

WHITE PAPER

Deep Learning Use Cases for Computer Vision

Six Deep Learning-Enabled Vision Applications in Digital Media, Healthcare, Agriculture, Retail, Manufacturing, and Other Industries

Published 2Q 2016

Published in Partnership with:



BRUCE DALEY
Principal Analyst

CLINT WHELOCK
Managing Director

SECTION 1

EMBEDDED VISION ALLIANCE



The [Embedded Vision Alliance](#), host of the Embedded Vision Summit, is a worldwide industry partnership that brings together technology providers who are enabling innovative and practical applications for computer vision. Its more than 50 members include suppliers of processors, sensors, software, and services. The mission of the Alliance is two-fold: 1) Inspire and empower product creators to incorporate visual intelligence into new products and applications, and 2) Enable Alliance member companies to accelerate success in computer vision by:

- Increasing their visibility among vision system and application developers
- Connecting them with the right ecosystem partners
- Delivering insights into market research, technology trends, and customer requirements

Founded in 2011, the alliance is led by Berkeley Design Technology, Inc. (BDTI), an independent technology analysis and engineering services firm. BDTI manages the Alliance on behalf of the member companies.

EMBEDDED VISION SUMMIT



The [Embedded Vision Summit](#) is *the* event for innovators who want to incorporate visual intelligence into products. The 2016 Summit, a three-day educational event, concentrates exclusively on deployable computer vision technology. The program is designed to inspire participants to use embedded vision technology and deep learning in new ways and to empower them with the know-how they need to integrate vision capabilities into products. The Vision Technology Showcase features demos, exhibits, and presentations on the enabling technologies.

SECTION 2

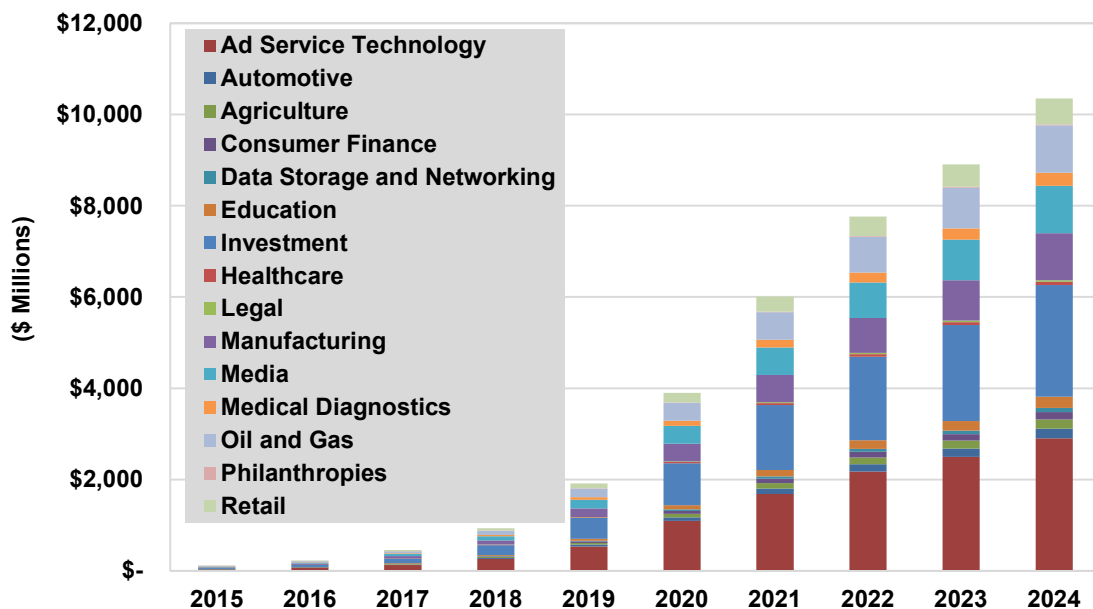
INTRODUCTION

The enterprise applications for deep learning have only scratched the surface of their potential applicability and use cases. Because it is data agnostic, deep learning is poised to be used in almost every enterprise vertical market, including agriculture, media, manufacturing, medical, healthcare, and retail, to name a few. Deep learning is particularly applicable to computer vision systems because it promises to be less costly, more accurate, and more reliable than traditional programming approaches.

Some of the most successful companies in the world have been early adopters of this technology. Although the enterprise market for deep learning is still small in relation to the total enterprise software sector, the variety, breadth, and scope of the applications that deep learning is being considered for suggests that a tremendous growth opportunity exists.

In this white paper, Tractica presents a diverse range of use cases for computer vision systems using deep learning. These examples provide a flavor for the level of deep learning activity that is currently taking place and the opportunities that it is creating for the future. Overall, Tractica forecasts that annual software revenue for enterprise applications of deep learning will increase from \$109 million in 2015 to \$10.4 billion by 2024, representing a compound annual growth rate (CAGR) of 65.8%. Note that this forecast pertains to the total market for deep learning software, rather than to the specific segment of the market related to deep learning-enabled computer vision.

Chart 2.1 Total Deep Learning Software Revenue by Industry, World Markets: 2015-2024



(Source: Tractica)

SECTION 3

USE CASES

3.1 STATIC IMAGE RECOGNITION, CLASSIFICATION, AND TAGGING

3.1.1 WHY THERE IS A NEED


Mobile devices and social media platforms have fueled an explosion of searchable images. Every day, Facebook users upload 350 million new photos, and Google is estimated to have indexed over 1 trillion images. Since over 50% of the human cerebral cortex is devoted to vision, images are a valuable commodity, but only if they can be found. Using human beings to recognize, classify, and tag static images is time-consuming, expensive, and tedious work that is highly error-prone.

3.1.2 HOW IT WORKS

Companies like Google and Facebook use deep learning to create static image recognition algorithms to identify objects within images. The objects are tagged with text that identifies what the object is called in human language. The tag is then indexed by grouping the object identifier and the location of the image with other objects with similar meaning. Depending on their business model, companies then associate the link with advertising related to the object. For example, when people search for an image of a cat, they are also served ads for cat food, kitty litter, or veterinarian services. Given their size and the number of advertisers they have, Google and Facebook can afford to hire the most accomplished deep learning talent on the planet and pay them handsomely. Organizations with fewer resources and less lucrative business models must find alternatives.

In addition to recognition and classification of still images, videos can also be analyzed utilizing similar techniques. While this is not as common today due to comparatively high computational demands, time and technology improvements will make automated video analytics more common in the future.

Figure 3.1 *Static Image Recognition, Classification, and Tagging Example*



(Source: Tractica)

MetaMind, an artificial intelligence (AI) technology startup, offers computer vision technologies to less profitable businesses in a wider variety of industries so that they can benefit from static image recognition. Customers can use MetaMind’s existing classifications to categorize and tag static images, or they can train MetaMind’s software to work with customized visual classifications by uploading their own examples. In addition, MetaMind’s

language technology enables customers to use machine learning to analyze language cues and context. Using a dynamic memory network, MetaMind hopes to train AI to mimic the human brain by processing sound, sight, and text at the same time so that it can answer complex questions that require synthesis and prioritization of simultaneous input.

3.1.3 BENEFITS

Google's mission is to organize the world's information and make it universally accessible and useful. Tagging images efficiently is a great step forward in organizing information. In addition, large-scale static image recognition can identify hidden patterns in data that are not obvious to the human mind. Much of the progress made in other use cases in this paper, such as improving and expanding crop yields in agriculture or manufacturing higher quality products, is a direct result of research in applying deep learning to static image recognition.

3.1.4 CHALLENGES

This technology is not a panacea. Deep learning algorithms will identify patterns that are statistically valid, but conceptually false. At their statistical heart, deep learning analyses will produce false positives and false negatives and make mistakes no human would ever make. Relying on these conclusions too heavily may result in false conclusions. In addition, humans can hack and game deep learning systems. Finally, widespread adoption of deep learning image recognition systems makes privacy more difficult to maintain.


3.2 DIGITAL RADIOLOGY ANALYSIS

3.2.1 WHY THERE IS A NEED

Radiologists are highly-trained, highly-skilled physicians; consequently, reimbursement for their services is one of the most expensive healthcare costs. As these costs soar in the developed world and societies in many large economies, including the United States, Japan, and Spain, face the challenge of an aging population, many stakeholders are looking for a way to shift from highly reimbursable, volume-driven diagnostic practices to less expensive, more effective methods. They would like to eliminate waste in the current system, reduce the number of mistakes, and compensate radiologists for complex medical diagnoses rather than for routine tasks.

3.2.2 HOW IT WORKS

Computer aided diagnosis (CAD) is already widely used in breast and lung nodule detection. Researchers and entrepreneurs are beginning to deploy deep learning to identify and analyze medical images from many sources, including X-rays, magnetic resonance imaging (MRI) scans, computed tomography (CT) scans, and 3D technology. Such systems are trained to identify and aggregate a large quantity of medical data and improve their pattern recognition over time with new images and other health data to better diagnose medical anomalies in images.

Figure 3.2 Digital Radiology Analysis Example


(Source: Tractica)

Building on the early successes of CAD, Enlitic, a deep learning startup, is training its neural net with a large data set of hospital and medical imaging company records of diagnostic images, lab results, doctors' notes, and claims records. Enlitic's goal is to find invisible patterns in the data that are not obvious to human cognition.

3.2.3 BENEFITS

Enlitic and other innovators envision speedier, more accurate, and cost-effective diagnoses. Large data opens the door to evidence-based medicine, and the technology may be of particular benefit in third world countries and rural areas where healthcare professionals are often hard to find. Also, by decreasing the cost of initial diagnostic procedures, this technology may enable more people to obtain the screenings that are essential to early diagnosis of medical conditions, improving the accuracy of the system and leading to evidence-based medicine.

3.2.4 CHALLENGES

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. The radiology community has adopted many advances in technology, but may be more reluctant to embrace deep learning tools that appear to usurp their role. Although the technology may eventually result in a lower demand for radiologists, there is likely to always be a demand for the eye, brains, skill, and experience of talented radiologists.

3.3 AGRICULTURAL CROP HEALTH ANALYSIS

3.3.1 WHY THERE IS A NEED

As the global population grows, food security and healthy crops are vital to avoiding large-scale human catastrophes, such as famine, social unrest, and violence, related to shrinking resources. Agriculture is one of the largest global economic sectors, employing over 1 billion people. Most of these people are subsistence farmers, which makes them and their families particularly vulnerable to crop failure because they are dependent on seasonal and yearly crops to survive. Even farmers and larger agricultural enterprises that are not at the subsistence level struggle to overcome the vagaries of nature and other forces that are destructive to their crops and livelihood. Farmers around the world are looking for better ways to identify crop problems and improve harvests. Much of the improvement of yields in developing economies has come from adopting techniques from the developed world.

3.3.2 HOW IT WORKS

Remote agricultural crop health analysis combines deep learning, cloud-based computing, and remote imaging devices to conduct satellite-based inventories of crops; identify soil health, wet spots in fields, or irrigation and waste runoff; and monitor crop canopies. The sensing is based on the fact that in the near infrared (NIR) part of the spectrum, healthy plants reflect up to 7 times more light than in the visible spectrum. GeoVisual Analytics, a Boulder-based technology startup, uses remote imaging to distinguish healthy plants (those reflecting relatively more NIR) from diseased or thirsty plants (those reflecting less NIR). As part of the segmentation and specialization capability this approach offers, GeoVisual can also use its technology to focus on higher-value specialty crops, such as fresh vegetables. Meanwhile, Caterpillar, the construction and mining equipment manufacturer, has developed the CAT S60 smartphone, which features a thermal camera. This device can be mounted on a truck, serving as a “poor man’s satellite” that is able to track crop problems through heat signatures. Similar smartphone technology with NIR cameras may be developed in the near future.

Figure 3.3 *Agricultural Crop Health Analysis Example*



(Source: Tractica)

2.3.3 BENEFITS

The power to analyze and predict crop yields at an early stage allows farmers and agronomists to develop effective short-term interventions. As more and more imagery is collected within and across regions, the hidden patterns that are only accessible to deep learning can also help inform large-scale, long-term crop strategies. The human and economic benefits of expanding food supplies and food security cannot be overestimated in a world where 70 million people are born each year.

3.3.3 CHALLENGES

Although researchers are trying to reduce the cost by employing reusable rockets, satellites remain expensive to launch. They currently cost between \$100 million and \$260 million (depending on the size of the payload). Satellites are also constrained by cloud cover. Drone technologies are controversial because of security and privacy concerns and are subject to restrictions by the Federal Aviation Administration (FAA) and other government agencies. Manned aircraft offer a viable platform for relatively inexpensive, broad area imagery, but for some applications, image resolution can be a limitation. Even a device like the Caterpillar S60 smartphone, which costs less than \$600, is prohibitively costly for subsistence farmers. As the technology spreads throughout the world, it may face even greater scrutiny and restrictions in regions with more military and political sensitivities than in the United States. As these technologies grow, they must include more advanced compliance mechanisms and have built-in security and privacy features in order to be widely adopted.

3.4 CLINICAL TRIAL MEDICATION COMPLIANCE

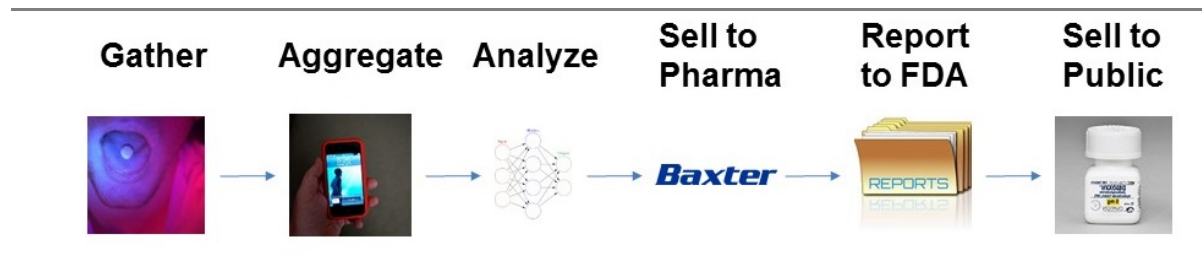
3.4.1 WHY THERE IS A NEED

Patients not taking the right medications, at the right time, and in the right doses is a major problem in medicine. The problem is particularly acute in clinical trials. Studies have found that many trial participants are non-compliant even when they are paid to be part of the study or are only taking placebos. If the trial participant fails to take the medication, then the result is likely to be a false negative. The financial consequences of a failed clinical trial can be enormous. In 2014, the Tufts Center for the Study of Drug Development estimated that developing a new prescription drug costs over \$2.5 billion and a single clinical trial may cost over \$100 million. Accurately predicting patient outcomes from clinical trials benefits the larger patient population, healthcare professionals, and the public, but it is only one example of medication non-compliance.

3.4.2 HOW IT WORKS

In response, innovative scientists and entrepreneurs are turning to biometrics, big data, remote communications, deep learning, and computer vision to increase medication compliance. One approach is a medication compliance app that uses computer vision and deep learning. It can be downloaded to certain mobile devices, which makes it an inexpensive option for the 200 million Americans who own smartphones. The app reminds participants when a medicine must be taken and uses facial recognition software to authenticate patients. The participant takes a picture of their pill so that it can be confirmed as the correct medication and takes another picture of the pill in their mouth. One of the companies using this approach, AiCure, is working to document its effectiveness even among complex populations, including elderly stroke patients, people suffering from schizophrenia, and participants in clinical trials for medications to prevent HIV.

Figure 3.4 Clinical Trial Medication Compliance Example



(Source: Tractica)

3.4.3 POTENTIAL PROBLEMS

Remote monitoring of medication compliance shares many of the security, privacy, and access challenges of the larger digital health sector. Mobile devices require protection from hacking and spoofing, including multiple security layers and monitoring redundancies. Personal information must comply with data privacy legislation. Network connections must be continuous enough to support extensive data transfer and integrity, including synchronizing information on the device during power outages. In addition, although two-thirds of Americans currently own smartphones, not all patients have access to mobile technology.

3.4.4 CHALLENGES

Beyond clinical trial participants, there are other groups of patients who are at risk of medical non-compliance, including seniors balancing taking many medications, people with mental illness who may intentionally avoid medications, and substance abusers. Widespread adoption could ultimately lead to improvements in patient engagement and monitoring of individual patients' reactions to medications.

3.5 CLOTHES AND ACCESSORIES SIZING AND FITTING


3.5.1 WHY THERE IS A NEED

During the dot-com boom, online clothing sales were predicted to grow to 40% to 50% of total sales. Although online sales of other kinds of merchandise, such as books, have reached 50% of the market in the past 15 years, the percentage of online clothing sales hovers around 20%. The difficulty in finding the correct size and fit is one of the primary reasons that consumers are reluctant to buy clothes online. And their concern is not groundless; sizing varies among clothing manufacturers, and it is difficult to ascertain fit from online images. Consequently, 30% to 40% of online clothing purchases are returned. In addition, certain kinds of clothes fitting, such as for shoes and bras, are difficult even in brick-and-mortar stores.

3.5.2 HOW IT WORKS

Computer vision has been used to address the problem of finding the right fit online by using mobile devices, digital scanning, and deep learning. For example, San Francisco-based ThirdLove offers a free app for iPhones that enables women to find the right bra from home. The app calculates a woman's measurements, cup size, shape, and fit by comparing the distance between body contours in front and side view photos that the woman takes while wearing a supportive bra and fitted top, standing before a mirror, and holding her iPhone camera at the level of her navel. The app bases the sizing on instantaneous calculations, and protects consumer privacy by not storing the images on the cloud or the iPhone.

Figure 3.5 Clothes and Accessories Sizing and Fitting Example



(Source: Tractica)

The Stockholm-based company Volumental offers computer vision applications for sizing shoes and eyewear. Its VANDRA hardware platform contains four depth cameras that make a 3D, volumetric scan of a foot in 2 seconds, including areas that retailers normally do not measure. The data is sent to a tablet, which a salesperson consults in order to make shoe recommendations from store inventory or other retailers. As another application of computer vision fitting, Volumental's VACKER platform uses 3D face scanning to assist opticians and customers to find the most flattering eyewear for different face shapes and sizes. While its technology is currently store-based, Volumental anticipates that depth cameras will

eventually be integrated into smartphones, tablets, and laptops in 2017, allowing customers to use VANDRA and VACKER for online purchasing as well.

3.5.3 CHALLENGES

Technologies that accurately assess clothing/accessory size and fit provide a more satisfying and individualized retail experience for customers, not to mention enhancing their overall comfort by ensuring that they wear clothes that flatter and suit them. At the same time, computer vision fitting cuts down on retailers' expenses related to product returns and customer complaints. By providing more seamless online shopping, computer vision fitting also supports online markets, expands customer choice, and will eventually provide access to personally tailored clothes.

3.5.4 CHALLENGES

Some of this technology is dependent on features like depth sensing cameras that are not yet available to ordinary digital consumers, prompting some analysts to say that the technology is "ahead of its time." Also, frustration is notoriously influential in deterring online customers from converting their interest in an item to a purchase, and smartphone apps can be buggy. Some people may have trouble holding phones at the right angles and following instructions. Older phones may not have the touch sensitivity that is essential to the fitting process. Nevertheless, ThirdLove's success, which includes return rates of less than 5% and one-third of bra sales resulting from the iPhone sizing app, demonstrates the vast potential of fitting technology.

3.6 INDUSTRIAL AUTOMATION QUALITY ASSURANCE

3.6.1 WHY THERE IS A NEED

Improving the quality of the parts that go into manufactured goods greatly benefits the consumer and the maker. Defective bolts and fasteners, which are typically made with high-quality metal, can cause manufacturers to build faulty and dangerous equipment. In addition to the damage to their brand, manufacturers can lose money when their products fail. For example, when defective bolts and fasteners pass quality inspections and are placed in automobiles, they can result in damage to the car and serious accidents. In 2012, the number of automobile recalls almost equaled the number of sales, and a high volume of warranty payments impacted automakers' profits. When Portland Bolt, a leading bolt manufacturing company, discovered that the metal bolts' grain lines were the source of the defective bolt problem, the challenge was to develop quality assurance mechanisms to consistently identify these tiny but destructive grain line flaws.

3.6.2 HOW IT WORKS

Portland Bolt improved the quality of its products by gathering data through a \$100 scanner purchased at Best Buy and creating a deep learning grain bolt analysis tool on the Amazon Web Services cloud platform. It scans its bolts and uses deep learning to compare them with cloud-based metal bolt grain images to find the grain line flaws. In the same vein, Sight Machine has developed a manufacturing analytics platform that gathers and analyzes quality assurance and efficiency data throughout production cycles and across multiple vendors for its manufacturing clients.

Figure 3.6 Industrial Automation Quality Assurance Example


(Source: Tractica)

3.6.3 BENEFITS

The metal bolt grain case study shows how dangerous and costly tiny manufacturing flaws can be. However, machine vision and large-scale data analysis, which cannot be completed by the human eye on an assembly line or in a quality control process, are able to “see” patterns that can predict flaws that would otherwise be invisible. In addition, web and mobile technologies, dashboards, and customized data modeling are able to share information about these patterns in order to educate workers, quality engineers, and managers about the most effective and efficient manufacturing and quality assurance processes.

3.6.4 CHALLENGES

Many manufacturers have experience with 30-year-old, manually programmed machine vision and quality control technology that is not connected to the Internet or the cloud. In order for them to recognize the potential of this new wave of vision technology and make progress in their quality control functions, they will have to overcome their natural cynicism and unlearn the earlier machine vision technology.

3.7 OTHER OPPORTUNITIES

3.7.1 EMOTION RECOGNITION

Human beings are social animals. As part of our neural wiring, humans assess each other’s emotional cues on a subconscious level without realizing it. However, some people, such as individuals with autism, have more trouble interpreting these cues. Some cues are difficult for everyone to read. Other cues are even intentionally misleading. Missing from most definitions of AI is this subconscious emotional intelligence. If AI is ever to successfully work and communicate with humans on a peer basis, it needs to take into account human emotion. Technology that helps people recognize and read emotions also has the potential to help humans communicate more effectively with each other.

Emotion recognition technology uses deep learning to identify images of prototypical facial expressions of such universal human emotions as joy, surprise, sadness, disgust, fear, and anger. Emotient, an AI startup that was recently acquired by Apple, has developed software named Facet that reads emotions in pictures or video frames with resolutions as low as 40 by 40 pixels. When Facet analyzes video sequences, it is able to track emotions over time and even catch the “micro-expressions” that humans cannot control and are generally unaware they are signaling.

3.7.2

NEUROMORPHIC COMPUTING

Just about any human intellectual activity is a candidate for the practical application of the current generation of AI. Organic intelligence, because it has evolved over billions of years of trial and error, has a much more complex electrochemical structure of neurons and synapses (to name but one of its many layers) than deep learning. In an attempt to overcome the binary limitations of the current generation of computers, a new approach to AI is being developed, to replicate the biological structure of organic nervous systems in hardware. This effort is known as neuromorphic computing.

IBM has developed a chip named TrueNorth that attempts to mimic the way human brains recognize visual patterns, relying on densely interconnected webs of transistors similar to the brain's neural networks. When it was first announced in *Science* in 2014, the chip had 4,096 processors, each of which contained 256 integrate-and-fire spiking neurons (with binary output) each with 256 inputs. Each processor can compute all 256 neurons 1,000 times per second. The overall peak performance was 4,096 times 256 times 256 times 1,000 = 266 GSops (billion synaptic operations per second). The power consumption was around 100 mW. Since then, the chip's performance has been improved and it recently went into testing at Lawrence Livermore Labs.

SECTION 4

COMPANY DIRECTORY

AiCure

902 Broadway, 6th Floor
New York, NY 10010, USA
www.aicure.com
+1.800.570.0448

Caterpillar

501 SW Jefferson Street
Peoria, IL 61614, USA
www.cat.com

Enlitic

10 Jackson Street
San Francisco, CA 94111, USA
www.enlitic.com

Facebook

1 Hacker Way
Menlo Park, CA 94025, USA
www.facebook.com
+1.650.543.4800

Google Inc.

1600 Amphitheatre Parkway
Mountain View, CA 94043, USA
www.google.com
+1.650.253.0000

GeoVisual Analytics

1525 Spruce Street, #200
Boulder, CO 80302, USA
www.geovisual-analytics.com
+1.303.955.1575

H2O.ai

1185 Terra Bella Avenue
Mountain View, CA 94043, USA
www.h2o.ai
+1.650.429.8337

MetaMind, Inc.

172 University Avenue
Palo Alto, CA 94301, USA
www.metamind.io

Sight Machine

300 Broadway, Suite 28
San Francisco, CA 94133, USA
www.sightmachine.com
+1.888.461.5739

ThirdLove

2325 3rd Street, Ste. 215
San Francisco, CA 94107, USA
www.thirdlove.com
+1.866.205.6833

Volumental

Drottning Kristinas väg 53
114 28 Stockholm, Sweden
www.volumental.com
+46.76.882.53.45

SECTION 5

ACRONYM AND ABBREVIATION LIST

Artificial Intelligence	AI
Compound Average Growth Rate	CAGR
Computed Tomography	CT
Computer Aided Diagnosis	CAD
Federal Aviation Association (U.S.)	FAA
Magnetic Resonance Imaging	MRI
Natural Language Processing	NLP
Near Infrared	NIR
Three-Dimensional	3D

SECTION 6

TABLE OF CONTENTS

Section 1	1
Embedded Vision Alliance	1
Embedded Vision Summit	1
Section 2	2
Introduction	2
Section 3	3
Use Cases	3
3.1 Static Image Recognition, Classification, and Tagging.....	3
3.1.1 Why there Is a Need.....	3
3.1.2 How It Works	3
3.1.3 Benefits.....	4
3.1.4 Challenges.....	4
3.2 Digital Radiology Analysis.....	4
3.2.1 Why there Is a Need.....	4
3.2.2 How it Works	4
3.2.3 Benefits.....	5
3.2.4 Challenges.....	5
3.3 Agricultural Crop Health Analysis.....	5
3.3.1 Why there Is a Need.....	5
3.3.2 How it Works	6
3.3.3 Benefits.....	6
3.3.3 Challenges.....	6
3.4 Clinical Trial Medication Compliance	7
3.4.1 Why there Is a Need.....	7
3.4.2 How it Works	7
3.4.3 Potential Problems	7
3.4.4 Challenges.....	8
3.5 Clothes and Accessories Sizing and Fitting.....	8
3.5.1 Why there Is a Need.....	8
3.5.2 How it Works	8
3.5.3 Challenges.....	9
3.5.4 Challenges.....	9
3.6 Industrial Automation Quality Assurance	9
3.6.1 Why there Is a Need.....	9
3.6.2 How it Works	9
3.6.3 Benefits.....	10
3.6.4 Challenges.....	10
3.7 Other Opportunities.....	10
3.7.1 Emotion Recognition	10
3.7.2 Neuromorphic Computing	11
Section 4	12
Company Directory	12
Section 5	13
Acronym and Abbreviation List	13
Section 6	14
Table of Contents	14

Section 7	16
Table of Charts and Figures	16
Section 8	17
Scope of Study	17
Sources and Methodology	17
Notes	18
Section 9	19
Additional Reading	19

SECTION 7

TABLE OF CHARTS AND FIGURES

Chart 2.1	Total Deep Learning Software Revenue by Industry, World Markets: 2015-2024.....	2
Chart 8.1	Tractica Research Methodology.....	18
Figure 3.1	Static Image Recognition, Classification, and Tagging Example.....	3
Figure 3.2	Digital Radiology Analysis Example	5
Figure 3.3	Agricultural Crop Health Analysis Example.....	6
Figure 3.4	Clinical Trial Medication Compliance Example	7
Figure 3.5	Clothes and Accessories Sizing and Fitting Example.....	8
Figure 3.6	Industrial Automation Quality Assurance Example	10

SECTION 8

SCOPE OF STUDY

This white paper covers the market for computer vision and deep learning technologies, providing real-world use cases of how they are used in various industry verticals. The verticals covered include agriculture, media, manufacturing, medical, healthcare, and retail.

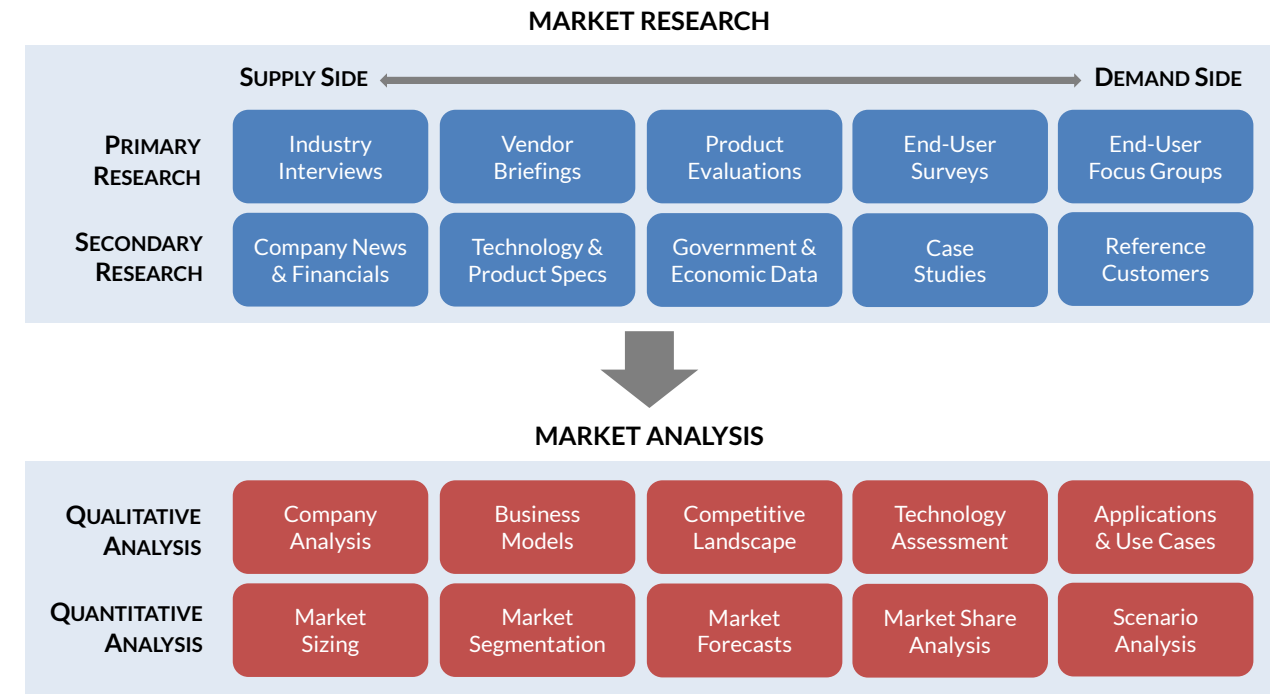
SOURCES AND METHODOLOGY

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 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 if errors in the firm's research are identified and verified, such errors are corrected promptly.

Chart 8.1 Tractica Research Methodology


(Source: Tractica)

NOTES

CAGR refers to compound annual growth rate, using the following 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 2015 U.S. dollars unless otherwise noted. Percentages may not add up to 100 due to rounding.

SECTION 9

ADDITIONAL READING

Tractica's Automation & Robotics advisory service examines global market trends and business models surrounding the growth of robotics and other automation technologies for consumer, commercial, and industrial applications. Research focus includes the development of AI and cognitive computing technologies, adoption patterns for automation and robotics technologies in key application areas, the dynamics of human interaction with these technologies, and assessment of economic and societal impacts of automation. Reports include detailed profiling of key industry players, new product introductions, and market forecasts by geography and segment.

Deep Learning for Enterprise Applications

Advertising Technology, Financial Services, Media, Manufacturing, Oil & Gas, Retail, and Other Enterprise Markets for Deep Learning Software and Systems

Published 4Q 2015

<https://www.tractica.com/research/deep-learning-for-enterprise-applications/>

This Tractica report examines the market for deep learning across 15 industry sectors, including 10-year revenue forecasts for the period from 2015 through 2024. The report forecasts software revenue, services revenue, and hardware sales, and provides segmentation by delivery method (cloud, device, and data center) as well as world region. The report also includes profiles of 20 key industry players in the global deep learning market.

3D Cameras and Sensors

3D Imaging Hardware and Software for Consumer & Mobile, Automotive, Robotics & Industrial, and Other Applications: Global Market Analysis and Forecasts

Published 4Q 2015

<https://www.tractica.com/research/3d-cameras-and-sensors/>

This Tractica report examines the market trends and technology issues surrounding 3D imaging technologies and presents 10-year forecasts for hardware and software during the period from 2015 through 2024. The report presents in-depth analysis of 3D cameras and sensors, as well as their enabling technologies, and also includes detailed profiles of 26 key industry players. Market forecasts, segmented by world region, include 3D imaging hardware unit shipments, device pricing and revenue, software licenses, and revenue for four key application markets.

Computer Vision Technologies and Markets

Hardware and Software for Automotive, Sports and Entertainment, Consumer, Robotics and Machine Vision, Surveillance and Security, and Medical Markets

Published 2Q 2015

<https://www.tractica.com/research/computer-vision-technologies-and-markets/>

This Tractica report examines the market dynamics, technology issues, and business opportunities surrounding computer vision and presents 5-year forecasts for computer vision hardware and software during the period from 2014 through 2019. Six application markets are analyzed in depth, and key industry players are profiled. Market forecasts, segmented by world region and vertical market, include hardware unit shipments, device pricing, and software revenue.

Artificial Intelligence for Enterprise Applications

Deep Learning, Predictive Computing, Image Recognition, Speech Recognition, and Other AI Technologies for Enterprise Markets: Global Market Analysis and Forecasts

Published 2Q 2015

<https://www.tractica.com/research/artificial-intelligence-for-enterprise-applications/>

This Tractica report examines the practical application of artificial intelligence within commercial enterprises. The technologies covered include cognitive computing, deep learning, machine learning, predictive APIs, natural language processing, image recognition, and speech recognition. The report presents 10-year forecasts for AI software, along with AI-driven services and hardware sales, for the period from 2015 through 2024. Forecasts are segmented by world region as well as by technology and industry application. The study also includes detailed profiles of more than 30 key industry players.

Published 2Q 2016

© 2016 Tractica LLC
1111 Pearl Street, Suite 201
Boulder, CO 80302 USA
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.