increase from \$0.81 worldwide in 2017 to \$52.63 million in 2025.

Table 2.95 Education for Autistic and Speech Deficient Children in Education, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.81	2.10	4.15	7.30	11.99	18.70	27.77	39.23	52.63	N/A

(Source: Tractica)

2.11.5 FOREIGN LANGUAGE TUTORING

Learning a foreign language is difficult. Even for the most apt language learners, to learn a language is to embrace another culture, grammatical universe, and idiomatic schema. While the best way to learn is to travel in-country and acculturate, the reality is most people learn foreign languages in the classroom. Computer-Assisted Language Learning (CALL) has been around for decades, but AI is introducing new efficiencies and personalization to these platforms.

With advancements in NLP and DL, AI is now infusing language learning. Personalized tutoring and adaptive learning are combined with increasingly sophisticated translation and interpretation algorithms. Recent advancements in DL and neural machine translation, pioneered by the University of Montreal, Stanford, Google, Baidu, and others have improved translation significantly in the last 5 years.

Duolingo is a downloadable software app (with web portal) that offers foreign language training using ML, NLP, and DL, as well as voice recognition. Users undergo exercises in which both translation and speaking are required to complete the course. Automatic speech recognition (ASR) requires that pronunciation, verb tenses, and syntax are correct and intelligible in order to advance to the next level. Duolingo recently introduced three new chatbots (German, French, and Spanish) to its platform as a real-time way for users to have conversations without feeling embarrassed. The bots have personas and even go by the names of Officer Ada, Renee the Driver, and Chef Roberto.

Tractica forecasts that the annual revenue for foreign language tutoring in education will increase from \$2.15 worldwide in 2017 to \$140.23 million in 2025.

Table 2.96 Foreign Language Tutoring in Education, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	2.15	5.60	11.06	19.45	31.95	49.82	73.99	104.54	140.23	N/A

(Source: Tractica)

2.11.6 SPOKEN FLUENCY EVALUATION

Reading and writing a language is one set of skills, but speaking and understanding with verbal fluency is an altogether different challenge. As language learners rely on conversation, oral, video, and audio practice, NLU and voice recognition offer new capabilities for spoken fluency evaluation. Specifically, spoken fluency evaluation software uses speech recognition, ML, and DL to understand how native speakers pronounce words and phrases (including variations by dialect), and where language learners may fall short.

These algorithms learn “acceptable” levels of recognition accuracies. Research in linguistics can also be analyzed by these models, such as findings that show that articulation rate and phonation-time ratio are good predictors of fluency.

Carnegie Speech is a software platform used by government agencies for teaching and assessing spoken language to non-natives to ensure maximum accuracy with minimal training time. Its software uses a combination of speech recognition and pinpointing technology that models each learner’s speaking characteristics and delivers personalized spoken-language training curriculum. To do this, the technology compares each user’s speaking profile against a composite statistical model of native speaker’s speech.

Duolingo, the language learning app outlined in Section 2.11.5 also incorporates speech evaluation into its exercises, including instant assessment. Users undergo exercises in which both translation and speaking are required to complete the course. ASR requires pronunciation, verb tenses, and syntax to be correct and intelligible in order to advance to the next level.

Tractica forecasts that the annual revenue for spoken fluency evaluation in education will increase from \$1.29 worldwide in 2017 to \$84.14 million in 2025.

Table 2.97 Spoken Fluency Evaluation in Education, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	1.29	3.36	6.64	11.67	19.17	29.89	44.39	62.72	84.14	N/A

(Source: Tractica)

2.11.7 TEXTUAL QUESTION ANSWERING

Personal tutors and one-to-one instruction are perhaps the most ancient form of pedagogy. Parent-child, master-apprentice, and teacher-student models for learning are inherently personalized given the nature of teaching, demonstrating, questioning, and answering in shared contexts.

In the field of computer science, textual question answering is the technique of extracting a sentence or text snippet from a document or database of information that responds directly to a specific query. In education, this constitutes the ability to simulate personal tutoring, as the system should be designed to provide an accurate and satisfactory explanation to any question. NLP, with ML and/or DL support these sorts of chatbot applications. In spoken contexts, these are often called dialog agents, or bots that learn to understand the meaning of an inquiry and provide answers rich with context.

Figure 2.14 SimCoach, a 3D Virtual Agent Interacts and Assists Military Personnel with Breaking Down Barriers



In 2009, the U.S. Defense Centers of Excellence (DCoE) for Psychological Health and Traumatic Brain Injury funded development for SimCoach to help veterans, service members, and family members access information, support, and resources available in areas like healthcare, life transitions, jobs, and community.

(Source: University of California Institute for Creative Technologies)

Despite the recent success of chatbots, the technology still has a way to go before AI-based Q&A reaches human levels of awareness, language mastery, or nuance. As more computational models assign mathematical values to the meanings of words and use them to successfully read text and derive meaning, AI's ability to wield language plus learning should improve.

Tractica forecasts that the annual revenue for textual question answering in education will increase from \$0.7 million worldwide in 2016 to \$586.72 million in 2025.

Table 2.98 Textual Question Answering in Education, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.70	9.69	24.16	46.98	82.06	134.31	208.97	309.96	437.60	586.72	111.3%

(Source: Tractica)

2.12 ENERGY

2.12.1 SATELLITE IMAGERY FOR GEO-ANALYTICS

Satellite imagery has long been a closed domain with high-resolution image databases only available to a select few companies and organizations, such as weather centers, government agencies, the military, and oil & gas companies. Being able to track changes on the ground from space has been vital for these industries, but required human analysis for years. Rapid increases in the availability and improvement in the level of detail of satellite imagery, and advancements in AI, CV, and DL have created new ways of identifying features, tracking changes, and extracting value from satellite imagery.

Apart from providing a way for humans to track the planet on a daily basis, this also means that image processing will have to be automated, in order to take advantage of this quick refresh rate and trove of imagery data. Collecting information through aerial imaging may be cheaper than a full networked sensor and connectivity implementation, for example. DL is

particularly helpful given it requires low or no feature engineering. Some basic challenges remain when it comes to weather, viewpoint, lighting, and atmospheric unpredictability. Applications in the energy sector are varied. Some include helping energy and utilities providers discover areas of new energy resources, better understanding the impacts of humans on the earth's resources, and measuring metal and commodity productions, as well as oil storage tanks that are not included in public records. More generally, satellite imagery can help track a bounded area with alerts and updates provided when something changes in that specific area, or for historical changes over said area. These are not just new applications, but new business models that provide country-wide, or object-specific analysis of satellite imagery to vertical markets.

PowerScout is using DL to analyze satellite data to detect and determine homes that would be likely candidates for solar panels given their positioning and exposure to light. This helps optimize sales and marketing costs associated with targeting the right potential buyers. The company has trained two neural networks to determine: 1) whether a home already has solar panels (or not), and 2) whether nearby vegetation would hamper installation or energy generation efforts. It is also developing an e-commerce site in order to use this data to let users run feasibility and estimated value and returns on solar panel installation for their homes. Based on "solar worthiness," the service then matches potential buyers with local installers, and offers tailored financing options. In the future, PowerScout hopes to use this data to optimize community solar sales by suggesting installations wherein multiple residents could take advantage.

Tractica forecasts that the annual revenue for satellite imagery for geo-analytics in energy will increase from \$36 million worldwide in 2016 to \$160.57 million in 2025.

Table 2.99 Satellite Imagery in Geo-Analytics in Energy, Annual Revenue, 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	36.00	38.62	42.45	48.01	55.99	67.23	82.68	103.20	129.28	160.57	18.1%

(Source: Tractica)

2.12.2 WEATHER FORECASTING

Weather monitoring and analysis is growing increasingly important as renewable energy industries grow. Energy forecasting has been part of utilities' operations and planning for over a century. As in other industries like airlines, CPG, oil & gas, and agriculture, the need for forecasting supply, demand, and prices has not changed, but the extent to which weather and consumption patterns are digitized has.

AI and sensor data from hundreds of thousands of sources collected and monitored in real-time (and over many years) are transforming the level of understanding and ability to forecast conditions. In addition to weather data, such engines combine streaming data from social feeds, news reports, transportation data, and historical data on storms or other weather events. While no one can ever fully predict the future, AI techniques apply reinforcement learning on past predictions and actual outcomes. By comparing predictions with accuracies, the model is able to learn and improve simulation capabilities, as well as forecast much further into the future. As energy suppliers aim for greater precision in all supply, demand, pricing, and distribution processes, for instance, the ability to accurately forecast and pinpoint environmental conditions is key.

AI is being used for weather forecasting in energy in a number of natural energy sectors, such as wind, solar, water, etc. Solar forecasts, which integrate weather patterns and solar production estimates, help grid managers predict how much solar energy will be produced across their system on a given day and allow utilities to better allocate resources and avoid the need to ramp up reserve power plants. Wind power forecasting helps wind farms estimate expected production, particularly important in areas of high seasonal variation. Utility companies could mine and model historical data of damage to power lines or telephone poles and then couple that information with hyper-local forecasts to better plan for how many repair crews would be needed and where. The applications are endless.

JEA, the water, sewer, and electric provider for Jacksonville, Florida, now uses an automated supervisory control system to optimize its pumping and distribution systems. The Optimized System Controls of Aquifer Resources (OSCAR) controls and adjusts the water system every minute by evaluating weather, historical water consumption, and other supervisory control and data acquisition (SCADA) data and uses DL and heuristic algorithms to generate the forecast to predict sub-grid hourly consumption. Based on the forecasting, operators have switched from reactive to proactive by ensuring, for example, that pumping is assigned to the water plant closest in proximity to the demand. Energy consumption is then minimized, while water generation is maximized during on-peak periods.

Tractica forecasts that the annual revenue for weather forecasting in energy will increase from \$1.78 million worldwide in 2016 to \$303.33 million in 2025.

Table 2.100 Weather Forecasting in Energy, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
(\$ Millions)	1.78	8.11	17.38	30.85	50.16	77.37	114.77	164.46	227.57	303.33	77.0%

(Source: Tractica)

2.13 FASHION

2.13.1 FASHION TREND PREDICTION

The ability to predict preferences, behaviors, and market movements has many mission-critical applications in defense, weather, finance, news, etc., but can also be applied to retail markets. One such area with high potential is in predicting fashion trends. The global fashion industry is valued at \$3 trillion, accounting for 2% of the world's gross domestic product (GDP), and employs millions of people internationally. It also encompasses a massive ecosystem of businesses, including designers, manufacturers, distributors, marketers, advertisers, etc. The challenge in this area is efficiently matching supply and demand, for all to benefit. Currently, fashion brands and retailers work with a limited amount of data to predict what products to order and when to discount or replenish them. If they predict wrong, the result is loss of income due to mark-downs, waste, and popular items selling out.

What has historically been developed based on traditional market research methods is now being assessed through diverse data streams fed to algorithms. By analyzing large amounts of data, such as the browsing and shopping history of every single one of a fashion brand's online customers, as well as those of its competitors, AI can tell a retailer how to align product drops to match demand, and even how to display products in a store to maximize sales.

Stitch Fix is an online platform that provides highly personalized styles to women and men. It provides its personal stylists with tools and technology to help hand-select clothing and accessories that fit shoppers' preferences, lifestyles, and body shapes. The company

introduced Deep Style to help collect and learn from product, customer, service, and workflow data in order to predict and deliver “just right” fashionable styles at the individual level. Specifically, it uses photographs to quantify the style and identify unique attributes of items in its collection. It is also using the model to associate one article of clothing with related accessories and color schema, as well as using all product and customer data to inform computer-generated clothing it can use to simulate new designs.

Figure 2.15 Stitch Fix Uses Deep Learning to Analyze Styles and Design New Clothing



In the first image, an actual shirt is analyzed in conjunction with other data sets to create a recommended shirt. In the second image, the model designs multiple “like” shirts, based on different variables for different user segments.

(Source: Stitch Fix)

Tractica forecasts that the annual revenue for fashion trend prediction in fashion will increase from \$3.52 million worldwide in 2017 to \$167.98 million in 2025.

Table 2.101 Fashion Trend Prediction in Fashion, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	3.52	8.83	16.76	28.35	44.78	67.13	95.86	130.23	167.98	N/A

(Source: Tractica)

2.14 FINANCE

2.14.1 AUTOMATED CREDIT SCORING

Credit rating agencies use credit scoring to monitor, penalize, and incentivize their customers to pay up. In the United States, one's FICO score is designed to indicate current financial circumstances and historical behavior demonstrating a willingness to pay off loans. When it comes to determining one's credit score, credit agencies use payment history, length of history, debt burden, types of credit, and recent credit searches.

Credit rating agencies are now beginning to explore AI, ML, and DL to aid in credit scoring, primarily to assess creditworthiness more precisely through more nuanced evaluations of data. Instead of looking at one or a few separate variables, AI engines help consider mitigating interactions between multiple variables. For instance, even if a consumer skipped payments on 2 debts within 24 months, but paid consistently for 12 months straight, and obtained new lines of credit, that may be weighted to mitigate the risk of the past missed payments. The other potential benefit is to consider people who might not have been able to get a score in the past, via traditional logistic regression-based scoring (which looks at credit history). The problem with using AI for credit scoring is one of transparency and decision accountability, particularly given regulations and the fact that these decisions can have very

real tangible benefits or setbacks in peoples' lives. Particularly in the case of neural networks, providing a "reason code"—an explanation of why they were denied credit—is opaque at best. How do we train a system to look at interactions with many variables, product one clear reason for declining credit, and enable it to articulate that reason?

Companies experimenting in this arena include large credit agencies, such as Equifax, Experian, and FICO, lenders like Elevate, and a host of fintech vendors like IDAnalytics, and others. All are working with regulators, lawyers, and compliance officers to ensure compliance.

Reference Section 2.14.7 for an overview of AI used for loan analysis.

Tractica forecasts that the annual revenue for automated credit scoring in finance will increase from \$2.14 million worldwide in 2016 to \$72.1 million in 2025.

Table 2.102 Automated Credit Scoring in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.14	3.02	4.47	6.83	10.64	16.58	25.43	37.78	53.62	72.10	47.8%

(Source: Tractica)

2.14.2 AUTOMATED REPORT GENERATION

Financial services generate reports for internal stakeholders, as required by auditors and regulators for compliance, as part of client programs, or even as formal products. Many financial functions remain reliant on manual processes, fragmented data, and legacy systems. Slow turnaround times, excessive effort spent on data collation and validation, and inconsistent reporting of results can ultimately create a variety of negative impacts and delays. As the amount of data flowing into and across organizations grows more and more massive, the problem is not just one of content distribution, but of the time it takes to comprehensively identify and organize insights that are useful and consumable.

AI is now a tool well suited for report generation. Using NLP, ML, and DL, in some cases, companies are using AI to collate reports far more rapidly than humans. Automated report generation tools generally support the following tasks:

- **Data Sourcing:** Identifies and extracts data from relevant internal and external sources, including industry news and reports, social media listening, and competitor intelligence.
- **Data Interpretation:** Upon consolidating data in standardized formats, the solution aligns the data in templates, codes, and prepares it for analysis using ML.
- **Data Analytics:** Defines business rules and correlation/causality at scale. With predictive modeling and data enrichment, solutions can run hundreds of "what if" scenarios and perform trend analysis
- **Narrative and Semantic Commentary:** Using NLP and generation, solutions can sometimes automate variance analysis and commentary writing in a systematic and structured way.

These capabilities allow for highly customizable automated report generation. With proper data inputs, both chief financial officers (CFOs) and financial services companies can use these tools to rapidly analyze revenue, market shares, trends, competitors, geographies, etc.

Financial services consulting and technology provider Synechron is working with financial institutions to read composite data sets and convert them to natural language advice through its AI platform, Neo. The platform can collect and understand hundreds of financial data records and generate easy-to-understand written summaries of the data.

Applying Neo to business processes enables firms to reduce operational risk and increase efficiency, obtain real-time data and insights for decision-making, and automate regulatory compliance and adherence to customer service-level agreements (SLAs) and customer experience. Synechron is somewhat unique in the natural language ecosystem in that it is something of a legacy financial services vendor, an early indication of the coming permeation of AI DNA-capabilities. Other companies in this space include Automated Insights, Narrative Science, Genpact, and a variety of other vertical specialists.

Tractica forecasts that the annual revenue for automated report generation in finance will increase from \$2 million worldwide in 2016 to \$187.19 million in 2025.

Table 2.103 Automated Report Generation in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.00	4.13	7.77	13.87	23.81	39.47	62.93	95.77	137.96	187.19	65.6%

(Source: Tractica)

2.14.3 BIOMETRIC IDENTIFICATION

The definition of identity is expanding to include digital biometrics. Just as manufacturers of mobile devices and IoT-enabled machinery and infrastructure have begun to use biometrics, such as fingerprints or facial recognition, to enable identity authentication, so are financial services companies.

ML and DL power identity authentication for biometrics-based recognition in areas like facial recognition, speech recognition, fingerprint recognition, retina recognition, and potentially other biometrics as well. They may be used as part of two- or three-factor authentication, and are generally seen as more secure and non-replicable than passwords. Apple and Google pioneered this with their respective fingerprinting technologies integrated with Apple Pay and Android and Samsung Pay. MasterCard recently launched MasterCard Identity Check, often referred to as “selfie-pay,” which lets users validate purchases by taking a photo of themselves or scanning their fingerprint at the time of purchase. To prevent spoofs, such as someone holding up a photo to the camera’s lens, MasterCard requires customers blink to confirm it really is their face. A host of other companies are working in this area, such as AimBrain, Signicat, Bytes Systems, Applied Recognition, and Sensory.

Tractica forecasts that the annual revenue for biometric identification in finance will increase from \$0.15 million worldwide in 2016 to \$43.29 million in 2025.

Table 2.104 Biometric Identification in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.15	0.63	1.46	2.86	5.17	8.81	14.28	21.94	31.80	43.29	88.1%

(Source: Tractica)

2.14.4

CONVERTING PAPERWORK INTO DIGITAL ASSETS

A vast range of businesses will benefit from converting paper documents and other unstructured data, such as email, PDFs, charts and graphs, into structured data for a wide range of uses, but particularly operational efficiency. To support converting paperwork into digital assets and data, NLP is combined with ML and sometimes DL to effectively “read” text. Using machines to handle the processing that humans have already done (when logging information on paper the first time) frees up employees to manage other more complex tasks.

App Orchid is a SaaS-powered company that combines NLP, Big Data, machine intelligence, and data science in one toolbox to help companies process and analyze structured and unstructured data for business intelligence. For insurance, App Orchid intends to help companies with actuarial analysis, catastrophe risk and damage analysis, targeted risk analysis, underwriting, claims processing, and fraud control by creating “a virtual repository of intelligence” and then “harness the power of all this information with artificial intelligence and machine learning to enable advanced predictive analysis simply by typing in a question.” As of the writing of this report, there is no evidence that the company has successfully secured an insurance customer in this regard.

Tractica forecasts that the annual revenue for converting paperwork into digital assets in finance will increase from \$11.98 million worldwide in 2016 to \$782.82 million in 2025.

Table 2.105 Converting Paperwork into Digital Assets in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	11.98	28.49	53.80	92.09	148.70	229.41	338.35	474.54	628.55	782.82	59.1%

(Source: Tractica)

2.14.5

PATIENT DATA PROCESSING

Insurance companies have adjacent uses for patient data, distinct from the healthcare industry. Patient data processing in the healthcare context is outlined in Section 2.17.16. Unlike in healthcare, insurance companies’ primary objective is to use and process patient data to model risk factors, improve insurance products, prevent losses and fraud, and reduce the amount of money they pay out.

Regardless of objective, patient data processing in both healthcare and insurance amount to Big Data and have, therefore, come into the realm of AI. In the near term, insurance companies also see potential in leveraging electronic health records and patient data to reduce fraud by detecting anomalies or patterns associated with fraudulent activity. Algorithms can also help speed up claims processing, by automatically assessing the severity of a claim and predicting costs from historical data, sensors, images, etc. Despite significant regulatory, political, and commercial hurdles, using electronic health records, in addition to a range of other data sets, represents an opportunity for insurance companies to better target customers with the coverage they need.

Fitsense.io is a company focused on personalizing insurance products by using app and device data. It has built a data aggregation platform that integrates, processes, and securely stores data across various channels (e.g., wearables, biometrics, health apps, demographic data, etc.). It uses ML and NLP to model and interpret raw data into specific customer and risk profiles, then leverages that data to help insurance companies design and substantiate

new insurance products and services. In particular, the app enables insurance companies to offer their own white-labeled self-quantification, health management, and incentive programs. Other companies developing in this space include DreamQuark, Big Cloud Analytics, and Cognicore, which offers a chatbot assistant for complaints and claims resolution and then uses interactions to improve insurance products and services.

Tractica forecasts that the annual revenue for patient data processing in finance will increase from \$1.7 million worldwide in 2016 to \$1.88 billion in 2025.

Table 2.106 Patient Data Processing in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.70	23.94	62.08	125.88	229.56	391.73	632.53	966.04	1,389.42	1,876.31	117.8%

(Source: Tractica)

2.14.6 EMPLOYEE EXPENSE MANAGEMENT

Companies spend tremendous amounts of time and costs on and sometimes even outsourcing employee expense management. The emergence of mobile brought numerous expense tracking apps capable of streaming expense inputs and monitoring, but AI presents new opportunities for companies to reduce time, efforts, and costs associated with managing just about every element of employee expenses.

AppZen is focused on a unique use case, reducing travel and expense costs for companies. AppZen uses NLP to automate expense report auditing and instantly detect fraud and compliance issues. The solution works like this: identify company expense policy, how much are you allowed to spend, etc., the audit process and the kind of issues the company typically sees is fed into the program using CV to scan credit card transactions, travel bookings, paperwork, and especially receipts. Semantic analysis is run on the documents. Then the engine looks at all the text and tries to figure out what it means. For example: what drinks on the receipt are alcoholic? Which charge is an in-room movie? Data is extracted and augmented using the model, and behavior is tracked by individual. The technology automatically detects accidental as well as intentional fraud, and provides real-time compliance to Internal Revenue Service (IRS) rules, Foreign Corrupt Practices Act (FCPA) regulations, and company policies. The solution also automatically audits and assigns risk scores to every expense, protecting the company from expense misuse and regulatory non-compliance. Because it can provide proof for enforcement and prove repeat offenders, it averages reductions in travel and expenses (T&E) by 2% to 5% according to AppZen. Comcast, Hitachi, Equinix, and other large enterprises are customers. Channel partners include Oracle, Concur, and NetSuite.

Tractica forecasts that the annual revenue for employee expense management in finance will increase from \$0.07 million worldwide in 2016 to \$29.48 million in 2025.

Table 2.107 Employee Expense Management in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.07	0.39	0.96	1.92	3.49	5.97	9.70	14.92	21.64	29.48	96.0%

(Source: Tractica)

2.14.7

LOAN ANALYSIS

Assessing, selecting, and underwriting loans is a complicated and historically very manual and costly process, in which an underwriter looks at a range of public documents and assesses risk and credit using FICO scores. But in the era of Big Data, and a proliferation of fintech companies, the current model for analyzing creditworthiness is being redefined. Not only are there hundreds of other data sets that creditors now find relevant—e-commerce transactions, first payment defaults, web browsing history, location data, to name a few—ML and AI can help process and learn from more data sets. Programs learn from correlations and surface patterns that may be relevant for assessing risk. These tools open up benefits and risks for both lenders and consumers.

Digital lenders pull in data, such as SAT scores, text-based punctuation behavior, licenses obtained, and even how many of their phone contacts have last names. In one cautionary example, such companies have found correlations between late-night internet use and bad loan repayment. That said, government research has found that FICO scores can place disadvantage on younger borrowers and people from other countries, as lower income and low or no credit applicants are targeted with higher interest loans. The fundamental question facing this market is, with the wealth of information about people on the internet, how does a company go about extracting the relevant information to best complement core financial data without inadvertently discriminating against or disenfranchising certain segments of people.

Companies like Underwrite.ai, Datanomers, and Upstart are all using ML and, in some cases, DL to pull in diverse data sets, analyze for correlations associated with lending risk, and reduce risk for lenders.

Tractica forecasts that the annual revenue for loan analysis in finance will increase from \$1.54 million worldwide in 2016 to \$54.68 million in 2025.

Table 2.108 Loan Analysis in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.54	2.21	3.30	5.09	7.98	12.49	19.21	28.60	40.64	54.68	48.6%

(Source: Tractica)

2.14.8

PERSONAL FINANCIAL ADVISOR

Managing one's personal finances can be an overwhelming task. While some people prefer to do it on their own, many opt to hire a professional financial advisor to help manage accounts, offer investment advice, flag questionable activity, and guide account holders through difficult financial processes or decisions.

AI-powered bots and digital assistants are now taking on various elements of personal[ized] financial advisory. So-called “finance bots” typically analyze data across multiple accounts, and identify areas to invest, areas to save, areas of spend, and offer advice. They often include other capabilities like custom alerts, money transfer capabilities, check deposit, FAQs, and customer support services. Several companies are experimenting with VDAs as wealth management assistants.

Indian bank Kasisto is a mobile-only bank, which has built MyKai, a financial assistance bot that helps users manage their money, track expenses (by merchants, timeframe, location, amount), set budgets, pay others, and analyze spending across other channels, such as

Facebook, Slack, SMS, etc. It also integrated with Tradelt, which works with Nasdaq, E-Trade, and other brokers; this allows the bot to monitor both market dynamics and assess the user's investment portfolio to suggest strategies related to holdings and shares.

Figure 2.16 Kasisto's MyKai, a Personal Financial Advisor Chatbot



(Source: Kasisto)

Enterprise VDA vendor OpenStream has a white-labeled virtual assistant called EVA. It is notably the engine behind a ground-breaking VDA from finUNO. finUNO's fin1 is an automated wealth advisor that helps users buy and sell investment products, track the specified market activity, and manage investment transactions in real time. fin1 can send reminders of balances and make suggestions based on previous activity or news.

Tractica forecasts that the annual revenue for personal financial advisor in finance will increase from \$0.37 million worldwide in 2016 to \$317.08 million in 2025.

Table 2.109 Personal Financial Advisor in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.37	3.86	9.93	20.21	37.11	63.84	103.99	160.26	232.62	317.08	111.8%

(Source: Tractica)

2.14.9 RISK ASSESSMENT AND COMPLIANCE

Financial institutions and other companies spend billions every year on fraud, counterfeit, and a host of other threats to credit, identity, and financial transactions. The World Bank estimates the amount of money laundered each year amounts to somewhere between \$2 trillion and \$3.5 trillion. To combat this tremendous problem, anti-money laundering (AML) compliance and penalties cost banks approximately \$18 billion annually. The size of this problem, coupled with the vast number of transactions and actors involved in international financial systems, not to mention the lack of data mutualization and duplicative efforts, means that today's AML procedures are extremely manual and labor intensive.

To combat these threats and help reduce costs, companies are developing AI-based methods for assessing customer risk and improving regulatory compliance. AI and DL are being applied to better learn legal requirements and variations, spot inconsistencies across millions of real-time transactions, identify policy violations, and mitigate against financial crimes, money, or trade-based laundering, and assess liquidity risk.

NextAngles is a venture within Mphasis Corporation, which provides AI tools to mitigate know your customer (KYC)-AML and trade-based money laundering (TBML)-related risks by using algorithms to rapidly review large amounts of documentation, detect suspicious or inconsistent information, run sentiment analysis, and automate transaction monitoring. It also provides solutions that equip financial crime investigators with tools for more rapid and proactive data compilation, consolidation, analysis, and inference drawing. The platform “learns” over time and offers easier scale through a conversion of structured English documents into executable rules.

RAGE Frameworks, an automation technology and services provider, is supporting clients with assessing daily business and credit risk. Companies can set up the system to assess dozens of business drivers to ascertain such risk. For instance, enterprises can use RAGE to monitor the internet, stock investments, competitive intelligence, market developments, and a range of risk factors, and then apply linguistic analysis and learning to identify and alert them of various business risks as they evolve in real time.

Insurance companies are also looking at AI and DL tools to assess policy-holder risk. The 2013 Monsanto-Climate Corporate acquisition leverages agricultural and weather data to support such a business model. It is using DL to determine the underwriting risk of selling farmers insurance against weather-related losses.

Tractica forecasts that the annual revenue for risk assessment and compliance in finance will increase from \$11.62 million worldwide in 2016 to \$197.55 million in 2025.

Table 2.110 Risk Assessment and Compliance in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	11.62	14.26	18.40	24.95	35.27	51.16	74.65	107.25	148.97	197.55	37.0%

(Source: Tractica)

2.14.10 TAX FILING AND PROCESSING

Filing taxes and processing returns is typically a stressful endeavor. The complexity of multiple returns, assets, dependents, and a litany of regulatory requirements renders it daunting for most individuals and small businesses. Even with the use of a professional accountant, errors can be costly and painful. In an effort to streamline the process, companies are exploring the use of AI for tax filing and processing. The idea is to use NLP, ML, and DL to more rapidly analyze tax filing data, tax codes, and deductions to reduce errors and ensure tax payouts are correct.

U.S.-based tax filing provider H&R Block introduced IBM Watson to its software in the 2017 personal tax filing season with the intent of helping tax professionals more easily identify deductions and credits. To develop the technology, it gave Watson tax data that included 74,000 pages of federal tax code and thousands of tax-related questions H&R Block had accumulated over 60 years, and worked with human tax professionals to refine questions and the interface. H&R Block CEO Bill Cobb told CNBC on June 14, 2017 that the addition

of Watson was one of the reasons the company beat Wall Street estimates for its most recent quarter. When onscreen, preparers and customers could see Watson at work. "We had...various call-outs of deductions and credits you could take," said Cobb, "What was also a very pleasant surprise was how much our tax pros got excited."

Intuit introduced ExpenseFinder as a feature for its tax services for the self-employed in December 2016. According to a company press release, ExpenseFinder "Finds deductible business expenses that self-employed may not know they can claim, saving them money. It proactively uncovers business expenses by securely gathering and automatically scanning bank accounts and credit card transactions and recommending potential deductible business expenses. Customers then confirm which expenses apply to their business to help them get every deduction." In the same timeframe, the company launched a natural language friendly search engine called ExplainWhy, which according to the company blog, can answer questions, giving users a personalized explanation of their tax deductions, credits, and refunds.

Tractica forecasts that the annual revenue for tax filing and processing in finance will increase from \$1.12 million worldwide in 2016 to \$108.02 million in 2025.

Table 2.111 Tax Filing and Processing in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.12	2.35	4.45	7.96	13.70	22.74	36.28	55.24	79.60	108.02	66.1%

(Source: Tractica)

2.14.11 TRANSACTION FRAUD DETECTION

Fraud and money laundering are extremely costly and difficult for financial institutions to identify, not to mention resolve. As methods and tools have adapted in the digital age, detecting transaction fraud is an ongoing priority (not unlike cybersecurity) as fraudsters constantly adjust their tools and methods in the digital age. While fraud detection software has been on the market for some time, approaches relying solely on historical data and business rules are insufficient to mitigate the evolving threat.

To detect new fraud schemes, AI, ML, NLP, and DL, are being explored in ways that do not solely rely on pre-programmed rules or models-based or historical data. The goal for such systems is to become self-learning, with models continuously updating individual profiles, threat profiles, payment methods, situations, behaviors, and other parameters. AI is also useful in helping process multiple data types, as new payment types and methods require flexibility in data processing. In addition, analyzing credit/debit card usage patterns and device access allows security specialists to identify points of compromise.

Brighterion supports AI-powered fraud detection and compliance solutions with iPrevent, a real-time behavioral profiling engine that uses multiple AI technologies and smart agents to identify and stop previously unknown fraud schemes. Profiles developed by the system include thousands of dimensions and behavioral characteristics, tracking where and when transactions take place and the context and methods for each transaction. Collectively, these feed risk scores and reason codes for each. iPrevent currently supports more than 6,200 transactions per second. Profiles are automatically updated in real time and new intelligence is associated across all relevant business lines.

MasterCard recently rolled out its own deployment, Decision Intelligence, which is designed to detect fraud and increase the accuracy of real-time card approvals, and false declines at check-out. The system leverages account information, such as customer value segmentation, location, merchant, device data, time of day, type of purchase, etc., and applies proprietary algorithms to provide a predictive score to the card issuer. Many other companies, such as Stripe and Feedzai, are working in this space.

Tractica forecasts that the annual revenue for transaction fraud detection in finance will increase from \$10.62 million worldwide in 2016 to \$364.33 million in 2025.

Table 2.112 Transaction Fraud Detection in Finance, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	10.62	15.05	22.35	34.29	53.51	83.55	128.31	190.77	270.90	364.33	48.1%

(Source: Tractica)

2.15 GAMING

2.15.1 CREATE DYNAMIC AND INTERACTIVE VIDEO GAME EXPERIENCES

Although gaming as an industry includes a wide range of games, including casinos, board games, fantasy sports, and beyond, the ability to create dynamic and interactive video game experiences is becoming a popular application for machine and deep learning.

Game developers are using neural networks to support all kinds of use cases, such as predicting player actions, inferring and recognizing player goals, developing adaptations to unpredictable player actions, and learning from simulated game environments. Another area is that of the creation and development of non-player/playable characters (NPCs), in which other characters in the game help play out the game's storyline and act according to pre-determined or responsive behavioral prompts. Game development company, Unity Technologies is using machine and deep learning to train NPCs overnight, freeing up human developers to work on the core gaming experience. Instead of creating NPCs with purely predetermined activities, they are using AI to train them to better support the storyline through more dynamic behaviors.

Perhaps one of the most interesting aspects of DL applications within the virtual gaming context is that companies are exploring it to support development of large-scale interactive gaming environments. A company called Improbable uses their SpatialOS platform to enable developers to build virtual worlds that can accommodate thousands of simultaneous players and scenarios at the same time. Its SpatialOS Innovation Games Program recently partnered with Google's Cloud platform to power massive networks of multiplayer online games in which virtual worlds are populated by thousands of players interacting with and changing a single dynamic environment, playing out in real-time over weeks.

Gaming environments are also crossing over from purely entertainment into industries as well, particularly in support of education, learning, and training students or employees on new skill sets. With the rise of VR and AR hardware, new form factors will support immersive training and gaming will becomeing more mobile and less fixed.

Tractica forecasts that the annual revenue for creating dynamic and interactive video game experiences in the gaming sector will increase from \$5.14 million worldwide in 2016 to \$290.3 million in 2025.

Table 2.113 Creating Dynamic and Interactive Video Game Experiences in Gaming, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	5.14	11.11	20.14	33.60	53.26	81.15	119.10	167.87	226.22	290.30	56.5%

(Source: Tractica)

2.16 GOVERNMENT

2.16.1 AGENT-BASED SIMULATIONS FOR DECISION-MAKING

Complex organizations like governments and businesses have long understood the importance of long-term strategy. Entire industries of strategic consulting firms, financial and industry analysts, and competitive intelligence brokers exist to help such organizations plan for future scenarios. In government contexts, strategic decision-making might include areas like national security, disaster response, (smart) city or utilities infrastructure, employment, policy decisions, and social or political campaigns. In any of these contexts, government planners are faced with the challenge of understanding highly complex systems and designing sophisticated technical schema, governance frameworks, and feasible outcomes, while balancing costs and “what-if” scenarios.

In perhaps one of the most alluring applications for AI, agent-based simulation for decision-making is useful in simulating and predicting the behavior of complex systems, where there are millions of individual entities or agents (humans, cars, viruses, etc.) that can have multiple dynamic characteristics. Each of the entities interacts with each other and behavior can be simulated using AI techniques like reinforcement learning to understand and plan for complex system benefits from simulation. In the past, developers and planners were limited by compute power, and the ability to scale or introduce new elements in real time. GPUs and high-speed processors are helping make virtual simulation possible.

The company Improbable is developing a simulation-building platform called SpatialOS, which allows companies to create large-scale virtual worlds in the cloud. While its platform has primarily been used by gaming companies, Improbable is also working with more government entities. Last year, the company worked with the British government to develop a simulation of the internet itself. It is now working in conjunction with governments, urban planners, and academics to develop elements like traffic patterns, energy consumption, waste management, and pollution in order to help governments with smart city planning. Improbable, which recently received \$500 million in funding from SoftBank, is also working on enabling digital simulations of economies, natural and biological systems, and virtual worlds of any type. According to Herman Narula, co-founder and CEO of Improbable, “Virtual worlds are going to be the playing ground where A.I. is going to evolve.”

Another company developing similar simulation technologies for smart cities, as well as autonomous cars, drones, and robotic toys and gaming, is Prowler.io.

Tractica forecasts that the annual revenue for agent-based simulations for decision-making in the government sector will increase from \$0.03 million worldwide in 2017 to \$3.15 million in 2025.

Table 2.114 Agent-Based Simulations for Decision-Making in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.00	0.03	0.09	0.19	0.34	0.58	0.94	1.47	2.20	3.15	185.9%

(Source: Tractica)

2.16.2 BEHAVIORAL ANALYTICS

Governments of all types have great interest in understanding their populace. Behaviors, in particular, signal insights that are often omitted, forgotten, or hidden from solicited feedback, such as surveys, censuses, or reporting. A range of methods and technologies is used to collect information about people's behaviors, including all manner of ML, DL, NLP, CV, and MR.

Given the astronomical size of data gathered about millions of people, it is not far-fetched to consider governments' abilities to track behavior of groups and individuals across the internet. Organizations like America's National Security Agency (NSA) or Britain's Government Communications Headquarters (GCHQ) have been shown to possess programs that can see thousands of different parameters, including both metadata and content.

PRISM, for example, was an alleged tool used by the NSA to collect private electronic data like emails, chats, voice over internet protocol (VoIP) call records, cloud-based files, and other digital interactions belonging to users of major internet services like Gmail, Facebook, Outlook, and others. With the rise of consumer smartphones and ubiquitous sensing technologies, such data are highly precise and more mobile, and if they can indeed be aggregated by government surveillance programs, offer the most intimate details about our behaviors. It is unlikely such tools would be commercialized and are likely to remain secret initiatives supporting specific intelligence agency objectives.

Tractica forecasts that the annual revenue for behavioral analytics in the government sector will increase from \$0.7 million worldwide in 2016 to \$21.69 million in 2025.

Table 2.115 Behavioral Analytics in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.70	0.97	1.39	2.05	3.09	4.70	7.11	10.62	15.45	21.69	46.4%

(Source: Tractica)

2.16.3 CONVERTING PAPERWORK INTO DIGITAL ASSETS

Government is notorious for paperwork. Indeed, there is a vast range of organizations that will benefit from converting paper documents and other unstructured data, such as email, PDFs, charts, and graphs, into structured data for a wide range of uses, but particularly operational efficiency. A Deloitte study found that, in federal government jobs, documenting and recording information consumer half a billion hours per year. AI can help to significantly reduce administrative tasks involved in converting paperwork and processing documentation. NLP, ML, DL, and even bots can be applied in these contexts to both capture paper documents and automate paperwork processing, such as data entry, filling in forms, invoicing, and automating reports on this information, as outlined in Section 2.14.2.

HyperScience converts paperwork into digital data for enterprises and government institutions. The platform can scan handwritten or typed text, and fill in all the data fields into the organization's (government or business) system, and then some. According to the company website, HyperScience claims its HS Evaluate tool "is designed for organizations which process complex claims or applications that require evaluation and contextual judgment. Our AI software can automatically review an extensive claim file, eliminate duplicate entries, assess eligibility, and then deliver precise adjudication decisions. HS Evaluate makes decisions based on the current context while factoring in millions of data points from past experience to create a more accurate picture, more comprehensively than a human could."

HyperScience's Tim Kalimov told Tractica how Evaluate could help a government institution. "The U.S. government processes a large number of disability claims per year. We can help people review these forms faster," said Kalimov, "Each case could be thousands of pages and a human has to go through that. So we can go through large case files and substantively eliminate duplicate content, to where the case is reduced by 20% to 30%."

Tractica forecasts that the annual revenue for converting paperwork into digital assets in the government sector will increase from \$11.98 million worldwide in 2016 to \$390.22 million in 2025.

Table 2.116 Converting Paperwork into Digital Assets in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	11.98	19.92	31.67	48.91	73.84	109.15	157.62	221.30	300.07	390.22	47.3%

(Source: Tractica)

2.16.4 CROWD ANALYTICS

Crowds and public gatherings typically offer lots of information, including cause, size, demographic, movement patterns, etc. In 2015, IHS estimated there were over 245 million operational cameras active globally. When it comes to digitally analyzing this information, many approaches fall short. For instance, crowd analytics may be effective in a pre-configured setting (e.g., a plaza), but analyzing crowd formations in new or unknown areas fails. As public and commercial infrastructures install more and more cameras, and as CV and recognition techniques advance, DL finds new application in analyzing crowds. Research institutions in China and India have been working to develop training data and DL solutions capable of estimating crowd density and dynamics.

Such government surveillance tactics are one of many with which DL researchers and innovators are experimenting and include object detection and identification, behavioral analytics, sentiment analysis, facial recognition, and even predicting social unrest or geopolitical events. In conjunction with surveillance footage, the sheer volume of video produced is fed into DL models to detect abnormalities, identify and trace moving objects, better manage large crowds, and ensure municipal and public safety.

Tractica forecasts that the annual revenue for crowd analytics in the government sector will increase from \$0.7 million worldwide in 2016 to \$17.89 million in 2025.

Table 2.117 Crowd Analytics in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.70	0.93	1.28	1.83	2.68	4.00	5.98	8.84	12.79	17.89	43.3%

(Source: Tractica)

2.16.5 DIALECT CLASSIFICATION

Dialect identification (DID) and classification is a subset of the general challenge of language identification (LID) by computers. “LID refers to the process of automatically identifying the language class for given speech segment or text document. DID is arguably a more challenging problem than LID, since it consists of identifying the different dialects within the same language class,” according to [“Automatic Dialect Detection in Arabic Broadcast Speech”](#), a research paper published in August 2016. For security agencies, being able to identify the dialect or accent that a speaker uses can help them better identify an individual.

Tractica forecasts that the annual revenue for dialect classification in the government sector will increase from \$0.17 million worldwide in 2017 to \$15.71 million in 2025.

Table 2.118 Crowd Analytics in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.17	0.46	0.93	1.69	2.88	4.68	7.32	10.97	15.71	N/A

(Source: Tractica)

2.16.6 DISASTER AND EMERGENCY MANAGEMENT

When disasters strike, they wreak havoc on communities. From weather, public health, terrorist attacks, or infrastructure failure, one of the essential roles of government is to help provide protection, evacuation, and relief. Depending on the type of emergency, disaster and emergency management have historically relied on industry experts (e.g., meteorologists, epidemiological researchers, military intelligence) and local law enforcement to assess threats, damage, and response plans. But in reality, this data is often reported after the fact, such as data from emergency rooms or urgent care centers. To better prepare for larger-scale disasters, governments set up agencies and budgets, and allocate funding for organizations like the Red Cross, which can help scale prediction, planning, and response efforts. AI is now a developing tool to aid in disaster and emergency management. The role of AI here is to help analyze more data from more sources for patterns or other notable signatures. Organizations are using NLP and DL to mine for trends across unstructured data sets.

In 2014, the Centers for Disease Control (CDC) Situational Awareness Branch conducted a study using AI software provider, Luminoso, designed to help predict disease outbreaks and help detect them in real time. Specifically, it wanted to more accurately predict the spread of the flu, but also demonstrate the value of analyzing unstructured data to ascertain the severity of a range of public health threats, such as Ebola, MERS, and other unknown diseases. The CDC and Luminoso developed a framework to correlate mentions of symptoms, detect changes in symptoms over time, and monitor core concepts and expressions that would discern the flu from other diseases. It analyzed Twitter feed data, as well as free text references, aligned with doctors’ and hospital reports. As more unstructured data were analyzed, it modified existing frameworks and models to further increase accuracy

in number of people who contracted the flu, the severity and duration of the current strain of flu, and the effectiveness of the flu vaccine. In the case of the Ebola outbreak, the problem was exacerbated by misinformation, misperception, and even conspiracy theories proliferated on social media. This data helped the CDC rapidly surface and plan effective responses to educate, inform, and reassure an anxious public.

Tractica forecasts that the annual revenue for disaster and emergency management in the government sector will increase from \$0.68 million worldwide in 2017 to \$62.09 million in 2025.

Table 2.119 Disaster and Emergency Management in Government, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
-	0.68	1.81	3.68	6.67	11.37	18.52	28.95	43.39	62.09	N/A	

(Source: Tractica)

2.16.7 FACIAL RECOGNITION

Facial recognition is a computer or machine's ability to identify or verify a person based on their facial characteristics. Computer applications use digital images, video frames, and video feeds to recognize people's faces. AI supports facial recognition through various ML and DL techniques, sometimes involving CV. Recognition algorithms are commonly divided into two main approaches:

- **Geometric:** Looks at distinguishing features (face, nose, shape of eyes)
- **Photometric:** Takes a statistical approach by processing an image into values, then eliminates variances by comparing the values with templates

Advancements in processing power and in other adjacent technologies have brought about complimentary techniques to enhance facial recognition. Some of these include:

- **3D Facial Recognition:** Using 3D sensors to capture information about shape, depth, lightfall
- **Skin Texture Analysis:** Uses image recognition to turns unique lines, spots into a mathematical space
- **Thermal Analysis:** Uses thermal cameras to detect head shape, while accessories such as glasses or make-up are undetected
- **Eye and Retina Recognition:** Detects unique features of a person's eyes
- **Emotion Recognition:** Facial expressions or physical features are analyzed against databases to determine the subject's disposition

Facial recognition is a verifiable biometric and useful in a variety of commercial applications, outlined in Section 2.9.7. In government contexts, uses for the technology typically center around security, law enforcement, fraud prevention, and identity authentication. Most of these applications work by using advanced cameras that reference images and footage collected against large databases of facial images.

Facial recognition technology has been synonymous with surveillance and security in the United States for years. In 2010, the Federal Bureau of Investigations (FBI) updated and enhanced its fingerprint database with other advanced biometrics. Called Next Generation

Identification, the database now holds about 50% of Americans' photographs. A recent oversight committee hearing presented analysis of the program, which troubled lawmakers, privacy groups, and even vendors of facial recognition technology. For example, the hearing reported that roughly 80% of photos in the FBI's network are non-criminal entries, including pictures from driver's licenses and passports. The algorithms used to identify matches are inaccurate about 15% of the time, and are more likely to misidentify black people than white people.

Meanwhile, Georgetown University recently published a report, the result of a year-long investigation into American police use of facial recognition technology. It found that merely having a state-issued driver's license or photo ID allows police to remotely search for and identify an individual's face from photos posted on social media without a warrant or court supervision. The report also found that police departments across the United States regularly use special smartphones and body cameras designed to capture faces, irises, and in some cases, DNA swabs of people stopped; data are pinged back to biometric databases, which return any identifying information or criminal records.

Other countries are grappling with the technology, in an array of efforts to weed out bad actors, while simultaneously quelling public and privacy groups' concerns of civil liberties.

- Canada recently revealed it used facial recognition technology to identify 15 suspects wanted on immigration warrants, who all used false identities to apply for travel documents. The technology was used to help locate and arrest those ineligible to stay in Canada as a result of being involved in terrorism, organized crime, or human rights violations.
- The Macau government is now using facial recognition at automated teller machines (ATMs) in order to prevent Chinese money laundering and anti-terrorism commonly tied to casinos.
- The Mexican government employed face recognition software to prevent voter fraud in its 2000 presidential election. Some individuals had been registering to vote under several different names, in an attempt to place multiple votes. By comparing new face images to those already in the voter database, authorities were able to reduce duplicate registrations.
- The Australian people and New Zealand t Services use SmartGate, an automated border processing system that uses facial recognition to compares individuals' faces with the image in the e-passport to verify identity and counter fraud.

Scores of other countries are at varying levels of regulation when it comes to how such systems are used, where images may be sourced, and what levels of consent are required.

Tractica forecasts that the annual revenue for facial recognition in the government sector will increase from \$0.17 million worldwide in 2017 to \$15.76 million in 2025.

Table 2.120 Facial Recognition in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.17	0.46	0.93	1.69	2.89	4.70	7.35	11.01	15.76	N/A

(Source: Tractica)

2.16.8

OBJECT DETECTION FOR SURVEILLANCE

Perhaps the most valuable use of AI in security is the use of object detection and classification, which takes sensor data, often from cameras, and then uses complex algorithms to classify these objects so that the AI system can then “learn” their characteristics, and recognize them in real time. The challenge is not in capturing images, as today’s HD cameras can present images in stunningly clear detail. However, in a moving environment, objects can appear to change size as a vehicle or camera approaches. The angle at which an object is viewed can also skew its appearance, and the presence of other factors (rain, bright sunlight, low lighting, glare, dirt, snow, or any other number of obstructions) can alter the appearance of an object, making it hard to accurately and consistently identify the object.

This is an area where machine vision and ML can provide invaluable support. By capturing a wide range of images of objects from a variety of vantage points, angles, and in different conditions, a repository of images that can be definitively classified as that object can be created, and used to “train” a ML system to identify and classify objects that resemble objects in the repository. By then assigning various other attributes to each object, such as whether the object is a person, car, animal, weapon, permanent, temporary, or capable of motion, the system can begin to develop logical rules on handling each object and the rules for dealing with them.

In government contexts, surveillance and closed-circuit television (CCTV) cameras can use object detection to learn patterns in an area, detect faces, gender, heights, mood, read license plates, and identify anomalies, potential threats, unaccounted for packages, etc. Following multiple camera feeds can track individuals’ movements over time and distances as well. Data feeds are typically stored in large databases that allow users to search for 15-second increments, such as: “everyone who entered X parking lot between 6 p.m. and 9 p.m. on Friday night” and conduct forensic analysis. It is perhaps worth noting that such systems have received tremendous pushback from civil liberties groups and privacy advocates, yet programs continue to grow.

A company called SCW sells such cameras to government agencies. 3-VR is another facial recognition provider in this space.

Tractica forecasts that the annual revenue for object detection for surveillance in government will increase from \$1.17 million worldwide in 2016 to \$193.64 million in 2025.

Table 2.121 Object Detection for Surveillance in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.17	3.33	6.91	12.75	22.08	36.67	58.82	91.12	135.78	193.64	76.5%

(Source: Tractica)

2.16.9

PREDICTING SOCIAL UNREST AND GEOPOLITICAL EVENTS

Governments have been developing programs and methods to monitor civilians for years and these programs arguably include some of the most sophisticated and comprehensive surveillance technologies around. Given the astronomical size of data gathered about millions of people, it is not far-fetched to consider governments’ abilities to perform simulations and scenario planning of any number of geopolitical events. Not only could organizations like America’s NSA or Britain’s GCHQ run Big Data analysis and predictive

analytics using data they collect, but they could enhance scenario parameters using data monitored from other countries as well, particularly countries experiencing certain sorts of social unrest or geopolitical dynamics. It is unlikely such tools would be commercialized and are likely to remain secret initiatives supporting specific intelligence agency objectives.

Tractica forecasts that the annual revenue for predicting social unrest and geopolitical events in the government sector will increase from \$0.7 million worldwide in 2016 to \$17.83 million in 2025.

Table 2.122 Predicting Social Unrest and Geopolitical Events in Government, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
0.70	0.93	1.28	1.82	2.68	3.99	5.96	8.82	12.75	17.83	43.2%	(Source: Tractica)

2.16.10 REAL-TIME VIDEO ANALYTICS

As camera technologies have sharpened the quality of video feed image precision, so too have analytics supporting such capture. As video feeds have expanded in volume, video analytics represent the only way to extract value in form of insights, patterns, action, from so much data.

AI is increasingly becoming a core enabler for video analytics, particularly for real-time analysis and action. DL, CV, and object and facial recognition enable accuracy and speed when it comes to analysis. DL also helps analyze and process multiple video and data streams and can help multiple systems communicate with each other. Common video analytics solutions may deploy various AI techniques to support the following areas:

- **Behavior Monitoring:** Motion detection, footfall or pedestrian traffic, facial detection, privacy masking, vandalism detection, theft or suspicious activity detection
- **People Monitoring:** People counting, people scattering, crowd analytics, line management
- **Vehicle Monitoring:** Vehicle classification, license plate monitoring, traffic monitoring, road monitoring
- **Device Monitoring:** Protection against tampering with camera, infrastructure, perimeter, or other intrusion

Use cases in government might include smart city security, law enforcement, intelligent transportation systems, public gatherings, etc. Essentially, video analytics technology helps security software “learn” what is normal so it can identify unusual, and potentially harmful activities. The technology requires operator feedback as pure object detection is insufficient. It is also not without some controversy as some warn of dangerous consequences when AI “decides” what or who looks suspicious.

A few companies in this space include Avigilon, which designs and manufactures video surveillance software and equipment, as well as elinfochips, which designs embedded software.

Tractica forecasts that the annual revenue for real-time video analytics in the government sector will increase from \$0.1 million worldwide in 2017 to \$9.01 million in 2025.

Table 2.123 Real-Time Video Analytics in Government, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
-	0.10	0.26	0.53	0.97	1.65	2.69	4.20	6.30	9.01	N/A	(Source: Tractica)

2.16.11 SENTIMENT ANALYSIS

Understanding the emotional context and drivers of citizens has been critical to the strategy of government and effective democracy since its earliest days. In recent times, research in the areas of public opinion, sentiment, and emotion in natural language texts, speech, music, and other media have grown under the umbrella of subjectivity analysis and affective computing.

The popularity of the internet and the rapid expansion of social media, as well as a wide variety of user-generated content, have become available online. Beyond research, an entire sentiment analytics software industry has emerged to support commercial efforts at understanding sentiment. But the major challenge has remained: how to process and organize at scale vast amounts of rich, unstructured user-generated content, such as open-ended text, audio, video, call logs, images, etc.

Many governments are using social media tracking to measure national and local sentiment. Over the past few years, sentiment analysis has been primarily used for messaging, policy positioning, and campaigning. For example, the Obama team used advanced sentiment analysis to gauge public opinion on policy announcements and campaign messages ahead of the 2012 presidential election. But sentiment analysis can also be used to identify trends and issues among constituents, and more importantly, push toward citizen-centric models where priorities are driven, even designed, according to citizen needs.

In Finland, the Veikkausm is a government-owned betting agency that runs 20 different gambling ventures in the country (e.g., a national lottery, scratch tickets, football pools, sports betting, etc.). It uses Big Data analytics, including AI-driven sentiment analysis, to identify those suffering from gambling addictions and ensure citizens are not compulsive or self-destructive in their gambling. To do this, it analyzes all gaming transactions, develops user profiles based on transactions, and scans customer service emails, social media data, etc. Running sentiment analysis across this data helps the Veikkausm determine tone, addictive tendencies, and potential problems, and cease marketing activity with certain profiles. In the future, the agency says it might intervene directly with addicts.

While using this information to bridge awareness gaps between vulnerable or concerned citizens and the government programs designed to help them present obvious benefits, sentiment analysis is not without controversial applications. One newspaper, *The Australian*, recently obtained a 23-page report in which Facebook had conducted deep sentiment analysis to discern when teenagers were suffering from depression, anxiety, stress, or defeat. In light of the revelations, Facebook acknowledged it had shared this information with advertisers, later saying it was a mistake, as policies prohibit advertisers from targeting based on emotional state. Outcry ensued as critics feared exploitation of vulnerable or insecure teens, or manipulation by targeting them with specific products like acne cream or make-up. Meanwhile, Facebook asserts that its role in taking interest in monitoring emotions

transcends its business model (advertising); particularly as the platform has seen an increasing number of young people expressing suicidal intentions, even attempting to broadcast suicides via Facebook Live. Facebook is using AI and pattern recognition to flag posts and review cases, offering help to users if deemed appropriate. While Facebook is a commercial enterprise, and not a government, it illustrates many of the same capabilities as mixed implications for AI-driven sentiment analysis.

Tractica forecasts that the annual revenue for sentiment analysis in the government sector will increase from \$0.7 million worldwide in 2016 to \$22.6 million in 2025.

Table 2.124 Sentiment Analysis in Government, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
	0.70	0.98	1.41	2.10	3.19	4.86	7.38	11.04	16.08	22.60	47.1%

(Source: Tractica)

2.16.12 SOCIAL MEDIA BOTS

There is no doubt that social media is a tool used to both proliferate and influence opinion. Its success is in no small part due to the ability for *anyone* to share their opinion with relatively low and sometimes no censorship. User-generated content is both the foundation and driver of social media consumption. We are drawn to platforms where our friends, family, and colleagues engage and we stay there based on the content they and others in our broader networks share. The emergence of social media bots has introduced an altogether external vector in the role, experience, and use of social media by governments.

The term “bots” is a generic term for small computer programs coded to detect and analyze certain inputs, and then trigger specific responses and outputs. In the case of social media bots, these can be programs that look, act, and speak like actual human users. The ability to deploy these bots at scale and design them to do, say, or engage with other users as a person or group of people would and about any topic or agenda (not to mention the general lack of regulation or legal precedent for bots) is why many have called social media bots “automated propaganda.”

In particular, individuals, political groups, or governments can create bots that *automatically*:

- Post content, associated with supporting or rejecting specific topics, campaigns, people
- Share or endorse content, associated with specific keywords, individuals, hashtags, sentiments, etc.
- Promote education, access to content, services, groups, events, etc.
- Interact one-to-one (at scale) with unique messages to thousands of other users at the same time
- Follow or “friend” specific people or campaigns
- Start new accounts, open new groups or threads
- Populate accounts with fake information (e.g., metadata, geography, GPS spoofing)

When bots (not humans) proliferate content, spread ideas, or influence language around campaigns, it threatens the very integrity of consensus and the will of the crowd. Such

“manufactured” consensus gives the illusion of significant popularity, even virality. When bots automatically post, share, or comment on stories that are one-sided or entirely fabricated, patently false narratives suddenly look as though they have been endorsed by hundreds of thousands of people.

The 2016 U.S. presidential election brought the use of social media bots for political purposes to the forefront. Not only has AI advanced the capabilities and sophistication of bots themselves, but as Oxford University’s report on Computational Propaganda Research Project states, “During the 2016 campaign, a bipartisan range of domestic and international political actors made use of political bots.” The same research study found that during the first and second presidential debates, a third of pro-Trump tweets and nearly a fifth of pro-Clinton tweets were generated by bot accounts. Bots were also used during the Brexit vote in the United Kingdom; in India’s 2014 elections; and by ISIS to amplify propaganda across thousands of [bot] accounts.

The question of social media bots and their influence remains unresolved, as measuring influence is notoriously difficult. What is clear is the impact that communities can have on human thought patterns: ethnic and cultural values, group think, echo chambers, and crowd consensus. In the age of social media bots, what requires greater transparency is the composition of communities themselves—who is human and what is a bot? Even ML algorithms designed to analyze thousands of features to detect “bot or not” are not 100% accurate, and unable to predict what they will do or say next.

Social media bots do not necessarily have to be tools for automated propaganda. In fact, bots can be helpful tools for disseminating and curating information, for tailoring content to the individuals’ specific interests or concerns, or even aiding individuals with accessing services or advice.

Tractica forecasts that the annual revenue for social media bots in the government sector will increase from \$0.03 million worldwide in 2016 to \$0.51 million in 2025.

Table 2.125 Social Media Bots in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.03	0.04	0.05	0.06	0.09	0.13	0.18	0.26	0.37	0.51	36.0%

(Source: Tractica)

2.16.13 STREET LIGHTING

Electricity costs associated with street lighting and municipal lighting infrastructure account for significant energy spending. For years, massive lighting infrastructure relied on pre-configured settings and timers running on energy sourced from the grid, so lights came on at night and went off in the mornings. Even the introduction of light-emitting diodes (LEDs) to programmable lighting management systems reduced urban electricity costs some 70%. As more sensors, cameras, and network connectivity are making their way into municipal infrastructure, even lightbulbs themselves, AI becomes a natural next step for smart(er) lighting. Networked lamps are core nodes in multi-functional communications of a smart city. Intelligent lighting lowers municipal electricity costs, enables demand-driven lighting, and reduces CO₂ emissions.

Despite the increase of smart lighting, many bulbs remain only switchable, and far from intelligent communicators. AI introduces new ways for lighting infrastructure and systems to

“learn” and adapt autonomously to environmental context. Learning will happen through sensing and combining other integrated data into logical conclusions: for example, using presence sensor data, such as motion, noise, pollution, etc., to control room temperature or building security. Using CV, object recognition, or even DL to mine large data sets for patterns in electricity demands, AI-powered lighting will enable more automatic/autonomous lighting, greater accuracy for energy allocation, variations tied to specific events, and the acceleration of cost saving associated with reduced emissions and electricity distribution.

The city of Glasgow is currently demonstrating intelligent street lighting in which its lighting network uses real-time data to improve lighting, safety, Wi-Fi, financials, and environmental impacts. Energy-efficient bulbs are capable of noise detection, air quality improvements, footfall detection, and even Wi-Fi-provisioning for city services and citizens. Lighting infrastructure is integrated with Glasgow’s Operations Center, which feeds real-time data in from other automated city systems for analysis, management, maintenance, and optimization over time.

Finnish company Helvar helps construction, real estate, and municipal clients develop AI-enabled lighting solutions. Helvar is developing self-learning algorithms to serve as “out-of-the-box” learning systems that integrate with and support other BASs. Specifically, its lighting systems will analyze data on behavior patterns and predictions to help designers improve building layouts. Self-learning algorithms in lighting systems can also help benefit maintenance programs by delivering automatic re-configurations or updates.

Tractica forecasts that the annual revenue for street lighting in the government sector will increase from \$0.71 million worldwide in 2017 to \$65.48 million in 2025.

Table 2.126 Street Lighting in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.71	1.91	3.88	7.04	11.99	19.53	30.53	45.75	65.48	N/A

(Source: Tractica)

2.16.14 TRAFFIC LIGHT MANAGEMENT

Traffic control systems typically combine traditional traffic lights with an array of sensors. Dynamic control signals adjust the timing and phasing of lights according to limits that are set in controller programming, typically updated once every 2 to 3 years. Nonetheless, traffic congestion still carries significant costs: an estimated \$121 billion a year in the United States alone, plus 25 billion kilograms of CO₂ emissions and 40% of urban drivers’ time spent idling, according to Carnegie Mellon.

As cities and municipal infrastructure become increasingly connected through sensors and data analytics, AI will become a critical tool to aid with learning from and better predicting traffic flow. DL is well suited for this use case, given the diverse and often unstructured and time-series data sets flowing in from a range of inputs influencing optimal lighting energy utilization. Pedestrian traffic, private, commercial, and public vehicle movement and concentration, weather, and municipal services are just some of the diverse and huge data sources analyzed to optimize traffic lights.

Carnegie Mellon and the city of Pittsburgh, Pennsylvania are developing an AI-enabled traffic management system in which signals communicate with each other to adapt to changing traffic conditions. The technology monitors vehicle numbers via fiber-optic video receivers

and makes real-time state changes with the objective of avoiding congestion and the amount of time vehicles spend idling. In pilot testing of the program, results raised eyebrows: it reduced travel time by 25%; idling time by more than 40%; and emissions by 21%.

Surtac is a startup born out of this project, which aims to commercialize the technology. Radar sensors and cameras at each light detect traffic, then use algorithms to develop and refine a “timing plan,” which helps move and route vehicles based on recommending the most efficient route. Every signal makes its own timing decisions, operating in a “decentralized” mode where nodes themselves learn and act. The system also sends data from signal to signal “downstream” so that other intersections can act accordingly. The long-term plan for Surtac is to communicate directly with cars. Some cars will be shipped with short-range radios by the end of 2017. With these feedback loops, drivers could know when lights are about to change, or be alerted to nearby traffic conditions.

Tractica forecasts that the annual revenue for traffic light management in the government sector will increase from \$1.43 million worldwide in 2017 to \$130.96 million in 2025.

Table 2.127 Traffic Light Management in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	1.43	3.83	7.76	14.08	23.99	39.06	61.06	91.51	130.96	N/A

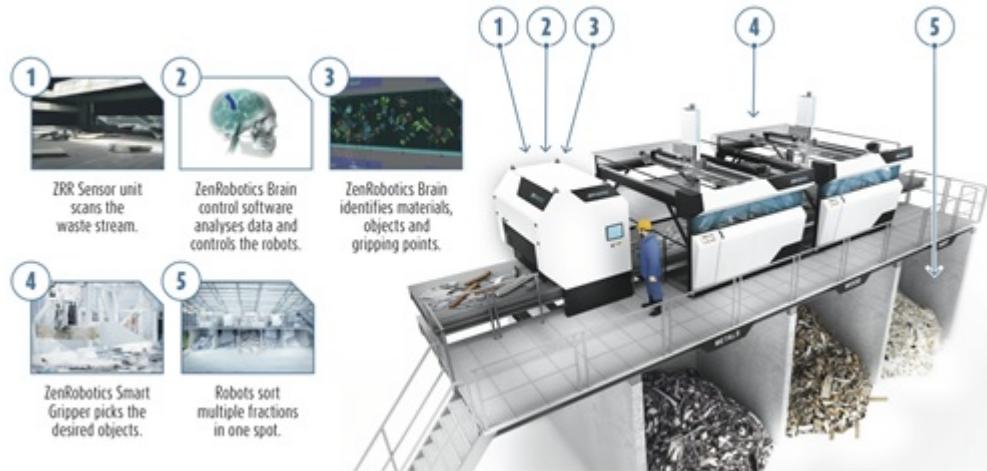
(Source: Tractica)

2.16.15 WASTE SORTING AND RECYCLING

Humans create a lot of waste. Of the 80% of recyclable waste produced every year, only 2% actually gets recycled. In the era of plastics, convenience, and hyper-consumerism, landfill waste is increasing greenhouse gas (GHG) emissions and contributing to climate change, contamination, risk to wildlife, and pollution. For businesses, costs of transport, disposal, and auditing reach into the millions, and are exasperated by variable packaging and recycling rules and regulations that vary by jurisdiction. AI is now being applied to this problem, using CV, robotics image and object recognition, and ML. In particular, current solutions are working on sorting waste, the greatest challenge businesses cite when it comes to waste management.

ZenRobotics uses robotics for waste separation with industrial robots powered by its software (ZenRobotics Brain), using CV, ML, and sensor data fusion for rapid sorting. Customers can select from common materials (e.g., metal, wood, cardboard) or the system can be trained for specific objects or new waste fractions. It uses infrared spectrum sensors, 3D sensor systems, hi-res gigabyte (GB) camera, imaging metal detector, and visual light spectrum sensor. Its ZRR2 unit, with two arms, conducts roughly 4,000 picks per hour. Multiple robots working together 24/7 process waste more quickly and accurately than humans.

Figure 2.17 Zen Robotics Waste Processing Workflow



(Source: ZenRobotics)

Startup Intuitive Robots is working on an automated waste sorting bin that uses AI to sort trash automatically. The device will use image recognition and DL in order to be able to identify any item of trash from any angle at different stages of decomposition. The technology, currently still in development, will identify items of trash as they are disposed and immediately sort them. As the system advances over time, performance should improve. The other cost efficiency of this product is its potential to eliminate the need for paid waste audits, given the ability for automated report generation, which could be submitted at any time. According to the company's founders, "the goal with this bin is to have 100% diversion rate: everything will be sorted instead of going to a landfill, helping both the environment and business to win."

Tractica forecasts that the annual revenue for waste sorting and recycling in the government sector will increase from \$0.05 million worldwide in 2017 to \$4.59 million in 2025.

Table 2.128 Waste Sorting and Recycling in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.05	0.13	0.27	0.49	0.84	1.37	2.14	3.21	4.59	N/A

(Source: Tractica)

2.16.16 WEATHER FORECASTING

Weather forecasting is part of government planning and resource allocation. Natural disasters can have catastrophic impacts on societies and national security. Weather forecasting helps governments support research into the impacts of climate change. AI and sensor data from hundreds of thousands of sources collected and monitored in real time (and over many years) are transforming the level of understanding and ability to forecast conditions. In addition to weather data, engines combine streaming data from social feeds, news reports, transportation data, and historical data on storms or other weather events.

While no one can ever fully predict the future, AI techniques apply reinforcement learning on past predictions and actual outcomes. By comparing predictions with accuracies, the model is able to learn and improve simulation capabilities, forecasting much further into the future.

AI can be used to perform weather pattern detection, such as cyclonic activity or other extreme weather events. NERSC has used CNNs to classify threatening climate events like cyclones. This work was performed on a CPU-only Cray XC30 supercomputer, where both the training and inference ran on the same platform, although some effort was involved in adapting the CNN algorithm to the climate data. The main goal for NERSC was to have a model learn the characteristics of a cyclone and classify it, an area where human decision-making variance is an issue. With the algorithm having between 80% and 90% accuracy in identifying extreme weather events, this is only the start and shows that AI techniques can be used for classification and identification of more complex weather systems and events.

Tractica forecasts that the annual revenue for weather forecasting in the government sector will increase from \$0.01 million worldwide in 2018 to \$0.34 million in 2025.

Table 2.129 Weather Forecasting in Government, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.00	0.00	0.01	0.02	0.04	0.06	0.10	0.16	0.24	0.34	87.7%

(Source: Tractica)

2.17 HEALTHCARE

2.17.1 AUTOMATED REPORT GENERATION

Businesses generate reports for a variety of reasons, from internal knowledge sharing and regulatory compliance to general operations, auditing, and accountability. In healthcare, there are thousands of different types of reports that must be generated, from emergency department patient flow to immunizations and the frequency of diagnoses and beyond. The primary driver in automated report generation is operational efficiency, allowing human employees to focus on more complex tasks.

Automated report generation frequently leverages natural language generation, in addition to NLP, ML, and/or DL. Automated report generation tools generally support the following tasks:

- **Data Sourcing:** Identifies and extracts data from relevant internal and external sources, depending on the application.
- **Data Interpretation:** Upon consolidating data in standardized formats, the solution aligns the data in templates, codes, and prepares it for analysis using ML.
- **Data Analytics:** Defines business rules and correlation/causality at scale. With predictive modeling and data enrichment, solutions can run hundreds of “what if” scenarios and perform trend analysis
- **Narrative and Semantic Commentary:** Using NLP and natural language generation, solutions can sometimes automate variance analysis and commentary writing in a systematic and structured way.

Several companies have emerged with a focus on automated report generation across multiple vertical markets.

3M is a multinational conglomerate that offers thousands of products in a range of industries, including medical products and electronics. 3M is focused on patient data processing and automated report generation healthcare-specific use cases. The company's NLP platform, bolstered by its 2010 acquisition of Cogent Systems, is used for computer-assisted coding (CAC) and clinical documentation improvement (CDI), especially in the healthcare industry. Its CodeRyte CodeMonitor provides an automated review of clinical documentation and compares the resulting evaluation and management (E/M) CPT codes to physician-assigned E/M CPT codes. Codes in agreement can go directly to a billing system. Its 360 Encompass System Professional uses NLP to provide auto-suggested codes for improved productivity, accuracy, and compliance for hospital-based professional fee coding.

Tractica forecasts that the annual revenue for automated report generation in healthcare will increase from \$0.52 million worldwide in 2016 to \$246.25 million in 2025.

Table 2.130 Automated Report Generation in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.52	5.66	13.44	25.04	41.98	66.02	98.72	140.75	191.03	246.25	98.4%

(Source: Tractica)

2.17.2 BIO-MARKER DISCOVERY

In biology, a biomarker or biological marker is a measurable indicator of some biological state or condition, typically used to assess biological, pathogenic processes, and/or pharmacological responses to therapies. Biomarkers can be cellular, biochemical, molecular, etc. Essential for unlocking precision medical treatment, biomarker discoveries are resulting in a paradigm shift in the treatment of specific diseases, particularly in oncology.

In the field of computational biology, AI and ML are now being applied to analyze huge clinical and genomic databases and identify relevant predictive biomarkers for specific types of diseases. Depending on the application, researchers are using a range of data sets (e.g., genomic data, gene expressions, proteomics, clinical data, etc.), and integrate signals and timeframes from this data to develop molecular profiles. In some cases, humans provide canonical disease or drug maps to cover various therapeutic areas and disease types. Hundreds of canonical pathways are analyzed and enriched to infer a disease or drug's mechanisms of actions (MOA). From here, molecular profiling data is fed into algorithms that use ML to identify biomarkers and/or drug sensitivity to specific biomarkers.

Certain biomarker solutions (tests) involve sets of genes and/or proteins in blood. This represents significant potential savings in diagnostic tests, given the relatively low costs and sophisticated resources of blood tests compared to other diagnostic methods like biopsies. This is a promising area of research for precision medicine as biomarkers are highly relevant for diagnoses, treatments, drug discovery, and molecular profiling. Furthermore, identifying risks for preventative care can save millions over time.

Lantern Pharma is working with Intuition Systems to drive precision oncology biomarker identification and drug discovery. Researchers from Lantern will use Intuition System's platform for Big Data analysis; analyzed data will also be co-related to patients' responses for Lantern's clinical stage drugs. The idea is to create molecular profiles based on biomarkers and patients' [favorable] responses to specific treatments. Targeting specific genomic profiles, stratifying and treating such profiles with a narrow scope of clinical trials represents a step toward precision therapies that are especially unmet in the cancer market.

Canadian company **Imagia** is building an artificial clinical intelligence platform to detect and quantify cancer changes. Its Deep Radiomics uses radiomics and DL to analyze clinical imaging data for biomarkers associated with cancer patient outcomes. The platform structures patient information, predicts the characteristics and genetic profiles for specific tumors, and provides evidence for patient prognosis, all of which is designed to trigger appropriately timed diagnostic and therapeutic procedures.

Tractica forecasts that the annual revenue for bio-marker discovery in healthcare will increase from \$1.63 million worldwide in 2017 to \$98.93 million in 2025.

Table 2.131 Bio-Marker Discovery in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	1.63	4.23	8.27	14.40	23.43	36.18	53.21	74.46	98.93	N/A

(Source: Tractica)

2.17.3 CLUSTERING AND PHENOTYPE DISCOVERY

In biology, phenotypes are the composite of an organism's observable characteristics, such as morphology, development, physiological properties, behavior, and manifestations of behavior. Phenotypes result from the expression of an organism's genotype and are influenced by environmental factors that impact both. Today, healthcare practitioners use broad, clinically-driven descriptions to classify phenotypes.

Another application for AI on the road to more personalized medicine is that of inferring precise phenotypic patterns from population-scale clinical data; in other words, allowing large clinical databases to be analyzed so as to precisely show what all phenotypes are and how they progress over time. Using unsupervised learning helps researchers identify patterns (features) that collectively form a compact and expressive representation of source data. Over time, researchers working in this space expect data-driven phenotypes to expose unknown disease variants and subtypes and other genetic associations.

Scientists from Tufts University and the University of Maryland, Baltimore County used AI to gain insight into the biophysics of cancer. Their ML platform predicted a trio of reagents, which the scientists said they would have never considered, which was able to generate a never-before-seen cancer-like phenotype in tadpoles. When treated by the unique set of reagents, pigment cells over the left eye converted to an invasive cancer-like form, while other areas of the tadpole, such as the right eye, remained normal. According to the study, "this is the first time an AI system has been used to discover the exact interventions necessary to obtain a specific novel result in a living organism... and could help human researchers in fields such as oncology and regenerative medicine control complex biological systems."

Tractica forecasts that the annual revenue for clustering and phenotype discovery in healthcare will increase from \$0.77 million worldwide in 2017 to \$46.71 million in 2025.

Table 2.132 Clustering and Phenotype Discovery in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.77	2.00	3.90	6.80	11.06	17.08	25.12	35.16	46.71	N/A

(Source: Tractica)

2.17.4 COMPUTATIONAL DRUG DISCOVERY

Drug discovery is the process by which new medications are discovered. The methods for drug discovery and pharmaceutical research and development (R&D) have largely centered around identifying the active ingredient from traditional remedies by simply by serendipity. Over the last hundred or so years, pharmacology as a field evolved as large chemical libraries and natural products and extract libraries were tested in intact cells or whole organisms to identify effects. Upon sequencing the human genome (which enabled rapid cloning and synthesis of large quantities of purified proteins), it has become common to use high-throughput screening of large compounds' libraries against isolated biological targets. Even still, new drug development costs run about \$2.6 billion per year, take as long as 14 years, and less than 10% of potential medications make it to market, according to Tufts University and the U.S. FDA.

AI offers new ways for researchers to leverage existing databases, develop new databases involving bigger and more diverse data, and to predict how molecules will behave and how likely they are to make a useful drug, thereby saving time and money on unnecessary tests. DL could help with drug development by finding patterns in sparse pathology data combined with large genomic data sets.

Many large pharmaceutical companies are partnering with AI drug discovery startups in a bid to reduce costs and time to market. GlaxoSmithKlein (GSK) recently announced a \$43 million partnership with Exscientia to search for drug candidates for up to 10 disease-related targets. Atomwise recently partnered with drug giant Merck and published first findings of Ebola treatment drugs last year. BenevolentAI is a British company focused on developing better drugs to target diseases of inflammation and neurodegeneration, and rare cancers. The idea is to use much of the dark data within pharma R&D organizations and apply vast data sets available on human health and biological systems to DL systems that learn and reason from interaction between human judgement and data. Numerous other companies are emerging in this space, such as Calico, Numerate, Globavir, NuMedi, twoXAR, and Cloud Pharmaceuticals.

Tractica forecasts that the annual revenue for computational drug discovery in healthcare will increase from \$3.91 million worldwide in 2016 to \$448.16 million in 2025.

Table 2.133 Computational Drug Discovery in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	3.91	11.35	23.09	41.30	68.89	109.46	166.70	243.13	338.45	448.16	69.4%

(Source: Tractica)

2.17.5 CONVERTING PAPERWORK INTO DIGITAL ASSETS

In heavily regulated industries, such as healthcare, documentation is not only required, it is essential for appropriate care and understanding individual medical histories and contexts.

Healthcare will benefit from converting paper documents and other unstructured data, such as email, memos, PDFs, charts, and graphs, into structured data for both operational efficiency and optimizing healthcare services people receive. AI can significantly reduce administrative tasks involved in converting paperwork and processing documentation. NLP, ML, DL, and even bots can be applied in these contexts to both capture paper documents and automate paperwork processing, such as data entry, filling in forms, invoicing, and automating reports on this information, as outlined in Section 2.14.2.

Nuance works to help reduce the time it takes to document medical interactions, by its own claim up to 45%. Its Computer Assisted Physician Documentation (CAPD) uses AI to provide clinical documentation improvement guidance within doctor workflows. The solution offers recommendations to speed up input and uses integrations to process reimbursements, compliance documentation, and so forth. It also offers a Computer-Assisted Clinical Documentation Improvement (CACDI) solution to analyze clinical information in search of areas that may require further clarification to accurately capture the severity or nuance of patients' issues.

Tractica forecasts that the annual revenue for converting paperwork into digital assets in healthcare will increase from \$5.39 million worldwide in 2016 to \$334.71 million in 2025.

Table 2.134 Converting Paperwork into Digital Assets in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	5.39	13.86	26.20	43.96	69.03	103.39	148.49	204.27	268.14	334.71	58.2%

(Source: Tractica)

2.17.6 FACIAL RECOGNITION

Facial recognition is a computer or machine's ability to identify or verify a person based on their facial characteristics. Computer applications use digital images, video frames, and video feeds to recognize people's faces. AI supports facial recognition through various ML and DL techniques, sometimes involving CV. Recognition algorithms are commonly divided into two main approaches:

- **Geometric:** Looks at distinguishing features (face, nose, shape of eyes)
- **Photometric:** Takes a statistical approach by processing an image into values, then eliminates variances by comparing the values with templates

Advancements in processing power and in other adjacent technologies have brought about complementary techniques to enhance facial recognition, outlined in Section 2.9.7. In healthcare contexts, uses for the technology often fall into the following categories:

- Identity authentication, including verification, security validation, anti-fraud
- Patient-doctor check-ins, via video conference
- Medical diagnostics via retina, skin, or other face-based image analysis
- Insurance and risk modeling
- Attendance, check-in/check-out for medical employees

FaceIn offers hands-free cloud-based software that supports physicians and staff members at Florida Heart & Vascular Associates to click-in and out via facial recognition. This

eliminates any fraudulent attendance and has the added benefit of minimizing germ transmission. Other companies, such as Compumatic, Fareclock, and uAttend, support facial recognition for employee time tracking.

Facial recognition for medical diagnostics and treatment also shows promise. In 2012, a team of researchers launched a 5-year National Institutes of Health (NIH)-funded project to determine whether pediatric patients' pain could be accurately measured by facial recognition software. The software was programmed to recognize 20 muscle movements known to indicate pain, then models were trained to measure pain based on images of patients in pain. Then it performs a regression analysis of the levels of intensity in those patients who displayed signs of pain. The software applies the data to measure pain in other subjects, using a 0-to-10 scale. The software's assessment of pain came closer than nurses' assessments to those self-reported levels. Emotient, the software used to underlie the study, was acquired by Apple in 2016.

Researchers from the National Human Genome Research Institute (NHGRI) recently published findings showing it successfully used facial recognition software to diagnose DiGeorge syndrome (or velocardiofacial syndrome), a rare genetic disease that affects Africans, Asians, and Latin Americans. The success of the technique is especially notable because the disease can manifest across multiple defects in the heart, cleft palate, learning issues, etc., making it difficult for clinicians to detect in diverse populations. Based on 126 individual facial features, researchers made correct diagnoses for all ethnic groups 96.6% of the time.

Tractica forecasts that the annual revenue for facial recognition in healthcare will increase from \$0.89 million worldwide in 2017 to \$42.29 million in 2025.

Table 2.135 Facial Recognition in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.89	2.22	4.22	7.14	11.27	16.90	24.13	32.79	42.29	N/A

(Source: Tractica)

2.17.7

GENOMIC DATA MAPPING AND ANALYSIS FOR PERSONALIZED HEALTHCARE AND PRECISION MEDICINE

As tools and research in modern healthcare have evolved, so have the industry's aspirations for more personalized medical services, products, and experiences. Perhaps no other greater advancement than the Human Genome Project, which mapped the human genome, has opened our eyes to the potential for precision medicine. But mapping was only the beginning; for the last 16 years, researchers have been working to analyze DNA at scale. But this too is in the early stages. "We have vast amounts of data; three billion data points per individual," explains Stephan Sanders, assistant professor at UCSF School of Medicine. "What we have less of is the other end: clean data of phenotypes or outcomes."

DL has been powering much of the ML-driven genomic data analysis to date, as researchers use it to explore areas such as gene splicing, epigenetics, and genetic causes for disease. Algorithms are fed thousands of chunks of DNA, along with proteins coded from those sequences, and after seeing thousands of these examples, AI helps assess variations, such as gene mutations, and detect when and why certain processes go astray.

One of the most well-known players in this space is Toronto-based startup Deep Genomics. Founded by Brandan Frey, the company is using DL to better understand how genomic realities, alterations, and variations across individuals and populations manifest as diseases. There is over a decade of work that has gone into Deep Genomics' DL algorithm development: it began by teaching the computer how to read the basic genetic code, associate sequences with the corresponding ribonucleic acid (RNA) and protein outputs, and from there have been working on identifying the manifestation of a disease. It has published research showing how DL can help identify patterns in DNA that might contribute to diseases like spinal muscular atrophy and nonpolyposis colorectal cancer.

In what Frey coins “a Google search engine for genomics,” the company is building out a database in which a user could eventually enter a combination of mutations found in a patient and the model should output the likelihood and severity of specific diseases. The company also aims to use genomic data to help with drug development that addresses the behavior of faulty genes. Its goal is to better understand diseases, disease mutations, and gene therapies, eventually using these findings to inform precision medicine and personalized therapies.

Longer-term applications for genomic data mapping and analysis have to do with incorporating genomic data into broader topologies for how genes interact with the environment. Researchers have some idea for how to integrate early findings and incorporate other contexts, such as biometric data, daily habits, and behavioral data, medications, and so forth. A company called iCarbonX is working towards this by offering a digital health management platform based on a combination of behavioral, biological, and psychological data. The company partnered with research institutions, pharmaceutical factories, hospitals, insurance companies, and other health management providers to offer a patient-facing platform called Meum. Meum runs on a massive database combining ‘panoramic’ life data, including genetics, molecular profiles, phenotype, time series interactions, and beyond to offer patients personalized solutions and recommendations in both medical and wellness areas like skincare, exercise, weight management, and so forth. The company is working towards diverse partnerships in order to build out even more personalized and precision and streamline medical research at scale.

Tractica forecasts that the annual revenue for genomic data mapping and analysis for personalized healthcare and precision medicine will increase from \$13.3 million worldwide in 2016 to \$207.8 million in 2025.

Table 2.136 Genomic Data Mapping and Analysis for Personalized Healthcare and Precision Medicine in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	13.30	17.37	23.53	32.71	46.12	65.15	91.03	124.30	164.09	207.80	35.7%

(Source: Tractica)

2.17.8 HOSPITAL PATIENT MANAGEMENT SYSTEM

Hospitals and medical clinics are not just managing staff, suppliers, and infrastructure, they are also held to high regulatory standards for managing patients' data for billing, treatment, checking-in, checking-out, what drugs or procedures are administered, etc. Ensuring the right patients receive the right information, direction, and treatment has serious consequences. While much of the investment and AI-driven development in patient records involves patient data processing, as outlined in Section 2.17.16, Tractica also identifies

hospital patient management as another area where AI can be applied. This refers to efficiently managing patients in and out of the hospital. Healthcare providers are under growing pressure to identify and enforce best practices that efficiently deliver high-quality care across an entire patient experience, the ability to learn from, and in certain cases, automate certain workflows. Often, this involves feeding algorithms electronic medical records (EMR), financial, and insurance data to analyze and predict outcomes of specific treatment or surgical procedures based on past contexts.

Ayasdi uses AI to analyze EMR and financial data across thousands of procedures and millions of patient events, unsupervised and semi-supervised learning to automatically resurface groups of similar patient procedures and recommend specific clinical pathways at the lowest costs for local patients. The company's solution also includes justifications with details about each input to its recommended pathways. These are compared to existing guidelines integrated into analytics to monitor adoption and adherence with standardized clinical pathways and identify new trends. Wellframe, Zephyr Health, and many others are working on various aspects of patient management system automation and analytics.

Tractica forecasts that the annual revenue for hospital patient management systems in healthcare will increase from \$1.62 million worldwide in 2016 to \$147.19 million in 2025.

Table 2.137 Hospital Patient Management Systems in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.62	4.67	9.29	16.16	26.21	40.46	59.83	84.73	114.50	147.19	65.1%

(Source: Tractica)

2.17.9

MARKET INTELLIGENCE FOR LIFE SCIENCES

As in most any business setting, life sciences companies need market intelligence and CRM tools to help wield Big Data associated with their prospects, customers, products, etc. As medicine has grown increasingly digital at every level, new markets are opening up to support the use of this data beyond medical treatment and diagnostic contexts. AI-powered market intelligence for life sciences is about using ML, NLP, and DL to power Big Data analytics for strategic marketing, business development, sales, and customer engagement. Insights mined from big disparate data sources do not just help with smarter prospecting or engagement, but AI is also powering recommendation engines, personalized communications, predictive customer insight bots, and advanced data visualization involving large amounts of disconnected data.

Zephyr Health is a data analytics company that uses global health data to support life sciences companies (pharmaceutical, biotech, medical device, diagnostic firms, etc.) with strategic engagement. The platform integrates thousands of data sets (e.g., geographic trends, hospital profiles, research programs, publications, physician affiliations, prescription behavior, channel preferences, account penetration, drug trials, competitor sales, etc.) to deliver highly-targeted insights and recommendations. It is using ML algorithms to power a variety of proprietary applications purpose-built for life sciences sales, marketing, customer engagement, and productivity enhancement. Some examples include product launch planning, identifying new market opportunities, accelerating field sales via targeted recommendations, prioritization of high potential customers and segments, and more.

Tractica forecasts that the annual revenue for market intelligence for life sciences in healthcare will increase from \$0.36 million worldwide in 2017 to \$16.96 million in 2025.

Table 2.138 Market Intelligence for Life Sciences in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.36	0.89	1.69	2.86	4.52	6.78	9.68	13.15	16.96	N/A

(Source: Tractica)

2.17.10 MEDICAL DIAGNOSIS ASSISTANCE

Since the dawn of medicine, education and clinical training have focused on how to diagnose and treat ailments. Medical diagnostics accounts for billions in spending every year, as doctors and patients pursue tests required to identify the issue before they can even begin treating it. In the United States alone, as many as 40,500 patients die annually in an intensive care unit (ICU) as a result of misdiagnosis, according to a 2012 Johns Hopkins study.

Penetration of AI into medical diagnostics has the potential to not only enhance doctors' speed, accuracy, and preventative strategies, but to advance society's collective understanding of the body and medical treatment. For centuries, doctors have been using one-on-one medical diagnoses by matching patients' symptoms to various lists, common effects, or frameworks associated with diseases. Seasoned doctors surely offer their experience, intuition, and extra training to diagnoses and treatment plans—expertise that algorithms may never quite match—but there remain tremendous errors in medical diagnoses, or diagnoses come too late. At least 80% of cancers could be effectively treated if detected earlier.

Approaching this problem using NLP and DL involves feeding medical records and images into neural networks and algorithms begin to detect patterns and abstractions, not just across symptom-disease associations, but across diseases, patients, geographies, environments, etc. The ability to take in, retain, analyze, and learn from so much diverse data simply transcends human capability and bandwidth. Doctors of all types will increasingly begin to leverage AI-generated inputs in their diagnostics.

Researchers from Sutter Health and the Georgia Institution of Technology demonstrated that, upon analyzing EHR using neural networks, they were able to predict heart failure as early as 9 months before doctors. Freenome is tackling the problem of cancer diagnosis by using DL to detect cell-free DNA sequencing of cancer in the blood. The model clustered characteristics by location, which helps scientists and doctors pinpoint where cancer is growing in the first place; a critical part of the puzzle. DeepMind Health in the United Kingdom has acquired data from the National Health Service (NHS) to allow its algorithms to look for early warning signs for specific conditions like Acute Kidney Injury (AKI).

A number of adjacent use cases will also frame AI's ability to aid in medical research, diagnostics, treatments, etc., with clustering and phenotype discovery, bio-marker discovery, treatment recommendations, genomic mapping, virtual assistants for patients, and beyond. It is also worth noting that a host of bioethical and ethical issues could arise, particularly around genomic targeting, less clinical drug testing, explainability of systems, inadvertent erroneousness, etc. Moreover, the standards in medicine are very high, which contributes to a bias within the profession against innovation. Even if these digital diagnostic tools are able to reach a 99.999% success rate, they will never be perfect, and mistakes due to false readings could lead to medical malpractice lawsuits and product liability issues.

Tractica forecasts that the annual revenue for medical diagnostic assistance in healthcare will increase from \$3.6 million worldwide in 2016 to \$180.58 million in 2025.

Table 2.139 Medical Diagnostic Assistance in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	3.60	6.61	11.34	18.63	29.65	45.83	68.62	99.04	136.95	180.58	54.5%

(Source: Tractica)

2.17.11 MEDICAL IMAGE ANALYSIS

Analyzing medical images was a task left to doctors and radiologists until only very recently. Instruments supporting medical imagery really only emerged in the 1940s and 1950s, starting with the camera. Historically, analyzing medical images has been difficult, highly prone to human error or oversight, and time-consuming and costly. Medical images like magnetic resonance imaging (MRIs), X-rays, computed tomography (CT) scans, and other diagnostic images are essential to better understanding and diagnosing a wide range of conditions. When it comes to diagnosing critical conditions, including cancer, neurodegeneration, and heart disease, the faster and smarter the speed, precision, and predictive capabilities, the better.

Analyzing images is a strong application for DL and CV within the realm of patient data processing. In particular, DL is now being applied to automate the analysis and increase accuracy, precision, and understanding of images down to the pixel. Some of the more common applications include:

- **3D Computer Vision:** Images are analyzed and highly detailed 3D models are then rendered.
- **Autograding of Eye Diseases:** Image recognition is able to detect specific kinds of eye diseases (e.g., macular degeneration, those associated with diabetes, etc.).
- **Detection and Segmentation of Radiology Images:** Millions of radiology images, often in 3D, are fed into neural networks, enabling them to segment them by organ, tumor, tissue, etc.

Enlitic uses DL networks that analyze medical imaging data, such as X-rays and MRIs, to identify even the smallest suspicious clues (e.g., tumors, hairline fractures, spots, etc.). Its networks increase diagnostic accuracy in less time and at a reduced cost compared to traditional diagnostic methods. Enlitic's software also allows comparison of an individual patient's radiological data with millions of other patients who received the same diagnosis in order to identify and track treatment outcomes for the most similar cases.

In June of 2017, Google's DeepMind announced a long-term project in which ML algorithms parse millions of retina and eye scans to tease out early warning signs that human doctors might otherwise miss. "There's so much at stake, particularly with diabetic retinopathy," says DeepMind co-founder Mustafa Suleyman. "If you have diabetes you're 25 times more likely to go blind. If we can detect this, and get in there as early as possible, then 98% of the most severe visual loss might be prevented."

Tractica forecasts that the annual revenue for medical image analysis in healthcare will increase from \$0.07 million worldwide in 2016 to \$1.523 billion in 2025.

Table 2.140 Medical Image Analysis in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.07	31.99	80.24	152.18	257.27	406.31	609.02	869.58	1,181.24	1,523.42	202.8%

(Source: Tractica)

2.17.12 MEDICAL TREATMENT RECOMMENDATION

In medicine, deciding on appropriate treatment options for patients is often a highly complex and risky endeavor. Doctors and internists must take into account a wide range of considerations not only in the manifestation of symptoms, but in a patient's environmental and genealogical contexts, relevant medical research, and impacts of treatment and/or medication on patients, not to mention risks of malpractice or unforeseen vectors. This is in the most basic sense what IBM calls "the doctor's dilemma—too much information." In addition to diagnostic assistance, AI is also helping doctors in determining optimal treatment plans. More and more medical and research institutions are leveraging NLP and DL for (big) data analysis in the name of driving faster and more precise treatment.

The Pediatric ICU of the Children's Hospital of Los Angeles is currently using recurrent neural network (RNN) and CNN DL to analyze 10 years of EHR, across 20,000 patients in order to simulate and develop better treatments, create illness profiles, and observe patient outcomes over time. "Our overarching goals are to keep more kids alive, to reduce the length of their stays as well as morbidities and ancillary effects," explains David Ledbetter, of the Children's Hospital of Los Angeles. "But we also aim to be an augmentation to doctors by mining for collective wisdom: Wisdom from over roughly 10,000 years' worth of patient data as well as by analyzing the state-of-the-art information to recommend personalized treatments for particular patients at particular points to optimize their outcome."

Longer-term applications involve using diverse data sets for medical research, drug and treatment development, and preventative care. Integrating patient data with its AI health tool enables IBM's Watson Health to mine patient data to find relevant facts about family history, current medications, or any pre-existing conditions, providing alerts or early warning signs through its system. IBM's Watson computer is currently in use by oncologists at Memorial Sloan-Kettering Cancer Center in New York. IBM's software draws from 600,000 medical evidence reports, 1.5 million patient records and clinical trials, and 2 million pages of text from medical journals to help doctors develop treatment plans tailored to patients' individual symptoms, genetics, and histories.

U.K.-based Babylon recently launched an AI-based app, into which users report (via text or speech) the symptoms of their illness into the ML/DL-fueled symptom checker to receive accurate medical advice. Unlike IBM's efforts, Babylon covers illnesses beyond cancer. The app mines a patient's history, genetics, behavior, biology, environmental circumstances, and checks them against disease databases, running analysis on hundreds of millions of combinations of symptoms. While current U.K. regulations prohibit AI from making diagnostics, the app recommends appropriate courses of action, typically including booking a doctor's appointment or considering over-the-counter medications. As of June 2017, for £4.99 (\$7.10) per month, its ~250,000 users can book appointments, and consult with one of about 100 doctors 12 hours a day, 6 days a week. Over time, the app aims to integrate with wearable devices to include ongoing/real-time data inputs around heart rate, sleep patterns, cholesterol levels, and other biometrics into its algorithms.

While these applications show great promise, particularly at scaling the ability to monitor in detail individuals' health and analyze massive amounts of data in seconds, a host of risks remain in the technology. Misdiagnosis, over-treating, under-treating, increased office visits, cognitive biases, cultural frictions, and even legal implications are just the tip of the iceberg. Nonetheless, these efforts may prove an important tool in long-term efforts to enable healthcare that is not just curative, but moving toward a more preventative model through precision medicine.

Tractica forecasts that the annual revenue for medical treatment recommendation in healthcare will increase from \$5.18 million worldwide in 2016 to \$303.78 million in 2025.

Table 2.141 Medical Treatment Recommendation in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	5.18	10.24	18.19	30.48	49.06	76.35	114.81	166.14	230.14	303.78	57.2%

(Source: Tractica)

2.17.13 MEDICATION COMPLIANCE FOR CLINICAL TRIALS AND GENERAL USAGE

Medical adherence, or complying with a clinical mandate to follow prescribed orders for treatment and rehabilitation, are notoriously difficult to enforce. Fundamentally, if researchers are not certain patients are taking the medications or partaking in essential treatments of the clinical study, they cannot know if the data collected from such tests is good data. It is estimated that medication adherence is around 50% per trial, and some 30% of clinical trials fail. What is estimated to be a \$300 billion challenge has been a core focus of the healthcare industry as it embraces new technologies. From direct mailers and text messages to mobile app notifications, and even wearable alerts, clinicians have been experimenting with new ways to ensure medication adherence for years.

AI offers another approach to medication adherence: by leveraging image recognition and basic smartphone functionalities, clinicians can offer an easy-to-use way for patients to register their adherence, while learning from trends in the data. AI could also be applied to analyzing data sets beyond the image recognition, pulling in biometric data from wearables, such as nutrition data or sleep data. As some 90% of patient behavior is unknown in outpatient settings, this offers the added benefit of visibility into outpatient behaviors beyond just taking medications.

AICure is focused squarely on tackling the issue of medical adherence in clinical trials, a \$15 billion problem for pharmaceuticals. It uses facial recognition and motion sensing technologies in mobile devices to visually confirm medication ingestion (by the right person at the right time). Using [Health Insurance Portability and Accountability Act (HIPAA) compliance] facial recognition, automatic medication identification (via image recognition), and real-time ingestion confirmation, the patient's experience is comparable to using a smartphone to take a picture. The app can be downloaded on any device, and customized to patient demographics, disease type, and communications preferences. Caretakers and medical providers have access to real-time analytics to ensure adherence, while algorithms help identify those at highest risk of poor adherence or need for hospitalization. The tool also helps prevent fraud and duplicate enrollments in clinical trials.

Tractica forecasts that the annual revenue for medication compliance for clinical trials and general usage in healthcare will increase from \$2.53 million worldwide in 2016 to \$66.09 million in 2025.

Table 2.142 Medication Compliance for Clinical Trials and General Usage in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.53	3.64	5.37	8.01	11.98	17.80	25.98	36.88	50.46	66.09	43.7%

(Source: Tractica)

2.17.14 METHODS FOR MONITORING VITALS

Medical data can be non-time series-based like patient notes, images, medications, allergies, and demographics, or it can be both multi-dimensional and time series data like heart rate, blood pressure, glucose levels, or other vitals. This kind of data has typically been gathered in hospitals, while medical professionals are monitoring or nearby, but advancements in wearables and in-home medical services are enabling new methods for monitoring vitals.

Using wearable data inputs, via bracelets, heart monitors, patches, sensor-enabled clothing, or other body sensors, data flowing from these devices can aid in new ways to monitor vitals, both from within hospitals and care facilities or remotely in patients' homes. Medical providers can leverage tools that use AI and ML to both analyze multi-dimensional time series data, and identify anomalies. While the "next step" capability to deploy or execute some sort of treatment or medical intervention remotely is possible, it remains early days, given the nascence of the technology and potential for false positive. In the meantime, patients can enjoy greater flexibility, while medical providers can access more visibility, particularly into outpatient vitals.

ZOLL's LifeVest is a wearable defibrillator offered to patients at risk of sudden cardiac arrest (SCA). The vest continuously monitors patients' cardiac state via electrocardiogram (ECG), photoplethysmogram (PPG), body temperature, blood pressure, galvanic skin response (GSR), and heartrate sensors. If a life-threatening arrhythmia signature is detected by onboard algorithms, the vest itself delivers a remote treatment resuscitation to restore a normal heart rhythm.

Tractica forecasts that the annual revenue for methods for monitoring vitals in healthcare will increase from \$1.18 million worldwide in 2017 to \$71.51 million in 2025.

Table 2.143 Methods for Monitoring Vitals in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	1.18	3.06	5.98	10.41	16.94	26.15	38.46	53.82	71.51	N/A

(Source: Tractica)

2.17.15 MINING, PROCESSING, AND MAKING SENSE OF CLINICAL NOTES

A significant portion of health record data, particularly context-rich physician and nurse's notes, lab reports, and discharge summaries, which are collectively called clinical notes, is unstructured text. Clinical notes are a sort of subset of broader medical-related paperwork, but constitute a host of challenges to processing digitally to extract longer-term value. NLP with ML and DL software is being used to mine clinical notes to extract and connect information to provide much improved patient analytics and a more comprehensive view of

healthcare. Often, notes are audio transcribed, and involve critical details that are essential for diagnosis and improving personalized care.

Research initiatives and some commercial deployments are underway. However, the path to clinical note mining will be challenging. In a paper published by Notre Dame researchers in 2H 2016, [*Mining The Clinical Narrative: All Text Are Not Created Equal*](#), researchers noted this about current clinical note mining initiatives:

Most of the systems have been developed specifically for specialized applications and for limited domains. Recent work has attempted to expand the scope of these techniques through the utilization of linguistic tools such as improved lexicon, and complex grammars. However foundational work done by Harris has already established the existence of what are known as sublanguages: "specialized domains that exhibit specialized constraints due to limitations of the words and relations of the subject matter.

The assumption that all text extracted from the EMR can be consumed and analyzed in the same manner, regardless of its source, is limiting. The NLP techniques, on which these multi-source systems are based, process data in a statistical manner, thus their ability to produce reliable output is highly dependent on the underlying data. It then stands to reason: if the sources of clinical text are in some way fundamentally different, no high-level linguistic tool will provide an accurate or effective model.

Some companies have commercial offerings in the space, including Hindsait, CloudMedx, and CareCentra.

Tractica forecasts that the annual revenue for mining, processing, and making sense of clinical notes in healthcare will increase from \$0.58 million worldwide in 2017 to \$24.40 million in 2025.

Table 2.144 Mining, Processing, and Making Sense of Clinical Notes in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.58	1.42	2.66	4.43	6.90	10.19	14.34	19.20	24.40	N/A

(Source: Tractica)

2.17.16 PATIENT DATA PROCESSING

The generation, input, processing, analysis, security, compliance, and utilization of patient data create massive challenges to healthcare organizations the world over. Moving patients through hospital systems requires and generates massive amounts of paperwork, documents, and data. This is, of course, not to mention the ability to learn from and use such data predictively. Indeed, the greatest challenges when it comes to patient data center around processing and analytics at scale.

The complexity, dysfunction, overwhelming amount of unstructured data, and lack of standards in the world's healthcare systems make patient data processing a ripe application for AI. Patient data processing will make use of both ML and DL in combination with NLP. When it comes to **administrative processing**, more and more medical institutions are leveraging DL for data analysis in the name of driving faster and more precise treatment. Institutions conducting medical research are also taking advantage of these techniques for

clinical processing of patient data. DL is also being applied to analyze medical images and can aid doctors as they analyze images.

SyTrue uses NLP and ML to support patient data processing. The tool helps integrate disparate sources of information to produce a comprehensive and “deep-dive” view of patient groups in physician practices, hospitals, and other healthcare institutions. The first phase of data processing is cleaning, wherein it uses ML to identify zones of information, such as a document header, footer, etc. Then SyTrue uses a semantic rules engine for a second phase of cleaning. “U.S. healthcare has 20 different coding schemes, so you need something to know how to sort that out. NLP will naturally extract data in context. Then you can add a rules base or ML approach to it, and then after that you could apply DL,” said Kyle Silvestro, the company’s CEO, chairman, and founder. SyTrue is developing models to support medical diagnostic assistance and treatment recommendations, as the goal is to offer a “longitudinal view,” powering data processing across every healthcare interaction. Xerox, Conduit, TransPortal, Abed-Graham Healthcare Strategies, and Neo4J are among the company’s clients and partners.

The Pediatric ICU of the Children’s Hospital of Los Angeles is currently using RNN and CNN DL to analyze 10 years of EHR, across 20,000 patients in order to simulate and develop better treatments, create illness profiles, and observe patient outcomes over time. “Our overarching goals are to keep more kids alive, to reduce the length of their stays as well as morbidities and ancillary effects,” explains David Ledbetter, of the Children’s Hospital of Los Angeles. “But we also aim to be an augmentation to doctors by mining for collective wisdom: wisdom from over roughly 10,000 years’ worth of patient data as well as by analyzing the state-of-the-art information to recommend personalized treatments for particular patients at particular points to optimize their outcome.”

Much of the latest innovation focuses on using decision trees and neural networks around patient data to improve fraud detection, claims processing, scanning (analog or digital) patient records, marketing, behavioral analysis, and preventive insurance. Unlike fixed statistical models, dynamic models using AI adapt to shifting parameters, making areas like fraud detection and claims processing self-learning and far more cost-effective than current models.

Longer-term applications for DL and patient data involve using diverse data sets for medical research, drug and treatment development, and preventative care. Integrating patient data with its AI health tool enables IBM’s Watson Health to mine patient data to find relevant facts about family history, current medications, or any pre-existing conditions, providing alerts or early warning signs through its system. DeepMind Health in the United Kingdom has acquired data from the NHS to allow its algorithms to look for early warning signs for specific conditions like AKI.

Another adjacent application for patient data is in genomic data mapping and analysis, outlined in Section 2.17.7.

Tractica expects patient data processing using AI to become much more commonplace by 2025, unleashing creative ways of understanding patient groups and health conditions, identifying hidden efficiencies in healthcare, and improving precision medicine. Tractica forecasts that the annual revenue for patient data processing in healthcare will increase from \$9.74 million worldwide in 2017 to \$465.14 million in 2025.

Table 2.145 Patient Data Processing in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	9.74	24.46	46.41	78.49	123.99	185.88	265.44	360.62	465.14	N/A

(Source: Tractica)

2.17.17 PORTABLE AND LOW-COST ULTRASOUND DEVICE

Medical equipment is notoriously expensive, but sometimes emerging technologies converge to radically alter cost structures and access. Devices like ultrasounds have historically only been accessible to medical professionals, in hospital settings, and require extensive training. Today more than 60% of the world lacks access to medical imaging.

A portable and low-cost ultrasound device would help “democratize” access to one of the most important diagnostic tools in medicine. Such a tool is being developed by Butterfly Networks, which claims to use AI to reinvent the ultrasound by putting all typical ultrasound components on a single silicon chip. Onboard the chip are DL algorithms trained by ultrasound experts. The resulting imager will also be more portable than any existing ultrasound on the market today, and learning algorithms means the device will require far less training to use effectively and interpret results. The effort is led by Dr. Jonathan Rothberg, who has helped other medical device startups including Clarifi, RainDnace, Ion Torrent, CuraGen, and 454 Life Sciences.

Tractica forecasts that the annual revenue for portable and low-cost ultrasound devices in healthcare will increase from \$.02 million worldwide in 2016 to \$4.19 million in 2025.

Table 2.146 Portable and Low-Cost Ultrasound Devices in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.02	0.09	0.20	0.37	0.63	1.01	1.54	2.26	3.16	4.19	81.9%

(Source: Tractica)

2.17.18 PREDICTING ILLNESS AND PATIENT OUTCOMES

Beyond a healthy diet, good sleep, and exercise, there has historically been little focus on *preventing* illnesses as integrated into healthcare regimes. Instead, we react: a symptom appears and we respond. One of the overarching goals technological integration in healthcare, is to enable preemptive or preventative care, instead of reactive care. AI, in conjunction with numerous other technologies like mobile, wearables, voice interactions, video, social media, and genomic mapping, have the potential to help predict illnesses and patient outcomes. In both cases, an AI engine combined with extensive medical knowledge covering thousands of conditions, symptoms, findings, and cases could surface patterns that would be impossible for a human to detect. Solutions can offer preemptive ways to:

- Take action toward early intervention or treatment selection
- Design personalized medical policies based on probabilities
- Reduce treatment variation and improve outcomes
- Predict individual responsiveness to treatment in both R&D and post-market contexts

Even in the case of predicting patient outcomes, in which AI could help offer some statistically-based prediction on drug reactions, adherence, speed to recovery, etc. could offer opportunities to save costs, increase personalization, etc. Evidation Health analyzes health outcomes data, while generating real-world clinical and economic evidence in order to identify and deploy the most effective and efficient interventions for patients. Its approach uses ML to help support treatment and rehabilitation decisions with large empirical data sets, not just intuition or available tools. Another company, Counsyl, takes a very “preemptive” approach to predicting illness. It provides a platform for couples to submit their DNA prior to conceiving a baby. The platform offers probabilistic risk of 100+ health conditions that could be passed from parents to children.

The primary hurdle in this area is the time it takes to improve predictions based on feedback loops. As Christine Lemke, President and co-founder of Evidation Health explains, “Within seconds, Google knows whether its search engine prediction is correct. But in healthcare, the feedback loop—which is often measured in terms of impact on biometric or cost outcomes—can take years.” Second, such analytics, and a host of use cases outlined in this section signal shifting power dynamics between physicians and patients. If physicians do not readily adopt predictive analytics for fear of losing decision-making power or liability concerns, then patients will go straight to the algorithms to find their own answers. But consumers need to trust the accuracy and reliability of these recommendations.

Tractica forecasts that the annual revenue for predicting illness and patient outcomes in healthcare will increase from \$0.75 million worldwide in 2017 to \$45.49 million in 2025.

Table 2.147 Predicting Illness and Patient Outcomes in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.75	1.94	3.80	6.62	10.77	16.63	24.47	34.24	45.49	N/A

(Source: Tractica)

2.17.19 TEXT CLASSIFICATION AND MINING FOR BIOMEDICAL LITERATURE

According to the “[Biomedical Literature Mining for Biological Databases Annotation](#)” research paper:

In biomedical research, there are thousands of specialized data repositories, focusing on particular molecules, organisms or diseases, which offer sets of richly annotated records. To ensure data of the highest quality, manual data entry and curation (annotation) processes are generally performed on these databases. Database curators are domain experts who search biomedical research literature for facts of interest, and manually transfer knowledge from published papers to the database. This helps experts to consolidate data about a single organism or a single class of entity, often in conjunction with sequence information. Most importantly, this process makes the information searchable through a variety of automated techniques, given that the curators use standardized terminologies or ontologies. However, as the volume of biomedical literature increases, so does the burden of curation, making annotation databases incomplete and inconsistent with the literature. It has been shown empirically that manual annotation cannot keep up with the rate of biological data generation (Baumgartner et al., 2007)... This motivates the upsurge of interest in text mining techniques which enable various degrees of automation in the analysis of scientific literature, such as identification of named entities, classification of documents, extraction of relevant facts (i.e., relationships

between two or more named entities expressing a fact), and generation of hypotheses (Cohen & Hersh, 2005; Jensen et al., 2006; Krallinger et al., 2005).

Scholars have choices between four resources, including Microsoft Academic, Google Scholar, Baidu Scholar, and Paul Allen's Semantic Scholar. As of June, 2017, through the Allen Institute for Artificial Intelligence, a cross-platform sharing of metadata, user behavior data, and other resources will take place within the Open Academic Search (OAS) working group.

A commercial product called Qinsight by Quertle was launched in September 2016. The company claims the solution covers more than 40 million documents, including searching the full text of 10 million. The content "includes, essentially, all biomedical and biological journals, patent grants and applications, NIH grant applications, TOXLINE databases, AHRQ treatment protocols, and more," according to the press release.

Tractica forecasts that the annual revenue for text classification for biomedical literature in healthcare will increase from \$0.5 million worldwide in 2017 to \$30.47 million in 2025.

Table 2.148 Text Classification for Biomedical Literature in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.50	1.30	2.55	4.44	7.22	11.14	16.39	22.94	30.47	N/A

(Source: Tractica)

2.17.20 VIRTUAL ASSISTANTS FOR DOCTORS

Doctors have long relied on other colleagues—doctors, nurses, specialists, etc.—to offer feedback, validation, or recommendations to support decision-making. The stakes for informed decision-making in healthcare are among the highest in any industry, not just because of the critical (sometimes life or death) nature of many healthcare applications, but given the amount of data, regulations, providers, and money involved.

As the market for virtual assistants grows in brand and consumer contexts, similar techniques are starting to crop up in healthcare, supporting both doctors and patients. Using NLP, ML, and DL, as well as complementary techniques in voice recognition, image recognition, potentially even CV, the objective of virtual assistants is not to replace doctors, but to expedite decision-making by basing analysis on more data sources taking into account thousands of other cases, faster.

To support doctors and clinicians, medical diagnostic app Babylon (outlined in Section 2.17.12) plans to help doctors to "accurately identify the disease and the most appropriate treatment" through what it claims as the largest curated knowledge graphs of medical content. In addition to fueling the conversational interface, the NLP engine turns text and speech into structured data, transcribing consultations, and summarizing clinical records.

In 2014, speech technology provider Nuance launched a pilot called Florence, a virtual assistant for doctors. The software was designed to help physicians update EMR faster and to anticipate further actions based on prior work. Florence used speech recognition, allowing doctors to speak to the program instead of typing. It also helped to automate and streamline CPOE. According to Dr. Anthony Sagel, Chief Medical Officer at Landmark Hospitals, Florence reduced his time entering orders by 35% during his 6-month trial period. However, as of June 2017, Nuance has not launched or announced the commercial availability of

Florence.

Tractica forecasts that the annual revenue for virtual assistants for doctors in healthcare will increase from \$0.04 million worldwide in 2016 to \$19.01 million in 2025.

Table 2.149 Virtual Assistants for Doctors in Healthcare, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
0.04	0.50	1.17	2.15	3.54	5.47	8.04	11.26	15.01	19.01	99.7%	(Source: Tractica)

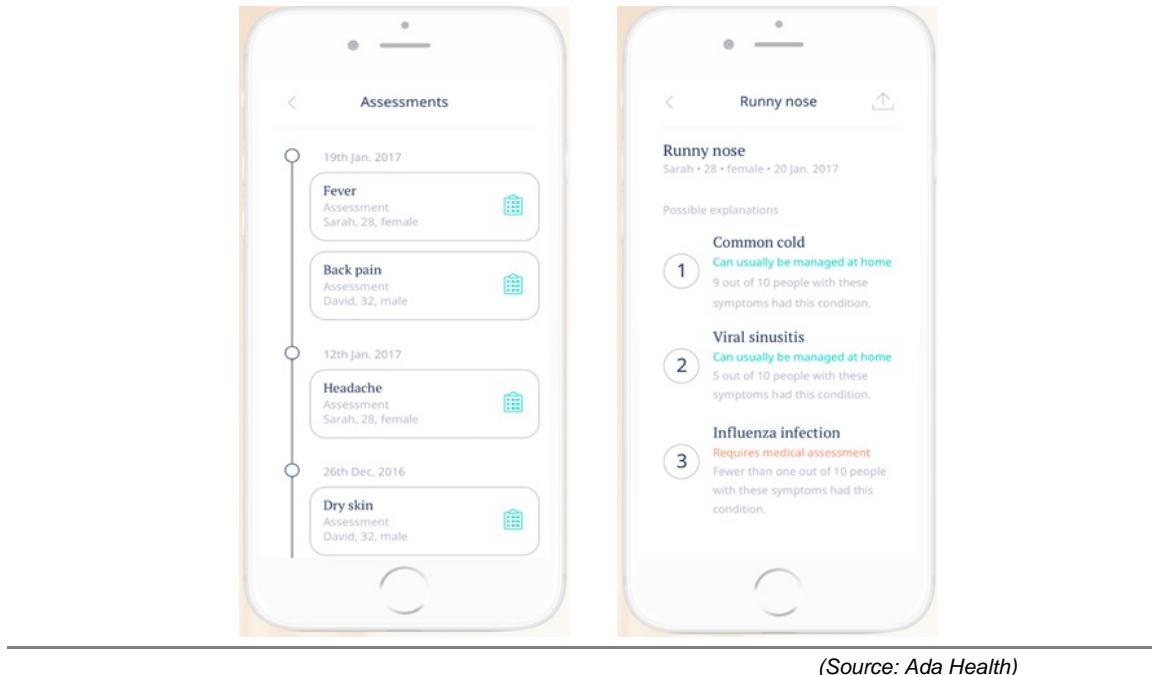
2.17.21 VIRTUAL ASSISTANTS FOR PATIENTS

Despite radical advancements in healthcare technologies, diagnostics, medicine, and treatments in the 20th century, the truth is millions of people have little to no relationship with their doctors or healthcare providers. Chronic cost and labor constraints in the healthcare industry have limited one-to-one relationships and the quality of care possible, as many doctors are required to see dozens (or more) patients in a single day.

Virtual assistants could provide some relief to this, as AI can be trained to mine large data sets and deliver advice, triage questions, promote medication adherence, or facilitate appointment scheduling for individual patients. Using NLP, DL, ML, and potentially CV (using patients' mobile device cameras for instance), virtual assistants are not likely to replace human doctors, but can scale their ability to provide guidance.

Ada offers a "telemedicine" mobile app-based virtual assistant that uses ML, NLP, and image recognition to support patients' understanding (current regulations prohibit formal AI-based diagnostics) and the ability to self-care. Patients can fill out a robust personal assessment, which becomes more personalized with each interaction, ask questions, indicate symptoms and severity, and even chat live with doctors and schedule appointments.

Figure 2.18 Ada Health App Delivers Virtual Assistance for Patients



(Source: Ada Health)

The app, which recently announced integration with Amazon's Alexa, also offers doctor assistance; the company claims doctors love it because it can collect important details they might miss or patients might forget to mention. Since its launch, the app has successfully diagnosed both common and rare conditions, and because its continuous training includes real human doctors, it pools shared expertise. The broader objectives are to help ensure patients can be proactive with their healthcare, avoiding unnecessary visits, while also making more informed decisions when symptoms need doctors' attention. Either way, time is saved, as the app serves as a pre-screening agent and helps create a digital paper trail prior to consultation.

Tractica forecasts that the annual revenue for virtual assistants for patients in healthcare will increase from \$13.33 million worldwide in 2016 to \$1.244 billion in 2025.

Table 2.150 Virtual Assistants for Patients in Healthcare, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	13.33	42.64	85.80	148.65	238.51	363.32	529.55	738.60	982.75	1,243.60	65.5%

(Source: Tractica)

2.18 INFORMATION TECHNOLOGY

2.18.1 AUTOMATED CODE DEVELOPMENT

Let not the irony be lost that those who develop code and software, including myriad AI technologies, may one day be displaced by AI itself. Given the potential for greater scale, fewer errors or bugs, and improved speed and costs associated with development, numerous researchers are working on AI-powered solutions to automatically develop code.

DL, in particular, is designed to extract knowledge from large datasets in order to apply its learnings to specific applications or situations.

Researchers from Microsoft and the University of Cambridge have recently developed a system called DeepCoder that uses a program synthesis technique. Program synthesis pieces together lines of code taken from existing software; a neural network is trained to predict properties of the program by looking at lists of outputs and inputs from each code fragment and optimizing which pieces of code were needed to achieve the most desired result. In effect, this technique is merely automating what human coders do, only the AI can search more thoroughly and attempt to assimilate pieces of code in more configurations than a human might think to try. As the system learns which combinations work and which fail, it is constantly learning from its experience—something that takes many years for humans. Furthermore, the system uses ML to mine databases of source and sort code fragments based on its prioritization of usefulness. Future versions of DeepCoder, could handle more complex and useful (if still tedious) tasks, but auto-generating programs that scrape information from websites or categorize photos save time.

Gamalon is developing probabilistic programming techniques to facilitate learning from less data. Gamalon calls its technique Bayesian program synthesis, named after Thomas Bayes, an 18th century mathematician. Instead of specific variables, the code uses probabilities to refine predictions based on experience. This requires fewer examples (less training) to make determinations, but Gamalon has also built the program to re-write its code and refine its “knowledge” and adjust probabilities as new examples are provided.

These examples offer an impressive glimpse into automated code generation. Such an approach, according to the DeepCoder team, could democratize coding altogether, enabling non-coders to build programs simply by describing their idea to the system and letting the AI handle the rest. The technology has a long way to go before any job displacement or at-scale adoption. For one thing, current solutions like DeepCoder are only capable of handling challenges limited to about five lines of code (not complex code of multiple lines). The objective for now is merely to free coders from more tedious tasks so they can focus on more sophisticated development, while significantly decreasing the time it takes to develop code.

Tractica forecasts that the annual revenue for automated code development in the information technology (IT) sector will increase from \$6.51 million worldwide in 2016 to \$717.49 million in 2025.

Table 2.151 Automated Code Development in Information Technology, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	6.51	16.18	32.02	57.45	97.31	157.88	246.28	368.62	527.12	717.49	68.6%

(Source: Tractica)

2.18.2 COMPUTER-AIDED DESIGN

The discipline of design is as old as human creativity itself, as exemplified throughout history in the success (or failure) of tools, architecture, cities, infrastructure, transportation, homes, and just about every commercial endeavor. And as design tools have evolved over the years, particularly with computer programs and design software, humans have still remained a constant in the process. The advent of AI in design marks a milestone in design technology; one in which the role of humans may fundamentally change. ML and DL are beginning to pervade design tools. Instead of human designers ideating and creating concepts

themselves, new tools are emerging that ingest vast data inputs to recommend optimal designs given the parameters.

Autodesk, a 3D design software provider, is now using AI and DL to expedite the design process. From large architecture and structures, down to the smallest parts, bits, and screws, its “generative design” platform uses a similar technique to product recommendation engines wherein users input certain criteria, specifications, dimensions, and ideas, and the model returns suggestions based on those inputs. The tool takes into account these inputs, as well as broader data around structural integrity, biological comparisons, weight, estimated costs to build, and other custom parameters. Stanley Black & Decker, a household and industrial tool manufacturer, recently used Autodesk’s generative design tool to analyze and recommend a design for a crimper, a small tool used to hang electrical lines. It set parameters, such as weight, dimensions, and costs to manufacture, and after two weeks, the model produced about 100 unique design suggestions. In the end, the team chose the design with the best compromise of weight and costs to manufacture.

While computer or AI-generated design shows tremendous promise, in accelerating ideation, prototyping, and product innovation, the technology also faces barriers before widespread adoption. First, it requires extensive computing power to crunch through data and highly unique parameters; second, this makes it very expensive for most businesses, although Autodesk is actively working on improving algorithms to bring these costs down; and third, computer-generated designs must be reliable and durable under stress tests, including those printed using 3D printers. Each application and the mechanisms involved to prototype such designs introduce new considerations for manufacturers and designers. This is an added challenge to building trust and adoption among stakeholders. For now, many adopters of this technology are using it primarily for software/web development or for smaller parts or bits, rather than large or mission-critical assets.

Tractica forecasts that the annual revenue for computer aided design in the IT sector will increase from \$0.04 million worldwide in 2016 to \$5.82 million in 2025.

Table 2.152 Computer Aided Design in Information Technology, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.04	0.11	0.24	0.45	0.77	1.27	1.98	2.98	4.27	5.82	75.7%

(Source: Tractica)

2.18.3 MOBILE APPLICATION DEVELOPMENT

Since the birth of mobile devices and particularly the smartphone revolution, mobile app development has been a booming area of skill development. The interplay between AI-driven mobile app development and the mobile market itself is somewhat circular: AI has and will continue to have a significant influence on the market as mobile leaders like Apple and Google lead the AI market. Mobile app development spurs AI development and vice versa.

Reinforcement learning itself is very useful for enabling apps to learn from different user(r) dynamics and preferences and determine how to optimize. ML is also a common tool for app developers to assess what features, functions, and overall directions the app could go. AI does this by storing and analyzing user and behavior data in order to improve engagement. Some tools allow mobile app developers to determine where to focus more development resources to improve workflows, layout, or even to boost user engagement and retention. ML is also used to enhance mobile app search analysis, wherein specific techniques (e.g.,

image recognition tied to product inventories or sentiment analysis) are leveraged to perform optimal search given the context of the app. Various ML and DL frameworks can also help power specific capabilities, such as recommendation engines or image recognition. Finally, AI has and will continue to power numerous adjacent capabilities that collectively evolve mobile app capabilities. Some examples include voice recognition, image recognition, object recognition, gesture recognition, AR, CV, language translation, chatbots, photo organizing, etc.

While AI will continue to power mobile app development, and likely the evolution of mobile app capabilities, it is also essential that human developers play a core role. So much depends on both the integrity of the data fed to algorithms, and the quality and adaptability of the algorithms over time—to audit for bias, to account for new contexts or market developments, and to adhere to potential changes in regulatory compliance, for example.

A number of companies are focused on AI-powered mobile app development, including GoodWorkLabs, Krify, Diffco, TechUgo, Vital.ai, AppSquadz Technologies, and many others.

Tractica forecasts that the annual revenue for mobile application development in the IT sector will increase from \$0.24 million worldwide in 2016 to \$63.22 million in 2025.

Table 2.153 Mobile Application Development in Information Technology, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
	0.24	1.08	2.47	4.71	8.23	13.59	21.42	32.27	46.33	63.22	85.4%

(Source: Tractica)

2.18.4

NETWORK/INFORMATION TECHNOLOGY OPERATIONS MONITORING AND MANAGEMENT

Network and IT operations monitoring and management is a foundation of any IT department, and consists of a wide range of tools and capabilities for managing the provisioning, capacity, performance, and availability of computing, application, and networked environments. These environments face increasing constraints, particularly legacy monitoring tools and services that simply cannot scale to address the complexity of dynamic architectures and outputs.

As these environments are fundamentally about both security and automation, IT management service providers are exploring how and where to leverage AI to support these tasks. The idea is to use AI to help IT infrastructure become more self-healing via predictive and preemptive maintenance; to focus automation on outcomes rather than just specific tasks. As IT environments grow more complex and unpredictable, algorithms are viewed as an essential supplement, if not requirement, to monitor, learn from potential vulnerabilities, and, in some cases, automatically execute on quality assurance, network and storage optimization, anomaly detection, and asset life cycle management. The benefit of these tools is not only the potential to eradicate human errors—often the primary cause of major IT issues—but to reduce costs associated with the many laborious and meticulous tasks of IT operations.

Enterprise IT operations and network environments are effectively critical infrastructure for the businesses and services that rely on them. While adoption is moving rapidly in this area, trust, security, and reliability are paramount. For those working in this arena, the longer-term

goal is not to replace humans, but for humans to spend their time telling the software what to do, while the details and execution are determined by the AI tools.

Moogsoft focuses developing automation and analytics solutions for IT operations and DevOps environments. The company uses ML, DL, clustering algorithms, and entropy calculation to identify the root causes of any issue, cluster them into actionable situations, surface situations, catalog and analyze previous issues, and enable early detection of future issues. DevOps and IT operations teams can see in real time the nature of issues customers are experiencing and manage issues via socialized workflows promptly. The company claims its tools have helped clients reduce the mean time to restore (MTTR) by up to 60%. Logz.io, Infosys, Instart Logic, and Akamai, among others, are companies using AI for IT operations management.

Tractica forecasts that the annual revenue for network/IT operations monitoring and management in the IT sector will increase from \$0.13 million worldwide in 2016 to \$142.8 million in 2025.

Table 2.154 Network/Information Technology Operations Monitoring and Management in Information Technology, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.13	2.02	5.14	10.19	18.15	30.28	48.03	72.61	104.49	142.80	116.9%

(Source: Tractica)

2.18.5

SIMULATING WORLDS FOR ARTIFICIAL INTELLIGENCE TRAINING

In order to simulate the functionality of software, developers and designers used to have to undergo significant, costly, and sometimes risky prototyping periods. Extensive testing, evaluation, re-configuring, re-testing, and repeat, often in conjunction with manually processed data sources (e.g., road data, safety compliance, etc.) was the status quo in order to advance features, functions, and designs to a point of reliability, security, and scale. AI is influential in this area, particularly as it can power very precise and highly programmable environments that can be used to simulate worlds for AI training. The benefits to testing in simulated worlds are manifold. One, costs are often lower as environments can be programmed with many (one day infinite) variables and parameters, so that a wide range of scenarios can be incorporated, learned, and tested over and over.

Using games like Pac-Man, chess, or other board games as a way to test and train AI systems (through reinforcement learning) has helped accelerate algorithms and model development for years, but only recently has the focus turned to training AI for real-world applications. The technique is driven by a reward function, only instead of points as rewards in a game, reward functions in the physical world might be a vehicle stopping for a dog or a robot successfully picking up a cup. Beyond gaming, simulated environments for AI training are gaining fast traction in robotics and autonomous vehicle development, but are also applicable in IT environments for testing, security patching, in training for AR or VR environments, for employee training, and beyond.

OpenAI, an AI research foundation, recently unveiled Universe, an open-source “digital playground” where developers can virtually test and train AI using games, apps, and websites. Universe contains thousands of environments with an expanding catalog of everything from space to biological science apps. The software also enables “transfer learning,” in which an agent takes what it has learned in one application and applies it to

another, enabling what OpenAI calls “general-purpose” knowledge about the world. This is a small but significant step toward more generalized AI, outlined in Section 2.5.11.

Tractica forecasts that the annual revenue for simulating worlds for AI training in the IT sector will increase from \$0.19 worldwide in 2017 to \$14.21 million in 2025.

Table 2.155 Simulating Worlds for AI Training in Information Technology, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.00	0.19	0.50	1.00	1.80	3.00	4.77	7.22	10.40	14.21	154.3%

(Source: Tractica)

2.18.6 SOFTWARE CODE ERROR CHECKING

Finding errors in software code has historically been a matter of developers reviewing or writing their own test code to find any bugs, or more often, encountering bugs, system glitches, or failures after deployment. Automatic bug-repair, patch-generation, or program-repair all fall under this use case. AI is being used to address this problem as it can be trained to understand specific programming languages, and then generalize from patterns. Automated bug-fixing emerged in the early 2000s and typically involved patch generation, where programs were analyzed and candidate patches were derived using ML or genetic programming; then patches were validated against specifications or a test version of the program. This approach was demonstrated in 2015 by MIT researchers using AI to automatically fix software bugs by replacing faulty lines of code with lines that worked from other programs. The researchers used a set of successful human-generated patches obtained from open-source software repositories, trained the model to identify and rank candidate patches based on likelihood of success, and then validated them against a suite of test cases.

DiffBlue, a company spun out of Oxford University’s incubator, uses AI to develop a mathematical model of any code base that can check for and correct faulty code. It has trained its software to “understand” code enough to serve the repetitive and labor-intensive tasks of testing code for bugs, as well as to automatically flag exploitable bugs and generate tests for those. It is also working on a refactoring product, which mines for and re-writes bad or out-of-date code. The startup currently works with Java and C, but plans to expand to others. It is currently used by all major banks in the United Kingdom. The longer-term goal of the company is to enable illiterate people to program.

Tractica forecasts that the annual revenue for software code error checking in the IT sector will increase from \$2.45 million worldwide in 2016 to \$506.29 million in 2025.

Table 2.156 Software Code Error Checking in Information Technology, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.45	9.19	20.31	38.24	66.41	109.30	171.96	258.73	371.19	506.29	80.8%

(Source: Tractica)

2.18.7 WEBSITE CREATION

In the early days of the web, website design used to be a far more labor-intensive task, requiring extensive programming. As the internet has increased in size, reach, and standardization, website development has become somewhat more democratized. Sites like Wordpress, Wix, Squarespace, Weebly, and others have helped anyone create a decent looking site at a low cost. Businesses and enterprises still spend significant resources on website creation, integrations, search engine optimization (SEO), and ongoing development and maintenance.

Meanwhile, website development tools are exploring the use of AI to enable easier, more personalized website creation. The idea is that users input certain information such as their business name, location, based preferences, and various algorithms mine images, text, and millions of websites or web design interactions to find optimal design patterns. Many question the true artistic capability of such systems, particularly when developers' abilities to customize AI-generated sites is limited. Given the importance of owned web real estate to businesses, there may be limits to these tools in the early days.

Wix, a do-it-yourself (DIY) website development company uses AI to recommend specific website features, layout, text, images, buttons, and aesthetic based on users' needs, location, and business. Its algorithms also locate content from around the web to add custom design elements. Users answer a short questionnaire and Wix's Artificial Design Intelligence (ADI) service suggests designs based on inputs. The company developed the algorithms by mining data from more than 86 million user interactions.

The Grid is a website design platform that takes any piece of content and builds a responsive design website around it. Using image recognition, automated cropping, algorithmic palette, and typography selection, the platform's objective is automation, automatically updated and reformatting every time new content is added. The site currently offers little in the way of custom development. Like Wix ADI, The Grid starts by asking users a few simple questions, and then determines optimal combinations of branding, layout, design, and content.

Tractica forecasts that the annual revenue for website creation in the IT sector will increase from \$0.2 million worldwide in 2016 to \$57.85 million in 2025.

Table 2.157 Website Creation in Information Technology, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
0.20	0.97	2.24	4.28	7.50	12.41	19.58	29.51	42.38	57.85	87.7%	(Source: Tractica)

2.19 INVESTMENT

2.19.1 ALGORITHMIC TRADING STRATEGY PERFORMANCE IMPROVEMENT

Every day, computers perform billions of calculations and make millions of electronic trades. Algorithmic trading, sometimes called "algo-trading," has been part of automating investment for years. Algorithms create rough schedules for when, how many shares, and at what price to buy or sell, and follow schedules accordingly; when changes in the market occur, the algorithm checks if the situation is applicable and does or does not trigger execution. The most common application of algo-trading is to enhance trading strategies, including arbitrage, intermarket spreading, market making, and speculation.

New advancements in AI and DL are being applied to improve strategy and performance. In this context, neural networks can uncover complex patterns, trends, and relationships unable to be detected by humans in high-input/high-speed environments. The idea is that, just as DL successfully identifies particular features in common to cat images, it may be able to identify particularly lucrative features of stocks as well.

Goldman Sachs, Bridgewater Associates, Cerebellum Capital, Euclidean, Man (AHL) Group, and a number of other established investment hedge fund firms are investigating how and where they can apply DL. Meanwhile, a host of startups like Sentient Technologies, Clone Algo, Neurensic, Alpaca, and Binatix are working on using AI and DL to improve or automate investment and trading as well.

Aidya is a Hong Kong-based investment company applying evolutionary programming, chaotic dynamics, and probabilistic knowledge to algo-trading. The system ingests a range of inputs, such as price and volumes from around the world, news in numerous languages across multiple sources, and macroeconomic and company accounting data, and studies how multiple factors within these data sources have interrelated historically.

Given the high stakes, experts point to a number of remaining challenges in DL and AI-enabled algo-trading, namely around the limitations of models to fully regard (or disregard) noise, random vectors, and high uncertainty prevalent in financial markets. Furthermore, the very commoditization of such algorithms would erode their competitive predictability, until, that is, algorithms themselves advance in evolutionary computation.

Tractica forecasts that the annual revenue for algorithmic trading strategy performance improvement in investment markets will increase from \$45.66 million worldwide in 2016 to \$2.014 billion in 2025.

Table 2.158 Algorithmic Trading Strategy Performance Improvement in Investment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	45.66	66.43	102.57	164.98	270.46	441.78	703.03	1,066.98	1,518.99	2,013.53	52.3%

(Source: Tractica)

2.19.2 FINANCIAL SEARCH ENGINE

Investment professionals spend a great deal of time on research, sifting through digital financial documents, such as public securities filings, but also data from market research firms, blogs, press releases, conference call transcripts, investor presentations, news media, and thousands of other online sources. A growing area is in the use of search engines specifically focused on financial data across disparate domains. These engines are using NL and NLP to pull in and index thousands of disparate sources, which helps significantly reduce the amount of time spent on financial research by investment professionals.

One such company catering to this market is AlphaSense. Founded in 2011, AlphaSense accesses more than 1,000 sell-side research providers and 35,000 public companies to pull financial data and then adds NLP, ML, and DL to create an intelligent search that understands financial language, including a range of synonyms. The tool is currently used by 450 customers, including JP Morgan and Credit Suisse. The company landed a new round of funding in March 2016 of \$33 million, bringing its funding total to \$35 million. Uberpile also provides similar services.

Tractica forecasts that the annual revenue for financial search engines in investment markets will increase from \$0.27 million worldwide in 2016 to \$16.33 million in 2025.

Table 2.159 Financial Search Engines in Investment, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
0.27	0.43	0.73	1.23	2.09	3.49	5.62	8.59	12.29	16.33	57.8%	

(Source: Tractica)

2.19.3

MARKET INTELLIGENCE AND DATA ANALYTICS FOR INVESTMENT

Since the earliest days of commercial investment, financial analysts and researchers have been integral to tracking spending, competitive forces, consumer trends, events, and other relevant factors for investment decision-making. As investing and trading have grown more and more data-intensive, ML, NLP, and particularly DL are highly sought after tools for market intelligence and data analytics. These tools typically mine enormous multi-dimensional data sources and sets to surface trends, notify users of opportunities, or enable query-based recommendations. AI may eventually replace financial analysts, as humans can take hours to collect and analyze data, take time to learn and evolve, are susceptible to emotions like greed or fear, and are slow to adapt to changing market conditions. AI takes seconds, can correct its errors in minutes, and adapt quickly and without human input.

Discover Patterns uses DL to support investors and analysts with trend detection and decision-making. Its Integrated Network Reality model analyzes big unstructured data streams, using context engines, billions of agent discoveries, and analyst support to discover and then track emerging themes across markets, looking at competitive patterns that are either evolving or dissipating. Themes are mapped to investments with industry context engines pulling in information across diverse areas. The system is designed for immediate analyst or client consumption. An analyst might start with an idea (e.g., robotics), at which point related themes are displayed and further themes are recommended. Analysts can continue to drill into specific industries, subsectors, geographies, companies, technologies, interfaces, and trends themselves. Tracking, movements, social feeds, and a variety of other parameters can be customized and visualized over time.

Most large investment institutions, such as Goldman Sachs, have built or acquired companies that conduct AI-based market intelligence and data analytics, as such tools are competitive in and of themselves. Another provider is a company called Othoz.

Tractica forecasts that the annual revenue for market intelligence and data analytics for investment will increase from \$0.03 million worldwide in 2016 to \$5.83 million in 2025.

Table 2.160 Market Intelligence and Data Analytics for Investment in Investment, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
0.03	0.09	0.19	0.37	0.68	1.18	1.95	3.03	4.37	5.83	76.5%	

(Source: Tractica)

2.19.4

SATELLITE IMAGERY FOR GEO-ANALYTICS

Satellite imagery has long been a closed domain with high-resolution image databases only available to a select few companies and organizations, such as weather centers, government agencies, the military, and oil & gas companies. Being able to track changes on the ground from space has been vital for these industries, but required human analysis for years. Rapid increases in the availability and improvement in the level of detail of satellite imagery, and advancements in AI, CV, and DL have created new ways of identifying features, tracking changes, and extracting value from satellite imagery.

Apart from providing a way for humans to track the planet on a daily basis, this also means that image processing will have to be automated, in order to take advantage of this quick refresh rate and trove of imagery data. Collecting information through aerial imaging may be cheaper than a full networked sensor and connectivity implementation, for example. DL is particularly helpful given it requires low or no feature engineering. Some basic challenges do remain when it comes to weather, viewpoint, lighting, and atmospheric unpredictability.

In the investment space, satellite images are being used to forecast growth by analyzing real estate, construction, energy resources, retail parking lots, etc. More generally, satellite imagery can help track a bounded area with alerts and updates provided when something changes in that specific area, or for historical changes over said area. These are not just new applications, but new business models that provide country-wide, or object-specific analysis of satellite imagery to vertical markets.

Orbital Insight uses CV and DL to take millions of geospatial images and provide insights based on these images. Providing, for example, the relative count and distribution of cars parked in a retailer's parking lot offers retailers insights into traffic patterns and inputs to forecast traffic over time. It provides similar services, via satellite, to measure crude oil stored in containers and assess oil supply in real time. This information is also factored into retail index for investors.

Tractica forecasts that annual revenue for satellite imagery for geo-analytics in investment markets will increase from \$0.09 million worldwide in 2016 to \$0.62 million in 2025.

Table 2.161 Satellite Imagery for Geo-Analytics in Investment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.09	0.10	0.11	0.13	0.16	0.21	0.28	0.38	0.49	0.62	24.0%

(Source: Tractica)

2.20

LEGAL

2.20.1

AUTOMATED REPORT GENERATION

Law firms and legal entities generate reports for internal stakeholders, as parts of client engagements, or even as formal products. Report generation is important across areas like billing, accounting, and case management. As the amount of data flowing into and across organizations grows, the problem is not just one of content distribution, but of the time it takes to comprehensively identify and organize insights that are useful and consumable.

AI is well suited for report generation. Using NLP, ML, and DL in some cases, companies are using AI to collate reports far more rapidly than humans. AI-generated reports can surface relevant metrics, tables, and charts, and generate multiple paragraphs of narrative.

Automated report generation tools generally support the following tasks:

- **Data Sourcing:** Identifies and extracts data from relevant internal and external sources, including industry news and reports, social media listening, and competitor intelligence
- **Data Interpretation:** Upon consolidating data in standardized formats, the solution aligns the data in templates, codes, and prepares it for analysis using ML.
- **Data Analytics:** Defines business rules and correlation/causality at scale. With predictive modeling and data enrichment, solutions can run hundreds of “what if” scenarios and perform trend analysis
- **Narrative and Semantic Commentary:** Using NLP and NLG, solutions can sometimes automate variance analysis and commentary writing in a systematic and structured way

RAVN offers a cognitive computing platform for enterprise search and reporting. In the legal space, it supports contract analysis, due diligence, cost projections, and analytics.

Tractica forecasts that the annual revenue for automated report generation in legal will increase from \$2.17 million worldwide in 2016 to \$693.43 million in 2025.

Table 2.162 Automated Report Generation in Legal, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.17	11.36	26.56	51.11	89.72	148.54	234.51	353.59	507.96	693.43	89.8%

(Source: Tractica)

2.20.2 CONTRACT ANALYSIS

In the United States, the legal profession and judicial process consume a larger share of GDP than in any other advanced economy. Massive costs associated with this sector, and its inherently language- and document-based structures and elements, render it ripe for ML and efficiency improvements enabled through algorithms. In particular, the contract analysis and document review and discovery processes are labor- and time-intensive. It is estimated that between 30% and 50% of the time at a company's legal firm or in-house legal department is spent on contract analysis.

Companies are using AI for legal contract analysis for various legal areas, such as due diligence, general commercial compliance, lease abstraction, real estate, corporate organization, and others. Having AI tools work on contract analysis presents an opportunity to have legal support workers work up the stack and spend more time on higher ticket value items like client recommendations. Across broader applications, judicial systems can apply DL to analyzing millions of individual cases and decades of case law to predict outcomes for future cases and accelerate case resolutions (both domestic and international cases) in court.

Companies like Kira, Beagle, Legal OnRamp, Adnsensa, Seal Software, eBrevia, and Luminance also offer solutions in contract analysis and document discovery. San Francisco-based Judicata is using DL to find patterns in unstructured legal text to convert to structured data, and hence organize the entire body of case law into a legal “genome” in order to predict the outcome of future cases.

Tractica forecasts that the annual revenue for contract analysis in legal will increase from \$16.69 million worldwide in 2016 to \$957.16 million in 2025.

Table 2.163 Contract Analysis in Legal, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	16.69	29.90	51.25	85.23	138.21	218.48	335.37	496.92	706.06	957.16	56.8%

(Source: Tractica)

2.20.3

LEGAL DOCUMENT REVIEW AND RESEARCH

A significant portion of the practice of law is based on understanding precedent, so lawyers of all types spend a great deal of time, an estimated 30% to 40% of their time, in legal research, looking for case law to support arguments and positions they take. In private practice, the time they spend is no longer billable to the client, so there is significant incentive for legal research to be conducted more efficiently.

Case law in the United States alone encompasses the federal system, all 50 states, more than 3,000 counties, and an untold number of cities and municipalities. Some of this data is structured, some is not. Research tools, such as Westlaw and LexisNexis, have begun to transition their database queries from keyword search to natural language, but the more significant advancement has been with case law analytics.

Casetext focuses on the legal research use case. Casetext has built a platform, CARA, that has ingested millions of legal cases and articles. Unlike popular legal databases like Westlaw and LexisNexis, CARA does not use keyword search. Customers scan in the documents they are working on and the system then provides the cases that are relevant. CARA was launched in September 2016 and current customers include some of the world's largest law firms (DLA Piper and Quinn Emanuel) and many smaller firms. Other legal document review and research companies using AI include Fastcase, Ravel Law, and ROSS.

Tractica forecasts that the annual revenue for legal document review and research in legal will increase from \$13.12 million worldwide in 2016 to \$604.02 million in 2025.

Table 2.164 Legal Document Review and Research in Legal, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	13.12	21.55	35.09	56.56	89.94	140.41	213.84	315.25	446.49	604.02	53.0%

(Source: Tractica)

2.21

LIFE SCIENCES

2.21.1

CREATE SYNTHETIC LIFE FORMS

Can we use AI to create life? So much depends on the definition of life. While we are far from using AI to grow a brain or a beating heart, there are developments underway that are working toward using ML and DL to synthetically create life forms. In these early days of AI, various computational biology and chemical manufacturing applications are experimenting with ways to use M to manipulate simple organisms and assist in biolab automation.

Zymergen is using AI for microbe strain optimization. It produces industrial chemicals from single-cell organisms (microbes) that can be used in a variety of materials and parts. Industrial chemicals are used in everything from soap to car parts to paint, but has historically depended on petroleum. What Zymergen does is use algorithms to develop microbes that can serve the same material function, but are not from petroleum. Microbes are versatile and Zymergen reprograms the genetic DNA so that the microbes churn out raw material byproducts, which are then used commercially. Zymergen leverages AI in two ways: in algorithms that sort through millions of different genetic combinations to produce the best chemical for the application; and in robotics, where a robotic workforce is used to assemble DNA, stir liquids, and aid in experiments. The company has raised over \$170 million in funding.

Another company, Gingko Bioworks, also uses AI microbe strain automation, and supports applications such as flavor and fragrance in the food industry, probiotics for soldiers susceptible to stomach bugs, ingredients for cosmetics, pharmaceuticals, and more.

Tractica forecasts that the annual revenue for creating synthetic life forms in life sciences will increase from \$6.63 million worldwide in 2016 to \$343.38 million in 2025.

Table 2.165 Creating Synthetic Life Forms in Life Sciences, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	\$6.63	\$10.65	\$17.15	\$27.57	\$44.08	\$69.72	\$108.48	\$164.84	\$242.67	\$343.38	(Source: Tractica)

2.22 LOGISTICS

2.22.1 DEMAND FORECASTING FOR WAREHOUSE AND SUPPLY CHAIN

The ability to understand product demand has a direct impact on the economic viability of any business. In the past, companies have been reactive or merely formulaic in their approach to gauging supply orders to fulfill demand. With the advent of ML and, increasingly, DL, companies are able to analyze, learn from, forecast, and predict demand with far greater accuracy and with regard to a wider range of forces. This is an area where AI will intersect with the IoT in that it will incorporate sensor data and radio frequency identification (RFID) tag monitoring, as well as support modifications at each level of supply chain. For example, DL can be applied to improve parts and labor sourcing, channel optimization, product inventory, quality assurance, fraud, risk modeling, weather forecasting, and predictive maintenance to support reliable supply chain operations.

Companies like SupplyMind, Epicor, and DemandWorks provide software supporting this use case.

Tractica forecasts that the annual revenue for demand forecasting for warehousing and supply chain in logistics will increase from \$1.83 million worldwide in 2016 to \$39.85 million in 2025.

Table 2.166 Demand Forecasting for Warehousing and Supply Chain in Logistics, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.83	2.43	3.35	4.77	6.95	10.21	14.93	21.42	29.80	39.85	40.8%

(Source: Tractica)

2.22.2 MACHINE/VEHICULAR OBJECT DETECTION/IDENTIFICATION/AVOIDANCE

Perhaps the most valuable use of AI in vehicles is the use of object detection and classification, which takes sensor data, often from cameras, and then uses complex algorithms to classify these objects so that the AI system can then “learn” their characteristics, and recognize them in real time.

The challenge is not in capturing images, as today’s HD cameras can present images in stunningly clear detail. However, in a moving environment, objects can appear to change size as a vehicle or camera approaches. The angle at which an object is viewed can also skew its appearance, and the presence of other factors (rain, bright sunlight, low lighting, glare, dirt, snow, or any other number of obstructions) can alter the appearance of an object, making it hard to accurately and consistently identify the object.

This is an area where machine vision and ML can provide invaluable support. By capturing a wide range of images of objects from a variety of vantage points, angles, and in different conditions, a repository of images that can be definitively classified as that object can be created, and used to “train” a ML system to identify and classify objects that resemble objects in the repository. By then assigning various other attributes to each object, such as whether the object is informational like a sign, whether or not it is permanent or temporary like a road barrier, or whether or not it has the capability of motion and how it typically moves, the system can begin to develop logical rules on handling each object and rules for dealing with them.

In logistics, object detection, avoidance, and identification will serve a number of applications including, but not limited to autonomous trucks, drone delivery, autonomous forklifts, or other indoor machines that move about, robotics, visual inspection, workplace safety, etc. Reference Section 2.23.2 for a description of this use case in manufacturing.

Tractica forecasts that the annual revenue for machine/vehicular object detection/identification/avoidance in logistics will increase from \$7.68 million worldwide in 2017 to \$584.65 million in 2025.

Table 2.167 Machine/Vehicular Object Detection/Identification/Avoidance in Logistics, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	7.68	20.45	41.13	73.72	123.44	196.17	296.94	427.62	584.65	N/A

(Source: Tractica)

2.22.3 LOCALIZATION AND MAPPING

As the movement of goods across the supply chain undergoes radical transformation, relying less on human labor and more on machines, ML, DL, and CV are becoming central technology enablers for robotics and autonomous machines to reliably move goods about.

Localization and mapping concerns the need and computational ability to simultaneously construct maps of the immediate environment, while updating both the agent's position on that map and movement therein. In the context of logistics, localization and mapping is a core technique for autonomous movement of cars, trucks, drones, or any other autonomous vehicle. Whether managing inventory in a warehouse environment, shipping goods via autonomous trucks, or drone-based delivery systems, the localization and mapping technology is essential so goods, infrastructure, people, and revenue are not damaged.

Amazon's delivery robots, which are the product of its acquisition of Kiva Systems, currently number well over 45,000 robots working alongside 23,000 people across 20 fulfilment centers, support repetitive and physically demanding tasks involved in inventory movement and warehouse management. These robots rely on localization and mapping to move around without damaging their surroundings or goods. Previously, Amazon workers would have to search shelves for products needed to fulfil each order; now robots, roughly the size of a footstool, carry shelves around seamlessly based on order needs and rearrange shelves in tightly packed rows. This increases efficiency of how warehouse surface area is used as well as the speed of fulfilment.

Tractica forecasts that the annual revenue for localization and mapping in logistics will increase from \$6.48 worldwide in 2017 to \$493.34 million in 2025.

Table 2.168 Localization and Mapping in Logistics, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	6.48	17.25	34.70	62.21	104.17	165.53	250.57	360.84	493.34	N/A

(Source: Tractica)

2.22.4 SATELLITE IMAGERY FOR GEO-ANALYTICS

Satellite imagery has long been a closed domain with high-resolution image databases only available to a select few companies and organizations, such as weather centers, government agencies, the military, and oil & gas companies. Being able to track changes on the ground from space has been vital for these industries, but required human analysis for years. Rapid increases in the availability and improvement in the level of detail of satellite imagery, and advancements in AI, CV, and DL have created new ways of identifying features, tracking changes, and extracting value from satellite imagery.

Apart from providing a way for humans to track the planet on a daily basis, this also means that image processing will have to be automated, in order to take advantage of this quick refresh rate and trove of imagery data. Collecting information through aerial imaging may be cheaper than a full networked sensor and connectivity implementation, for example. DL is particularly helpful given it requires low or no feature engineering. Some basic challenges do remain when it comes to weather, viewpoint, lighting, and atmospheric unpredictability.

In logistics, applications generally involve monitoring the volumes and production of goods. For example, shipping companies (and investment firms) can now count the number of ships arriving and leaving ports to gauge trade volume of a country. More generally, satellite imagery can help track a bounded area with alerts and updates provided when something changes in that specific area, or for historical changes over said area. These are not just new applications, but new business models that provide country-wide, or object-specific analysis of satellite imagery to vertical markets

Tractica forecasts that the annual revenue for satellite imagery for geo-analytics in logistics will increase from \$0.32 million worldwide in 2016 to \$2.37 million in 2025.

Table 2.169 Satellite Imagery for Geo-Analytics in Logistics, Annual Revenue, 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
0.32	0.36	0.42	0.50	0.63	0.81	1.06	1.41	1.85	2.37	25.2%	(Source: Tractica)

2.22.5 SUPPLY CHAIN AND LOGISTICS (FREIGHT TRANSPORT, RETAIL)

Supply chain and logistical operations are undergoing radical transformation. The more data surrounding not only the movement of materials and goods, but the businesses, customers, machinery, infrastructure, and economic forces influencing their movement, the more supply chains are becoming automated. AI will play a significant impact in supply chain automation, although it is one of a number of technological innovations imposing radical changes. IoT, 3D printing, blockchain, on-demand services, and autonomous vehicles will also transform accelerate automation. The impact of AI lies across just about every part of the supply chain, and is fundamentally about helping organizations analyze the massive amounts of data being generated for operational and service efficiency, and to improve working capital management. As visibility into current and future scenarios grows, AI will also power more agility in product development, reducing time to market and increasing risk planning capacity. Below are some of the areas where AI will power supply chain automation for logistics:

- **Customer Demand Forecasting and Analysis:** Companies can use predictive analytics, ML, and DL to analyze, learn from, forecast, and predict demand with greater accuracy and with regard to a wider range of forces
- **Supply Chain Operations and Execution:** Factories are becoming more automated at every level, using AI to power robotics, autonomous machinery, anomaly detection, predictive maintenance, etc. Siemens is working toward developing an entirely automated and self-organizing manufacturing plant wherein demand and order information would be automatically processed as work orders, which would be fed into the production execution.
- **Supplier Management and Customer Management:** As companies leverage virtual conversational agents to handle more robust customer service tasks, these will be integrated with product inventory, procurement, enterprise resource planning (ERP), CRM, and other operational systems. IPSoft's Amelia agent is used by an oil & gas company for answering and triaging invoicing questions from its suppliers.
- **Logistics and Warehousing:** Manufacturers and distributors can use AI to analyze and automate manufacturing, storage, and movement of materials and goods in warehouse environments. Examples include CV-enabled cameras, robotics, or autonomous forklifts.
- **Logistics and Transportation:** Companies across the supply chain can use AI to help handle domestic and international movement of goods; Sensor data integration (e.g., RFID, quick response (QR) code, beacons) can provide greater accountability for tracking shipments, changes in conditions, arrivals to ports, etc.

- **Procurement:** Increased real-time visibility and classification of spending helps improve compliance adherence and drive cost reductions, even automate certain tedious procurement activities. The Singaporean government is trialing the use of AI to prevent procurement fraud by analyzing procurement requests, HR and financial data, tender approvals, workflows, and non-financial data to identify corruption or negligent practices.
- **New Product Development:** Companies will be able to use AI to power more rapid product development in a number of ways: leveraging 3D digital models of products for more rapid prototyping; using labor sourcing data to optimize design integrity; incorporating UX and behavioral analytics for faster innovation; and reducing time to market through automated orders processing.

Over time, Tractica expects automated supply chains will increasingly move toward connectivity, creating an ecosystem of multiple chains—manufacturing, agricultural, distribution, retail, finance, transportation—that enable seamless movement and visibility of products and information from one side to the other.

Tractica forecasts that the annual revenue for supply chain and logistics will increase from \$6.04 million worldwide in 2016 to \$132.44 million in 2025.

Table 2.170 Supply Chain and Logistics in Logistics, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	6.04	8.01	11.06	15.79	23.04	33.89	49.58	71.16	99.03	132.44	40.9%

(Source: Tractica)

2.22.6 WEATHER FORECASTING

Logistics companies benefit from the ability to forecast weather events as foresight can help ensure minimal disruption to supply chain, ordering components ahead of time, wear and tear on vehicles, moving assets to a safer location, identifying alternative sources of supply or alternative routing for fleets, or potentially signaling when to evacuate employees. In the United States alone, the cost of weather-related delays in the freight industry was estimated at \$8.7 billion (an estimated 1.6% of the total estimated freight market) in 2012, according to the U.S. Department of Transportation.

AI and sensor data from hundreds of thousands of sources collected and monitored in real time (and over many years) are transforming the level of understanding and ability to forecast conditions. In addition to weather data, engines can combine streaming data from social feeds, news reports, transportation data, and historical data on storms or other weather events. While no one can ever fully predict the future, AI techniques apply reinforcement learning on past predictions and actual outcomes. By comparing predictions with accuracies, the model can learn and improve simulation capabilities, and forecast much further into the future.

IBM Watson's Supply Chain Risk Insights platform combines AI-powered weather forecasting with Big Data analysis supporting procurement, demand management, manufacturing, and supply chain risk management. Weather forecasts are integrated into the platform so logistics and supply chain companies can more quickly anticipate, create contingency plans, and safely deal with extreme weather. When Hurricane Patricia, the second most-intense tropical cyclone on record, was about to strike Mexico, weather

forecasting powered by IBM's cognitive supply chain helped its Guadalajara production center prepare. Although the storm struck north of the city, IBM evacuated the center as a precaution and made contingency plans immediately. Some inbound shipments were routed to the United States and then shipped back to Mexico after the hurricane passed.

A company called Riskpulse offers futures traders information about how inclement weather could disrupt the supply chain in the short term.

Tractica forecasts that the annual revenue for weather forecasting in logistics will increase from \$0.02 million worldwide in 2016 to \$5.98 million in 2025.

Table 2.171 Weather Forecasting in Logistics, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.02	0.10	0.23	0.44	0.78	1.28	2.02	3.05	4.38	5.98	87.7%

(Source: Tractica)

2.23 MANUFACTURING

2.23.1 3D PRINTING ARM CONTROL

While 3D printing still faces a number of challenges before it achieves widespread industrial or consumer adoption, the industry has grown rapidly over the last 5 years. Materials, waste reduction, compliance, and reliability remain important areas for the 3D printing industry to overcome, but AI is beginning to address core challenges associated with efficiency.

Using NL and DL, 3D printing arm control offers a significant advancement in one of the most challenging areas the technology supporting the market faces: time to print. As algorithms learn parameters over time, they could improve and suggest new materials or even structures to achieve the same or similar design integrity. Reference Section 2.18.2 for more on how AI is impacting design.

Ai Build is developing algorithms specifically for 3D printing and learning. The idea is to use AI to understand and learn optimal parameters (e.g., best material, design, extrusion thickness, cost mitigation, and time required) for constructing specific products in the most efficient way possible. It is developing a robotic arm capable of printing very large modular structures with high-resolution finishing, using a variety of polylactic acid (PLA) and acrylonitrile butadiene styrene (ABS)-based materials. This also helps with the fundamental problem of robots not being able to see during production; Ai Build's robotic arm will be equipped with sensors allowing it to track what it does, see how it moves, and find and correct any errors.

Tractica forecasts that the annual revenue for 3D printing arm control in manufacturing will increase from \$3.53 million worldwide in 2016 to \$246.41 million in 2025.

Table 2.172 3D Printing Arm Control in Manufacturing, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	3.53	8.64	16.47	28.33	45.90	71.04	105.12	147.98	196.85	246.41	60.3%

(Source: Tractica)

2.23.2

MACHINE/VEHICULAR OBJECT DETECTION/IDENTIFICATION/AVOIDANCE

The ability to “see” in factory and manufacturing settings is very often what has defined quality for parts and products produced. In the past, precision of parts and elements relied on humans, and, later, heavy machinery for preconfigured repetitive evaluation and sorting.

With advances in machine and CV, which are becoming DL enabled, the ability to more accurately and precisely detect and identify specific features automates tasks like fault detection, failure type detection, visual inspection, inventory monitoring, product testing, workplace safety, video analytics, and potentially additive manufacturing techniques like 3D printing in the long term. Meanwhile, these techniques are infusing industrial robots, and also support safer working environments for human employees. As data from thousands of cases flows in, neural networks help the robots quickly learn and predict thousands of non-automated manufacturing tasks.

Another impact of DL-based object detection is flexibility and reduced downtime. Replacing a robot or machine on the production line is costly, slow, and requires downtime to calibrate. A project by the University of Nottingham in England created smart algorithms to help machines self-optimize during start-up, which achieved 50% reductions in ramp-up time. Rethink Robotics is using object detection classification and avoidance techniques to have its manufacturing robots, Baxter and Sawyer, perform safely in the presence of humans. In the industrial and manufacturing robotics space, collaborative robots like Baxter and Sawyer are becoming much more prevalent, with leading robot manufacturers like ABB, Kuka, and Yaskawa all using CV-based object avoidance techniques to have robots work safely alongside humans.

Tractica forecasts that the annual revenue for machine/vehicular object detection/identification/avoidance in manufacturing will increase from \$3.04 million worldwide in 2016 to \$280.05 million in 2025.

Table 2.173 Machine/Vehicular Object Detection/Identification/Avoidance in Manufacturing, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	3.04	9.01	18.09	31.77	51.97	80.75	119.66	168.46	223.96	280.05	65.3%

(Source: Tractica)

2.23.3

PREDICTIVE MAINTENANCE

As manufacturing produces more and more digital replications of physical assets like parts, machines, vehicles, equipment, and even in process manufacturing environments, new capabilities around monitoring, learning, and predicting maintenance needs continue to emerge. In industrial environments, this trend has been evolving alongside the ML field for over three decades, through various configurations of data mining, case-based reasoning, knowledge-based systems, genetic algorithms, and, increasingly, neural networks. But aging infrastructures, increased digitization and threat vectors have caused maintenance loads to grow in scope.

In addition to techniques like sequence analysis, which can be used to understand failure patterns and follow-on failures, ML and DL are now being used to perform predictive models or recurrent event models. These models support learning and prediction around specific functional failures, as well as optimization of parallel systems. They also help with

maintenance scheduling by identifying appropriate expertise, prioritization, and scheduling based on risk.

A number of startup and data analytics companies are working on predictive maintenance using AI in the fields of manufacturing, aerospace, and automotive. For many current and “next generation” products, specifically in IoT-enabled devices and machinery, robots, and autonomous transportation vehicles and machines, these capabilities will be manufactured into devices themselves to streamline the configuration of predictive maintenance programs.

Konux specializes in predictive maintenance in industrial and transportation applications. One of its customers is German railway company Deutsche Bahn, which uses KONUX to both digitize and run preventative maintenance around switching, the component critical to on-time operations and dispatching. Taking a “holistic” approach, it monitors a wide range of infrastructural systems, physical assets, and movements and uses ML to detect issues and anomalies and plan maintenance or other switches in advance. The model is built to both predict future errors and eliminate downtime altogether.

ABB's RobotStudio is a simulation tool used to optimize robotic parts manufacturing in real time. The system uses offline programming using real-time data to test alternatives without shutting down production. Using RobotStudio, engineers can simulate real-world situations to identify problems with a new robot design. If problems exist, they can be addressed before the tool is actually made of steel and iron. RobotStudio is also used to create macros and to model welding, gluing, and image processing procedures.

Other companies working in this area include DataRPM, Machina Metrica, Pivotal, Falkonry, Simularity, Tellmeplus, and Augury, among many others. In some cases, this use case can enable new business models, wherein predictive modeling is used to support industrial IoT pricing models, such as supporting and replenishing assets based on wear and tear.

Tractica forecasts that the annual revenue for predictive maintenance in manufacturing will increase from \$0.43 million worldwide in 2016 to \$347.21 million in 2025.

Table 2.174 Predictive Maintenance in Manufacturing, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.43	6.48	16.05	31.00	53.84	87.53	134.71	196.22	269.29	347.21	110.6%

(Source: Tractica)

2.23.4 REAL-TIME VIDEO ANALYTICS

As camera technologies have sharpened the quality of video feed image precision, so has analytics supporting such capture. As video feeds have expanded in volume, video analytics represents the only way to extract value in form of insights, patterns, action, from so much data. AI is increasingly becoming a core enabler for video analytics, particularly for real-time analysis and action. DL, CV, and object and facial recognition enable accuracy and speed when it comes to analysis. DL also helps analyze and process multiple video and data streams and can help multiple systems communicate with each other. Common video analytics solutions may deploy various AI techniques to support the following areas:

- **Behavior Monitoring:** Motion detection, footfall or pedestrian traffic, facial detection, privacy masking, vandalism detection, theft or suspicious activity detection

- **People Monitoring:** People counting, people scattering, crowd analytics, line management
- **Vehicle Monitoring:** Vehicle classification, license plate monitoring, traffic monitoring, road monitoring
- **Device Monitoring:** Protection against tampering with camera, infrastructure, perimeter, or other intrusion

Use cases in manufacturing might include facilities security, employee tracking, operations monitoring, equipment, machinery, or cargo protection, theft or tampering prevention, intelligent logistics systems, compliance reporting, object tracking, etc. Essentially, video analytics technology helps security software “learn” what is normal so it can identify unusual and potentially harmful activities. The technology requires operator feedback as pure object detection is insufficient.

DVTEL’s ioimage product supports Marine Container Services’ 130,000-square foot facility plus 6,000 surrounding acres with advanced video analytics. The system protects some \$5 million dollars in goods stored, as vehicles are constantly entering and exiting the facility. Intelligent video surveillance analytics detects for specific events and provides 24/7 remote monitoring. Users define detection parameters and, once a threat is detected, a central station is notified and an operator manages a real-time response. Other companies in this space include 3VR, Cisco, Avigilon, and SightLogix.

Tractica forecasts that the annual revenue for real-time video analytics in manufacturing will increase from \$2.69 million worldwide in 2017 to \$126.44 million in 2025.

Table 2.175 Real-Time Video Analytics in Manufacturing, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	2.69	6.79	13.01	22.21	35.36	53.15	75.48	100.86	126.44	N/A

(Source: Tractica)

2.23.5

LOCALIZATION AND MAPPING

As manufacturing processes undergo radical transformation, relying less on human labor and more on machines, ML, DL, and CV are becoming central technology enablers for robotics and autonomous machines to reliably move about in manufacturing environments. Localization and mapping concerns the need and computational ability to simultaneously construct maps of the immediate environment, while updating both the agent’s position on that map and movement therein. In the context of manufacturing, localization and mapping is a core technique for autonomous movement of robots, cars, trucks, drones, or any other autonomous machine that moves. In manufacturing, localization and mapping may be used to expedite numerous tasks that fixed robots cannot, such as moving parts or products from one station to another.

ASTI is developing automated guided vehicles (AGVs) and mobile robots for use in factories and warehouses akin to automated forklifts, stackers, and pallet trucks. These machines are equipped with lasers and sensors to localize themselves, and are used for a variety of manufacturing and logistics-related tasks requiring no human intervention. Examples include moving large parts (e.g., 30-ton airplane parts); moving food (e.g., breads from station to station); automated battery changing; automatic intermediate storage; and more. The longer-term goal of the company is to free manufacturers from traditional factory assembly lines

and use automated vehicles to move parts (and processes) around. The company operates in more than 15 countries and counts the PSA Group Ltd., the manufacturer of Peugeot and Citroen cars, drug-maker GlaxoSmithKline Plc., Pepsi, Proctor & Gamble, Grupo Bimbo, SAD, and Spanish food-maker Campofrio Food Group SA among its clients.

Tractica forecasts that the annual revenue for localization and mapping in manufacturing will increase from \$0.30 million worldwide in 2017 to \$15.27 million in 2025.

Table 2.176 Localization and Mapping in Manufacturing, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.30	0.76	1.47	2.53	4.06	6.18	8.88	12.02	15.27	N/A

(Source: Tractica)

2.23.6 SENSOR DATA FUSION IN MACHINERY

Sensor data fusion is the technique used to aggregate, or “fuse together” multiple sensor data feeds and other data feeds in order to ascertain a more complete or multi-dimensioned picture of operations. The resulting multi-dimensional data offers less uncertainty than if the data feeds were viewed individually. In manufacturing, sensor data fusion concerns the ability for manufacturers to monitor machines and equipment and make sure they are functioning properly and will not fail. Using AI and DL for sensor data fusion is most advanced in automotive applications, as it is essential for minimizing risks or failure in cars, particularly automated car. But beyond auto manufacturers, sensor data fusion also applies to manufacturing equipment and machinery, so that multiple feeds (like temperature, vibration, cameras, or tension) could be monitored as a “whole” picture in order to reduce downtime, preemptively order parts, alert stakeholders, make environmental changes, etc.

Tractica forecasts that the annual revenue for sensor data fusion in machinery in manufacturing will increase from \$2.67 million worldwide in 2017 to \$137.59 million in 2025.

Table 2.177 Sensor Data Fusion in Machinery in Manufacturing, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	2.67	6.82	13.18	22.73	36.57	55.62	79.97	108.25	137.59	N/A

(Source: Tractica)

2.23.7 VOICE/SPEECH RECOGNITION

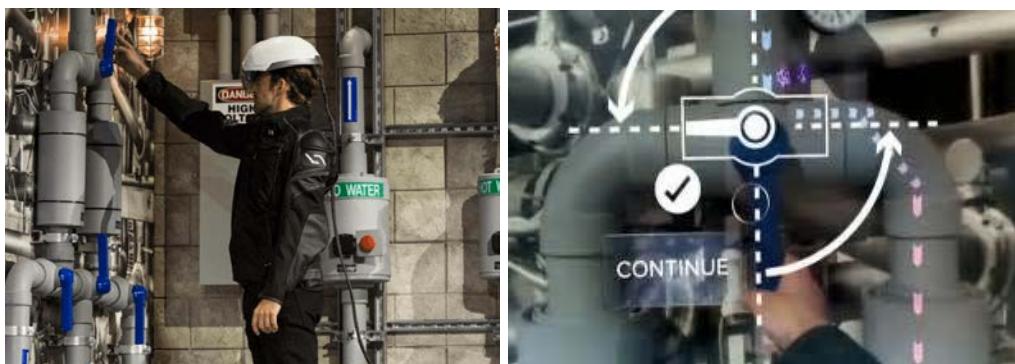
Until recently, voice and speech recognition were hardly a viable mode of interaction with computers or machines, not to mention dialog to which critical functions would be ascribed. In many industrial environments, resources are limited, time equals money, and errors can be costly. Finding economical, reliable methods to streamline employee tasks presents companies with an opportunity to gain a competitive advantage. Recent advancements in speech recognition technologies have led to a surge in development. Today, voice control represents a rapidly growing trend, as it vies to become the primary user interface in specific hands-free environments.

Manufacturing is one such industry in which voice/speech recognition will help technicians more easily give commands and control machines. Factory floor technicians or machinists can now use voice and speech recognition to do certain tasks (e.g., start, stop, move, bring

up an interface, conduct and log inspections, deliver status reports, ask for remote assistance on issues, authenticate security credentials, etc.). While this functionality may be built into industrial machinery, the more common platform for voice recognition in manufacturing environments is through wearable headsets.

DAQRI is a wearable smart helmet used in heavy industrial environments, such as manufacturing, telecom or utilities repair, logistics, and other factory settings. The helmets are powered with industrial inertial measurement units (IMU), 360° field of view via multiple navigation cameras, thermal imaging, and eye and speech recognition. Users interact with and control AR-overlaid data visualization, guided work instructions, and remote service access via voice recognition and gaze tracking.

Figure 2.19 DAQRI's Smart Helmet Combines Voice Recognition and Augmented Reality for Real-Time Work Instructions



(Source: DAQRI)

Other headsets providing this capability include GoogleGlass, RealWear, Honeywell's Vocollect product, and others.

Voice recognition in manufacturing and other industrial environments represents a promising improvement in interface and potential cost savings in the time required to carry out certain tasks. Factors like impaired speech, lack of accuracy, errors in data input or data processing, computing costs, potential battery life, or connectivity constraints will curb market growth in certain use cases and sectors. Please see Tractica's report on [Speech and Voice Recognition](#) report for a deeper analysis of use cases.

Tractica forecasts that the annual revenue for voice/speech recognition in manufacturing will increase from \$0.02 million worldwide in 2016 to \$0.69 million in 2025.

Table 2.178 Voice/Speech Recognition in Manufacturing, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.02	0.03	0.05	0.07	0.11	0.17	0.26	0.38	0.53	0.69	48.4%

(Source: Tractica)

2.24 MEDIA AND ENTERTAINMENT**2.24.1 ALGORITHMIC NEWS STORIES**

In the digital age of more content and access to information than ever before, news publishers and media outlets are exploring all manner of format, distribution, channel, and data collection. Most content generation using online data has been compiled using structured data, in addition to traditional journalistic data gathering and reporting. With advancements in NLP, ML, and DL, media outlets are now beginning to tap into vast amounts of unstructured data to automatically generate news stories. Some of the same techniques used in automated report generation, as in financial services, are used to compile algorithmic news stories.

Automated Insights is a company focused on self-service natural language generation for automated report generation and algorithmic news stories. Clients use the company's platform to automate writing using a proprietary platform called Wordsmith. In 2014 alone, the Automated Insights platform generated over 1 billion pieces of targeted, personalized content with a team of 50 employees. Specific markets resonating with the company are e-commerce, financial services, and business intelligence. To use the Automated Insights Wordsmith platform, clients upload their data using a comma separated values (CSV) format or the company's API. Next, clients design an article using Wordsmith's editor. Clients can control length, tone, and variability, so every article or report "is unique, compelling and individually personalized." The last step includes the creation of client-specific narratives.

The Associated Press is one of the first newsrooms to have an automation editor to oversee automated articles; the company said it would boost its output of quarterly earnings stories fifteen-fold, celebrating that the technology would allow journalists to do more journalism and less data crunching. Specific use cases for the Automated Insights system include crime trends, sales summary, election results, portfolio summary, real estate descriptions, workout recaps, salesforce reports, product descriptions, airline delays, and account summaries.

The Washington Post debuted its Heliograph software for the Rio games during the summer of 2016, using it to update medal counts and victories. It was used more extensively during the U.S. presidential election in late 2016 to handle simpler stories while human journalists focused on the election.

Startup Wibbitz creates video content from text articles for media companies, such as USA Today, TMZ, and Time. Video content for the company is licensed from Reuters and Getty Images. In April 2017, Reuters launched News Tracer.

According to Reginald Chua, executive editor for editorial operations in a Thomson Reuters company blog post, News Tracer:

...is a capability we've developed...that finds events that are breaking on Twitter... and assigns them a newsworthiness score so you can focus on the things that are important, and the real magic of it is that it then gives a confidence score about how likely it is that those events are true. This is really critical because the landscape of news has changed dramatically. One thing we found is that it's been very good at finding certain types of events much more quickly than many mainstream news organizations are able to do. It was ahead on the Brussels Airport bombing by several minutes. It was ahead on the Chelsea bomb by again several minutes. Since we started keeping analytical records about a year ago, Reuters News Tracer has beaten global news outlets in breaking over 50 major news stories. This has given our Reuters journalists anywhere from an 8- to 60-minute head start."

Tractica forecasts that the annual revenue for algorithmic news stories in media & entertainment will increase from \$5.25 million worldwide in 2016 to \$16.31 million in 2025.

Table 2.179 Algorithmic News Stories in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	5.25	5.66	6.16	6.78	7.57	8.60	9.94	11.67	13.81	16.31	13.4%

(Source: Tractica)

2.24.2 AUDIO AND VIDEO MINING

While a significant portion of digital text is structured data, the vast majority of audio and video content is unstructured data. Innovators in such areas as sales performance and marketing are beginning to convert audio and video into structured data and leverage it. Organizations in the media space, often driven by advertising-based business models, are also getting on board. As an extension of image recognition and analysis, AI is now also being used by organizations to aid in audio and video mining. In a media context, speech and voice recognition can be mined for specific moments, such as a user posting a video about a product. In an advertising context, AI can be used to transcribe, identify keywords, and mine audio, video footage, or online media. DL can also be applied here for auto-generated speech-to-text transcription.

DeepGram aids media organizations in rapid search and discovery of specific clips from past archives of video and audio. The tool helps journalists conduct research more rapidly for sourcing relevant soundbites and for searching through audio and video streams. It offered journalists free access to the tool in the weeks leading up to the U.S. 2016 presidential election. The company provides an API that allows users to apply audio and video mining to calls, meetings, podcasts, video clips, and lectures, and then rapidly search them. Other companies in this space include Veritone, Tagasauris, Valossa, and Yactraq. Some companies in the video analytics space will begin to offer new intelligent solutions for security and public safety.

Tractica forecasts that the annual revenue for audio and video mining in media & entertainment will increase from \$1.15 million worldwide in 2016 to \$450.81 million in 2025.

Table 2.180 Audio and Video Mining in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.15	7.11	16.99	32.94	58.05	96.31	152.24	229.71	330.14	450.81	94.2%

(Source: Tractica)

2.24.3 FILM SCENE STRUCTURE

Although arguably one of humanity's most creative endeavors, developing film concepts, scenes, and potentially even screenplays is an area where media companies and creatives are experimenting with AI.

The screenplay for a recent film called *Sunspring* was "written" entirely by AI. (The AI named itself Benjamin.) Using a type of RNN known as long short-term memory (LSTM), often used for text recognition, researchers from NYU's Film School fed the model dozens of sci-fi screenplays from the 1980s and 1990s. Over time, the model learned common sentence

formations and word associations, as well as the typical elements of a screenplay like stage directions, character details, and intonations. While the screenplay and subsequent film received high praise at Sci-Fi London's 48-Hour Film Challenge (for which it was developed), Benjamin was but an experiment, and only able to develop screenplays based on other content, not its own authentic voice.

Meanwhile, to kick off the 2017 Sundance Film Festival, actress Kristin Stewart published a paper with Cornell University that demonstrated the use of neural style transfer to recreate scenes from her short film *Come Swim*, in the impressionistic painting style that inspired the film.

Tractica forecasts that the annual revenue for film scene structure in media & entertainment will increase from \$0.29 million worldwide in 2017 to \$21.93 million in 2025.

Table 2.181 Film Scene Structure in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.29	0.77	1.54	2.77	4.63	7.36	11.14	16.04	21.93	N/A

(Source: Tractica)

2.24.4 FONT RECOGNITION AND SUGGESTIONS

Just as designers employ all manner of imagery and colors, fonts are also important tools for delivering consistent and appropriate tone, feel, and message. To identify a font, designers have historically relied on people that charge fees and take time to determine fonts.

Recently, the enterprise design platform giant, Adobe (or more specifically an intern with Adobe), developed a DL-based font recognition system called DeepFont. Similar to Google's reverse image search, where a user can link to or upload an image and image recognition detects and spits out other links with the same image, DeepFont scans images and spits out the fonts from the image. Adobe now ships DeepFont within its Photoshop and Typekit products.

Tractica forecasts that the annual revenue for font recognition and suggestions in media & entertainment will increase from \$13.77 million worldwide in 2016 to \$184.57 million in 2025.

Table 2.182 Font Recognition and Suggestions in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	13.77	17.75	23.71	32.46	45.01	62.47	85.69	114.80	148.65	184.57	33.4%

(Source: Tractica)

2.24.5 GESTURE RECOGNITION

Another key area in which AI is required to bring full functionality to a technology is with gesture recognition. The ability to accurately track and recognize gestures made by humans (who, by their nature, are not capable of repeating a gesture repeatedly using the exact same speed, position, or trajectory) requires algorithms that can account for these variances, as well as understand context. In media environments, gesture recognition can be used to

facilitate screenless or touchless/hands-free interactions in physical environments, such as in gaming, in retail, with robots, or in cars. Gesture recognition can also leverage 3D sensing technology and be used to interact with TVs, displays, and mobile devices like kiosks, robots, tablets, or smartphones. The best way to think about opportunities for gesture recognition is to think of what would naturally be better controlled with hands, body positions, eyes, facial expressions, etc.; for example, to point to a product on a display or on a shelf, or to swipe right or left if standing in front of a connected fitting mirror. Gesture control is particularly useful for retailers because data collected informs metrics, such as length of engagement, products viewed, and product popularity, and if the product is integrated with a loyalty program or facial recognition, it would supplement individual profiles.

Fluid Motion offers gesture walls, which are giant screens displaying information about a store's products or current campaign, and customers can navigate products or information using their hands. Fluid Motion offers the technology for both in-store engagement and in store windows, in order to facilitate engagement with passersby and when the store is closed. Customers can even select products from the store window and make a purchase without entering the store.

As a mode of interaction that does not require text, but may be more suitable in certain environments compared to voice, gesture recognition introduces new cost savings and potential revenue generation opportunities across a number of industries. Some companies supporting gesture recognition in retail include Leap Motion, GestureTek, Omek, GestSure, and Thalmic Labs.

Tractica forecasts that the annual revenue for gesture recognition in media & entertainment will increase from \$0.06 million worldwide in 2016 to \$10.17 million in 2025.

Table 2.183 Gesture Recognition in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.06	0.19	0.41	0.77	1.34	2.20	3.46	5.20	7.46	10.17	78.6%

(Source: Tractica)

2.24.6 HUMAN EMOTION ANALYSIS

It is no secret that humans are emotional creatures, often motivated more by emotion than pragmatism when making consumption decisions. Economists, content creators, and advertisers have understood this for years. Given the commonality of the advertising-based business model for online content, media companies have never been more concerned with staying emotionally in-tune with consumers.

Although computers are far better at calculating statistical probabilities than anything resembling emotion, developers are working to train models to recognize, categorize, and tag human emotions so that algorithms can make decisions based on such categorizations. Techniques could involve CV, DL, or NLP, or even robotics depending on the use case. While this is an emerging and controversial area of AI, early studies show computers are very adept at identifying human emotions. As a result, more companies are turning to AI to aid in the quest to better understand, predict, advertise, and display based on human emotions.

Kairos offers an emotion analysis API, which pulls in data from video content and offers the following metrics:

- Attention measurement via total attention time, number of blinks, glances, and attention span
- Facial expression detection via detecting joy, disgust, fear, smiles, frowns, anger, and surprise
- Emotion detection via mining for positive, negative, and neutral sentiments
- Ethnicity detection via “understanding the diversity of the human face”
- Gender and age detection via assigning probability scores to each detected face

These insights are used by media and content producers to test people’s responses to new or current programming. Reference Section 2.2.2 for a discussion on human emotion analysis used in advertising.

Tractica forecasts that the annual revenue for human emotion analysis in media & entertainment will increase from \$0.51 million worldwide in 2017 to \$38.79 million in 2025.

Table 2.184 Human Emotion Analysis in Media & Entertainment, World Markets: 2016-2025

Units (\$ Millions)	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016- 2025)
-	0.51	1.36	2.73	4.89	8.19	13.02	19.70	28.37	38.79	N/A	(Source: Tractica)

2.24.7 MUSIC PRODUCTION AND GENERATION

Musicians and artists have been using software to aid in music production for decades, not only to create sounds used in musical arrangements, but to aid in the editing and mastering process. As software tools have grown in sophistication, an emerging industry of artificial intelligence in music (AIM) has emerged with applications in the area of music composition, performance, theory, and digital sound processing. Algorithms are typically fed large amounts of audio data and learn rhythmic, tonal, melody, instrumental, lyrical, or other data, then produce their own enhancements when interacting with a human player or, in some cases, generate their own arrangements altogether. AI can also be used to analyze data across multiple formats, for example MP3s, PDFs, Wav., etc.

FlowComposer is a tool and part of FlowMachines research project developed with Sony Computer Science Laboratory that uses AI to convert musical styles into computational objects, which then applies melodies and harmonies. FlowComposer is fed a catalog of music, and then used as a collaborative tool by musicians, producing various melodic and harmonic sequences, interacting with the musician, and editing them. The tool relies on Markov chains, which describes a system in terms of states and probabilities of moving between states. The FlowMachine research project also developed DeepBach, a neural network system that produces harmonization in Bach style.

Another company, Jukedeck, uses AI to compose music, which it then sells to media producers and content creators. It is using DL to learn from, compose, and adapt music. Jukedeck saves such content producers, be they YouTube creators or brand publishers, money on royalties (the songs are royalty-free... today), and songs are generated based on user-defined preferences for mood, style, tempo, and length. And oMax learns in real-time the typical features of a musician’s style and plays along in “machine-sounding” tones interactively to mirror the player’s style. Google and IBM are also developing in this area.

Tractica forecasts that the annual revenue for human emotion analysis in media & entertainment will increase from \$0.51 million worldwide in 2017 to \$38.79 million in 2025.

Table 2.185 Human Emotion Analysis in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.51	1.36	2.73	4.89	8.19	13.02	19.70	28.37	38.79	N/A

(Source: Tractica)

2.24.8 NEWS AND FEED CURATION FOR CONSUMERS

The method in which news consumers go about consuming their news is increasingly digital, particularly through social media. In 2016, more than 62% of American adults got their news from social media, according to Pew Research. In an age of more information and more access to more information than ever before, the role NLP and DL play in news delivery is a growing, controversial, and important one.

For news accessed on digital channels (e.g., news sites), search and advertising play a significant role in click-throughs and what content is served up in what order, sometimes even using different headlines for the same article. For news accessed via social media, content is typically driven by advertising and users' social graphs. The application of neural networks and search and content curation in social media introduce a host of efficiencies, but also significant issues.

The contention lies in two areas. First, advertising revenue is a fundamentally different incentive than public awareness. Social networks are transforming into publishing networks, with paid social media content on the rise. This means that access to news is at the mercy of advertising models. Many questions around this remain unanswered, such as how information, especially news, should or will be prioritized or deprioritized in favor of advertising, and who or what decides and monitors this. The second area to note is that of the [in]ability to fully explain neural networks' decision making. Two people can search the same query and receive entirely different search results; while this may be ideal in one scenario ("pizza near me"), it is problematic in other scenarios wherein informed decision-making has grave implications.

Tractica forecasts that the annual revenue for news and feed curation for consumers in media & entertainment will increase from \$2.37 million worldwide in 2016 to \$714.01 million in 2025.

Table 2.186 News and Feed Curation for Consumers in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.37	11.84	27.50	52.77	92.53	153.09	241.59	364.18	523.09	714.01	88.5%

(Source: Tractica)

2.24.9 SIMULATING CROWDS

Crowd simulation is the process of using simulation software to train "agents" to interact in a scene. AI is used to simulate large crowds primarily in gaming and media applications. In **gaming**, certain environments may need to have large numbers of people or other

characters that possess certain functions within the game. In **media production**, such as movies or TV shows, MR can be used to coordinate crowd movement, given rules of spatial proxemics, and human territories, and automatically generate “ambient” interactions, (i.e., those happening in the background) like responding to newcomers or an event.

Unity3D is a game development platform that provides crowd simulation to gaming companies.

Tractica forecasts that the annual revenue for simulating crowds in media & entertainment will increase from \$9.2 million worldwide in 2016 to \$327.36 million in 2025.

Table 2.187 Simulating Crowds in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	9.20	13.85	21.24	32.89	50.94	78.15	117.68	172.21	242.73	327.36	48.7%

(Source: Tractica)

2.24.10 SOCIAL MEDIA PUBLISHING AND MANAGEMENT

Since the emergence of social media, a vast array of tools has emerged in order to help companies effectively identify, monitor, engage, and learn from user-generated content related to their markets. In the case of media and entertainment companies, social media is a tool to listen to fans and consumers, offer bite-size and curated content, and an additional content distribution channel.

ML, DL, and NLP are growing rapidly as tools for mining big unstructured data sets (e.g., social media posts, comments, reddit threads, online communities, etc.). In addition to many of the use cases outlined in Section 2.7.15, brands and publishers are using AI to support social media publishing and management in the following ways:

- Detect what consumers/readers/viewers/fans want by analyzing unstructured data, monitoring sentiment, content engagement, trends (more rapidly than teams of people)
- Determine what colors, images, text, hashtags, and other elements resonate most with specific audiences
- Recommend optimal spend for each post
- Bots can source and score well performing content, then recommend new content to post based on scores of past content
- Automatically tag, classify mentions, photos, logo placements, and brand mentions using image recognition
- Offer alerts, information, reminders to check out new campaigns, etc.
- Communicate in multiple languages

Cortex offers AI-driven social media marketing for brand publishers. The tool uses NLP and ML to break down social media content into component parts—deployment data (frequency, day, time); content data (colors, keywords, emojis, subject, hashtags); performance data (likes, comments, shares, click-thrus, etc.); and promotion data (ad spend). AI then analyzes these components to find patterns across all of those categories, competitors, etc. It also automates social media publishing by suggesting optimized editorial calendars and strategic

suggestions based on that data. Cortex works with the band Maroon 5 to offer intelligence and scale to their social media efforts. Maroon 5 has a presence across numerous social media channels and these profiles rely on constant streams of authentic content designed to inspire band love, interactions, and ticket and album sales. In the first four months of Maroon 5's engagement with Cortex, it increased fan engagement on Instagram by 98%, and 39% on Facebook. It also uses Cortex to inform ongoing content development and distribution strategies.

Tractica forecasts that the annual revenue for social media publishing and management in media & entertainment will increase from \$0.61 million worldwide in 2016 to \$1.175 billion in 2025.

Table 2.188 Social Media Publishing and Management in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.61	16.07	41.75	83.33	148.84	248.75	394.85	597.30	859.82	1,175.26	131.7%

(Source: Tractica)

2.24.11 VIDEO EDITING

As video development and production has grown more digital, the tools for editing video have as well. AI is now playing an increasing role in video editing, currently assisting in the process, with fully AI-edited videos on the horizon. Using a mix of ML, image, and object recognition, most video editing identifies action or highlights based on defined and learned parameters. Snapchat is a very simple but real-time use of facial recognition to support its filters. In some cases, as in Apple's Live Titles feature, its video editing app, Clips, uses speech recognition to insert captions into videos.

An app called Antix supports AI-powered video editing for GoPro footage by automatically identifying the "most exciting" highlights from the footage. Antix controls the GoPro wirelessly and analyzes sensor data in the user's smartphone to record motion metadata for the clip, along with video from the wireless feed from the camera. It autotags "exciting" content in real time. The company works with consumers, as well as professional athletes and brand content producers.

Magisto provides AI-powered video editing geared toward consumers for social video posting. The company uses algorithms to analyze the content and action of the video, identify the "most compelling" parts, and make edits in the appropriate places, and stitches them together for seamless auto-editing, even recommending complementary music selections. Users can then select graphical themes like "warm and fuzzy," at which point the algorithm takes the theme into account and adjusts stylistic options for title, special effects, soundtrack, etc. Shred is another company supporting AI-enabled video editing.

Tractica forecasts that the annual revenue for video editing in media & entertainment will increase from \$0.5 million worldwide in 2017 to \$37.73 million in 2025.

Table 2.189 Video Editing in Media & Entertainment, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.50	1.32	2.65	4.76	7.97	12.66	19.16	27.60	37.73	N/A

(Source: Tractica)

2.25 OIL, GAS, AND MINING

2.25.1 AUTOMATED REPORT GENERATION

Oil and gas and energy companies generate reports for internal stakeholders, as required by auditors and regulators for compliance, or as parts of client programs. Many functions remain reliant on manual processes, fragmented data, and legacy systems. Slow turnaround times, excessive effort spent on data collation and validation, and inconsistent reporting of results can ultimately create a variety of negative impacts and delays. As the amount of data flowing into and across organizations grows more and more massive, the problem is not just one of content distribution, but of the time it takes to comprehensively identify and organize insights that are useful and consumable.

AI is now a tool well suited for report generation. Using NLP, ML, and DL, in some cases, companies are using AI to collate reports far more rapidly than humans. Automated report generation tools generally support the following tasks:

- **Data Sourcing:** Identifies and extracts data from relevant internal and external sources, including industry news and reports, social media listening, and competitor intelligence.
- **Data Interpretation:** Upon consolidating data in standardized formats, the solution aligns the data in templates, codes, and prepares it for analysis using ML.
- **Data Analytics:** Defines business rules and correlation/causality at scale. With predictive modeling and data enrichment, solutions can run hundreds of “what if” scenarios and perform trend analysis.
- **Narrative and Semantic Commentary:** Using NLP and natural language generation, solutions can sometimes automate variance analysis and commentary writing in a systematic and structured way.

App Orchid is a SaaS-powered company that combines NLP, Big Data, machine intelligence, and data science in one toolbox to help companies process and analyze structured and unstructured data for business intelligence. App Orchid aids energy clients as part of its business and claims it can “improve grid reliability, increase efficiency and provide superior customer service” by tapping analytical data from smart meters, controls, and power line sensors, as well as unstructured utility-related data in reports, emails, white papers, field assessments, customer interactions, and personnel observations. In January 2017, the company announced a contract to provide such services to Energinet.dk, the Danish Transmission System Operator to manage the country’s energy grid. Another company, P2Energy, specializes in reporting for oil and gas and energy companies.

Tractica forecasts that the annual revenue for automated report generation in oil & gas will increase from \$10 million worldwide in 2016 to \$437.93 million in 2025.

Table 2.190 Automated Report Generation in Oil & Gas, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	10.00	18.99	32.15	51.26	78.66	117.28	170.35	240.86	330.43	437.93	52.2%

(Source: Tractica)

2.25.2 OIL PRODUCTION OPTIMIZATION

Geophysical feature detection is a critical part of the workflow in the oil and gas industry. Seismic surveys are carried out in the exploratory phase and during various other phases, from planning to field characterization before and during oil production. Once the data is gathered, the seismic traces are then processed and analyzed by human experts. Typically, this process can take several months.

Recently, Shell and MIT partnered to use AI techniques to automate this process and improve workflow efficiencies. Using DL, the raw seismic traces were analyzed to discover and locate subsurface faults in the underground structure, which are likely to contain hydrocarbons, before running migration and interpretation models. While there are still challenges in training and computational requirements, the study proved that geophysical feature detection could be automated.

Oil exploration capital expenditure is estimated to be around \$100 billion per year, so any savings and efficiencies brought about by geophysical analysis is expected to be adopted widely across the oil and gas industry.

Tractica forecasts that the annual revenue for oil production optimization in oil & gas will increase from \$15.98 million worldwide in 2016 to \$944.23 million in 2025.

Table 2.191 Oil Production Optimization in Oil & Gas, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	15.98	35.47	64.01	105.46	164.91	248.67	363.80	516.75	711.04	944.23	57.3%

(Source: Tractica)

2.26 REAL ESTATE

2.26.1 REAL ESTATE DEVELOPMENT OPTIMIZATION

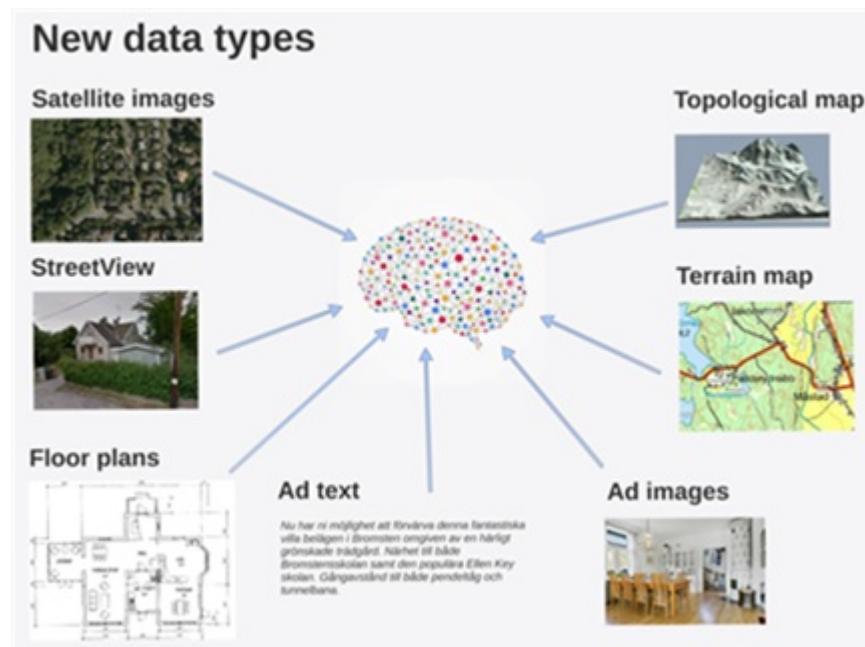
AI is being applied in real estate to better assess development opportunities. Using CV, developers are analyzing geographic images through drones and other technology to support models for valuation of properties and neighborhoods. In the past, property evaluation was one of the most important parts of a real estate broker's job. With hundreds of variables to analyze to correctly determine the value of any parcel of real property, predicting real estate values is a perfect use case for AI.

Seattle-based startup CityBldr has created a SaaS platform using AI to help determine the best use of all properties, to help developers find the most underutilized properties, and to help property owners understand the value of their properties as potential development sites by predicting what a developer might pay the property owner for the site's development value. The tool draws on 16 different public sources of data, including zoning codes, tax

history, transit, and parcel data, and generates 3 proprietary data sources. In the end, it analyzes more than 180 variables on 118 million U.S. properties to determine how plots of land can be improved to maximize their value.

Other applications include using DL to mine and extract key information from global contracts and real estate documents (e.g., leases, invoices, insurance policies, contracts, credit notes, etc.), sometimes in multiple languages. Data aggregation activities and particularly those involving unstructured data, such as in fiscal reporting, is another area where DL can be used to automate information extraction and expedite reporting. Leverton specializes in real estate data and document management, supporting large corporate real estate management companies.

Figure 2.20 Peltarion's Model Analysis of Millions of Data Points to Produce Real Estate Valuations



Swedish firm Peltarion uses DL to analyze data from historical data, object properties, demographic information, and nearby points of interest, using millions of data points to determine valuation. The platform is now used by most Swedish banks when administering home loans.

(Source: Peltarion)

Tractica forecasts that the annual revenue for real estate development optimization will increase from \$4.95 million worldwide in 2016 to \$440.74 million in 2025.

Table 2.192 Real Estate Development Optimization in Real Estate, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	4.95	15.09	30.95	55.19	90.90	140.84	205.74	282.54	363.94	440.74	64.7%

(Source: Tractica)

2.27 RETAIL**2.27.1 BEHAVIORAL ANALYTICS**

Gaining a deep understanding of customers has always been core to retailers' success, but the digital age has introduced a new universe of context, capabilities, and data available to ascertain such an understanding. Behavioral analytics used in retail to reveal insights around how consumers act and why; what motivates certain actions and when; and increasingly, where intention, consideration, interactions and transactions take place. What was born of e-commerce and online search, has been evolving to include real-world, in-store, mobile, in-home, and even in-car, as retailers collect and analyze as much data about their customers as possible.

As the volume and variety of consumer data has grown astronomically over the years, retailers, already struggling to connect e-commerce with in-store, are beginning to explore the use of AI for behavioral analytics. A range of methods and technologies are used to collect information about people's behaviors, including all manner of ML, DL, NLP, CV, and MR. Depending on the application, data may be processed in real time, draw on historical data, or over time; used for pre or post-purchase analysis; used for direct consumer interface or indirect analysis; or used for individual or group-level insights. Some applications include:

- **Segment Analysis:** Persona development; behavioral targeting
- **Individual Profiling:** Personalization; lead or retention prediction; behavioral targeting
- **Online/App Properties:** Site layout; UX and navigation; usage preferences
- **In-Store Properties:** Store layout; campaign displays; labor allocation
- **Product Recommendations:** If-then suggestions
- **Inventory Recommendations:** Predicting future sales trends; inventory needs, partner strategies
- **Security Risks:** Identifying or forecasting theft or fraud
- **AI Training:** Insights gathered to train models (e.g., speech, chatbots, etc.)
- **Staff Training:** Insights gathered to train employees (e.g., sales, support, etc.)

Security cameras with CV and video analytics, for example, are being used by retailers to monitor in-store and outside foot traffic, dwell time, and in-store engagement. Often, they are also used to identify suspect behavior and trigger intervention in cases of theft or crime.

Then there are applications like IBM Watson for Retail, which use Big Data analysis techniques to aid retailers with better engaging their customers in-store. Working with Sensitel, an IoT data analytics provider, the companies help retailers deliver smarter and more personalized shopping experiences. Specifically, Sensitel analyzes sensor, Wi-Fi, camera, digital displays, and other device data; monitoring movements and recognizing faces of shoppers. IBM helps enable the backend and data scientist workbench tools and libraries for high-speed analysis and development. Retailers can now monitor the location, facial expressions, and patterns of shoppers in real time, and direct staff to aid accordingly. For instance, ITM is working with a shoe store chain retailer to test and analyze how many minutes sales staff should wait before approaching customers.

Experimentations with different policies and approaches is helping the chains develop best practices. Another project conducted with a large multi-story retailer used behavioral analysis

to discover that very few customers were visiting the third floor of the building—an insight that can be optimized in a variety of ways. These techniques also help retailers identify improvements to in-store layouts, optimize labor allocation, inventory, and incorporate contextual signals useful for personalized customer experiences.

Tractica forecasts that the annual revenue for behavioral analytics in retail will increase from \$.06 million worldwide in 2016 to \$6.33 million in 2025.

Table 2.193 Behavioral Analytics in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.06	0.15	0.31	0.55	0.93	1.49	2.29	3.37	4.73	6.33	69.3%

(Source: Tractica)

2.27.2 CLOTHES SIZING AND FITTING

Sometimes it can seem a futile effort to search for the perfect fitting clothing, not to mention that tastes, fashions, and body types also change over time. A variety of AI techniques are being applied to help address this struggle. Some techniques, such as using image recognition and ML to “learn” one’s typical styles, colors, and fits are focused primarily on predicting fashion tastes and suitable outfits. A service like Thread, for example, asks users to submit photos of themselves, alongside their measurements, images of other clothes in their wardrobe, and budget constraints as inputs for their algorithms to curate (alongside human stylists) suggestions across a database containing thousands of pieces of clothing.

Other techniques use ML and CV and 3D scanning to automatically obtain measurements simply by having a shopper stand in front of a camera. In real time, such an application could assess measurements against a database of clothing and match an individual’s body shape and sizing with size profiles associated with specific shirts, pants, belts, etc.

Body Labs offers a horizontal solution that uses CV and 3D body motion tracking to map and model the body. Its solution is explored across a variety of applications, including companies working to enable individualized fitting and sizing for retail.

Tractica forecasts that the annual revenue for clothes sizing and fitting in retail will increase from \$11.66 million worldwide in 2016 to \$224.3 million in 2025.

Table 2.194 Clothes Sizing and Fitting in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	11.66	15.30	20.91	29.50	42.45	61.51	88.54	124.93	170.80	224.30	38.9%

(Source: Tractica)

2.27.3 CROWD ANALYTICS

Retailers have been monitoring and analyzing customer movement in B&M environments for years, whether through video cameras, beacons, or other sensing technology. AI introduces new capabilities to crowd analytics that involve image recognition and learning in CV-enabled cases, as well as in using neural networks to analyze, learn from, and predict information about traffic patterns, in-store displays, energy allocation, etc.

ShopperTrak is using DL to pose hypothetical foot traffic scenarios by inputting data from specific days, product launches, promotions, weather patterns, or other contexts to model and predict foot traffic, both digital and B&M. The model also uses back propagation to train itself over millions of simulations, by running predictions comparing outcomes to actual data, and then making adjustments accordingly.

Herta Security is a Barcelona-based company that uses DL for intelligent video analytics in malls, sports stadiums, airports, banks, and other retail environments. The company tracks and matches faces instantly and its system can be used to identify shoplifters and notify security personnel within 7 seconds.

Tractica forecasts that the annual revenue for clothes sizing and fitting in retail will increase from \$11.66 million worldwide in 2016 to \$224.3 million in 2025.

Table 2.195 Clothes Sizing and Fitting in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	11.66	15.30	20.91	29.50	42.45	61.51	88.54	124.93	170.80	224.30	38.9%

(Source: Tractica)

2.27.4

INTELLIGENT CUSTOMER RELATIONSHIP MANAGEMENT SYSTEMS

CRM systems have been helping organizations track and make sense of customer sales, marketing, and support interactions for years. What was born primarily as a sales tracking tool has expanded, with the advent of digital and social media, into robust platforms designed to unify insights around broader customer interactions and transactions, beyond just sales. Functionality tends to support at least four areas: contact management, customer acquisition, sales, and customer service. The goal of these systems is to facilitate “a single 360° view” of any individual customer, although this has been easier said than done, given the complexity of integrating online and offline customer profiles and behaviors. This is particularly challenging for legacy retailers with B&M environments.

AI is now infusing all aspects of CRM systems, and CRM more broadly. When it comes to **contact management**, companies are using ML and DL to mine large data sets for cleanliness and data integrity, purge bad data, help process incomplete contacts, suggest those to de-duplicate, etc. AI can be used to suggest potential contacts worth outreach as well. This is a particularly useful tool for sales enablement and **customer acquisition**. When it comes to sourcing, analyzing, prioritizing, and predicting prospective customers, AI is being applied for predictive lead scoring, suggested prioritization for sales outreach, and to optimize related sales workflows. ML and DL, in conjunction with NLP, are being applied for content curation and strategic outreach, wherein models process large data sets and then recommend specific content, offers, and outreach that may most resonate with particular kinds of prospects or customers.

AI-enabled CRMs are also helping companies assess which customers could be the most profitable and likely to respond to sales outreach. AI is also being used for **sales enablement, even predictive sales**. Similar to predictive or proactive customer service, AI can help scale sales agents read, triage, and respond to inbound prospects; analyze and predict the most appropriate action to take based on behavior and conversion trends; and even to filter, score, and prioritize similar leads. AI models take into account customer trends, but some companies, such as AgilOne, also fuse CRM data with external data from news, social media, weather, etc. to come up with sales leads and predictive pitches.

In retail environments, point of sale (POS) systems that collect contact and/or mobile data from in-store customers can be integrated with CRM, as can tablets or mobile devices used in-store to engage with customers. The plethora of emerging connected infrastructure in retail environments, such as connected displays, interactive mirrors, beacons, kiosks, and even robots, can also contribute insights around customer engagement, product interest, and other context to CRM tools. In effect, AI can be applied to every touchpoint—from e-commerce site layout to in-store personalization to augmenting loyalty.

Finally, the post-purchase phase of the customer life cycle is being enhanced by AI-enabled CRM systems as well. **Customer service**-related use cases enhance efficiencies on both enterprise and consumer sides. For consumers, the benefit should be more pain-free support experiences, void of redundant conversations or repetitive troubleshooting, and even delight through preemptive service actions. When tools like chatbots are effective, they can save customers time and energy. On the enterprise side, call centers and service agents are using AI to automate simple Q&A through chatbots; to automate triage and service escalation, activity capture, case classification, recommended responses, etc. AI is also increasingly used by service organizations to more efficiently allocate resources.

Many CRM providers, such as Salesforce.com, SugarCRM, Capillary Technologies, Infer, and AcuteIQ, provide AI-powered services across these four areas. Start-up ChiliData offers an AI-powered CRM system specifically for small to medium-sized business (SMB) clothing retailers. The company supports integration of multiple data streams to identify trends, develop unique client profiles, implement self-learning segmentation across customer demographics, automate communications for birthdays, sales, and surveys, develop key performance indicators (KPIs) for each customer, and offer an out-of-the-box Facebook chatbot for the brand. It also uses ML to automate product photo tagging using a visual recognition engine and identifying trends.

Tractica forecasts that the annual revenue for intelligent CRM systems in retail will increase from \$3.81 million worldwide in 2016 to \$40.27 million in 2025.

Table 2.196 Intelligent Customer Relationship Management Systems in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	3.81	4.49	5.51	7.03	9.29	12.57	17.20	23.40	31.19	40.27	30.0%

(Source: Tractica)

2.27.5 PREDICTIVE ANALYTICS FOR RETAIL

Retail is big business and includes a lot of Big Data. Bridging digital and physical worlds ("brick with click") demands analysis of extremely diverse and often unstructured data sets. Customer transaction and CRM data, browsing history, location data, sensor data, weather data, social media data, ad data, data from conversational commerce, and data across multiple websites are just a handful of the data retailers are mining to deliver highly personalized ads, product recommendations, marketing materials, purchasing options, campaigns, etc.

One of the greatest impacts of AI capabilities in retail is in supporting better analytics and real-time personalization through prediction and pattern detection. This includes customer-facing applications, such as those listed above, as well as a range of operational, logistical, and even legal ones, including but not limited to:

- Identification of individual store trends (e.g., unmet demand)
- Customer churn prediction
- Store layout and inventory efficacy
- Fraud detection
- Anomaly detection

These tactics, often used in conjunction with ML, can be extremely powerful for retailers to better understand and predict every aspect of their businesses. Retailers must be highly cognizant of consumer privacy protections, particularly when involving facial recognition and/or third-party data sets to identify individuals. For example, some fast food restaurants today are using vision-based AI to reliably read license plates of cars passing their franchise locations, then combining this data with public third-party data to relate the license plate information to an individual, and based on that input and their recognized movement patterns, create hyper-personalized marketing communications.

Tractica forecasts that the annual revenue for predictive analytics in retail will increase from \$10.47 million worldwide in 2016 to \$287.79 million in 2025.

Table 2.197 Predictive Analytics for Retail in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	10.47	15.07	22.23	33.30	50.10	74.92	110.20	157.76	217.76	287.79	44.5%

(Source: Tractica)

2.27.6 SENTIMENT ANALYSIS

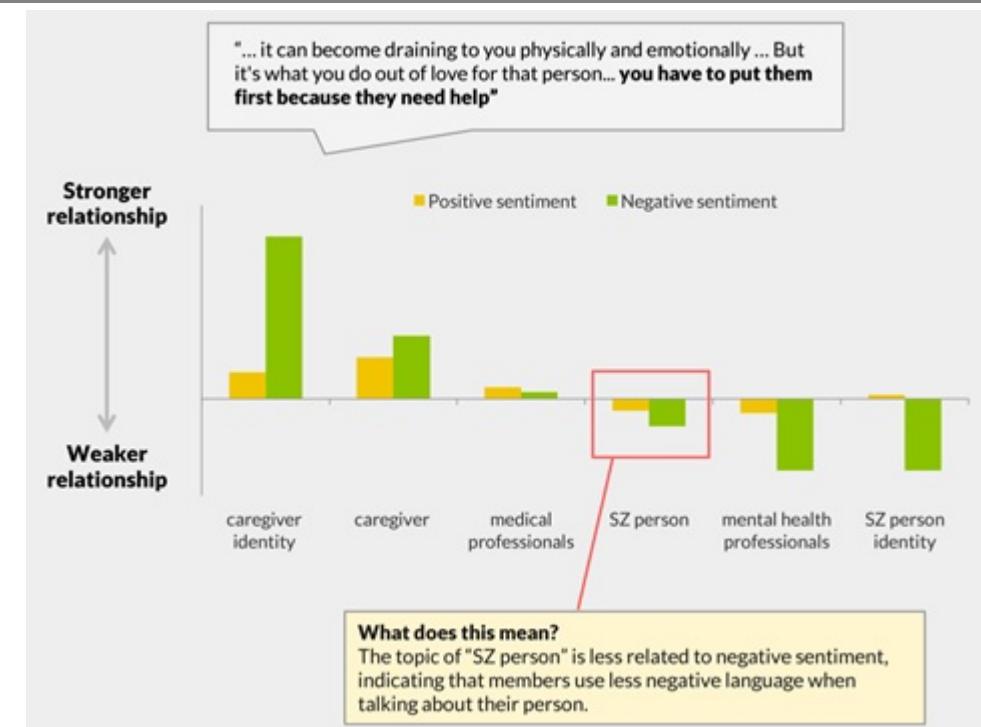
Understanding the emotional context of buyers has long been a tactic for retailers, just as it is in any social interaction. As commerce made its way into the digital realm, sentiment analysis, in which structured data is mined to understand shoppers' feelings and experience, has grown into its own industry. Sentiment analysis can be very useful for gaining an overview of public opinion, ideation, or feedback on a given topic.

Common approaches for measuring brand sentiment include the net promoter score (NPS), up/down votes, emojis, basic Likert scales, etc. Traditional sentiment analysis systems analyze text to return the sentiment as positive, negative, neutral, or mixed, based on dictionaries of positive and negative words, and define patterns that describe how to combine these words to form positive and negative phrases. Even when analyzing text like hashtags, however, these techniques often miss key insights or confuse sentiment for slang.

AI and NLP are now enhancing sentiment analysis by capturing and understanding the unstructured, more nuanced, and qualitative feedback, not just the best fitting response in a multiple choice. This data is combined with structured data sets for advanced analytics to surface trends. For example, retailers can track social media sentiment analysis, then using NLP, dig deeply into the rich nuances of comments and feedback. The ability to see beyond simple happy-neutral-angry or like-dislike then allows retailers to plan and act according to far more nuanced categories, personas, product lines, or campaigns. The majority of data available to most organizations is "dark," unstructured, and unused, but potentially full of valuable insights, so AI can be used as a tool to shed light on sentiments found in call logs, emails, transcripts, video and audio data, etc.

A large pharmaceutical company interested in optimizing C-Space, its online community of caregivers for people with schizophrenia recently partnered with AI software company Luminoso to better understand major issues these caregivers face and how to provide them with better resources and communications. Together, they used NLP and deep analytics on vast amounts of rich, but disparate and unstructured data, pulling together content from online communities, online discussion boards, multiple research projects, photo collages, and open-ended responses from surveys. Luminoso's software vectorized the data, meaning it effectively turns the text into mathematical vectors, then maps unstructured data based on relationships between topics and ideas. They uncovered a number of key themes and associations about the emotional composition of caretakers, their struggles, concerns, resource needs, and how they change over time. The pharmaceutical company also used the findings to improve community management, messaging, and support services.

Figure 2.21 Luminoso's Analytics Found Caregivers of Schizophrenic Patients Are Harder on Themselves than Any Other Group



(Source: Luminoso)

Tractica forecasts that the annual revenue for sentiment analysis in retail will increase from \$34.96 million worldwide in 2016 to \$1.322 billion in 2025.

Table 2.198 T Sentiment Analysis in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	34.96	55.82	88.66	139.66	217.37	332.45	496.25	717.33	996.37	1,322.13	49.7%

(Source: Tractica)

2.27.7

SUPERMARKET SHELF ANALYTICS

Retailers apply various technologies like weights, beacons, cameras, or other sensors to their in-store environments to help with inventory tracking, theft prevention, and other efficiencies. AI is now infusing many of these technologies and the software that supports them in order to automate learning and prediction. Data flowing from supermarket shelves carries important context for retailers, such as inventory, traffic flow, purchase frequency, linger time, etc. This data feed represents one of many valuable data sources that will drive a future trend of automated check-outs, where technology supports seamless “just-walk-out” purchase experiences. Whether through sensors, CV, or robotics, more and more AI-supported retail applications are incorporating shelf analytics.

Focal Systems is using CV and DL to drive more automated in-store experiences. The company offers a tablet that enables real-time out-of-stock detection, digital advertising opportunities, way-finding, and real-time promotions for shoppers. Shelf analytics are one feed that will help streamline B&M operations by eliminating lengthy product searches, checkout lines, out-of-stock items, and voice or tap-enabled shopper questions. A number of in-store robots coming into the market also include the ability to capture and analyze shelf data. Examples include Simbe Robotics and Intel's Tally robot and Fellow Robot's Navii robot used in Lowe's.

Tractica forecasts that the annual revenue for supermarket shelf analytics in retail will increase from \$0.09 million worldwide in 2016 to \$1.39 million in 2025.

Table 2.199 Supermarket Shelf Analytics in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.09	0.11	0.14	0.20	0.28	0.39	0.56	0.78	1.06	1.39	36.1%

(Source: Tractica)

2.27.8

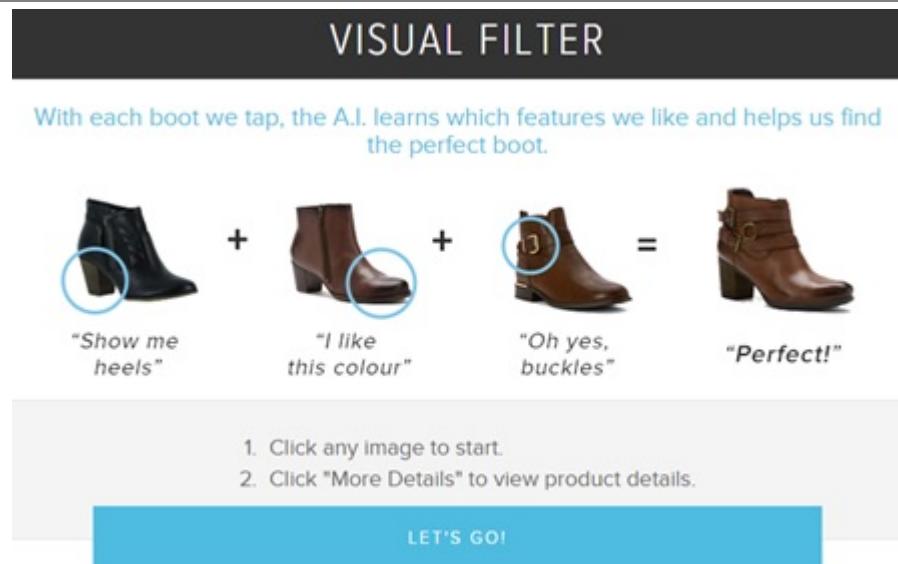
VISUAL SEARCH-BASED E-COMMERCE

The ability to see a product before buying it online is obviously essential to e-commerce, but it is nothing new. The problem is not one of seeing a product per se, rather one of reckoning text-based search inquiry with discovering exactly what a shopper wants. Shoppers spend hours and hours whittling down a short list by using filters, but this remains inefficient for both consumer and retailer.

What AI brings to visual search is image recognition in conjunction with NLP and ML. This allows e-commerce sites to respond in near-real time based on user inputs. Whether a user inputs their preferences via text or voice, or simply clicks on an image, that interaction can begin to train an algorithm to curate and display products that match what the user wants.

Sentient Technologies specializes in mining images and customer interactions to support this use case. Its customer, Shoes.com, is currently using DL to reflect recommended products based on what customers select in a 20-questions-inspired workflow. E-shoppers click the shoes most similar to what they are seeking, and the platform serves up personalized recommendations based on those initial selections. Instead of forcing shoppers to type out specific product feature requests (e.g., knee-high boot with 2-inch thick heel and pointy toe in red), the model uses image recognition to streamline the shopping process.

Figure 2.22 Shoes.com Partners with Sentient to Power Image-Based Search and Product Recommendations



(Source: Sentient Technologies)

Similar techniques are also applied to entire websites, with Sentient powering customers like Sunglass Hut and Cosabella with specific advertising images, applying A/B testing of aesthetics and workflows, and then feeding the model and, therefore, learning from all associated customer, image, ad, and behavioral data.

Tractica forecasts that the annual revenue for visual search-based e-commerce in retail will increase from \$4.41 million worldwide in 2016 to \$749.88 million in 2025.

Table 2.200 Visual Search-Based e-Commerce in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	4.41	15.97	34.49	63.61	108.30	174.80	269.75	398.14	560.38	749.88	76.9%

(Source: Tractica)

2.27.9 WEATHER FORECASTING

Retail companies benefit from the ability to forecast weather events as foresight helps model demand, supply, inventory, merchandizing, and commodities futures. (Reference Section 2.22.6 for an overview in the logistics space.) While retailers have understood the importance of weather forecasting on operational and supply chain efficiencies, more and more marketers are awakening to its importance in customer-facing contexts. Understanding how different types of weather impact store traffic, website traffic, sales in specific categories, weather-related promotions, weather-appropriate activities or resources, pricing, etc. helps retailers move both strategically and quickly.

AI and sensor data from hundreds of thousands of sources collected and monitored in real time (and over many years) are transforming the level of understanding and ability to forecast conditions. In addition to weather data, such engines combine streaming data from social

feeds, news reports, transportation data, and historical data on storms or other weather events. While no one can ever fully predict the future, AI techniques apply reinforcement learning on past predictions and actual outcomes. By comparing predictions with accuracies, the model is able to learn and improve simulation capabilities, as well as forecast much further into the future.

Retailers are working with IBM Watson and its recently acquired The Weather Company to deliver weather forecasting in support of communications, inventory optimization, and purchase rates, and to develop specific triggers based on how weather impacts customer behaviors. "We work with a national craft store brand to put together a "weather strategy," explains Paul Walsh, Director of Weather Strategy at IBM Global Business Services. "We essentially did an analysis to understand what type of weather had the biggest impact on traffic. First, we found that it was a local effect. For example, a rainy day in LA impacts consumer behavior differently than in Seattle. Or, say, snow in Atlanta impacts consumers differently than snow in Denver. This is important because you have to analyze the data at a local level in order to determine its impact. We then put in place a plan where the craft store could use weather predictions to alter their advertising message. For example, if it was going to be rainy, they would run ads that highlighted indoor crafting solutions. This, then, enabled consumers to be empowered rather than victimized by the weather."

Tractica forecasts that the annual revenue for weather forecasting in retail will increase from \$0.01 million worldwide in 2017 to \$0.72 million in 2025.

Table 2.201 Weather Forecasting in Retail, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.00	0.01	0.03	0.06	0.10	0.17	0.26	0.38	0.54	0.72	84.1%

(Source: Tractica)

2.28 SPORTS

2.28.1 ATHLETE FITNESS, SLEEP MONITORING, AND PERFORMANCE OPTIMIZATION

Unlike consumers or hobbyists using fitness trackers to count steps or track altitude, professional athletes use wearables and other technologies for high-precision monitoring to improve every element of their training and performance. Sleep, nutrition, training load, stress, travel, environment, interactions, movement, biometrics, wellness, injuries, and recovery are just some of the areas athletes are monitoring their data for performance optimization. As the sources and volume of data have expanded, along with sensor technologies powering wearables and cameras, AI has become a critical tool for athletes, coaches, managers, and medical providers to better understand and make decisions around all that data.

A leader in this area is an Australian company, Catapult Sports. Catapult offers both hardware (wearable) and software technologies that assess all aspects of a player's performance lives. The company works with thousands of players and hundreds of teams to design and scientifically validate metrics for performance optimization. Working with over 90 universities worldwide, Catapult has helped pioneer GPS tracking for team sports, having developed metrics around Inertial Movement Analysis and global navigation satellite system (GNSS) monitors to provide micromovement information that GPS and cameras are unable to capture and improve personal accuracy. Catapult has built extensive algorithms for tracking movement and positions unique to individual positions across over 30 sports, from bowlers

to quarterbacks.

One of Catapult's clients is the Houston Rockets, a National Basketball Association (NBA) team. Players wear devices that monitor the athlete's metabolic and musculoskeletal loads. These are used as part of periodization models, and use algorithms to determine for both individuals and the team optimum times and approaches for physical preparation and who is most physically ready at any given time. The Rockets analyze seasons upon seasons of player and game data, and are now able to take a proactive approach to physical demands, load levels, intervention strategies, rehabilitation, and readiness, and even identify the most debilitating physical setback a team could endure based on real-time health monitoring. A variety of other companies are working in this space, such as Zebra, Zephyr Technology Corp., and STATS LLC.

Clearly, some players and organizations push back on the privacy and discriminating impacts of such technology. As wearable data is used for selection, cited during contract negotiations, or potentially in legal contexts, many unions only allow the use of such devices during practices and not games. Although owners push back considering hefty sums paid out to professionals.

Tractica forecasts that the annual revenue for athlete fitness, sleep monitoring, and performance optimization in sports will increase from \$20.78 million worldwide in 2016 to \$507.16 million in 2025.

Table 2.202 Athlete Fitness, Sleep Monitoring, and Performance Optimization in Sports, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	20.78	40.43	69.26	110.14	165.08	233.56	310.84	388.16	455.82	507.16	42.6%

(Source: Tractica)

2.28.2 BIOMARKER-BASED ATHLETE PERFORMANCE OPTIMIZATION

As wearables have begun to pervade both consumer fitness and professional sports, the deeper question of how to leverage what data has advanced what performance optimization looks like. Professional athletes and coaches are now using advanced analytics, often powered by AI, to drive performance optimization based on biometrics like speed, distance, heartrate, history of injury, nutrition, and sleep, as well as by using genetic data and biomarkers. Alongside these other critical data sets, biomarkers found in blood can indicate general wellness, response to injury, and performance readiness. They can also enhance accuracy for nutrition personalization.

Orreco combines sports science and data science to offer professional athletes a platform for individualized performance optimization. Leveraging both individual data inputs and scientific research in biomarker analysis, GPS analytics, training load optimization, sleep analysis, performance nutrition, recovery protocols, and overtraining proclivity, Orreco uses IBM Watson to apply a range of data science techniques (DL, statistical modeling, predictive analytics, game theory, etc.) to offer individualized insights and strategies for refining performance at every level. DL, in particular, is used to discover previously unknown relationships between diverse data sets. The company works alongside athletes, as well as coaches and medical staff to provide hyper-personalized metrics and solutions.

Tractica forecasts that the annual revenue for biomarker-based athlete performance optimization in sports will increase from \$0.01 million worldwide in 2016 to \$26.55 million in 2025.

Table 2.203 Biomarker-Based Athlete Performance Optimization in Sports, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.01	0.55	1.46	2.91	5.10	8.18	12.16	16.81	21.73	26.55	133.9%

(Source: Tractica)

2.28.3 GAME OUTCOME PREDICTIONS FOR BETTING

Predicting the outcome of sports, casino, or other games is a business worth an estimated \$700 billion to \$1 trillion a year. Businesses and individual fans alike place significant money on who will win in just about every type of sporting event, from football to horse racing to poker. Entire adjacent industries benefit from game outcomes, including sports tourism, sporting goods manufacturing, advertisers, content creators, recruiters, etc. AI is being applied in this context to augment game outcome predictions by analyzing diverse data sets and past historical data to “learn from” and more accurately predict who will win, and when.

Vantage Sports has partnered with New Data Sports in developing an algorithm for pick forecasting. The company tracks obscure metrics that other industry trackers like the NBA do not measure, such as whether a pass was made to an open shot, or the number of times players contest shots (i.e., put their hands in the shooter’s face to block the shot). The company tracks dozens and dozens of metrics around player behavior for every team. Vantage Sports recently conducted DL analysis on its data versus public data and found that insights analyzed from its data yielded a 54% positive rate for prediction compared to 49% of correct predictions using public data.

Tractica forecasts that the annual revenue for game outcome predictions for betting in sports will increase from \$9.52 million worldwide in 2017 to \$469.34 million in 2025.

Table 2.204 Game Outcome Predictions for Betting in Sports, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	9.52	25.51	51.14	89.91	144.49	214.85	297.05	384.19	469.34	N/A

(Source: Tractica)

2.28.4 SPORTS STATISTICS ANALYSIS AND SEARCH

Millions of people worldwide follow sports. But for many in the digital age, following sports is more than watching games; rather, it is an ongoing and up-to-the second analysis of individual players, coaches, and managers; performances, predictions, and wagers; and even fantasy sports. Beyond fans, pro-teams and leagues use analytics to aid in all kinds of decisions, from building sponsorships to predicting injuries. AI for sports statistics analysis and search has numerous applications, mostly involving ML, NLP, and sometimes DL and CV. Given the vast amounts of data generated by each player, training session, game, and team, ML and DL are being explored as a means to make smarter training regimens, proclivity to stress or injury, and nutrition. Other data, such as fan engagement and composition, and geo-mapping or beacon data help feed recommendation engines to

improve fan experience, build loyalty, develop campaigns, personalize marketing, and predict ticket sales and renewals.

To support fans' ability to keep up with all the highlights of the 2017 Wimbledon, IBM Research and IBMiX used the Cognitive Highlights solution to assess all footage and automatically produce highlight packages for each match, then distributed that across digital platforms as soon as the play closes. To support these intelligent highlights over the course of the 13-day tournament, IBM used CV, combined with information from an on-court statistician pulling in data from an array of sensors tracking distance, serves, and speed. Using audio and video footage gathered from previous championships, the system was trained to recognize fan reactions and incorporate this and player reactions into the model.

Meanwhile, European soccer teams are putting wearables on players to track a range of biometrics, speed, location, training performance, and beyond. In basketball and baseball, many unions are against wearables on players; however, leagues like the NBA place SportVU cameras in the rafters, to collect player data via CV. Major League Soccer (MLS) teams in the United States have been using the Adidas mCoach Elite system for data collection and analysis. MLS teams in both the United States and Canada have been slowly creating a giant data warehouse for fan data from all of its 19 soccer clubs and countless CRM and ticketing systems. The data is used to create personalized campaigns and increase sales and loyalty across millions of fans.

For fans, an app called Statmuse offers an anytime real-time chatbot that users can ask and receive different stats for different players.

Tractica forecasts that the annual revenue for sport statistics analysis and search in sports will increase from \$0.77 million worldwide in 2016 to \$5.69 million in 2025.

Table 2.205 Sports Statistics Analysis and Search in Sports, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	0.77	0.90	1.10	1.39	1.81	2.38	3.10	3.94	4.82	5.69	24.9%

(Source: Tractica)

2.28.5 SPORTS TEAMS PLAYERS SELECTION

Selecting players in professional sports team is a careful process, as recruiting managers and coaches weigh many variables when searching for talent, chemistry, and increasing team success. The use of data in sports team player selection was made famous in the movie *Moneyball*, which tells the story of how the Oakland As used data analytics to select a winning team in 2002.

AI is taking the use of data for player selection a step further. Given the vast amounts of data generated by each player, training session, game, and team, ML and DL are being explored as a means to predict player outputs. Attributes include things like body composition, flexibility, anaerobic and aerobic power, visual tests, attention tests, interpersonal tests, psycho-motor skills, skill assessments, practice and game performances, etc. Applications vary widely and include the use of neural networks to predict training loads, injuries, individual performance, coachability, team performance, etc. Coaches can also leverage this to supplement decision-making around which players should play on a certain day depending on the opposition, health, position, and best team configuration.

The Australian Football League (AFL) has done significant research in this area and employs the use of neural networks to aid in player selection. WaiverWire.com uses ML to power a fantasy football drafting tool that recommends optimal player selection at each decision point during the draft. The tool combines proprietary ML models for point projections and qualitative adjustments based on content analysis to account for variables like projected performance, team changes, other players in the pool, or other factors. ESPN offers a similar tool for its March Madness bracket pools.

Tractica forecasts that the annual revenue for sports team player selection in sports will increase from \$0.02 million worldwide in 2017 to \$0.82 million in 2025.

Table 2.206 Sports Team Player Selection in Sports, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.02	0.04	0.09	0.16	0.25	0.37	0.52	0.67	0.82	N/A

(Source: Tractica)

2.29 TELECOMMUNICATIONS

2.29.1 PREDICTIVE MAINTENANCE

Current telecommunications services typically rely on heavy equipment, machinery, transformers, lines, boxes, poles, and a range of other infrastructure to maintain connectivity, reliability, and security. When parts go down, costs incurred are manifold: costs of machines, costs of maintenance required (i.e., labor, emergency rates), costs of downtime, and (often untold) costs of customer frustration and loss, especially when customers are businesses. The ability to manage so much capital outlay is critical. As in other industries like manufacturing, oil & gas, transportation, and beyond, telecommunications companies are also applying AI to predictive maintenance.

DataRPM provides predictive maintenance solutions for telecommunications and automotive suppliers. Its platform connects diverse sensor data across other data sets into a data lake where it “cleanses” the data, uses feature engineering and clustering, and runs multiple ML iterations and combinations of sensors to detect anomalies and suspect patterns associated with machine failure. Using derived patterns, the platform then generates labeled training data, enhanced by user validation, to enable faster predictions. Over time, this builds an ensemble of predictive models for manufacturing configurations to minimize any future failures. The company recently worked with a large telecom provider and was able to identify 85% of customers likely to be affected by set-top box failures—a staggering 40% of overall customers—in addition to the reasons behind failure. Using a combination of attributes that affect set-top boxes, such as manufacturers, software version, and time since install, it ran some 50,000 predictive models, selected the one with the best prediction accuracy, and used the analytics to target those customers with the highest likelihood of box issues with proactive retention marketing campaigns. In total, this accounted for some 36% of the company’s customer base.

Chinese telecommunications giant Huawei recently partnered with GE's industrial internet cloud platform Predix to support connectivity between industrial assets (Huawei's edge computing IoT (EC-IoT) and cloud applications. This partnership is geared specifically to allowing real-time machine health monitoring, data analysis and perception, and smart maintenance decision-making particularly in environments with bandwidth constraints.

Tractica forecasts that the annual revenue for predictive maintenance in telecommunications will increase from \$9.81 million worldwide in 2016 to \$285.47 million in 2025.

Table 2.207 Predictive Maintenance in Telecommunications, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	9.81	15.59	24.31	37.32	56.33	83.29	119.97	167.12	223.53	285.47	45.4%

(Source: Tractica)

2.29.2

PREVENTION AGAINST CYBERSECURITY THREATS

Cybersecurity represents one of the greatest threats to telecommunications providers, particularly as the volumes of data and complexity of IT networking infrastructure grow. All computer databases are, to some extent, vulnerable to being hacked. Today's devices, machines, and vehicles (aerial and otherwise) have more control units, computing power, lines of code, and wireless connections with the outside world than ever before—this renders them both more “intelligent” in connectivity, but also more vulnerable to hackers. In these scenarios, ML and DL are used to aid in learning from threats and predicting optimized protection for all types of telecommunications infrastructure, assets, and networks. Specifically, telecom companies can leverage ML, DL, and MR to review massive amounts of data to detect suspicious behavior, foresee equipment failure or downtime, identify threat types and profiles, and protect confidential information.

AI development is now targeting how to respond to cyberattacks on networks, working to quickly block suspicious communications and analyze malicious behavior and software—tasks still often allocated to humans. When under attack, the system will be able to identify the entry point and stop the attack, as well as patch the vulnerability.

DarkTrace is a startup in this space that aspires to mimic the human immune system in its response to security threats. Its Enterprise Immune System technology can detect previously unidentified anomalies and potential threats in real time, which other legacy approaches either fail to see or take longer to eradicate. By applying its unsupervised ML system, DarkTrace claims it has identified 30,000 previously unknown threats in over 2,400 networks, including zero-days, corporate espionage, IoT hacks, criminal campaigns, insider threats, and more stealth attacks. The company works with telecom providers, such as British Telecom, Telstra, and T-Mobile, to protect highly complex networks and datasets containing confidential information.

Reference Section 2.7.11 for an overview of cybersecurity tools used in enterprise settings.

Tractica forecasts that the annual revenue for prevention against cybersecurity threats in telecommunications will increase from \$7.23 million worldwide in 2016 to \$452.61 million in 2025.

Table 2.208 Prevention Against Cybersecurity Threats in Telecommunications, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	7.23	16.55	30.65	51.67	82.39	125.95	185.21	261.40	352.54	452.61	58.4%

(Source: Tractica)

2.29.3

IMPROVE CUSTOMER EXPERIENCE MANAGEMENT

In telecommunications, CEM refers to managing the telecommunications experience and quality of service. While there are a number of overlaps in functionality with traditional and “intelligent” CRM systems, outlined in Section 2.7.9, CEM is used in telecommunications to support various elements of customer experience that most CRM systems are not tooled to handle. Examples include auto-adjusting network parameters, service quality detection, website quality detection, and addressing network performance or security needs in real time.

Mobile and network service providers are now leveraging AI in a number of ways that both enhance customer experience and help automate quality of service. Advanced chatbots and virtual agents can use ML and NLP to handle support interactions via SMS or other messenger platforms and make necessary changes or updates. Network and device data can be used to predict and preemptively execute provisioning or other automation to optimize reliability. Real-time rating, charging, and mediation capabilities can streamline billing processes. Ongoing qualitative and quantitative customer interactions, requests, complaints, service logs, and cross-channel portals can be analyzed using ML, NLP, and DL to detect trends or performance issues across demographics, devices, time, or location. With integrations across CRM, operations tools, call center solutions, social analytics, etc., AI can help CEM systems convert interactions into insights across the entire customer and device life cycles.

Nokia offers an AI-enabled CEM tool for telecom providers, which offers cognitive analytics supporting customer insights, crowd insights, and marketing insights, and integrates this across access analytics, care and provisioning, payment, traffic monitoring, security, customer care, and device management.

Tractica forecasts that the annual revenue for improving CEM in telecommunications will increase from \$25.80 million worldwide in 2016 to \$660.91 million in 2025.

Table 2.209 Improving Customer Experience Management in Telecommunications, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	25.80	39.10	59.20	89.18	132.97	195.10	279.60	388.24	518.20	660.91	43.4%

(Source: Tractica)

2.29.4

FRAUD MITIGATION

Telecommunications fraud is the theft of telecommunication services or the use of telecommunication service to commit other forms of fraud. Fraud primarily occurs to a company with a weak defenses or poorly protected telecom infrastructure. Billing systems, VoIP, voice technologies, and network vulnerabilities can be exploited to gain access. Detecting transaction fraud is an ongoing priority (not unlike cybersecurity) as fraudsters constantly adjust their tools and methods in the digital age. While fraud detection software has been on the market for some time, approaches relying solely on historical data and business rules are insufficient to mitigate the evolving threat. According to the Communications Fraud Control Association (CFCA) 2011 Global Fraud Loss Survey, the CFCA estimates that telecom fraud costs the industry over \$40 billion annually.

To detect fraud schemes in telecom, AI, ML, NLP, and DL are being explored in ways that do not solely rely on pre-programmed rules or models based on historical data. The goal for such systems is to become self-learning, where models continuously update individual profiles, threat profiles, payment methods, situations, behaviors, and other parameters. AI is also useful here in helping process multiple data types, as new payment types and methods require flexibility in data processing. In addition, analyzing credit/debit card usage patterns and device access allows security specialists to identify points of compromise.

Meanwhile, the IoT and other technologies will create a very large number of endpoints for telecoms. With 5G, the nature of the network will change (i.e., high-density urban deployments). This will create a large number of patterns, difficult for humans to ascertain. Because AI is self-learning, it will also be used to help manage these patterns, fraudulent or otherwise.

Tractica forecasts that the annual revenue for fraud mitigation in telecommunications will increase from \$2.77 million worldwide in 2016 to \$122.33 million in 2025.

Table 2.210 Fraud Mitigation in Telecommunications, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	2.77	5.28	9.06	14.70	22.95	34.64	50.55	71.00	95.47	122.33	52.3%

(Source: Tractica)

2.29.5

INTELLIGENT CUSTOMER RELATIONSHIP MANAGEMENT SYSTEMS

CRM systems have been helping organizations track and make sense of customer sales, marketing, and support interactions for years. What was born primarily as a sales tracking tool has expanded, with the advent of digital and social media, into robust platforms designed to unify insights around broader customer interactions and transactions, beyond just sales. Functionality tends to support at least four areas: contact management, customer acquisition, sales, and customer service. The goal of these systems is to facilitate “a single 360° view” of any individual customer, increase customer share and retention, reduce churn, and raise revenue. In the telecom and mobile operator space, there are extensive opportunities given the near-constant connectivity and data flow involved in customer relationships.

In telecom environments, AI can be applied to nearly every aspect of the relationship, from e-commerce site layout to personalized notifications and promotions, to real-time responses to service requests, to augmenting loyalty and reducing churn. As 5G grows in availability, competitive pressures will force more dynamic customer relationships.

AI is now infusing all aspects of CRM systems, and CRM more broadly. When it comes to **contact management**, companies are using ML and DL to mine large data sets for cleanliness and data integrity, purge bad data, help process incomplete contacts, suggest those to de-duplicate, etc. AI can be used to suggest potential contacts worth outreach as well. This is a particularly useful tool for sales enablement and **customer acquisition**. When it comes to sourcing, analyzing, prioritizing, and predicting prospective customers, AI is being applied for predictive lead scoring, suggested prioritization for sales outreach, and to optimize related sales workflows. ML and DL, in conjunction with NLP, are being applied for content curation and strategic outreach, wherein models process large data sets and then recommend specific content, offers, and outreach that may most resonate with particular kinds of prospects or customers.

AI-enabled CRMs are also helping companies assess which customers could be the most profitable and likely to respond to sales outreach. AI is also being used for **sales enablement, even predictive sales**. Similar to predictive or proactive customer service, AI can help scale sales agents read, triage, and respond to inbound prospects; analyze and predict the most appropriate action to take based on behavior and conversion trends; and even filter, score, and prioritize similar leads. Not only do AI models take into account customer trends, but some companies, such as AgilOne, fuse CRM data with external data from news, social media, weather, etc. to come up with sales leads and predictive pitches.

Finally, the post-purchase phase of the customer life cycle is being enhanced by AI-enabled CRM systems as well. **Customer service**-related use cases enhance efficiencies on both enterprise and consumer sides. For consumers, the benefit should be more pain-free support experiences, void of redundant conversations or repetitive troubleshooting, and even delight through preemptive service actions. When tools like chatbots are effective, they can save customers time and energy. On the enterprise side, call centers and service agents are using AI to automate simple Q&A through chatbots; to automate triage and service escalation, activity capture, case classification, recommended responses, etc. AI is also increasingly used by service organizations to more efficiently allocate resources.

Many CRM providers, such as Salesforce.com, SugarCRM, Capillary Technologies, Infer, and AcuteIQ, provide AI-powered services across these four areas.

Tractica forecasts that the annual revenue for intelligent CRM systems in telecommunications will increase from \$34.56 million worldwide in 2016 to \$862.39 million in 2025.

Table 2.211 Intelligent CRM Systems in Telecommunications, Annual Revenue, 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	34.56	51.89	78.09	117.16	174.25	255.23	365.37	506.98	676.37	862.39	43.0%

(Source: Tractica)

2.30 TRANSPORTATION

2.30.1 MACHINE/VEHICULAR OBJECT DETECTION/IDENTIFICATION/AVOIDANCE

Perhaps the most valuable use of AI in vehicles is the use of object detection and classification, which takes sensor data, often from cameras, and then uses complex algorithms to classify these objects so that the AI system can then “learn” their characteristics, and recognize them in real time.

The challenge is not in capturing images, as today’s HD cameras can present images in stunningly clear detail. However, in a moving environment, objects can appear to change size as a vehicle or camera approaches. The angle at which an object is viewed can also skew its appearance, and the presence of other factors (rain, bright sunlight, low lighting, glare, dirt, snow, or any other number of obstructions) can alter the appearance of an object, making it hard to accurately and consistently identify the object. This is an area where machine vision and ML can provide invaluable support. By capturing a wide range of images of objects from a variety of vantage points, angles, and in different conditions, a repository of images that can be definitively classified as that object can be created, and used to “train” a ML system to identify and classify objects that resemble objects in the repository.

By then assigning various other attributes to each object, such as whether the object is informational like a traffic sign, whether or not it is permanent or temporary like a road barrier, or whether or not it has the capability of motion and how it typically moves, the system can begin to develop logical rules on handling each object and the rules for dealing with them.

Tractica forecasts that the annual revenue for machine/vehicular object detection/identification/avoidance in transportation will increase from \$3.04 million worldwide in 2017 to \$123.68 million in 2025.

Table 2.212 Machine/Vehicular Object Detection/Identification/Avoidance in Transportation, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	3.04	7.49	13.93	23.09	35.74	52.51	73.49	97.84	123.68	N/A

(Source: Tractica)

2.30.2 PREDICTING TRAFFIC DENSITY

One of the biggest applications for AI in the transportation space is in predicting traffic density and flow. In an age of smart and autonomous vehicles, accuracy in traffic flow is a key enabler of broader intelligent transportation systems. Improving accuracy in these systems is key for traffic operational efficiency, reducing carbon emissions, alleviating traffic congestion, helping road users make better decisions, and improving overall municipal efficiency.

DeepDrive is a research project out of the University of California at Berkeley that is currently using deep reinforcement learning in conjunction with microsimulation methods to support urban traffic optimization for future smart cities and smart cars. Using diverse data streams from both cars and cities (i.e., GPS-based measurements of speeds and delays, sensor data, static flow measurements with magnetic loops, Bluetooth re-identification, odometer data, and raw video feed), the project aims to support optimization and shared learning for autonomous vehicle traffic patterns. In other studies, even more data inputs from crowdsourcing and social media are fed into models.

Tractica forecasts that the annual revenue for predicting traffic density in transportation will increase from \$6.77 million worldwide in 2016 to \$439.02 million in 2025.

Table 2.213 Predicting Traffic Density in Transportation, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	6.77	15.82	29.50	49.91	79.72	122.00	179.51	253.45	341.90	439.02	59.0%

(Source: Tractica)

2.30.3 SENSOR DATA FUSION IN MACHINERY (SHIPS, UNMANNED SHIPS)

Sensor data fusion is the technique used to aggregate, or “fuse together” multiple sensor data feeds and other data feeds in order to ascertain a more complete or multi-dimensioned picture of operations. The resulting multi-dimensional data offers less uncertainty than if the data feeds were viewed individually.

In transportation, sensor data fusion concerns the ability for operators and controllers to monitor vehicles, ships, and equipment—manned or unmanned—and make sure they are functioning properly and will not fail. Using AI and DL for sensor data fusion is most advanced in automotive applications, as it is essential for minimizing risks or failure in cars, particularly automated cars. But beyond vehicle manufacturers, sensor data fusion could also be used for maritime surveillance, to detect abnormal behavior, or potentially even in port environments to help automate arrivals. As in other sectors, it is useful for monitoring a “whole” picture in order to reduce downtime, preemptively order parts, alert stakeholders, make routing or environmental changes, etc.

Tractica forecasts that the annual revenue for sensor data fusion in machinery in transportation will increase from \$1.45 million worldwide in 2017 to \$66.59 million in 2025.

Table 2.214 Sensor Data Fusion in Machinery in Transportation, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	1.45	3.61	6.80	11.44	17.99	26.86	38.22	51.77	66.59	N/A

(Source: Tractica)

2.30.4 LOCALIZATION AND MAPPING

As transportation processes undergo radical transformation, relying less on human operation and more on machines, ML, DL, and CV are becoming central technology enablers for autonomous vehicles to reliably move about in the world. Localization and mapping concerns the need and computational ability to simultaneously construct maps of the immediate environment, while updating both the agent's position on that map and movement therein. AI systems can provide that visibility via a model through two variables: an unknown variable, which is the location of the car, and observations about the car's location based on the sensor inputs at that given time. The AI component takes these two variables and, based on a randomized algorithm that repeatedly samples possible scenarios, returns a best estimate for where the vehicle currently is situated. These models can be refined over time by also incorporating HD, 3D maps, which provide more accuracy than typical 2D maps provided by Google and others. In the context of transportation, localization and mapping is a core technique for the autonomous movement of cars, trucks, ships, or any other autonomous machine that moves.

The V-Charge project is a collaboration between the EU Consortium of Volkswagen, ETH Zurich, Bosch, the University of Oxford, and others that aims to advance the development and adoption of autonomous electric vehicles to reduce traffic congestion and parking inefficiencies, and global CO₂ emissions in order to meet sustainable development goals. An essential part of this work involves enabling autonomous indoor navigation without modifications to infrastructure (e.g., parking lots or buildings); precise environmental perception and control in order to automate parking in tight spaces; and scheduling algorithms for parking spot and charging station assignment so as to make drop-off and pick-up seamless. All of these require high-precision localization and mapping in order to achieve full electric vehicle autonomy, and the objectives of the project.

Another project for the first Autonomous Bus program in the United States was recently announced by Proterra, the University of Nevada, Reno, and the Living Lab Coalition. Unlike other pilots to date, this project will pilot real road conditions from a public transit systems perspective, dynamic and dense navigation environments, and require quick emergency responses. Part of the project also involves refinement of robotic perception algorithms

required to respond to cues from multimodal sensors and localization and mapping optimization.

Tractica forecasts that the annual revenue for localization and mapping in transportation will increase from \$0.84 million worldwide in 2017 to \$33.18 million in 2025.

Table 2.215 Localization and Mapping in Transportation, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	0.84	2.05	3.80	6.27	9.67	14.17	19.79	26.29	33.18	N/A

(Source: Tractica)

2.30.5 VEHICLE NETWORK AND DATA SECURITY

As the transportation industries develop more connected and autonomous vehicles, they grapple with the nightmarish threat of cyber-hacking or terrorism of its fleets. Today's vehicles have more control units, computing power, lines of code, and wireless connections with the outside world than ever before, which is why vehicles of the future are cause for great security concerns. A recent study by Munich Re, the world's second-largest reinsurer, found that 55% of corporate risk managers surveyed named cybersecurity as their top concern for autonomous vehicles.

Even today, many systems within vehicles are separated so as to avoid penetration scenarios, where malicious actors enter through one system and attack another. There are two broad areas of vulnerability: network security, including command and control systems, databases, and communications (which all rely on network security); and platform security, including operational systems, engineering plants, and applications. Then there remains the constant internal threat, in the event an employee knowingly or unknowingly uploads malware into a critical system. Data security of drivers and their devices also cannot be ignored. There are also threats along the ecosystem: traffic controls, mobile devices, in-vehicle Wi-Fi, third-party vendors, etc. As manufacturers and operators gain increasing visibility into fleets of machines, sensors, data, and networks simultaneously open up new vulnerabilities and new security methods.

AI can be applied in an IoT security context, in which various techniques like ML, sensor data fusion, DL, CV, and MR can be used to enhance machine and device security by monitoring sensor and environmental data, analyzing systems and anomalous events, and acting accordingly. AI could pull in data from vehicles in transport, detect a new threat, and automatically issue the appropriate updates to every other vehicle's software for real-time defense intelligence. The AI could also update maps of where threats were and automatically reroute both manned and unmanned vehicles around them.

Trend Micro is a Chinese cybersecurity firm that specializes in connected office, connected car, and home protection. Its Smart Protection Network monitors and uncovers threat information from hundreds of millions of sensors, files, IPs, URLs, and mobile apps, and identifies over 500,000 new threats every single day! The company uses ML to proactively block new threats before they emerge.

Tractica forecasts that the annual revenue for vehicle network and data security in transportation will increase from \$1.86 million worldwide in 2017 to \$82.14 million in 2025.

Table 2.216 Vehicle Network and Data Security in Transportation, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	-	1.86	4.61	8.66	14.49	22.67	33.67	47.66	64.20	82.14	N/A

(Source: Tractica)

2.30.6 WEATHER FORECASTING

Transportation companies benefit from the ability to forecast weather events, as foresight can help ensure minimal disruption to transportation systems, supply chain movements, and wear and tear on vehicles, identifying alternative routing for fleets, or potentially signaling when to evacuate people. In the United States alone, the cost of weather-related delays in the freight industry was estimated at \$8.7 billion (an estimated 1.6% of the total estimated freight market) in 2012, according to the U.S. Department of Transportation.

AI and sensor data from hundreds of thousands of sources collected and monitored in real time (and over many years) are transforming the level of understanding and ability to forecast conditions. In addition to weather data, such engines combine streaming data from social feeds, news reports, transportation data, and historical data on storms or other weather events. While no one can ever fully predict the future, AI techniques apply reinforcement learning on past predictions and actual outcomes. By comparing predictions with accuracies, the model is able to learn and improve simulation capabilities, as well as forecast much further into the future.

In a smart city context, weather forecasting would be included in one of a huge array of data sources feeding a central brain for automated “intelligent” municipal services. Traffic signal control systems, street lighting, environmental pollution, CCTV systems, parking systems, de-icing systems for roads and bridges, and a host of other systems benefit from forecasting weather (and potential damages or wear).

For other types of transportation, NASA’s National Center for Atmospheric Research (NCAR) has been working to predict areas of turbulence, both in clear skies and within storms, and across remote areas of oceans, by applying AI to satellite data and computer-generated models of weather.

Tractica forecasts that the annual revenue for weather forecasting in transportation will increase from \$1.18 million worldwide in 2016 to \$202.22 million in 2025.

Table 2.217 Weather Forecasting in Transportation, World Markets: 2016-2025

Units	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	CAGR (2016-2025)
(\$ Millions)	1.18	5.39	11.76	21.24	35.11	54.77	81.52	115.91	157.05	202.22	77.0%

(Source: Tractica)

SECTION 3

RECOMMENDATIONS AND CONCLUSIONS

3.1

RECOMMENDATIONS

AI has the potential to disrupt numerous industries, workflows, jobs, and mechanisms for knowledge generation and sharing. In the era of Big Data and digitization, data availability, integrity, and standardization are paramount to successful AI deployments. Tractica recommends that businesses interested in exploring AI begin piloting projects. Begin with the following steps:

- **Invest in Understanding:** Invest in time, guidance, and talent to educate internal stakeholders and leadership about AI, particularly areas of application, differentiation, cost efficiencies, new revenue opportunities, security, overhype, controversy, and risk.
- **Define the Problem:** Begin not with AI, but with current pain points and problems. Know your highest-impact decision bottlenecks. AI and DL, in particular, are best applied to very specific questions and scoped problems, rather than general issues or experiments.
- **Prioritize Data Integrity:** Regardless of familiarity with AI, all enterprises should be prioritizing data cleansing, standardizing, consolidation, and formalizing processes to maintain and optimize data integrity across internal and external data sources. Part of this can and sometimes should involve user engagement, as users themselves can help train AI models.
- **Build Talent and Collaborations:** Tap into open-source communities, consortia, partnerships, universities, etc., in order to foster collaborative ideation and development for DL initiatives.
- **Monitor, Manage, and Secure:** Enterprises must constantly monitor and provide ongoing maintenance to AI models, as well as to other relevant operational analytics. For applications, set, monitor, and evolve KPIs, and assess risks. Formalize relevant security requirements, such as identity authentication, access controls, auditing, and privacy assessments, related to both model development and performance.
- **Provide Training, Support, and Communications:** It is also essential to coordinate necessary training and communications plans for the role of ML in employee, partner, and end-user workflows and experiences.

3.2

CONCLUSION

AI and the combination of technologies therein enable new capabilities and ways of thinking, both for machines and humans. Despite its potential, and perhaps because of its nature, the technology is also subject to overhype, oversell, under-delivery, and controversy. As we teach machines to perceive and think, it is critical that we design, build, apply, and scale mindfully, with individual and institutional regard for risks, unintended consequences, societal benefits, and human empowerment.

SECTION 4

ACRONYM AND ABBREVIATION LIST

Acrylonitrile Butadiene Styrene	ABS
Acute Kidney Injury	AKI
Advanced Driver Assistance System	ADAS
All-Terrain Vehicle	ATV
Amazon Web Services	AWS
Anti-Money Laundering	AML
Application Programming Interface	API
Artificial Intelligence	AI
Artificial Intelligence in Music	AIM
Augmented Reality	AR
Australian Football League	AFL
Autism Spectrum Disorder	ASD
Automated Guided Vehicles	AGV
Automated Teller Machine	ATM
Automatic Speech Recognition	ASR
Average Order Value	AOV
Brick and Mortar	B&M
Building Automaton System	BAS
Centers for Disease Control	CDC
Centimeter	cm
Central Processing Unit	CPU
Chief Executive Officer	CEO
Chief Financial Officer	CFO
Click-Thru Rate	CTR
Clinical Documentation Improvement	CDI

Closed Circuit Television.....	CCTV
Comma-Separated Values.....	CSV
Communications Fraud Control Association.....	CFCA
Compound Annual Growth Rate	CAGR
Computed Tomography	CT
Computer-Assisted Clinical Documentation Improvement	CACDI
Computer-Assisted Coding	CAC
Computer-Assisted Language Learning	CALL
Computer-Assisted Physician Documentation.....	CAPD
Computer Vision.....	CV
Consumer Packaged Goods	CPG
Convolutional Neural Network	CNN
Cost-per-Acquisition.....	CPA
Cost-per-Click	CPC
Cost-per-Lead	CPL
Customer Experience Management	CEM
Customer Relationship Management	CRM
Deep Learning.....	DL
Deep Neural Network	DNN
Defense Advanced Research Projects Agency	DARPA
Defense Centers of Excellence (U.S.)	DCoE
Deoxyribonucleic Acid.....	DNA
Dialect Identification	DID
Do-It-Yourself	DIY
Electrical Control Unit.....	ECU
Electrocardiogram.....	ECG
Electronic Health Records.....	HER
Electronic Medical Records.....	EMR

Enterprise Resource Planning	ERP
European Union	EU
Evaluation and Management	E/M
Expert Personal Shopper	XPS
Federal Bureau of Investigations (U.S.)	FBI
Field Programmable Gate Array	FPGA
Food & Drug Administration (U.S.)	FDA
Foreign Corrupt Practices Act	FCPA
Galvanic Skin Response	GSR
Geographic Information System	GIS
Gigabyte	GB
Global Navigation Satellite System	GNSS
Global Positioning System	GPS
Government Communications Headquarters (U.K.)	GCHQ
Grand Theft Auto	GTA
Graphics Processing Unit	GPU
Greenhouse Gasses	GHG
Gross Domestic Product	GDP
Head Mounted Display	HMD
Health Insurance Portability and Accountability Act	HIPAA
Hewlett Packard Enterprise	HPE
High Definition	HD
High-Performance Embedded Computing	HPEC
Human Resources	HR
Identification	IP
Inertial Measurement Unit	IMU
Information Technology	IT
Intensive Care Unit	ICU

Interactive Voice Response	IVR
Internal Revenue Service (U.S.)	IRS
Internet Protocol.....	IP
Internet of Things	IoT
Key Performance Indicator.....	KPI
Know Your Customer	KYC
Language Identification.....	LID
Light Detection and Ranging.....	LIDAR
Light-Emitting Diode.....	LED
Long Short-Term Memory	LSTM
Machine Learning.....	ML
Machine Reasoning	MR
Major League Soccer (U.S.).....	MLS
Magnetic Resonance Imaging	MRI
Massachusetts Institute of Technology	MIT
Massively Multiplayer Online.....	MMO
Mean Time to Restore.....	MTTR
Mechanism of Action.....	MOA
National Basketball Association.....	NBA
National Center for Atmospheric Research (NASA)	NCAR
National Energy Research Computing Center.....	NERSC
National Health Service (U.K.).....	NHS
National Human Genome Research Institute	NHGRI
National Institutes of Health (U.S.).....	NIH
National Security Agency (U.S.)	NSA
Natural Language Processing.....	NLP
Natural Language Understanding	NLU
Net Promoter Score	NPS

Network Service Provider.....	NSP
Neural Information Processing Systems.....	NIPS
Non-Player Character	NPC
Open Academic Search	OAS
Partners for Advanced Transportation Technology	PATH
Personal Computer	PC
Photoplethysmogram	PPG
Point of Sale.....	POS
Public Relations	PR
Questions & Answers.....	Q&A
Quick Response	QR
Radio Frequency Identification	RFID
Recurrent Neural Networks.....	RNN
Red, Green, Blue	RGB
Research and Development.....	R&D
Return on Investment.....	ROI
Return-Oriented Programming.....	ROP
Ribonucleic Acid.....	RNA
Search Engine Optimization	SEO
Service Level Agreement	SLA
Short Message Service (Text)	SMS
Simultaneous Localization and Mapping	SLAM
Single Instruction Multiple Data	SIMD
Small to Medium-Sized Business	SMB
Software-as-a-Service.....	SaaS
Software Development Kit.....	SDK
Stock Keeping Unit.....	SKU
Sudden Cardiac Arrest.....	SCA

Supervisory Control and Data Acquisition	SCADA
Synthetic Environment for Analysis and Simulations (U.S. Department of Homeland Security).....	SEAS
Television	TV
Three-Dimensional.....	3D
Travel and Expenses	T&E
Two-Dimensional	2D
Trade-Based Money Laundering	TBML
Unmanned Aerial Vehicle.....	UAV
User Experience.....	UX
Vice President.....	VP
Virtual Digital Assistant	VDA
Visual Media Reasoning	VMR
Virtual Reality	VR
Voice over Internet Protocol.....	VoIP

SECTION 5

TABLE OF CONTENTS

SECTION 1	2
Executive Summary.....	2
1.1 Introduction.....	2
1.2 Artificial Intelligence Expands Across Industries.....	2
1.3 Market Forecast	4
1.3.1 Total Revenue for Artificial Intelligence.....	5
1.3.2 Top 10 Use Cases for Artificial Intelligence	6
SECTION 2	7
Artificial Intelligence Use Cases.....	7
2.1 Overview	7
2.2 Advertising.....	7
2.2.1 Ad Insertions into Images and Video	7
2.2.2 Human Emotion Analysis	8
2.2.3 Interactive Window Displays	8
2.2.4 Performance Reporting and Analytics of Ad Campaigns	10
2.2.5 Querying Image Content	11
2.2.6 Static Image Recognition, Classification, and Tagging.....	11
2.2.7 Targeted Advertising Using Multi-Domain Customer Data (Social, Web, Context)	12
2.2.8 Video Content Analysis	13
2.2.9 Voice/Speech Recognition	14
2.3 Aerospace	15
2.3.1 Localization and Mapping (Aircraft and Drones)	15
2.3.2 Machine/Vehicular Object Detection/Identification/Avoidance (Aircraft, Drones, Satellites)	16
2.3.3 Predictive Maintenance (Aircraft, Drones, Satellites).....	17
2.3.4 Sensor Data Fusion in Machinery (Aircraft, Drones, Satellites).....	18
2.3.5 Swarming Drones	18
2.3.6 Vehicle Network and Data Security (Aircraft, Drones, Satellites)	19
2.3.7 Weather Forecasting	20
2.4 Agriculture	21
2.4.1 Food Safety	21
2.4.2 Livestock Management	21
2.4.3 Machine/Vehicular Object Detection/Identification/Avoidance	22
2.4.4 Satellite Imagery for Geo-Analytics	23
2.4.5 Sensor Data Analytics	24
2.4.6 Sensor Data Fusion in Machinery	25
2.4.7 Localization and Mapping.....	25
2.4.8 Weather Forecasting	26
2.4.9 Weed Identification.....	27
2.5 Automotive	27
2.5.1 Automated On-Road Customer Service.....	27
2.5.2 Building Generative Models of the Real World	28
2.5.3 Driver Face Analytics and Emotion Recognition	29
2.5.4 Gesture Recognition.....	30
2.5.5 Machine/Vehicular Object Detection/Identification/Avoidance	31
2.5.6 Personalized Services in Cars	31
2.5.7 Truck Platooning	32
2.5.8 Predicting Demand for On-Demand Taxis	33

2.5.9	Predictive Maintenance	33
2.5.10	Sensor Data Fusion in Machinery	35
2.5.11	Simulating Worlds for Artificial Intelligence Training	35
2.5.12	Surge Pricing for On-Demand Taxis	36
2.5.13	Localization and Mapping.....	37
2.5.14	Vehicle Network and Data Security.....	37
2.5.15	Virtual Testing and Simulation for Racing Cars	39
2.6	Building Automation	39
2.6.1	Building Automation and Energy Management.	39
2.7	Business Services.....	40
2.7.1	Agent-Based Simulations for Decision-Making	40
2.7.2	Audio and Video Mining	41
2.7.3	Automated Report Generation	42
2.7.4	Automated Workforce Scheduling.....	43
2.7.5	Chatbot-Based Brand/Service Interactions	44
2.7.6	Chatbot-Based E-Commerce and Sales	45
2.7.7	Crowdsourced Market Research.....	47
2.7.8	Enterprise Chatbots for Productivity and Collaboration	47
2.7.9	Intelligent Customer Relationship Management Systems	48
2.7.10	Intelligent Recruiting and Human Resources Systems.....	49
2.7.11	Prevention Against Cybersecurity Threats.....	50
2.7.12	Procurement Management.....	51
2.7.13	Project and Stakeholder Management.....	52
2.7.14	Real-Time News Analysis and Competitive Intelligence	53
2.7.15	Social Media Publishing and Management	54
2.7.16	Travel Concierge and Booking Services	55
2.7.17	Workflow and Project Management	56
2.8	Construction	57
2.8.1	Satellite Imagery for Geo-Analytics	57
2.9	Consumer.....	57
2.9.1	Automated Tour Guide and Itinerary Service.....	57
2.9.2	Building Generative Models of the Real World	58
2.9.3	Calendar, Meeting, Event Scheduling, and Reminders	59
2.9.4	Child Behavioral Analytics.....	59
2.9.5	Computer-Aided Art.....	60
2.9.6	Contextual Intelligence for Mobile	61
2.9.7	Facial Recognition.....	62
2.9.8	Language Translation Services.....	64
2.9.9	Local Search and Discovery.....	65
2.9.10	Movie Recommendations.....	66
2.9.11	Music Recommendations	67
2.9.12	Machine/Vehicular Object Detection/Identification/Avoidance	67
2.9.13	Personalized Health, Fitness, and Wellness Improvement.....	68
2.9.14	Predictive Typing Assistant	70
2.9.15	Product Recommendations	71
2.9.16	Relationships and Matchmaking	72
2.9.17	Search Engine Queries	72
2.9.18	Smart Oven Control with Food Recognition.....	73
2.9.19	Social Media Feed Curation.....	74
2.9.20	Static Image Recognition, Classification, and Tagging	76
2.9.21	Text-Based Automated Bots	76
2.9.22	Travel Concierge and Booking Service.....	77
2.9.23	Voice/Speech Recognition	78

2.10	Defense	80
2.10.1	Agent-Based Simulations for Decision-Making	80
2.10.2	Localization and Mapping (Aircraft and Drones)	80
2.10.3	Machine/Vehicular Object Detection/Identification/Avoidance (Defense Aircraft and Drones)	81
2.10.4	Predictive Maintenance (Defense Aircraft, Drones, Satellites)	82
2.10.5	Prevention Against Cybersecurity Threats	82
2.10.6	Satellite Imagery for Geo-Analytics	83
2.10.7	Sensor Data Fusion in Machinery (Defense Aircraft, Drones, Satellites)	84
2.10.8	Swarming Drones	85
2.10.9	Vehicle Network and Data Security (Defense Aircraft, Drones, Satellites)	86
2.11	Education	86
2.11.1	Personalized Tutoring and Adaptive Learning	86
2.11.2	Automated CliffsNotes, Study Notes, and Quiz Generators	87
2.11.3	Automated Grading of Tests	88
2.11.4	Education for Autistic and Speech Deficient Children	89
2.11.5	Foreign Language Tutoring	90
2.11.6	Spoken Fluency Evaluation	90
2.11.7	Textual Question Answering	91
2.12	Energy	92
2.12.1	Satellite Imagery for Geo-Analytics	92
2.12.2	Weather Forecasting	93
2.13	Fashion	94
2.13.1	Fashion Trend Prediction	94
2.14	Finance	95
2.14.1	Automated Credit Scoring	95
2.14.2	Automated Report Generation	96
2.14.3	Biometric Identification	97
2.14.4	Converting Paperwork into Digital Assets	98
2.14.5	Patient Data Processing	98
2.14.6	Employee Expense Management	99
2.14.7	Loan Analysis	100
2.14.8	Personal Financial Advisor	100
2.14.9	Risk Assessment and Compliance	101
2.14.10	Tax Filing and Processing	102
2.14.11	Transaction Fraud Detection	103
2.15	Gaming	104
2.15.1	Create Dynamic and Interactive Video Game Experiences	104
2.16	Government	105
2.16.1	Agent-Based Simulations for Decision-Making	105
2.16.2	Behavioral Analytics	106
2.16.3	Converting Paperwork into Digital Assets	106
2.16.4	Crowd Analytics	107
2.16.5	Dialect Classification	108
2.16.6	Disaster and Emergency Management	108
2.16.7	Facial Recognition	109
2.16.8	Object Detection for Surveillance	111
2.16.9	Predicting Social Unrest and Geopolitical Events	111
2.16.10	Real-Time Video Analytics	112
2.16.11	Sentiment Analysis	113
2.16.12	Social Media Bots	114
2.16.13	Street Lighting	115
2.16.14	Traffic Light Management	116

2.16.15	Waste Sorting and Recycling	117
2.16.16	Weather Forecasting	118
2.17	Healthcare	119
2.17.1	Automated Report Generation	119
2.17.2	Bio-Marker Discovery	120
2.17.3	Clustering and Phenotype Discovery	121
2.17.4	Computational Drug Discovery	122
2.17.5	Converting Paperwork Into Digital Assets	122
2.17.6	Facial Recognition	123
2.17.7	Genomic Data Mapping and Analysis for Personalized Healthcare and Precision Medicine	124
2.17.8	Hospital Patient Management System	125
2.17.9	Market Intelligence for Life Sciences	126
2.17.10	Medical Diagnosis Assistance	127
2.17.11	Medical Image Analysis	128
2.17.12	Medical Treatment Recommendation	129
2.17.13	Medication Compliance for Clinical Trials and General Usage	130
2.17.14	Methods for Monitoring Vitals	131
2.17.15	Mining, Processing, and Making Sense of Clinical Notes	131
2.17.16	Patient Data Processing	132
2.17.17	Portable and Low-Cost Ultrasound Device	134
2.17.18	Predicting Illness and Patient Outcomes	134
2.17.19	Text Classification and Mining for Biomedical Literature	135
2.17.20	Virtual Assistants for Doctors	136
2.17.21	Virtual Assistants for Patients	137
2.18	Information Technology	138
2.18.1	Automated Code Development	138
2.18.2	Computer-Aided Design	139
2.18.3	Mobile Application Development	140
2.18.4	Network/Information Technology Operations Monitoring and Management	141
2.18.5	Simulating Worlds for Artificial Intelligence Training	142
2.18.6	Software Code Error Checking	143
2.18.7	Website Creation	144
2.19	Investment	144
2.19.1	Algorithmic Trading Strategy Performance Improvement	144
2.19.2	Financial Search Engine	145
2.19.3	Market Intelligence and Data Analytics for Investment	146
2.19.4	Satellite Imagery for Geo-Analytics	147
2.20	Legal	147
2.20.1	Automated Report Generation	147
2.20.2	Contract Analysis	148
2.20.3	Legal Document Review and Research	149
2.21	Life Sciences	149
2.21.1	Create Synthetic Life Forms	149
2.22	Logistics	150
2.22.1	Demand Forecasting for Warehouse and Supply Chain	150
2.22.2	Machine/Vehicular Object Detection/Identification/Avoidance	151
2.22.3	Localization and Mapping	151
2.22.4	Satellite Imagery for Geo-Analytics	152
2.22.5	Supply Chain and Logistics (Freight Transport, Retail)	153
2.22.6	Weather Forecasting	154
2.23	Manufacturing	155
2.23.1	3D Printing Arm Control	155

2.23.2	Machine/Vehicular Object Detection/Identification/Avoidance	156
2.23.3	Predictive Maintenance	156
2.23.4	Real-Time Video Analytics	157
2.23.5	Localization and Mapping	158
2.23.6	Sensor Data Fusion in Machinery	159
2.23.7	Voice/Speech Recognition	159
2.24	Media and Entertainment	161
2.24.1	Algorithmic News Stories	161
2.24.2	Audio and Video Mining	162
2.24.3	Film Scene Structure	162
2.24.4	Font Recognition and Suggestions	163
2.24.5	Gesture Recognition	163
2.24.6	Human Emotion Analysis	164
2.24.7	Music Production and Generation	165
2.24.8	News and Feed Curation for Consumers	166
2.24.9	Simulating Crowds	166
2.24.10	Social Media Publishing and Management	167
2.24.11	Video Editing	168
2.25	Oil, Gas, and Mining	169
2.25.1	Automated Report Generation	169
2.25.2	Oil Production Optimization	170
2.26	Real Estate	170
2.26.1	Real Estate Development Optimization	170
2.27	Retail	172
2.27.1	Behavioral Analytics	172
2.27.2	Clothes Sizing and Fitting	173
2.27.3	Crowd Analytics	173
2.27.4	Intelligent Customer Relationship Management Systems	174
2.27.5	Predictive Analytics for Retail	175
2.27.6	Sentiment Analysis	176
2.27.7	Supermarket Shelf Analytics	178
2.27.8	Visual Search-Based E-Commerce	178
2.27.9	Weather Forecasting	179
2.28	Sports	180
2.28.1	Athlete Fitness, Sleep Monitoring, and Performance Optimization	180
2.28.2	Biomarker-Based Athlete Performance Optimization	181
2.28.3	Game Outcome Predictions for Betting	182
2.28.4	Sports Statistics Analysis and Search	182
2.28.5	Sports Teams Players Selection	183
2.29	Telecommunications	184
2.29.1	Predictive Maintenance	184
2.29.2	Prevention Against Cybersecurity Threats	185
2.29.3	Improve Customer Experience Management	186
2.29.4	Fraud Mitigation	186
2.29.5	Intelligent Customer Relationship Management Systems	187
2.30	Transportation	188
2.30.1	Machine/Vehicular Object Detection/Identification/Avoidance	188
2.30.2	Predicting Traffic Density	189
2.30.3	Sensor Data Fusion in Machinery (Ships, Unmanned Ships)	189
2.30.4	Localization and Mapping	190
2.30.5	Vehicle Network and Data Security	191
2.30.6	Weather Forecasting	192

SECTION 3	193
Recommendations and Conclusions.....	193
3.1 Recommendations	193
3.2 Conclusion.....	193
SECTION 4	194
Acronym and Abbreviation List.....	194
SECTION 5	200
Table of Contents.....	200
SECTION 6	206
Table of Charts and Figures.....	206
SECTION 7	212
Scope of Study	212
Sources and Methodology	212
Notes	213

SECTION 6

TABLE OF CHARTS AND FIGURES

Chart 1.1	Artificial Intelligence Software Revenue, World Markets: 2016-2025	5
Chart 1.2	Cumulative Artificial Intelligence Software Revenue, Top 10 Use Cases, World Markets: 2016-2025	6
Chart 7.1	Tractica Research Methodology.....	213
Figure 2.1	Nike Window Display Gammifies Shoppers' Interactions Using Motion Detection.....	9
Figure 2.2	Drones Use Mapping and Localization to Fly Indoors.....	15
Figure 2.3	BMW's 2016 7-Series Incorporates Gesture Recognition for Six Commands.....	30
Figure 2.4	Predictive Maintenance Dashboard for Connected Cars.....	34
Figure 2.5	Sample Sales Summary Populated by Artificial Intelligence.....	42
Figure 2.6	Mark Zuckerberg on Messenger Business at F8 Conference in 2016.....	44
Figure 2.7	Retailers are Integrating with Facebook Messenger App to Tie E-Commerce Directly to Facebook Experience.....	46
Figure 2.8	AI Painter, a Neural Network that Renders Photos as Paintings	61
Figure 2.9	Netflix Uses Artificial Intelligence for Personal Homepage Optimization and A/B Testing for Page Generation	66
Figure 2.10	Google's Smart Reply Offers Auto-Generated In-Context Responses.....	70
Figure 2.11	June Oven Uses Image Recognition to Identify, Automate, and Optimize Cooking	74
Figure 2.12	DARPA's Swarm of Drones Simulates Group Formations over California.....	85
Figure 2.13	Milo, a Humanoid Robot, Helps ASD Students Identify Human Emotions	89
Figure 2.14	SimCoach, a 3D Virtual Agent Interacts and Assists Military Personnel with Breaking Down Barriers	92
Figure 2.15	Stitch Fix Uses Deep Learning to Analyze Styles and Design New Clothing.....	95
Figure 2.16	Kasisto's MyKai, a Personal Financial Advisor Chatbot	101
Figure 2.18	Zen Robotics Waste Processing Workflow.....	118
Figure 2.19	Ada Health App Delivers Virtual Assistance for Patients.....	138
Figure 2.20	DAQRI's Smart Helmet Combines Voice Recognition and Augmented Reality for Real-Time Work Instructions.....	160
Figure 2.21	Peltarion's Model Analysis of Millions of Data Points to Produce Real Estate Valuations	171
Figure 2.22	Luminoso's Analytics Found Caregivers of Schizophrenic Patients Are Harder on Themselves than Any Other Group	177
Figure 2.23	Shoes.com Partners with Sentient to Power Image-Based Search and Product Recommendations.....	179
Table 2.1	Ad Insertions into Images and Video in Advertising, World Markets: 2016-2025	7
Table 2.2	Human Emotion Analysis in Advertising, World Markets: 2016-2025.....	8
Table 2.3	Interactive Window Displays in Advertising, World Markets: 2016-2025	9
Table 2.4	Performance Reporting and Analytics of Ad Campaigns in Advertising, World Markets: 2016-2025	10
Table 2.5	Querying Image Content in Advertising, World Markets: 2016-2025.....	11
Table 2.6	Static Image Recognition, Classification, and Tagging in Advertising, World Markets: 2016-2025	12
Table 2.7	Targeted Advertising Using Multi-Domain Customer Data in Advertising, World Markets: 2016-2025	13

Table 2.8	Video Content Analysis in Advertising, World Markets: 2016-2025.....	14
Table 2.9	Voice/Speech Recognition in Advertising, World Markets: 2016-2025.....	15
Table 2.10	Localization and Mapping in Aerospace, World Markets: 2016-2025.....	16
Table 2.11	Machine/Vehicular Object Detection/Identification/Avoidance in Aerospace, World Markets: 2016-2025.....	17
Table 2.12	Predictive Maintenance in Aerospace, World Markets: 2016-2025	17
Table 2.13	Sensor Data Fusion in Aerospace, Machinery World Markets: 2016-2025	18
Table 2.14	Swarming Drones in Aerospace, World Markets: 2016-2025	19
Table 2.15	Vehicle Network and Data Security in Aerospace, World Markets: 2016-2025	20
Table 2.16	Weather Forecasting in Aerospace, World Markets: 2016-2025	20
Table 2.17	Food Safety in Agriculture, World Markets: 2016-2025	21
Table 2.18	Livestock Management in Agriculture, Annual Revenue, 2016-2025	22
Table 2.19	Machine/Vehicle Object Detection/Identification/Avoidance in Agriculture, World Markets: 2016-2025	23
Table 2.20	Satellite Imagery for Geo-Analytics in Agriculture, World Markets: 2016-2025	24
Table 2.21	Sensor Data Analytics in Agriculture, World Markets: 2016-2025	25
Table 2.22	Sensor Data Fusion in Machinery in Agriculture, World Markets: 2016-2025	25
Table 2.23	Localization and Mapping in Agriculture, World Markets: 2016-2025.....	26
Table 2.24	Weather Forecasting in Agriculture, World Markets: 2016-2025	27
Table 2.25	Weed Identification in Agriculture, World Markets: 2016-2025	27
Table 2.26	On-Road Customer Service in Automotive, World Markets: 2016-2025.....	28
Table 2.27	Building Generative Models of the Real-World in Automotive, World Markets: 2016-2025.....	29
Table 2.28	Driver Face Analytics and Emotion Recognition in Automotive, World Markets: 2016-2025	29
Table 2.29	Gesture Recognition in Automotive, World Markets: 2016-2025	31
Table 2.30	Machine/Vehicle Object Detection/Identification/Avoidance in Automotive, World Markets: 2016-2025	31
Table 2.31	Personalized Services in Cars in Automotive, World Markets: 2016-2025.....	32
Table 2.32	Truck Platooning in Automotive, World Markets: 2016-2025	33
Table 2.33	Predicting Demand in On-Demand Taxis in Automotive, World Markets: 2016-2025	33
Table 2.34	Predictive Maintenance Taxis in Automotive, World Markets: 2016-2025.....	34
Table 2.35	Sensor Data Fusion in Machinery in Automotive, World Markets: 2016-2025	35
Table 2.36	Simulating Worlds for AI Training in Automotive, World Markets: 2016-2025	36
Table 2.37	Surge Pricing for On-Demand Taxis in Automotive, World Markets: 2016-2025.....	37
Table 2.38	Localization and Mapping in Automotive, World Markets: 2016-2025	37
Table 2.39	Vehicle Network and Data Security in Automotive, World Markets: 2016-2025	38
Table 2.40	Virtual Testing and Simulation for Racing Cars in Automotive, World Markets: 2016-2025	39
Table 2.41	Building Automation and Energy Management, World Markets: 2016-2025	40
Table 2.42	Agent-Based Simulations for Decision-Making in Business, World Markets: 2016-2025....	41
Table 2.43	Audio and Video Mining in Business, World Markets: 2016-2025	42
Table 2.44	Automated Report Generation in Business, World Markets: 2016-2025	43
Table 2.45	Automated Workforce Scheduling in Business, World Markets: 2016-2025	44
Table 2.46	Chatbot-Based Brand/Service Interactions in Business, World Markets: 2016-2025.....	45
Table 2.47	Chatbot-Based e-Commerce and Sales in Business, World Markets: 2016-2025	46
Table 2.48	Crowdsourced Market Research in Business, World Markets: 2016-2025.....	47
Table 2.49	Enterprise Chatbots for Productivity and Collaboration in Business, World Markets: 2016-2025	48
Table 2.50	Intelligent CRM Systems in Business, World Markets: 2016-2025	49
Table 2.51	Intelligent Recruiting and Human Resources Systems in Business, World Markets: 2016-2025	50
Table 2.52	Prevention Against Cybersecurity Threats in Business, World Markets: 2016-2025.....	51
Table 2.53	Procurement Management in Business, World Markets: 2016-2025.....	52
Table 2.54	Project and Stakeholder Management in Business, World Markets: 2016-2025.....	53

Table 2.55	Real-Time News Analysis and Competitive Intelligence in Business, World Markets: 2016-2025	54
Table 2.56	Social Media Publishing and Management in Business, World Markets: 2016-2025.....	55
Table 2.57	Travel Concierge and Booking Services in Business, World Markets: 2016-2025.....	56
Table 2.58	Workflow and Project Management in Business, World Markets: 2016-2025	56
Table 2.59	Satellite Imagery for Geo-Analytics in Construction, World Markets: 2016-2025	57
Table 2.60	Automated Tour Guide and Itinerary Services in Consumer, World Markets: 2016-2025....	58
Table 2.61	Building Generative Models of the Real World in Consumer, World Markets: 2016-2025 ..	59
Table 2.62	Calendar, Meeting, Event Scheduling, and Reminders in Consumer, World Markets: 2016-2025	59
Table 2.63	Child Behavioral Analytics in Consumer, World Markets: 2016-2025.....	60
Table 2.64	Computer-Aided Art in Consumer, World Markets: 2016-2025	61
Table 2.65	Contextual Intelligence in Mobile in Consumer, World Markets: 2016-2025	62
Table 2.66	Facial Recognition in Consumer, World Markets: 2016-2025.....	64
Table 2.67	Language Translation Services in Consumer, World Markets: 2016-2025	65
Table 2.68	Local Search and Discovery in Consumer, World Markets: 2016-2025	65
Table 2.69	Movie Recommendations in Consumer, World Markets: 2016-2025.....	66
Table 2.70	Music Recommendations in Consumer, World Markets: 2016-2025.....	67
Table 2.71	Machine/Vehicle Object Detection/Identification/Avoidance in Consumer, World Markets: 2016-2025	68
Table 2.72	Personalized Health, Fitness, and Wellness Improvement in Consumer, World Markets: 2016-2025	70
Table 2.73	Predictive Typing Assistants in Consumer, World Markets: 2016-2025	71
Table 2.74	Product Recommendations in Consumer, World Markets: 2016-2025.....	71
Table 2.75	Relationships and Matchmaking in Consumer, World Markets: 2016-2025	72
Table 2.76	Search Engine Queries in Consumer, World Markets: 2016-2025	73
Table 2.77	Smart Oven Control with Food Recognition in Consumer, World Markets: 2016-2025.....	74
Table 2.78	Social Media Feed Curation in Consumer, World Markets: 2016-2025	75
Table 2.79	Static Image Recognition, Classification, and Tagging in Consumer, World Markets: 2016-2025	76
Table 2.80	Text-Based Automated Bots in Consumer, World Markets: 2016-2025	77
Table 2.81	Travel Concierge and Booking Services in Consumer, World Markets: 2016-2025.....	78
Table 2.82	Voice/Speech Recognition in Consumer, World Markets: 2016-2025	79
Table 2.83	Agent-Based Simulations for Decision-Making in Defense, World Markets: 2016-2025	80
Table 2.84	Localization and Mapping in Defense, World Markets: 2016-2025.....	81
Table 2.85	Machine/Vehicle Object Detection/Identification/Avoidance in Defense World Markets: 2016-2025	82
Table 2.86	Predictive Maintenance in Defense, World Markets: 2016-2025	82
Table 2.87	Prevention Against Cybersecurity Threats in Defense, World Markets: 2016-2025	83
Table 2.88	Satellite Imagery for Geo-Analytics in Defense, World Markets: 2016-2025	84
Table 2.89	Sensor Data Fusion in Machinery in Defense, World Markets: 2016-2025	85
Table 2.90	Swarming Drones in Defense, World Markets: 2016-2025	85
Table 2.91	Vehicle Network and Data Security in Defense, World Markets: 2016-2025.....	86
Table 2.92	Personalized Tutoring and Adaptive Learning in Education, World Markets: 2016-2025 ...	87
Table 2.93	Automated Cliffs Notes, Study Notes, Quiz Generators in Education, World Markets: 2016-2025	88
Table 2.94	Automated Grading of Tests in Education, World Markets: 2016-2025.....	88
Table 2.95	Education for Autistic and Speech Deficient Children in Education, World Markets: 2016-2025	90
Table 2.96	Foreign Language Tutoring in Education, World Markets: 2016-2025	90
Table 2.97	Spoken Fluency Evaluation in Education, World Markets: 2016-2025	91
Table 2.98	Textual Question Answering in Education, World Markets: 2016-2025.....	92
Table 2.99	Satellite Imagery in Geo-Analytics in Energy, Annual Revenue, 2016-2025.....	93

Table 2.100	Weather Forecasting in Energy, World Markets: 2016-2025.....	94
Table 2.101	Fashion Trend Prediction in Fashion, World Markets: 2016-2025	95
Table 2.102	Automated Credit Scoring in Finance, World Markets: 2016-2025	96
Table 2.103	Automated Report Generation in Finance, World Markets: 2016-2025.....	97
Table 2.104	Biometric Identification in Finance, World Markets: 2016-2025	97
Table 2.105	Converting Paperwork into Digital Assets in Finance, World Markets: 2016-2025	98
Table 2.106	Patient Data Processing in Finance, World Markets: 2016-2025	99
Table 2.107	Employee Expense Management in Finance, World Markets: 2016-2025.....	99
Table 2.108	Loan Analysis in Finance, World Markets: 2016-2025	100
Table 2.109	Personal Financial Advisor in Finance, World Markets: 2016-2025	101
Table 2.110	Risk Assessment and Compliance in Finance, World Markets: 2016-2025	102
Table 2.111	Tax Filing and Processing in Finance, World Markets: 2016-2025	103
Table 2.112	Transaction Fraud Detection in Finance, World Markets: 2016-2025	104
Table 2.113	Creating Dynamic and Interactive Video Game Experiences in Gaming, World Markets: 2016-2025	105
Table 2.114	Agent-Based Simulations for Decision-Making in Government, World Markets: 2016-2025	106
Table 2.115	Behavioral Analytics in Government, World Markets: 2016-2025	106
Table 2.116	Converting Paperwork into Digital Assets in Government, World Markets: 2016-2025	107
Table 2.117	Crowd Analytics in Government, World Markets: 2016-2025	108
Table 2.118	Crowd Analytics in Government, World Markets: 2016-2025	108
Table 2.119	Disaster and Emergency Management in Government, World Markets: 2016-2025	109
Table 2.120	Facial Recognition in Government, World Markets: 2016-2025	110
Table 2.121	Object Detection for Surveillance in Government, World Markets: 2016-2025	111
Table 2.122	Predicting Social Unrest and Geopolitical Events in Government, World Markets: 2016-2025	112
Table 2.123	Real-Time Video Analytics in Government, World Markets: 2016-2025	113
Table 2.124	Sentiment Analysis in Government, World Markets: 2016-2025	114
Table 2.125	Social Media Bots in Government, World Markets: 2016-2025	115
Table 2.126	Street Lighting in Government, World Markets: 2016-2025	116
Table 2.127	Traffic Light Management in Government, World Markets: 2016-2025.....	117
Table 2.128	Waste Sorting and Recycling in Government, World Markets: 2016-2025	118
Table 2.129	Weather Forecasting in Government, World Markets: 2016-2025	119
Table 2.130	Automated Report Generation in Healthcare, World Markets: 2016-2025	120
Table 2.131	Bio-Marker Discovery in Healthcare, World Markets: 2016-2025.....	121
Table 2.132	Clustering and Phenotype Discovery in Healthcare, World Markets: 2016-2025.....	122
Table 2.133	Computational Drug Discovery in Healthcare, World Markets: 2016-2025	122
Table 2.134	Converting Paperwork into Digital Assets in Healthcare, World Markets: 2016-2025....	123
Table 2.135	Facial Recognition in Healthcare, World Markets: 2016-2025	124
Table 2.136	Genomic Data Mapping and Analysis for Personalized Healthcare and Precision Medicine in Healthcare, World Markets: 2016-2025.....	125
Table 2.137	Hospital Patient Management Systems in Healthcare, World Markets: 2016-2025.....	126
Table 2.138	Market Intelligence for Life Sciences in Healthcare, World Markets: 2016-2025	127
Table 2.139	Medical Diagnostic Assistance in Healthcare, World Markets: 2016-2025	128
Table 2.140	Medical Image Analysis in Healthcare, World Markets: 2016-2025	129
Table 2.141	Medical Treatment Recommendation in Healthcare, World Markets: 2016-2025	130
Table 2.142	Medication Compliance for Clinical Trials and General Usage in Healthcare, World Markets: 2016-2025.....	131
Table 2.143	Methods for Monitoring Vitals in Healthcare, World Markets: 2016-2025	131
Table 2.144	Mining, Processing, and Making Sense of Clinical Notes in Healthcare, World Markets: 2016-2025	132
Table 2.145	Patient Data Processing in Healthcare, World Markets: 2016-2025	134
Table 2.146	Portable and Low-Cost Ultrasound Devices in Healthcare, World Markets: 2016-2025	134

Table 2.147	Predicting Illness and Patient Outcomes in Healthcare, World Markets: 2016-2025	135
Table 2.148	Text Classification for Biomedical Literature in Healthcare, World Markets: 2016-2025	136
Table 2.149	Virtual Assistants for Doctors in Healthcare, World Markets: 2016-2025	137
Table 2.150	Virtual Assistants for Patients in Healthcare, World Markets: 2016-2025	138
Table 2.151	Automated Code Development in Information Technology, World Markets: 2016-2025	139
Table 2.152	Computer Aided Design in Information Technology, World Markets: 2016-2025	140
Table 2.153	Mobile Application Development in Information Technology, World Markets: 2016-2025	141
Table 2.154	Network/Information Technology Operations Monitoring and Management in Information Technology, World Markets: 2016-2025.....	142
Table 2.155	Simulating Worlds for AI Training in Information Technology, World Markets: 2016-2025	143
Table 2.156	Software Code Error Checking in Information Technology, World Markets: 2016-2025	143
Table 2.157	Website Creation in Information Technology, World Markets: 2016-2025	144
Table 2.158	Algorithmic Trading Strategy Performance Improvement in Investment, World Markets: 2016-2025	145
Table 2.159	Financial Search Engines in Investment, World Markets: 2016-2025	146
Table 2.160	Market Intelligence and Data Analytics for Investment in Investment, World Markets: 2016-2025	146
Table 2.161	Satellite Imagery for Geo-Analytics in Investment, World Markets: 2016-2025	147
Table 2.162	Automated Report Generation in Legal, World Markets: 2016-2025.....	148
Table 2.163	Contract Analysis in Legal, World Markets: 2016-2025.....	149
Table 2.164	Legal Document Review and Research in Legal, World Markets: 2016-2025	149
Table 2.165	Creating Synthetic Life Forms in Life Sciences, World Markets: 2016-2025.....	150
Table 2.166	Demand Forecasting for Warehousing and Supply Chain in Logistics, World Markets: 2016-2025	151
Table 2.167	Machine/Vehicular Object Detection/Identification/Avoidance in Logistics, World Markets: 2016-2025	151
Table 2.168	Localization and Mapping in Logistics, World Markets: 2016-2025.....	152
Table 2.169	Satellite Imagery for Geo-Analytics in Logistics, Annual Revenue, 2016-2025	153
Table 2.170	Supply Chain and Logistics in Logistics, World Markets: 2016-2025	154
Table 2.171	Weather Forecasting in Logistics, World Markets: 2016-2025	155
Table 2.172	3D Printing Arm Control in Manufacturing, World Markets: 2016-2025.....	155
Table 2.173	Machine/Vehicular Object Detection/Identification/Avoidance in Manufacturing, World Markets: 2016-2025.....	156
Table 2.174	Predictive Maintenance in Manufacturing, World Markets: 2016-2025	157
Table 2.175	Real-Time Video Analytics in Manufacturing, World Markets: 2016-2025	158
Table 2.176	Localization and Mapping in Manufacturing, World Markets: 2016-2025	159
Table 2.177	Sensor Data Fusion in Machinery in Manufacturing, World Markets: 2016-2025	159
Table 2.178	Voice/Speech Recognition in Manufacturing, World Markets: 2016-2025	160
Table 2.179	Algorithmic News Stories in Media & Entertainment, World Markets: 2016-2025.....	162
Table 2.180	Audio and Video Mining in Media & Entertainment, World Markets: 2016-2025.....	162
Table 2.181	Film Scene Structure in Media & Entertainment, World Markets: 2016-2025	163
Table 2.182	Font Recognition and Suggestions in Media & Entertainment, World Markets: 2016-2025	163
Table 2.183	Gesture Recognition in Media & Entertainment, World Markets: 2016-2025	164
Table 2.184	Human Emotion Analysis in Media & Entertainment, World Markets: 2016-2025	165
Table 2.185	Human Emotion Analysis in Media & Entertainment, World Markets: 2016-2025	166
Table 2.186	News and Feed Curation for Consumers in Media & Entertainment, World Markets: 2016-2025	166
Table 2.187	Simulating Crowds in Media & Entertainment, World Markets: 2016-2025.....	167

Table 2.188	Social Media Publishing and Management in Media & Entertainment, World Markets: 2016-2025	168
Table 2.189	Video Editing in Media & Entertainment, World Markets: 2016-2025.....	169
Table 2.190	Automated Report Generation in Oil & Gas, World Markets: 2016-2025	170
Table 2.191	Oil Production Optimization in Oil & Gas, World Markets: 2016-2025	170
Table 2.192	Real Estate Development Optimization in Real Estate, World Markets: 2016-2025.....	171
Table 2.193	Behavioral Analytics in Retail, World Markets: 2016-2025	173
Table 2.194	Clothes Sizing and Fitting in Retail, World Markets: 2016-2025	173
Table 2.195	Clothes Sizing and Fitting in Retail, World Markets: 2016-2025	174
Table 2.196	Intelligent Customer Relationship Management Systems in Retail, World Markets: 2016-2025	175
Table 2.197	Predictive Analytics for Retail in Retail, World Markets: 2016-2025.....	176
Table 2.198	T Sentiment Analysis in Retail, World Markets: 2016-2025	177
Table 2.199	Supermarket Shelf Analytics in Retail, World Markets: 2016-2025	178
Table 2.200	Visual Search-Based e-Commerce in Retail, World Markets: 2016-2025.....	179
Table 2.201	Weather Forecasting in Retail, World Markets: 2016-2025	180
Table 2.202	Athlete Fitness, Sleep Monitoring, and Performance Optimization in Sports, World Markets: 2016-2025	181
Table 2.203	Biomarker-Based Athlete Performance Optimization in Sports, World Markets: 2016-2025	182
Table 2.204	Game Outcome Predictions for Betting in Sports, World Markets: 2016-2025	182
Table 2.205	Sports Statistics Analysis and Search in Sports, World Markets: 2016-2025	183
Table 2.206	Sports Team Player Selection in Sports, World Markets: 2016-2025	184
Table 2.207	Predictive Maintenance in Telecommunications, World Markets: 2016-2025.....	185
Table 2.208	Prevention Against Cybersecurity Threats in Telecommunications, World Markets: 2016-2025	185
Table 2.209	Improving Customer Experience Management in Telecommunications, World Markets: 2016-2025	186
Table 2.210	Fraud Mitigation in Telecommunications, World Markets: 2016-2025.....	187
Table 2.211	Intelligent CRM Systems in Telecommunications, Annual Revenue, 2016-2025	188
Table 2.212	Machine/Vehicular Object Detection/Identification/Avoidance in Transportation, World Markets: 2016-2025.....	189
Table 2.213	Predicting Traffic Density in Transportation, World Markets: 2016-2025	189
Table 2.214	Sensor Data Fusion in Machinery in Transportation, World Markets: 2016-2025.....	190
Table 2.215	Localization and Mapping in Transportation, World Markets: 2016-2025	191
Table 2.216	Vehicle Network and Data Security in Transportation, World Markets: 2016-2025	192
Table 2.217	Weather Forecasting in Transportation, World Markets: 2016-2025.....	192

SECTION 7

SCOPE OF STUDY

This report examines the practical use cases and applications of AI within commercial enterprises, governments, and consumer markets. This report is a qualitative compendium to the [Artificial Intelligence Market Forecasts](#) report, in which Tractica quantitatively assesses the opportunity for AI across 29 industries using the same use case taxonomy. This report focuses on use cases only, and does not include comprehensive coverage of any one technology, application, industry, or quantitative analysis.

Within that scope, the report provides descriptions, industry context, examples, and revenue forecasts for each use case in each industry. The report also considers common themes across the broader AI market that will impact adoption.

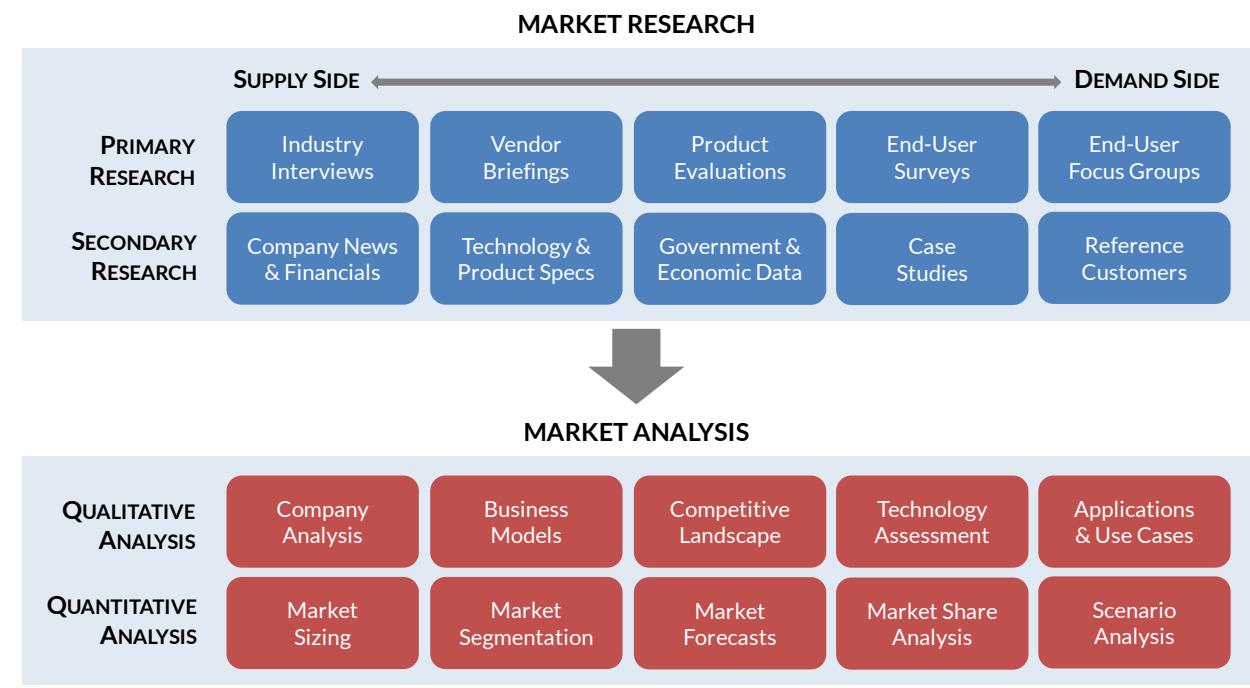
SOURCES AND METHODOLOGY

Tractica is an independent market research firm that provides industry participants and stakeholders with an objective, unbiased view of market dynamics and business opportunities within its coverage areas. The firm's industry analysts are dedicated to presenting clear and actionable analysis to support business planning initiatives and go-to-market strategies, utilizing rigorous market research methodologies and without regard for technology hype or special interests including Tractica's own client relationships. Within its market analysis, Tractica strives to offer conclusions and recommendations that reflect the most likely path of industry development, even when those views may be contrarian.

The basis of Tractica's analysis is primary research collected from a variety of sources including industry interviews, vendor briefings, product demonstrations, and quantitative and qualitative market research focused on consumer and business end-users. Industry analysts conduct interviews with representative groups of executives, technology practitioners, sales and marketing professionals, industry association personnel, government representatives, investors, consultants, and other industry stakeholders. Analysts are diligent in pursuing interviews with representatives from every part of the value chain in an effort to gain a comprehensive view of current market activity and future plans. Within the firm's surveys and focus groups, respondent samples are carefully selected to ensure that they provide the most accurate possible view of demand dynamics within consumer and business markets, utilizing balanced and representative samples where appropriate and careful screening and qualification criteria in cases where the research topic requires a more targeted group of respondents.

Tractica's primary research is supplemented by the review and analysis of all secondary information available on the topic being studied, including company news and financial information, technology specifications, product attributes, government and economic data, industry reports and databases from third-party sources, case studies, and reference customers. As applicable, all secondary research sources are appropriately cited within the firm's publications.

All of Tractica's research reports and other publications are carefully reviewed and scrutinized by the firm's senior management team in an effort to ensure that research methodology is sound, all information provided is accurate, analyst assumptions are carefully documented, and conclusions are well-supported by facts. Tractica is highly responsive to feedback from industry participants and, in the event errors in the firm's research are identified and verified, such errors are corrected promptly.

Chart 7.1 Tractica Research Methodology


(Source: Tractica)

NOTES

CAGR refers to compound average annual growth rate, using the formula:

$$\text{CAGR} = (\text{End Year Value} \div \text{Start Year Value})^{(1/\text{steps})} - 1.$$

CAGRs presented in the tables are for the entire timeframe in the title. Where data for fewer years are given, the CAGR is for the range presented. Where relevant, CAGRs for shorter timeframes may be given as well.

Figures are based on the best estimates available at the time of calculation. Annual revenues, shipments, and sales are based on end-of-year figures unless otherwise noted. All values are expressed in year 2017 U.S. dollars unless otherwise noted. Percentages may not add up to 100 due to rounding.

Published 3Q 2017

© 2017 Tractica LLC
1111 Pearl Street, Suite 201
Boulder, CO 80302
Tel: +1.303.248.3000
Email: info@tractica.com
www.tractica.com

This publication is provided by Tractica LLC ("Tractica"). This publication may be used only as expressly permitted by license from Tractica and may not otherwise be reproduced, recorded, photocopied, distributed, displayed, modified, extracted, accessed or used without the express written permission of Tractica. Notwithstanding the foregoing, Tractica makes no claim to any Government data and other data obtained from public sources found in this publication (whether or not the owners of such data are noted in this publication). If you do not have a license from Tractica covering this publication, please refrain from accessing or using this publication. Please contact Tractica to obtain a license to this publication.