Cognizant

Digital Systems & Technology

Machine Learning: The First Salvo of the AI Business Revolution

Advances in artificial intelligence, such as machine learning, promise new business-critical capabilities that create enhanced and differentiated customer experiences and market advantages. Here's a look at the thinking, tools and frameworks necessary to move forward in unlocking machine learning's vast potential.

Executive Summary

In the form of machine learning (ML), artificial intelligence (AI) is making its way into the enterprise, promising to transform many areas of the business including customer experience. Just as humans can learn and act from experience, machines can learn from data, identify patterns and make decisions. Furthermore, this cognitive ability of machines can improve dramatically over time. The higher intelligence demonstrated by machines has led to wide application across industries. Virtually every industry – such as banking and financial services, and healthcare and life sciences – is now leveraging ML's innate ability to revitalize core business processes and models.

From start-ups to leading players, organizations are interested in AI of all forms, with ML positioned as the best way forward to innovate and maximize business outcomes. International Data Corporation (IDC) estimates that spending on AI and ML will grow from $19.1 billion in 2018 to $52.2 billion by 2021.[1]

However, we see four key shortfalls that thwart ML adoption across businesses:

- A lack of infrastructure and of accessible data.
- A lack of understanding and skills concerning where and how to apply ML.
- A lack of a talent pool to collect the high-quality data necessary to train ML algorithms.
- A lack of understanding of how to build business models for ML.

This white paper details key dimensions of ML including the learning process, the developmental lifecycle and popular ML frameworks. It also showcases how ML is being used to create successful applications that are invoking new possibilities for the future. Finally, we offer guidance for enterprises seeking to leverage the technology to advance their digital business objectives. The Evolution of Machine Learning

The evolution of machine learning

From our perch, we see ML accelerating across enterprises, impacting fields as diverse as human communications, autonomous cars, fraud detection and disease diagnosis. With traditional computational approaches, algorithms are explicitly programmed to solve particular problems. In ML, however, the system can identify data patterns automatically and improve on experience without being explicitly programmed. Within a few years, ML could potentially become integral to far more effective and widely available applications that perform tasks on their own and create fundamentally new, profitable business models.

ML's increasing popularity is based on a variety of factors (see Figure 1):

- **Data availability:** The vast and growing amount of structured and unstructured data in various formats – such as plain text, raw images, video files and audio files – is fueling interest in ML. ML can learn from all these data sources and autonomously build applications.
- **Computing power:** Processing large volumes of data has become more affordable and cost-effective through cloud services. Efficient parallel computing is enabled by graphics processing unit (GPU) computing, one of the most pervasive, accessible and energy-efficient (i.e., faster and with less infrastructure) platforms for training ML models.
- **Tools and frameworks**: The availability of various open source frameworks, toolkits and libraries makes it easier to build, evolve, implement and scale the ML models using popular programming languages.

Machine learning's key drivers

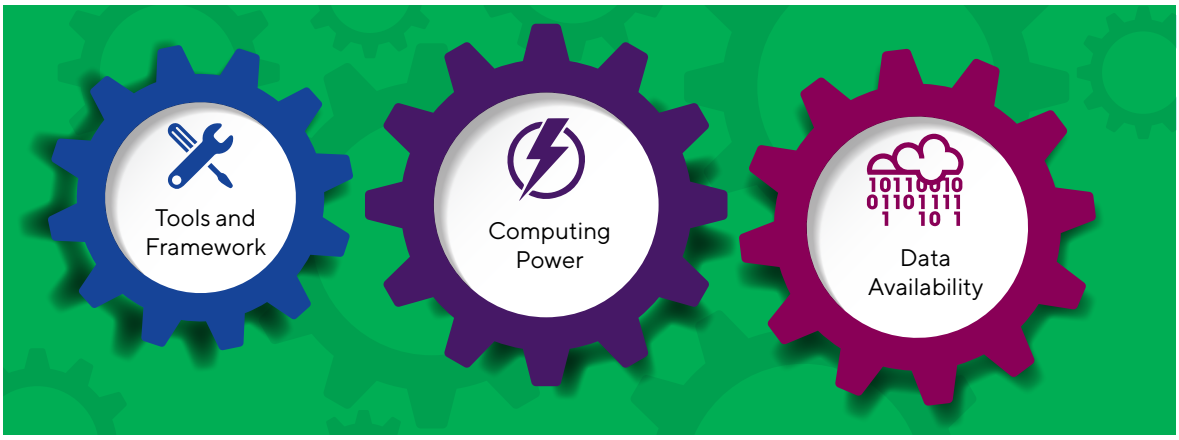

Figure 1

ML offers the potential for creating previously unattainable business platforms and solutions. Relevant use cases of ML at work include:

- Banks could apply biometrics for facial recognition and voice authentication to improve customer experience and security.
- Insurance companies could automatically recognize and assess vehicle damages, thereby lowering risk.
- Hospitals could create a library of fundus images to screen and detect diabetic retinopathy. (See our white paper, "How AI Enhances & Accelerates Diabetic Retinopathy Detection.")
- Retailers could enhance retail product metadata quality, and thereby help consumers make more effective buying decisions. (See our white paper, "Using AI to Enhance the Quality of Retail Metadata.")

ML is a constantly evolving field. Hence, organizations must stay ahead of how ML frameworks evolve, and continuously evaluate ML's potential impact on their businesses.

ML process/development lifecycle

ML is often an iterative process, and collecting, training and applying large volumes of data to develop suitable ML models is very challenging. The data used by an ML algorithm must contain the information about prior occurrences of a situation or condition, suitable to make predictions. When business conditions change, ML models can simply be retrained on new data, without the need to rewrite the instructions as you would for explicitly programmed systems. The essentials include:

- **The right questions:** In machine learning, asking the right questions pertaining to the identified business problem will help in understand the value to be generated by an ML solution.
- **The right data:** Having accurate and complete data sets will establish reliable inputs and outputs to train the model for an ML solution.
- **Success criteria:** The ML algorithms define the success criteria, which should be measurable on an ongoing basis. This enables the system to continually learn and adjust its algorithms to help the business meet its objectives.

We suggest the phased approach to the ML lifecycle depicted in Figure 2:

- **Deep analysis:** Establish a good understanding of the business, formulate the problem and define the success criteria. Identify the different data sources and influencers.

- **Algorithm selection:** Define the mathematical model that best represents the data distribution. Assuming the simplest case, where the data is distributed nearly linearly, the mathematical model would be y=f(x) where y is the output (dependent variable) and x is the input (independent variable).
- **Data preparation:** Understand the impact of the attributes on the desired outcome and synthesize new attributes as required. Curate and segregate data into training, validation and testing sets. Ensure that data distribution is consistent across training, validation and testing, and is representative of the full data set.
- **Training:** Iteratively refine the parameters of the mathematical model to best fit the training data. Use pre-built optimizers that can iteratively modify the model to deliver the right optimization and hyper parameters (e.g., a parameter whose value is set before the learning process begins).
- **Validation:** Validate whether the mathematical model defined is actually the best representation of the training data. Validate if the y=f(x) is the best representation of the data set, using the validation data set.
- **Testing:** Assess the model's quality by measuring prediction accuracy. Evaluate how close the prediction output is to the actual output for the given testing data set. Use the right metric to measure qualities such as precision, recall or F1 score.
- **Prediction:** Evaluate the mathematical model with the refined parameters for new data. Evaluate y=f(x) using different values of x, for new data (to include prediction timeline).
- **Model refresh:** If the nature of the underlying data changes, revisit the mathematical model and/or refine the parameters. If there are recurring instances of new data that doesn't fit y=f(x), revisit the equation itself or train with new data to further modify the model. Based on the nature of the underlying process, a refresh can be performed in real time or periodically (daily, weekly, monthly, etc.). The shorter the refresh cycle, the higher the cost.

ML development lifecycle

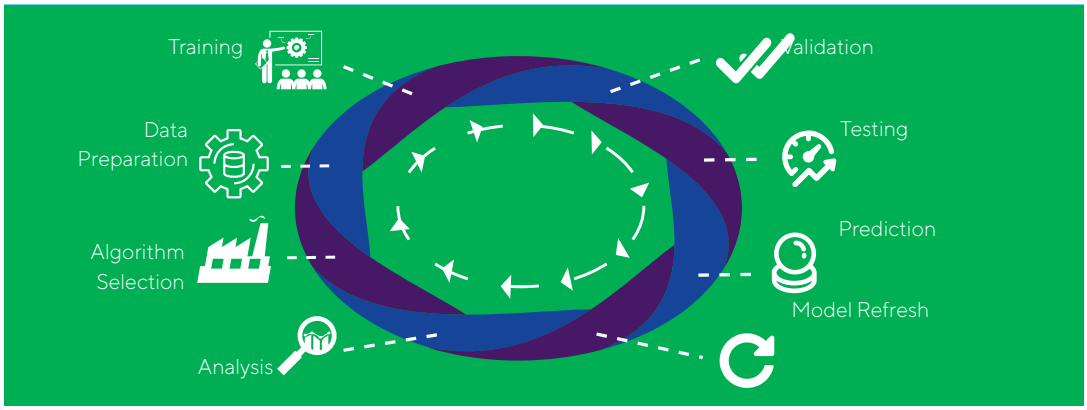

Figure 2

Machine learning involves mapping from a set of inputs to a set of outputs. Several examples are shown in Figure 3.

Sample ML application inputs & outputs

Input X	Output Y	Application
Patient Fundus Images	Detection	DR Screening
Retail Product Label Images	Extraction	Metadata Extraction
Face	Names	Face Recognition

Figure 3

At regular intervals, repeat the entire process to create a model that can be updated occasionally to adapt to changes. The learning curve will detail how the ML model is being developed over a period of time to achieve better business outcomes, using various popular ML frameworks such as TensorFlow, Theano, Keras, Microsoft Cognitive Toolkit, Caffe, Torch and Apache MXNet.

When delving into ML, choosing one framework from among the many choices can be daunting. There are different data sets, libraries, frameworks, toolkits and applications to contend with. It is thus wise to evaluate the options and become familiar with ML frameworks during the decision-making process, and make your choice accordingly.

ML frameworks

An efficient ML framework reduces the complexity of machine learning, making it accessible to more developers. An effective ML framework is:

- **Simple to understand and easy to code.**
- **Developer-friendly for building models.**
- **Easily optimized for high performance computing.**
- **Capable of automating the computation process.**

Examples of ML frameworks include the following:

- **TensorFlow** is an open-source software library that uses data flow graphs for numerical computation. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. The flexible architecture allows developers to deploy computation to one or more CPUs or GPUs in a desktop, server or mobile device with a single API.[2]

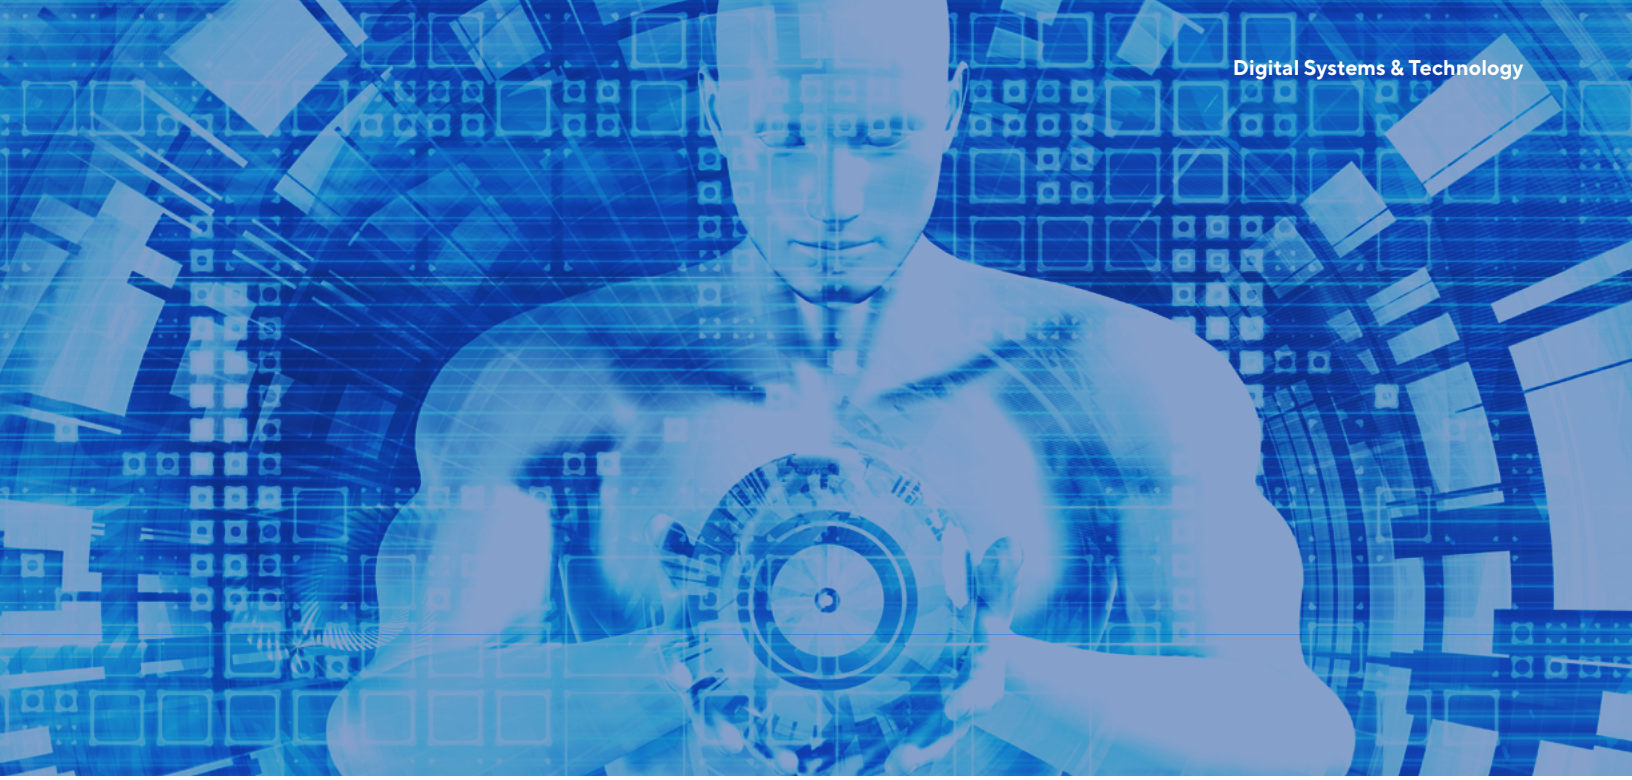

- **Theano** is a Python library that defines, optimizes and evaluates mathematical expressions with multidimensional arrays. Theano creates derivatives for functions with one or many inputs. It quickly evaluates expressions and detects and diagnoses many types of errors. Examples of some popular libraries are Lasagne, Blocks and Keras.[3]

- **Keras** is a high-level neural network set of APIs, written in Python and capable of running on top of TensorFlow, Microsoft Cognitive Toolkit or Theano. The ability to go from idea to result with the least possible delay is key to conducting proper research. Developers should use Keras if they need a deep-learning library to provide easy and fast prototyping (through user-friendliness, modularity and extensibility). It supports both convolutional networks (a class of deep, feed-forward artificial neural networks most commonly applied to analyzing visual imagery) and recurrent networks (a class of artificial neural network most commonly applied in handwriting or speech recognition), as well as combinations of the two. Importantly, it runs seamlessly on CPUs and GPUs.[4]

- **Microsoft Cognitive Toolkit** (previously known as CNTK) provides sophisticated algorithms and production readers for working with massive data sets. It trains and evaluates deep-learning algorithms faster than other toolkits, scaling efficiently in a range of environments – from CPUs to GPUs to multiple machines – while maintaining accuracy. Working with popular languages and networks such as C++ and Python, it empowers developers to customize any of the built-in training algorithms.[5]

- **Caffe** is a deep-learning framework made with expression, speed and modularity in mind. Models and optimization are defined by configuration, without hard coding, with an expressive architecture that encourages application and innovation. Caffe is perfect for research experiments and industry deployments which can use its power to process over 60 million images per day with a single NVIDIA K40 GPU. It was developed for computer vision/image classification by leveraging convolutional neural networks (CNNs).[6]

We recommend a quick comparison of popular ML frameworks across licensing, support roadmap, industry adoption, skills availability, feature-richness and interoperability.

- **Torch** is a scientific computing framework coded in a language called Lua. It has wide support for ML algorithms, and it's easy to read and understand. Torch comes with a large ecosystem of community-driven packages in ML, computer vision, signal processing, parallel processing, image, video, audio and networking. It also comes with a huge repository of sample code and guides to provide maximum flexibility and speed in building scientific algorithms while making the process extremely simple.[7]
- **Apache MXNet** is a modern, open-source, deep-learning framework that is used to train and deploy deep neural networks. It is scalable, allowing for fast model training. Also, it supports a flexible programming model and multiple languages – such as C++, Python, Julia, Matlab, JavaScript, Go, R, Scala, Perl and the Wolfram Language. Apache MXNet is a lean, flexible and ultra-scalable deep-learning framework that supports modern deep-learning models, including CNNs and long short-term memory networks (LSTMs).[8]

ML framework comparison

Before you get started, we recommend a quick comparison of popular ML frameworks across licensing, support roadmap, industry adoption, skills availability, feature-richness and interoperability. Licensing is a distinct issue among these ML frameworks. TensorFlow and Apache MXNet are licensed under Apache License 2.0, which is a permissive license whose main conditions require preservation of copyright and license notices. Theano, Torch and Caffe are licensed under a Berkeley Software Distribution (BSD) license, a permissive free software license that imposes minimal restrictions on the use and redistribution of the covered software. Keras is a permissive free software license originating at the Massachusetts Institute of Technology (MIT). Microsoft Cognitive Toolkit also uses the MIT licensing approach.

Figure 4, next page, displays ML framework ratings on industry adoption and skills availability, based on our experience working with ML algorithms as well as trends seen on Stack Overflow.[9]

Emerging ML prototypes

ML has the potential to redefine how business is conducted across industries. The following six use cases are examples of how ML will drive business value to enterprises.

Service desk optimization

Service desk or L1 technical support jobs (such as support ticket creation) contain rote and routine tasks that are increasing exponentially across industries. By 2021, more than 50% of enterprise spending will be on smart bots for optimization.[10] Smart speaker sales (like Amazon's Alexa) are projected to rise from $4.4 billion to $17.4 billion by 2022,[11] and New Voicebot Report says nearly 20% of U.S. adults own smart speakers.[12] Smart AI bots can replace the existing L1 support provided by human channels and by new voice sources such as Alexa and Google Home.

Solution: A user speaks to a smart AI bot through their choice of channel such as VoIP, a Chrome browser, Alexa or Google Home. The system converts the speech to text using Google ASR, understands the context and manages conversations using Dialogflow and NLP Engine and fetches/validates the information requested by the user from external resources such as knowledge repositories and databases. Once the system receives a response, it rephrases and converts the text into speech using Amazon Polly, a text-to-speech service that uses advanced deep learning technologies to synthesize human-sounding speech, and sends it back to the user. If the smart AI bot is unable to resolve the query, it automatically transfers the request to a human agent for resolution.

Benefit: Cutting down on resolution time produces operational cost savings. Smart AI bots reduce the call length from an average talk time (ATT) of three minutes to two minutes, based on our prototype. The smart AI bot enhances customer experience by offering round-the-clock connection with the contact center. It can help to reduce the number of full time equivalents (FTEs) for work order creation requests per month.

ML framework comparison

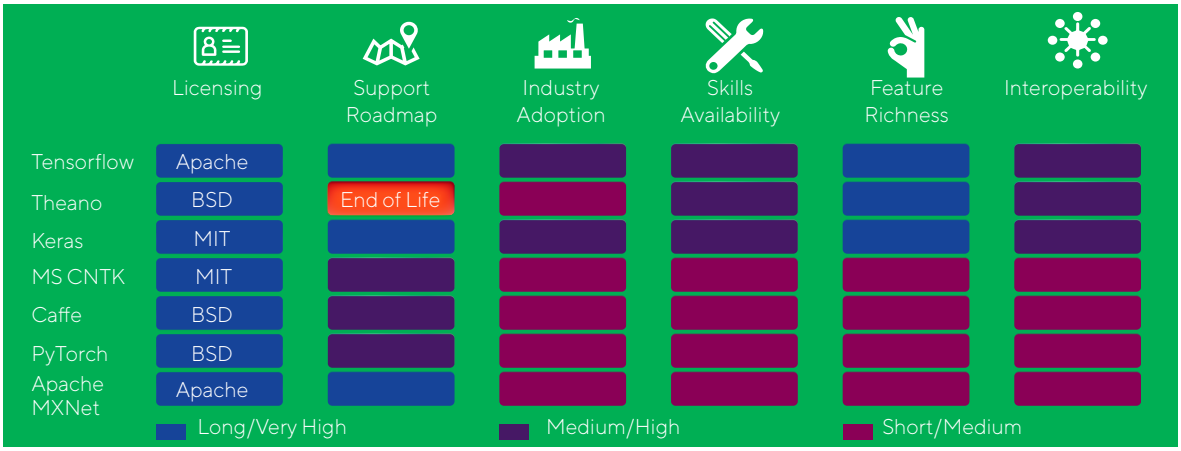

Figure 4

AI can leverage information to move assets to a safe-zone to prevent infection, and it can be used to take vulnerability scans.

Image de-duplication

Image/video duplication is a rapidly increasing challenge, as media content volume is already huge and growing rapidly across most organizations. With visual media, there can be multiple iterations of images varying only by format, size, quality, compression and image transformation. Media de-duplication can result in poor media search performance, excessive storage usage and potential fraud.

Solution: Computer vision with information retrieval technology can tackle image de-duplication. Here's how it works: Key features are extracted from media files that are robust to noise and transformations, and are indexed. A lightweight scalable search service is developed that identifies duplicates through a similarity score. The AI tool is integrated into the de-duplication service which is responsible for preventing duplicate media files from being uploaded to cloud storage. The solution comprises a two-engine model – one for querying and the other for indexing. The querying engine is used to check a file for duplication in the existing media corpus and the indexing engine is used to transform the media corpus into an N-dimensional feature space, and index the files for efficient querying.

Benefit: Automating the de-duplication of media files helps enterprises perform efficient querying – with an accuracy rate above 90% from our initial test results in identifying duplicate media files. The AI tool's elimination of manual effort will help enterprises save costs and efficiently process large amounts of content.

Vehicle damage assessment

Vehicle damage assessment is a tedious manual process, where an insurer must conduct a visual damage inspection to determine allocations for repair costs. The estimation process is labor-intensive and time-consuming as it requires domain experts to evaluate the damage. The inspection is done by a service center, which can be costly. There is a clear need to expedite vehicle damage assessment to speed up claim settlements and trim insurers' costs.

Solution: Use AI and ML technology like computer vision techniques and CNNs to automate vehicle damage detection by providing critical information to insurers about the magnitude of the damage.

Insurers can upload pictures of the damaged car, with variations, into an AI application that can access large data sets of various auto damage images. The uploaded images are then classified using the Keras and Theano framework, comparing the insurer's photos with thousands of other anonymized crash photos, thus enabling the ML system to detect and classify vehicle damage instantly.

Benefit: The AI application will help insurers to automate vehicle damage detection. This will save time in the event of an accident as well as yield cost savings by enabling insurers to process claims at faster rates.

Cybersecurity regulator

Cybersecurity Ventures has predicted that global ransomware damage costs will hit $11.5 billion by 2019.[13] One global ransomware attack was estimated to have spread across 150-plus countries and have infected more than 200,000 computers.[14] Such threats are unpredictable and uncontrollable. Yet enterprises need to fix these cyberattacks to protect customer data. AI and ML can augment security systems to better identify malicious activity and prevent cybercrime.

Solution: Traditional signature-based protection is to be replaced or enhanced by ML and AI, These systems can rapidly identify abnormal asset behavior, signs of infection, abnormal traffic and anomalies – all of which can then be automatically placed in a quarantined group and removed from network access. The next-generation AI firewall system constantly trains, and it identifies patterns of behavior and discriminates between normal and abnormal activities. Additionally, AI can leverage information to move assets to a safe-zone to prevent infection, and it can be used to take vulnerability scans.

Benefit: Applying ML and complex AI systems in cybersecurity is a reactive approach to analysis, notification and/or escalation of potential problems. An ML-based approach can help predict threats before potential data leaks occur.

Breast cancer diagnostics

Breast cancer is a leading cause of mortality for women in the U.S. The Breast Cancer Organization has estimated that 266,120 new cases of invasive breast cancer will be diagnosed in women in the U.S., and 63,960 new cases of non-invasive (in situ) breast cancer.[15] Today, breast cancer diagnosis is still a lengthy process, involving multiple tests and other processes. Hence there is a definite need for ML/AI to partially automate breast cancer diagnosis.

Solution: The AI system uses ML and computer vision technique to analyze images of tissues and predict which lesions are most likely to become cancerous. This approach has huge potential to reduce the specialist's workload in a typical pathology lab and to improve the incidence of early breast cancer detection.

Benefit: Early detection of cancer greatly increases the chances of successful treatment. The ML system ultimately could provide a more targeted approach in the breast cancer area and could also help women make more informed treatment decisions.

Advanced driver assistance systems

The advanced driver assistance systems (ADAS) component market will reach $67.43 billion by 2025.[16] ADAS are on the rise because of the combination of technological innovation and the strong need to improve vehicle automation and support self-driven vehicles. ADAS promise to make automobiles more intelligent, safer and more convenient.

Solution: Our Driver Eye is a key module of the driverless car. Its capacities are based on analyzing the real-time video input from mounted vision sensors on vehicles. They include: detecting traffic signals, recognizing roadside signboards (mandatory and cautionary), detecting lane markers, and identifying stop lines and zebra crossings. The Driver Eye prototype relies on one or more camera sensors mounted on the vehicle to provide vision-based driver assistance for an enriched driving experience and enhanced safety on the road. Our Driver Eye platform is built as an exclusive ML subsystem within an autonomous vehicle.

Benefit: Our solution functions as a vehicle's extended vision allowing the driver (or car) to take decisions based on traffic conditions. Cameras are low-cost vision sensors that can serve as the basis for multiple driver-assistance applications. Our vision is to bring the Driver Eye technology to the masses to enrich the driving experience, to enhance safety on the road and specifically to address the low-cost automotive safety application market.

The way forward under the influence of AI

As the future unfurls, machines will learn to do more and more. Choosing the right problem and solving it using ML can tip the scales and create value for any business. Organizations will gravitate to systems of intelligence to take advantage of ML's power to ensure market relevance by creating superior solutions that either anticipate or more effectively serve customers' needs, wants and desires.

Yet, key challenges remain. They include:

- **Data security:** Enterprises must provide access to their data sets to implement ML correctly and efficiently, but for security reasons the data remains tightly secured and inaccessible.
- **Infrastructure:** This is a must to test the various ML models with different tools. A lack of infrastructure may undermine frequent tests in developing the best possible outcomes.
- **Talent pool:** Creating an ML model requires domain expertise and technologists who can solve the problem – but there are few experts to be found in the industry.

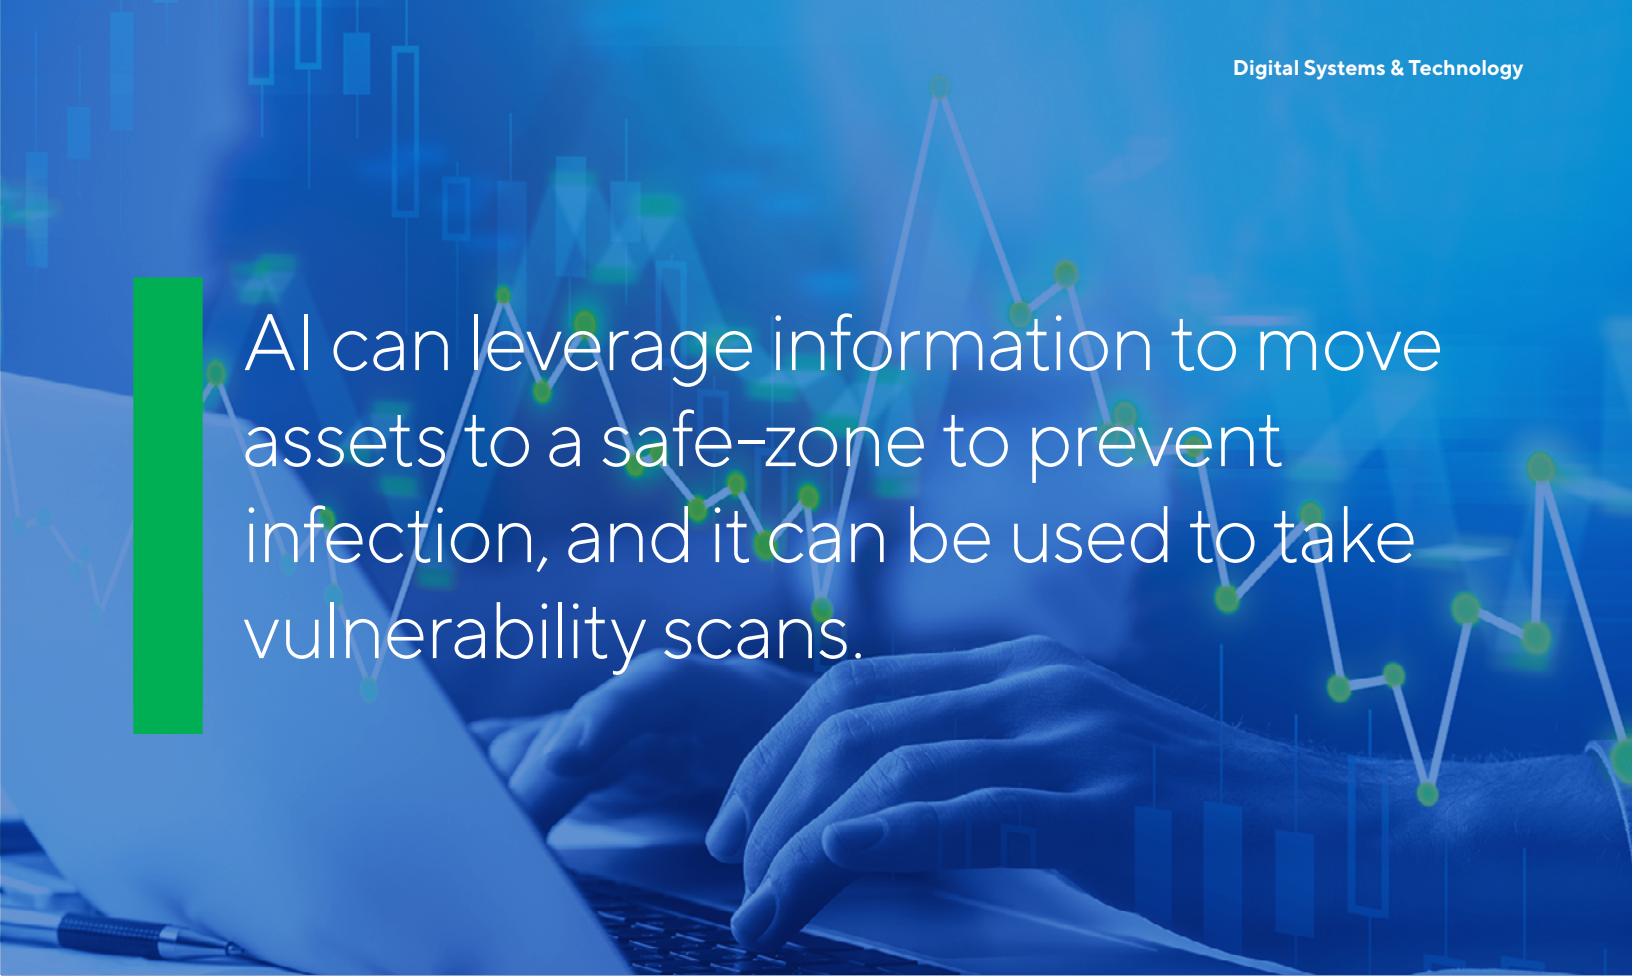

AI can leverage information to move assets to a safe-zone to prevent infection, and it can be used to take vulnerability scans.

- **Business model:** Implementing ML requires organizations to change their mindset, as well as implement a proper strategy for managing the change. There are few tools, frameworks and best practices to help guide business leaders forward.

If applied right, ML has boundless potential to improve business practices, increase return on investment (ROI), and satisfy customers with future-ready platforms and solutions. As a result, we recommend the following:

- **Start from business decisions.** These should be the key strategic decisions that affect the business, and the related metrics that need improvement.
- **Identify an appropriate business problem.** Then solve it with the right ML framework and packages.
- **Build business-problem-specific data sets.** The data set should be labeled (by structured or unstructured data) and aligned with the problem being solved.
- **Pilot with parallel runs.** Conduct and compare ML solution outcomes with those enabled by human decision-makers. Assess and iteratively improve the performance/accuracy of the ML solution.
- **Scale the ML solution.** Build out the solution with necessary hardware/software architecture and with associated infrastructure.
- **Manage change.** Institute a broad change management program to shift the internal decision-making mindset.

Endnotes

1 Press Release, IDC, March 22, 2018, https://www.idc.com/getdoc.jsp?containerId=prUS43662418.

2 "TensorFlow," https://www.tensorflow.org/.

3 "Theano," http://deeplearning.net/software/theano/.

4 "Keras," https://keras.io/.

5 "The Microsoft Cognitive Toolkit," https://www.microsoft.com/en-us/cognitive-toolkit/.

6 "Caffe," http://caffe.berkeleyvision.org/.

7 "Torch," http://torch.ch/.

8 "Apache SINGA," https://singa.incubator.apache.org/en/index.html.

9 "Developer Survey Results 2017," stackoverflow, https://insights.stackoverflow.com/survey/2017#technologies-and-occupations.

10 Kasey Panetta, "Gartner Top Strategic Predictions for 2018 and Beyond," Gartner, October 3, 2017, https://www.gartner.com/smarterwithgartner/gartner-top-strategic-predictions-for-2018-and-beyond/.

11 Jitesh Ubrani & Adam Wright, "Worldwide Quarterly Smart Home Device Tracker," IDC, March 29, 2018, https://www.idc.com/getdoc.jsp?containerId=prUS43701518.

12 Bret Kinsella, "Smart Speaker Consumer Adoption Report 2018," Voicebot, March 7, 2018, https://voicebot.ai/2018/03/07/new-voicebot-report-says-nearly-20-u-s-adults-smart-speakers/.

13 Steve Morgan, "Global Ransomware Damage Costs Predicted To Hit $11.5 Billion By 2019," November 14, 2017, https://cybersecurityventures.com/ransomware-damage-report-2017-part-2/.

14 "WannaCry ransomware attack," Wikipedia, https://en.wikipedia.org/wiki/WannaCry_ransomware_attack.

15 "U.S. Breast Cancer Statistics," https://www.breastcancer.org/symptoms/understand_bc/statistics.

16 "ADAS Market Size Worth $67.43 Billion By 2025," Grand View Research, February 2018, https://www.grandviewresearch.com/press-release/global-advanced-driver-assistance-systems-adas-market.

About the authors

Mahesh Balaji

Head, Cognitive Computing & Data Sciences Lab, Cognizant Global Technology Office

Mahesh Balaji heads Cognitive Computing and Data Sciences Lab within Cognizant's Global Technology Office. In this role, he is responsible for driving fundamental advancements in next-generation workloads, enabled by machine learning and deep learning. Mahesh is also responsible for nurturing technology thought leadership. He has 20 years of IT industry experience, 14 years of which with Cognizant. Mahesh's top priorities are to stimulate innovation in cognitive areas of information technology and to bring those cognitive innovations into the marketplace quickly; to apply these innovations to help our clients succeed; and to identify and nurture new and future areas for investment and growth. He has an MS in software engineering from the University of Houston, and a bachelor's degree in engineering from Madras University, India. Mahesh has multiple publications and patents to his credit and about 30-plus industry accreditations at strategic and tactical levels. He can be reached at Mahesh.Balaji@cognizant.com | www.linkedin.com/in/maheshbalaji/.

Ramakrishnan Viswanathan

Manager of Business Development, Cognitive Computing & Data Sciences Lab, Cognizant Global Technology Office

Ramakrishnan Viswanathan is a Manager of Business Development in the Cognitive Computing and Data Sciences Lab within Cognizant's Global Technology Office. In this role, he focuses on AI/ML services and emerging technologies. In addition to his time spent in GTO and Cognizant's Application Value Management Practice, Ram has over 14 years of experience in presales, strategic partnerships, business development and client relations. He works closely with business leaders to develop technology accelerators for the future. Ram has an executive program degree in sales and marketing (EPSM) from Indian Institute of Management, Calcutta (IIMC). He has also published compelling points of view on AI and cognitive technology. Ram can be reached at Ramakrishnan.Viswanathan3@cognizant.com | www.linkedin.com/in/ramakrishnanviswanathan.

Acknowledgments

The authors would like to thank Aan Singh Chauhan, Cognizant SVP & Chief Technology Officer within the company's Global Technology Office, and Alan Alper, Cognizant's Editorial Director, for their support and guidance on this paper.

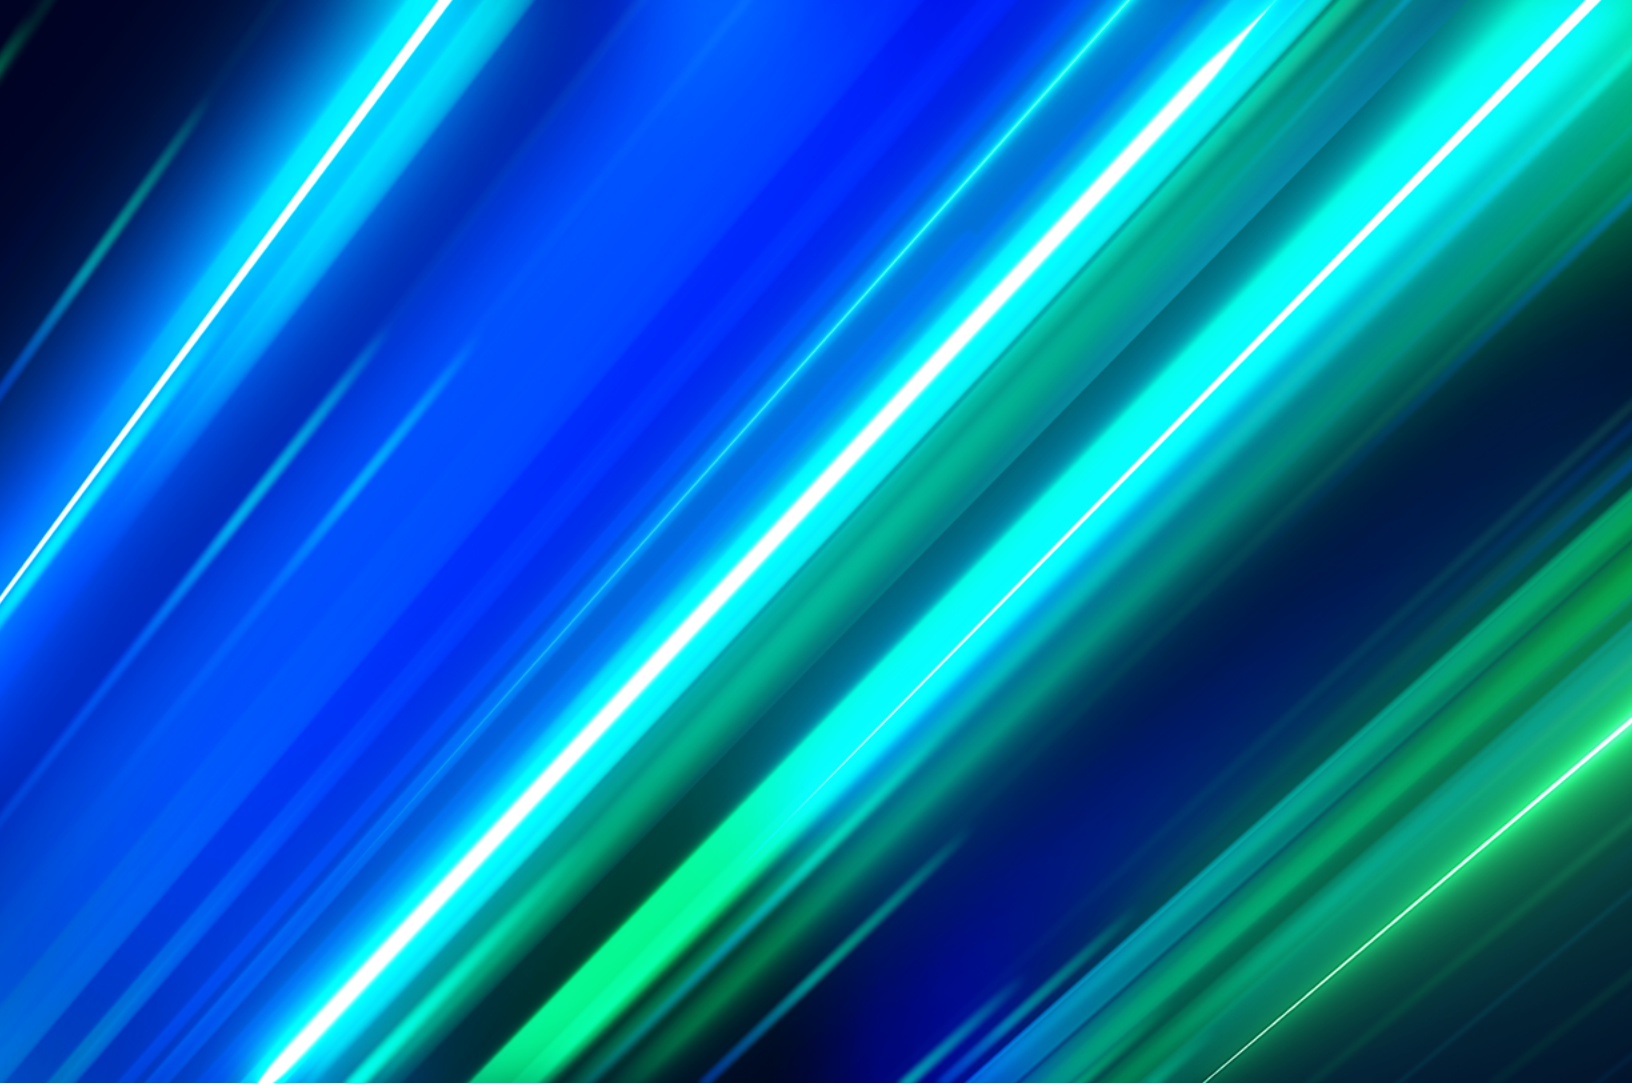

About Cognizant's Cognitive Computing & Data Sciences Lab

The Global Technology Office (GTO) is the core technology organization of Cognizant, with a mission to power and accelerate our capabilities to harness transformative technologies that enable our people, customers and processes to navigate the shift in the work ahead. As part of GTO, the Cognitive Computing & Data Sciences (CDS) Lab's vision is to explore emerging and cognitive technology areas in artificial intelligence, machine learning, natural language processing, voice/speech recognition and computer vision. The CDS lab builds innovative, industry-specific cognitive platforms and solutions for digital business transformation.

About Cognizant

Cognizant (Nasdaq-100: CTSH) is one of the world's leading professional services companies, transforming clients' business, operating and technology models for the digital era. Our unique industry-based, consultative approach helps clients envision, build and run more innovative and efficient businesses. Headquartered in the U.S., Cognizant is ranked 195 on the Fortune 500 and is consistently listed among the most admired companies in the world. Learn how Cognizant helps clients lead with digital at www.cognizant.com or follow us @Cognizant.

Cognizant

World Headquarters

500 Frank W. Burr Blvd.
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277

European Headquarters

1 Kingdom Street
Paddington Central
London W2 6BD England
Phone: +44 (0) 20 7297 7600
Fax: +44 (0) 20 7121 0102

India Operations Headquarters

#5/535 Old Mahabalipuram Road
Okkiyam Pettai, Thoraipakkam
Chennai, 600 096 India
Phone: +91 (0) 44 4209 6000
Fax: +91 (0) 44 4209 6060

TL Codex 3697.02