ARTIFICIAL INTELLIGENCE OF THINGS

Secure connectivity beyond today's IoT

CONTENTS

ARTIFICIAL INTELLIGENCE: BEYOND THE HYPE

If there is one true indicator to measure the disruptiveness of a **new technology**, it's certainly the public outpouring of fear and suspicion. If we use societal angst as a measure, the current renaissance of **artificial intelligence** (AI) is a good candidate for groundbreaking technological disruption.

AI will change life as we know it, as Elon Musk, Bill Gates, Stephen Hawking and other great minds have told us. The widespread anxiety about the **harmful consequences of AI applications** is not an unparalleled reaction to technological change but rather an expression of the societal unease that commonly precedes the changes associated with new technologies and the vast potential that comes with them.

We'd like to use the current debate as an opportunity to present an industry view on AI and its subsets machine learning (ML) and deep learning (DL). We examine how the shift of high performance processing from the cloud to the edge of networks has enabled the Internet of Things (IoT) to thrive, and how this paradigm shift has laid the foundations for AI to unfold its true potential.

And we're looking beyond today's IoT, toward a future where smart connected devices not only talk with each other but where they use artificial intelligence to interact with each other on our behalf. This new global fabric of artificially intelligent things one day will be known as the AIoT, **the Artificial Intelligence of Things**.

In the process, we will discuss the tremendous potential of these technologies, but also look at their constraints, and reflect on the real threats certain AI applications can pose to IoT security, and how to counter them.

WHERE WE STAND TODAY

As a discipline of mathematics – and, to a certain extent, of philosophy – AI lived in the shadows for more than six decades before public interest suddenly soared in the present period. One reason for the current publicity is that, for a long time, considerations on AI applications were purely theoretical, or at least science fiction. For artificial intelligence use cases to become real in the present IoT environment, three conditions had to be fulfilled:

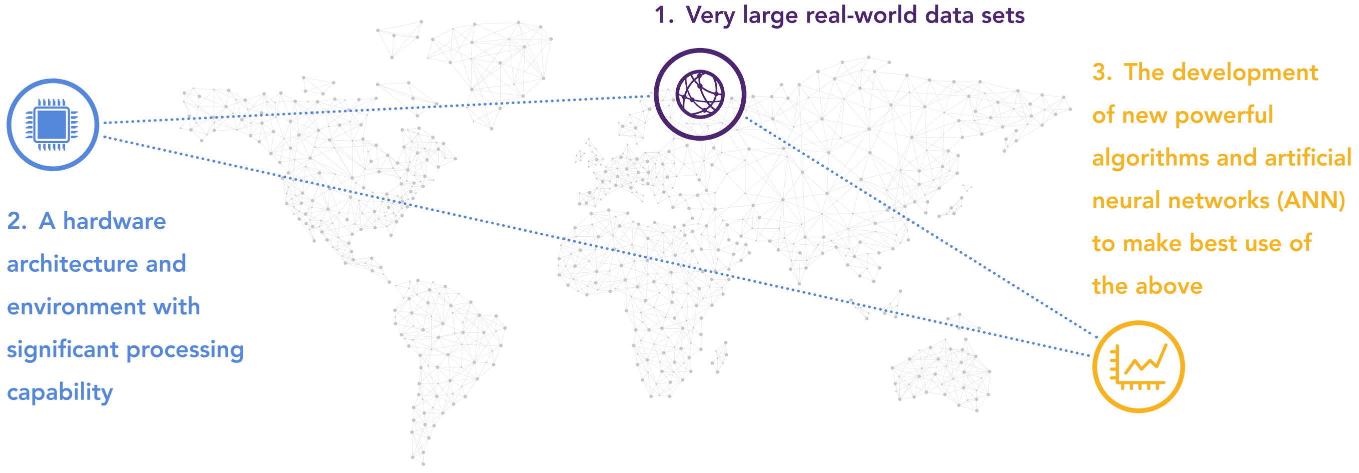

1. Very large real-world data sets

2. A hardware architecture and environment with significant processing capability

3. The development of new powerful algorithms and artificial neural networks (ANN) to make best use of the above

It is apparent that the two latter requirements depend on each other, and that the breakthroughs in deep neural nets could not have occurred without a significant increase in processing power. As for the input: large data sets of every quality – vision, audio, and environmental data – are being generated by an increasing number of embedded IoT devices. Today, the flow of data grows exponentially.

In fact, annual data generation is expected to reach 44 zettabytes (one zettabyte is 1 billion terabytes) by 2020, which translates to a compound annual growth rate (CAGR) of 141% over five years. Just five years after that, it could reach **180 zettabytes**.

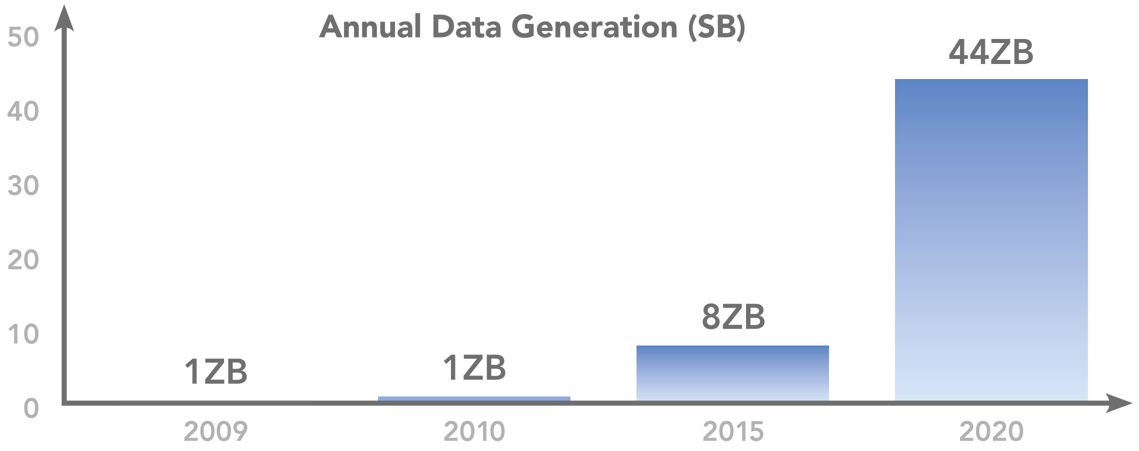

Massive growth in the amount of unstructured data being created by the increasingly connected devices, machines, and systems globally.

Starting around 2015, when multicore application processors and graphics processing units (GPUs) became widely available, we've also commanded the tools to cope with these amounts of data. Parallel processing became a much faster, cheaper, and more powerful business. Add fast, abundant storage and more powerful algorithms to sort and structure that data, and suddenly there is an environment in which AI can prosper and thrive.

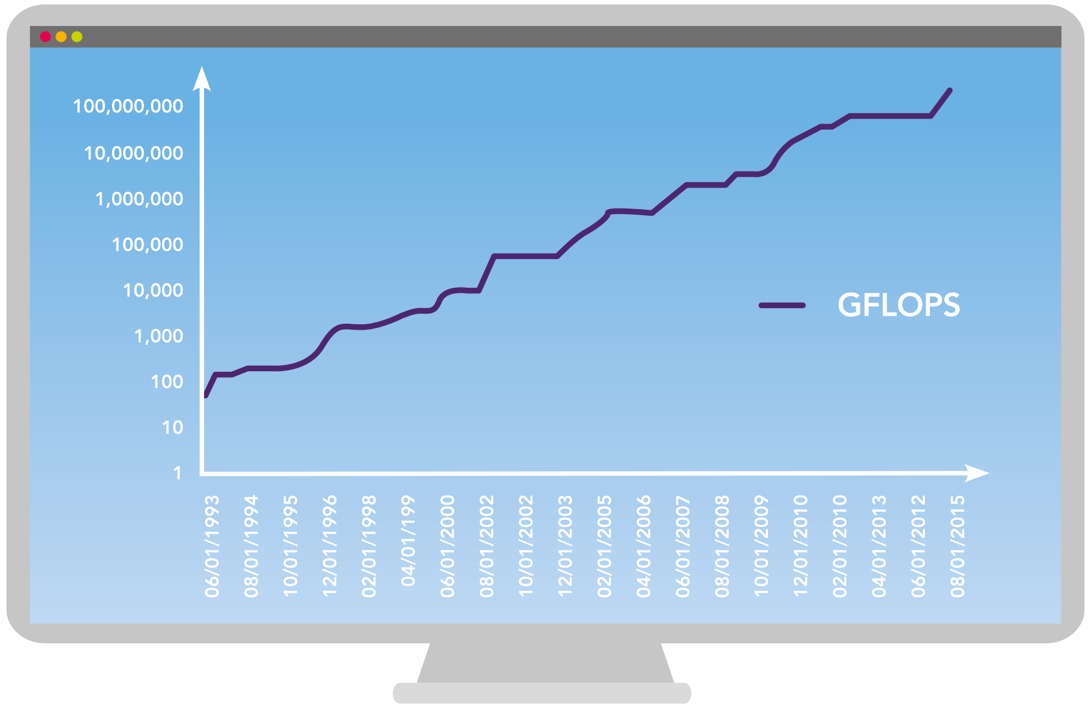

Exhibit 2: Raw compute performance of global supercomputers, measured in GFLOPs, has increased exponentially since 1993. Rpeak GFLOPS #1 ranked global supercomputers on the Top 500 list.

Still, in 2016 the EE Times senior researcher at Baidu's Silicon Valley AI Lab, Greg Diamos, commented on what AI lacked to live up to its true potential:

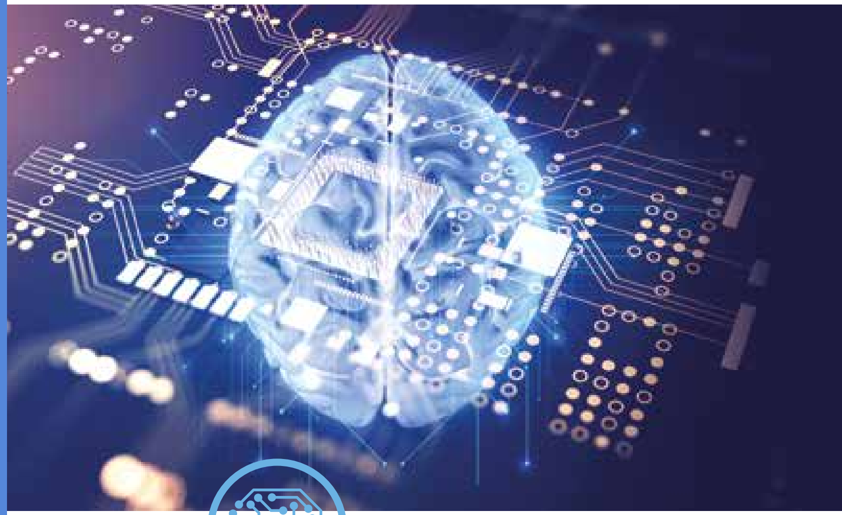

> "Today the job of training machine learning models is limited by compute, if we had faster processors we'd run bigger models… in practice we train on a reasonable subset of data that can finish in a matter of months. […] We could use improvements of several orders of magnitude – 100x or greater."
>
> - *Greg Diamos*

In 2018, neural network-trained AI voice recognition software is an integral part of a variety of consumer and industrial applications. Computing capability is increasing by roughly a factor of 10 each year, mainly driven by new classes of custom hardware and processor architecture. This computing boom is a key component in AI progress is helping make AI mainstream in the future.

WHAT IS AI?

By classic definition, **artificial intelligence** is a rather unspectacular affair. In his groundbreaking 1976 paper Artificial Intelligence: A Personal View, British neuroscientist and AI pioneer David Marr states: The goal of AI is to identify and to **solve useful information processing problems** and to give an abstract account of how to solve it, which is called a method.

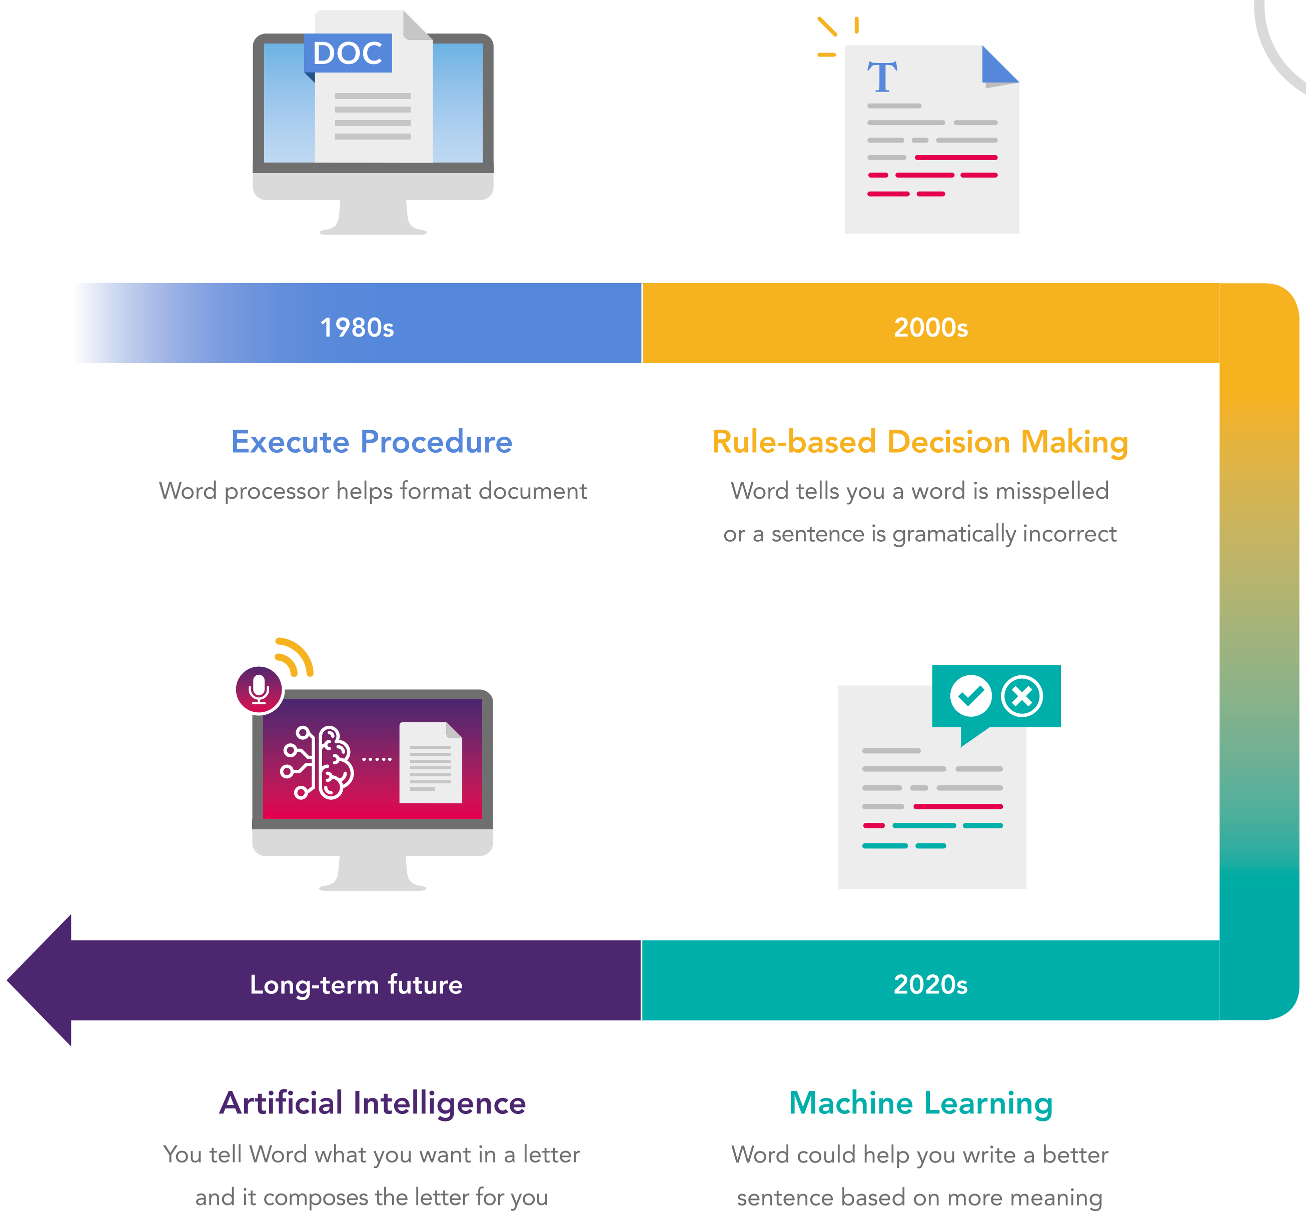

1980s

2000s

Execute Procedure

Word processor helps format document

Rule-based Decision Making

Word tells you a word is misspelled or a sentence is gramatically incorrect

Long-term future

2020s

Artificial Intelligence

You tell Word what you want in a letter and it composes the letter for you

Machine Learning

Word could help you write a better sentence based on more meaning

It's true that AI computing systems are vaguely inspired by the biological neural networks that constitute brains. However, it is a popular myth that AI looks to re-engineer the function of the human brain to enable machines to solve problems the way humans would. The neural network rather is a framework for many different machine learning algorithms to work together and process complex data inputs. The main deviation from biology is that ANN are focusing to perform specific tasks rather than universal problem solving and planning capabilities.

Compared with the scientific effort, today's industry's approach on AI is even more pragmatic. Rather than trying to achieve a replica of the human mind, current AI development uses human reasoning as a guide to provide better services or create better products. But how does that work? Let's have a look at the current approaches.

WHAT IS ML?

As a subset of artificial intelligence, machine learning uses statistical techniques to give computers the ability to learn without being explicitly programmed. In its most crude approach, machine learning uses algorithms to analyze data, and then make a prediction based on its interpretation.

The key element is that the machine is trained to learn from the data so that it can perform its given job. To achieve this, machine learning applies pattern recognition and computational learning theory, including probabilistic techniques, Kernel methods and Bayesian probabilities, which have grown from a specialist niche to become mainstream in current ML approaches. Rather than following static program instructions, ML algorithms operate by building a model from an example training set of input in order to make data-driven predictions expressed as outputs.

Computer vision is the most active and popular application field where ML is applied. It is the extraction of high-dimensional data from the real world to produce numerical or symbolic information – ultimately in the forms of decisions. However, until very recently, an extensive portion of hand coding was involved for machines to develop advanced pattern recognition skills. Human operators had to extract edges to define where an object begins and where it ends, apply noise removing filters or add geometrical information, e.g., on the depth of a given object. It turned out that even with advanced machine learning training software, it's not a trivial task for a machine to make real sense of a digital reproduction of its environment. This is where deep learning comes into play.

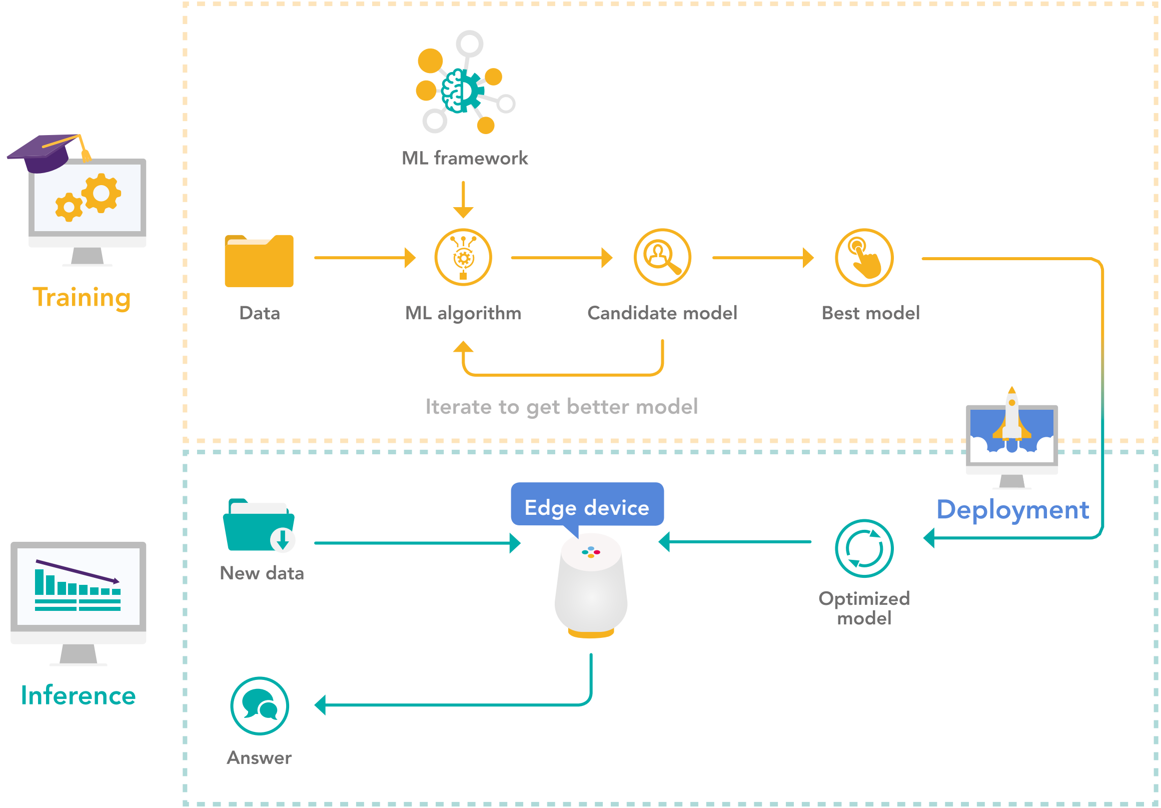

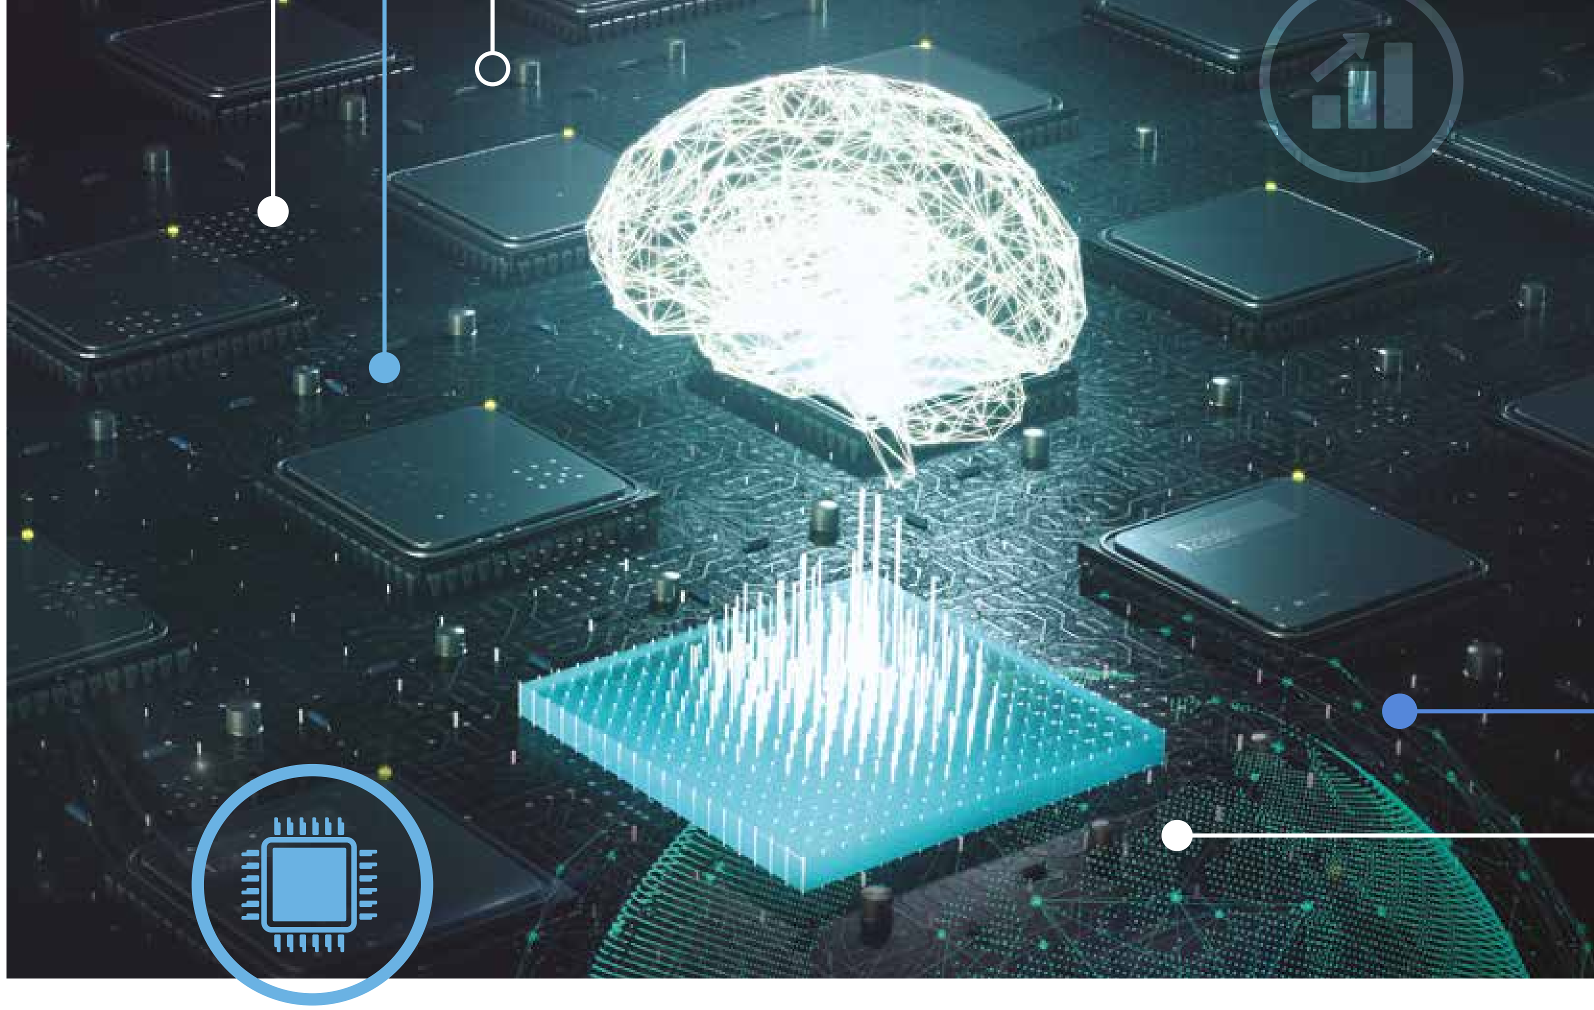

WHAT IS DL?

AIt's a decades' old idea that software can simulate a biological neocortex's array of neurons in an artificial "neural network." Deep learning algorithms attempt exactly that – to mimic the multilayer structure and functionality of the human neuronal network. In a real sense, a deep learning algorithm learns to recognize patterns in digital representations of sounds, images, and other data. But how?

With the current improvements in algorithms and increasing processing capacity, we can now model more layers of virtual neurons than ever before, and thus run models in much greater depth and complexity. For a long time, it was not feasible to use Bayesian techniques because, in order to calculate the evidence, it was necessary to perform probabilistic integrations by hand. Today, Bayesian deep learning is used in multilayer neural nets to tackle complex learning problems.

However, what we can do today still mostly falls into the concept of "narrow" or "weak AI" – technologies able to perform specific tasks as well as, or better, than humans can. For instance, AI technologies for image classification or face recognition perform certain facets of human intelligence, yet not the full spectrum, or even a combination of several human capabilities. A machine capable of performing a multitude of complex tasks, one that exhibits behavior at least as skillful and flexible as a human being, would be considered "strong AI." While the experts are divided over the question whether strong AI can ever be achieved, it doesn't stop them from trying.

SINCE 2013 INVESTMENT GROWTH IN AI HAS TRIPLED

The current need for more efficient systems for solving mathematical and computational problems becomes crucial, as the volume of relevant data is exponentially increasing. Consequently, the key players in the IT industry have focused on developing ICs and applications with the objective to position **AI at the core of the next-generation software technologies** in the market.

One solid indicator for the disruptive potential of the technology is the dimension of investment currently being channeled toward AI. According to McKinsey, $26 to 39 billion were invested in AI in 2016, most of it by tech giants such as Google™ and Baidu™. Since 2013, the external investment growth in artificial intelligence has tripled. In China, the year-on-year increase of AI investments is 76%. Alibaba™, the Chinese high-tech company, will invest $15 billion within the next five years in order to build up a worldwide AI technology network.

Alibaba will invest **$15 billion** in **AI** in the next 5 years

In 2016 **$26** to **$39 billion** was invested in **AI**

Noticeably so, investors are becoming increasingly attracted to startups around AI. Investments in these businesses have increased from $2.67 billion in 2014 to $5.02 billion in 2016. From 2016 to 2017, the global number of AI start-ups have increased about 141% and since 2016, more than 1,100 new AI companies raised equity funding.

Governments around the globe have started to see the relevance of the AI industry as well. The European Commission wants to put the EU in the top position of AI and declared to provide $1.7 billion for research and development of these technologies, a figure dwarfed by China's National Integrated Circuitry Industry Investment Fund that is targeting the development of AI technology with $47 billion. Several other countries have created national AI plans. The U.S. implemented a comprehensive AI research and development plan in May 2016, the UK launched a plan to improve access to data, AI skills, and AI research, and Canada announced a $125 million Pan-Canadian Artificial Intelligence Strategy.

$1.7 billion
for research and development

$47 billion
for development in AI technology

$125 million
Artificial Intelligence Strategy

ECONOMIC PROSPECTS ARE BRIGHT

Consequently, economic forecasters are in favor of the **emerging AI markets**. Research by PWC shows that the global GDP could be up to 14% higher in 2030 as a result of AI, which would be equivalent to an additional **$15.7 trillion**. "Of this, $6.6 trillion is likely to come from increased productivity and $9.1 trillion is likely to come from consumption-side effects."

Global GDP could be **14% higher** in 2030 as a result of **AI**

This would be the equivalent of **$15.7 trillion**

Looking at different countries, the greatest economic gains from AI is expected for China (26% boost to GDP in 2030) and North America (14.5% boost), equivalent to a total of $10.7 trillion, with China taking home $7 trillion of that total, dwarfing North America's $3.7 trillion in gains, and accounting for almost 70% of the global impact.

The effect of AI on the regional economy in Northern Europe is less, but still expected to be 9.9% of the GDP, equivalent to $1.8 trillion. In the UK, expectations are that AI could add an additional $814 billion (£ 630 billion) to the economy by 2035, increasing the annual growth rate of GVA from 2.5 to 3.9%.

Given the dependence on powerful computing hardware, it is no surprise that the strongest effects on the economy is expected in the IC sector. In the semiconductor industry alone, by 2025 AI-enabling components will have reached a global sales share of $60 billion. One of the fields with huge potential is the healthcare market, where AI can contribute to a growth of almost 40 % in the period from 2017 to 2027 and which is expected to reach $50 billion by 2027.

AI in the **healthcare market** will contribute to a **growth of almost 40%** by 2027

AI across industries in Germany could increase its **GDP 4%** by **2020**.

AI has grown by more than a 26% CQGR in the last two years across a diverse set of industries. From a strategic viewpoint, its biggest potential is seen in its complementary nature to the IoT: while the IoT component ensures a continuous supply of relevant data, the AI capability can serve as the system's inference engine, interpreting the data generated by the endpoints and driving their functions.

Merging both IoT and AI together in an integrated technology portfolio creates a powerful new platform for digital business value.

ADVANCING AI TO THE EDGE 1: HIGH-PERFORMANCE PROCESSING TO DO THE JOB

As we have seen, for AI to unfold its massive potential relies heavily on adequate hardware. Machine learning in particular, requires enormous **processing and storage capacity**. A training cycle for one of Baidu's speech recognition models for instance, requires not only four terabytes of training data, but also 20 exaflops of compute — that's equivalent to 20,000 quadrillion math operations per second — across the entire training cycle. Given its hunger for powerful hardware, it is no wonder that AI today is still mostly confined to data centers.

Uncoupling AI from the data centers and advancing it to the endpoints of the IoT will allow us to tap its full potential. This is exactly what NXP is about.

Let's take a closer look at the requirements. Today's IoT ecosystem has seen its very own disruption, namely the shift of data processing from the center of connected systems to the edge. Edge processing has taken the control of computing applications, data, and services away from some central nodes (the "core") to the periphery of the Internet. This is where IoT devices make contact with the physical world and where data comes in via various sensors (e.g. vision, voice, environmental). Processing this data at the edge, significantly decreases data volumes to be moved, thereby increasing privacy, reducing latency, and improving quality of service.

No longer relying on a central core also means the removal of a major bottleneck and potential single point of failure. Edge processing is based on distributed resources that may not be continuously connected to a network in such applications as autonomous vehicles, implanted medical devices, fields of highly distributed sensors, and a variety of mobile devices. To make use of AI in this challenging environment, an agile application that can retain learning and apply it quickly to new data is necessary. This capability is called inference: taking smaller chunks of real-world data and processing it according to training the program has done.

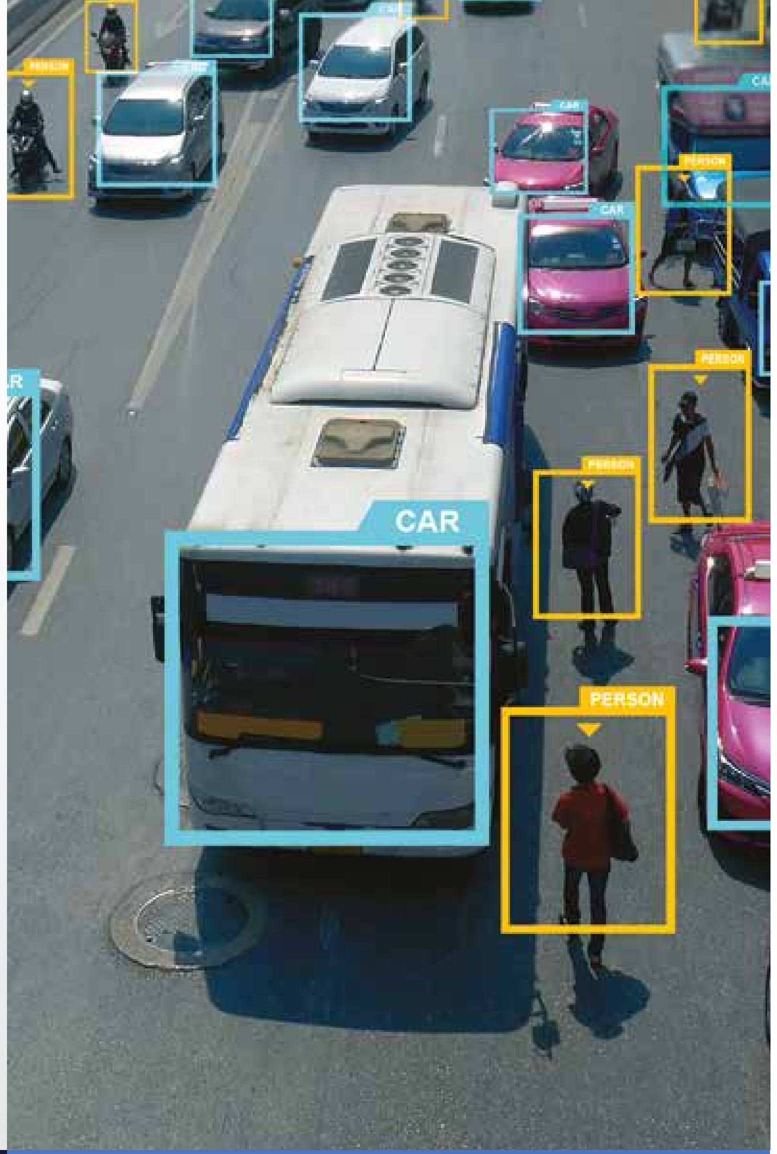

For inference to work in edge environments, processing architecture and hardware are required that are optimized and come with certain requirements on processing capacity, energy efficiency, security, and connectivity. NXP has established leadership in machine learning at the edge – particularly for the inferencing tasks – by developing an advanced portfolio of ICs that address the challenges of modern edge environments. In fact, we have been ranked as one of the world's top three artificial intelligence chipset companies.

For AI applications at the edge, the key design objective is to balance system costs with the end user experience. For example: food recognition in an ML-based microwave AI oven can happen in 1-2 seconds. Whereas, stop sign recognition, a pedestrian crossing, or the detection the detection of a drowsy driver's eyes closing behind the wheel in a vehicle require a much faster processing speed. For businesses to enhance their AI application portfolio, scalable processors and software enablement are key as they allow the developer to deploy the ideal IC for a broad range of specific AI applications.

The NXP portfolio covers almost the entire MCU and applications processors portfolio that is used in modern AI applications:

Product Name	AI Applications
i.MX 8M family	for advanced audio, voice and video, voice control, voice assistance
i.MX 8X family	safety certifiable and efficient performance for auto, industrial and consumer
i.MX 8 family	for advanced graphics, imaging and performance for auto, industrial and consumer
Crossover i.MX RT	for audio subsystems, consumer and healthcare, home and building automation, industrial computing, motor control and power consumption
i.MX 6, 7 & 8	for the consumer market and smart home, IoT cloud integration
Kinetis® and LPC MCUs with low-power Arm® Cortex	cores merge exceptional performance efficiency, memory scalability and integrated security. The on-chip hardware acceleration for symmetric cryptography reduces (CPU) loading, simplifies implementation, reduces software overhead, and allows the system to perform more efficiently
Layerscape® Products	the QorIQ® communications processor portfolio offers unmatched depth and breadth. With the addition of our next-generation QorIQ Layerscape series processors built on Arm core technology, the portfolio extends performance to the smallest form factor—from power-constrained networking and industrial applications to new virtualized networks and embedded systems requiring an advanced datapath and network peripheral interfaces
S32 MCUs and microprocessor units	for automotive and industrial applications provide the best architecture for performance and power efficiency. They're designed to address current and future connectivity, security, and safety challenges. Automotive systems for driver assistance and driver replacement are a primary domain of AI deployment. The S32 family of ICs provide domain specific solutions for inference for sensor signal processing, such as Radar, Optical sensors as well as powertrain, vehicle dynamics and vehicle network connectivity solutions

QUANTUM COMPUTING: THE MAGIC BULLET TO CREATE A STRONG AI?

To speed computation, quantum computers tap directly into an unimaginably vast **fabric of reality**—the strange world of quantum mechanics. Rather than store information using bits represented by 0s or 1s, as conventional digital computers do, quantum computers use quantum bits (qubits), to encode information as 0s, 1s, or both at the same time.

This superposition of states—along with the other quantum mechanical phenomena of entanglement and tunneling—enables quantum computers to manipulate enormous combinations of states at once. Given AI's hunger for fast processing, is quantum computing paving the path for strong AI? The answer is yes and no.

The truth is, over the past decade, quantum computing has moved from the realm of speculation and debates about its reality into demonstration of small prototypes. Not only has the concept been proven, but computing machines built with tens of qubits are already in operation. However, it is unclear if this scaling can continue unabated for quantum computers to reach or surpass the capabilities of classical computers. Fundamental questions related to the scaling of quantum computers exist that prevent scientists from predicting with any certainty whether much larger computers can ever be built.

An essential characteristic of quantum computers is that they lend themselves only to certain types of computational problems. Not all arbitrary algorithms used by ordinary computers are suitable for quantum computers. To date, only a handful of algorithms that can run on quantum computers have been developed. Among these algorithms is the Grover algorithm that can tremendously speed up an unstructured search of a very large amount of data and, thus, has tremendous potential for AI. If quantum computers ever get widely deployed, it's expected that they will be complementary to ordinary computers and not their replacements.

There have also been speculations and claims that quantum computers can break traditional cryptographic systems, such as public key algorithms (RSA, Diffie-Hellman, Elliptic Curve), and symmetric ciphers such as triple-DES and AES. Putting aside serious doubts about this claim (which arise from the probabilistic nature of quantum computers,) even if this breaking of cryptos becomes a reality it will be well in the future – somewhere around the 2030's.

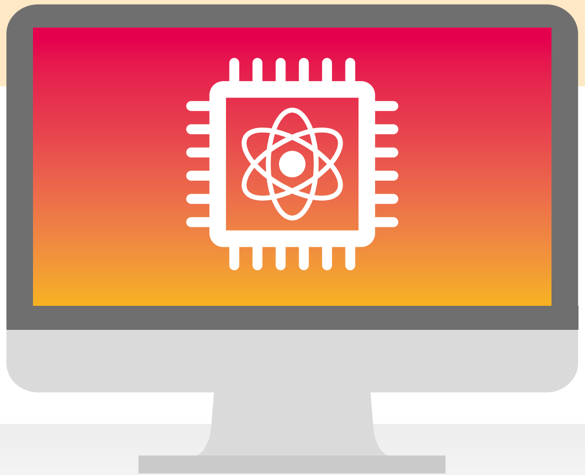

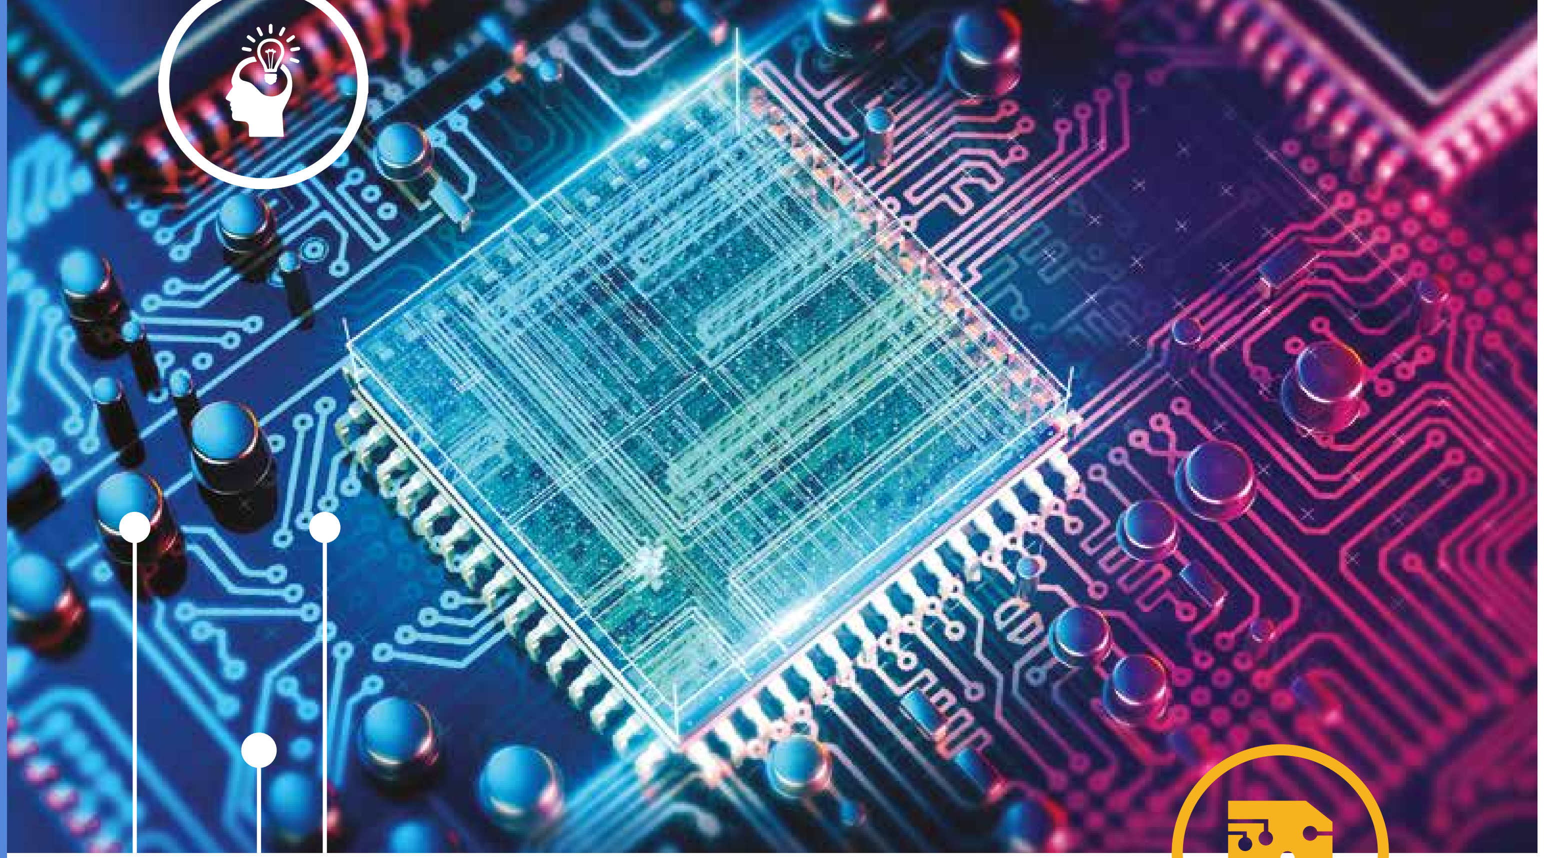

ADVANCING AI TO THE EDGE 2: DEDICATED MACHINE LEARNING ENVIRONMENT

To build **innovative AI applications** with cutting-edge capabilities, developers depend on a machine learning software environment that enables easy integration of dedicated functionalities into consumer electronics, industrial environments, vehicles, and other embedded applications, in general. But to roll out AI-based business models at a much broader scale and to make AI applications available to billions of end users across all verticals, the industry must first **overcome past limitations**.

This is why NXP has developed its machine learning hardware and software environment, which enables inference algorithms to run within the existing architecture. With its ML environment, NXP provides a variety of easy-to-implement ML functions on NXP's breadth of devices – from low-cost MCUs to breakthrough crossover i.MX RT processors and high-performance application processors. This ML environment provides turnkey enablement for choosing the optimum compute elements across the entire spectrum: from among Arm® Cortex M and A cores to high-performance GPUs and Digital Signal Processors (DSP) and custom accelerator architectures for advanced processing of data.

The NXP ML environment enables fast-growing machine learning use-cases in vision, voice, path-planning and anomaly detections, and the integration of platforms and tools for deploying machine learning models, including neural networks and classical machine learning algorithms, on those engines.

NXP Machine Learning environment enables:

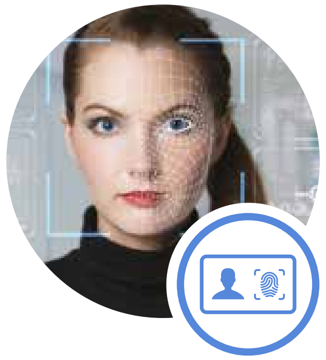

Vision

Voice

Path-Planning

Anomaly Detections

For customers integrating machine learning into their applications, the ML environment includes software tools that allow them to import their own pre-trained ML frameworks. By importing from their own open source software libraries of choice for high-performance numerical computations such as TensorFlow™ or Caffe models, designers can convert these models to optimized inference engines. Using the tools, a customer can train their models in the cloud, export these models, and deploy them inside NXP chips integrated into edge applications.

To support a broad range of customer needs, NXP has also created an expanding machine learning partner ecosystem to connect customers with complementary technology that can accelerate time-to-revenue with proven ML tools, inference engines, vertical applications, and design services.

And this is only the first step, as NXP is already working to integrate scalable artificial intelligence accelerators in its device that will boost machine learning performance by at least an order of magnitude. As a leader in the field of crypto acceleration technology, for NXP to do the same for AI functionality, this will further secure NXP's position as a leader in machine learning.

AI: IMPLICATIONS FOR IoT SECURITY

Every second, five new **malware variants are discovered**. Organizations across the globe are hit by one hundred previously unknown malware attacks every hour. And every day, one million new malicious files appear in the connected world.

With ever more devices and systems connected to the web, Cybercrime has become an increasing threat to our technological assets – and to the safety of our society as a whole.

At NXP, we built some of the most sophisticated secure devices in the world. We create countermeasures inside them to protect them against a broad range of logical and physical attacks, such as side-channel or template attacks. It's only a matter of time when hackers will rely on AI to extract secrets and critical information from secure systems, as it only enhances their "learning" capabilities. We must think about and check our defense mechanisms against these coming approaches.

Progress in artificial intelligence is closely related to the development of cyber threats. Machine learning proves to be a double-edged sword: While ML enables industry-grade malware detection programs to work more effectively, it will soon be used by the bad actors to enhance the offensive capabilities of their attacks. As a matter of fact, a group of researchers from the University of Amsterdam recently demonstrated how this can work. In their side-channel attack that leaked information out of the CPU's translation lookaside buffers (TLBs), the white-hat hackers used novel machine learning techniques to train their attack algorithm and bring it to a new level of performance. They are confident that machine learning techniques will improve the quality of future side-channel attacks.

In order to prevent the emergence of new, effective AI and ML techniques from changing the balance of power, we must focus on how to leverage artificial intelligence to improve system security and data privacy.

ML CAN ADD TO SYSTEM SECURITY

A good example of ML-based security is anomaly detection, where the system "detects" anomalous behavior or patterns in the data stream.

This process has been routinely applied to SPAM and malware detection in the past, but machine learning can be expanded to look for more subtle and complex anomalous behavior in a system. While monitoring and protecting from external threats is crucial for an effective system defense, few organizations are aware of inside threats. In a survey from Accenture in 2016, they found that two thirds of the surveyed organizations fell victim to data theft from inside the organization. In these instances, 91% reported that they did not have effective detection methods for identifying this type of threat. Machine learning can significantly aid in the development of effective, real-time profiling and anomaly detection capabilities, to detect and neutralize user-based threats from within the system.

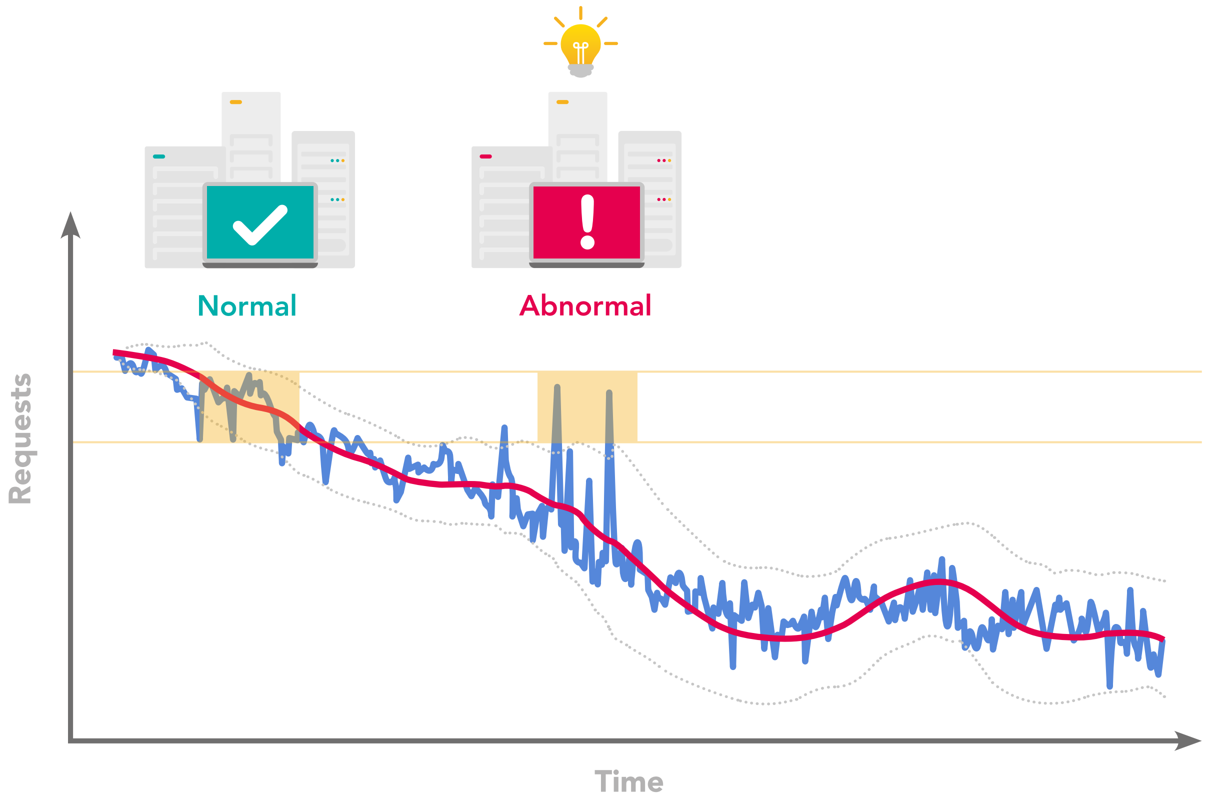

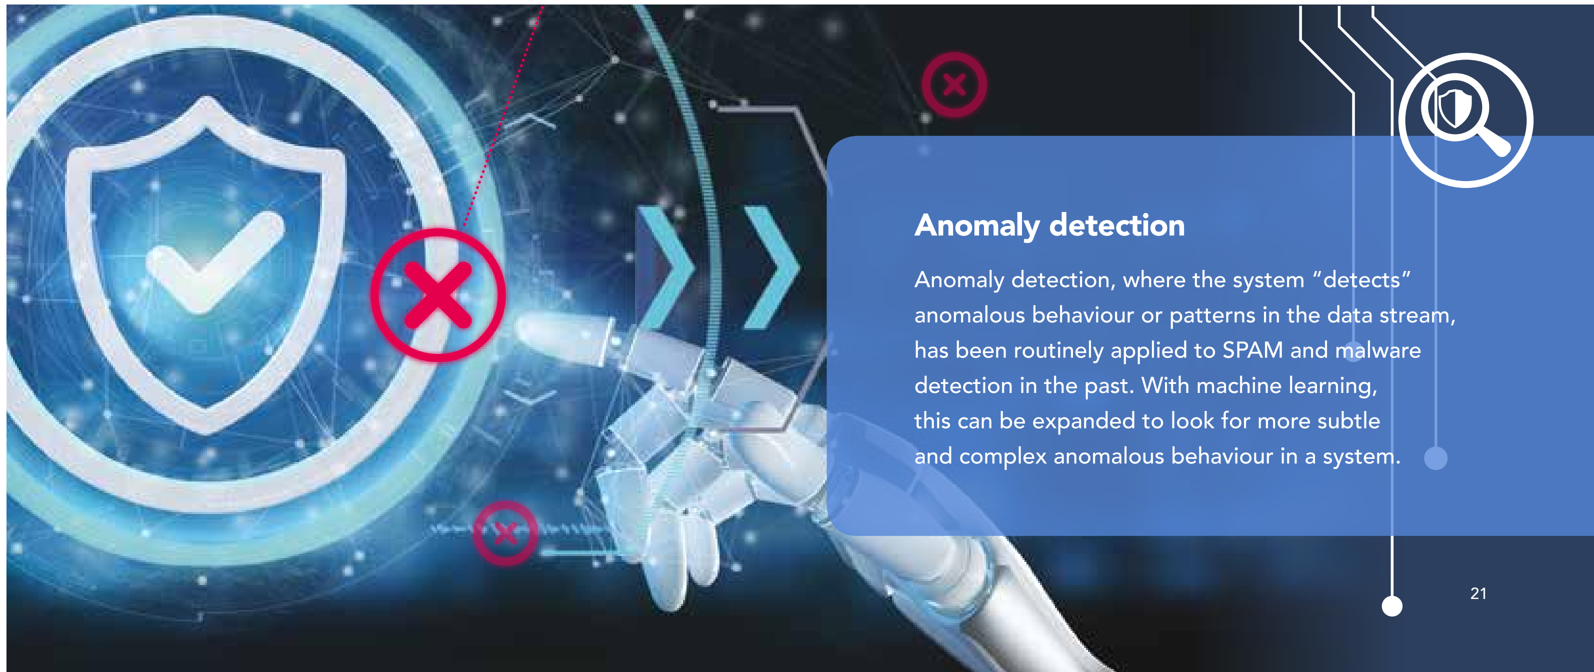

Anomaly detection

Anomaly detection, where the system "detects" anomalous behaviour or patterns in the data stream, has been routinely applied to SPAM and malware detection in the past. With machine learning, this can be expanded to look for more subtle and complex anomalous behaviour in a system.

Machine Learning can contribute to IoT Security – but machine learning itself must be secured.

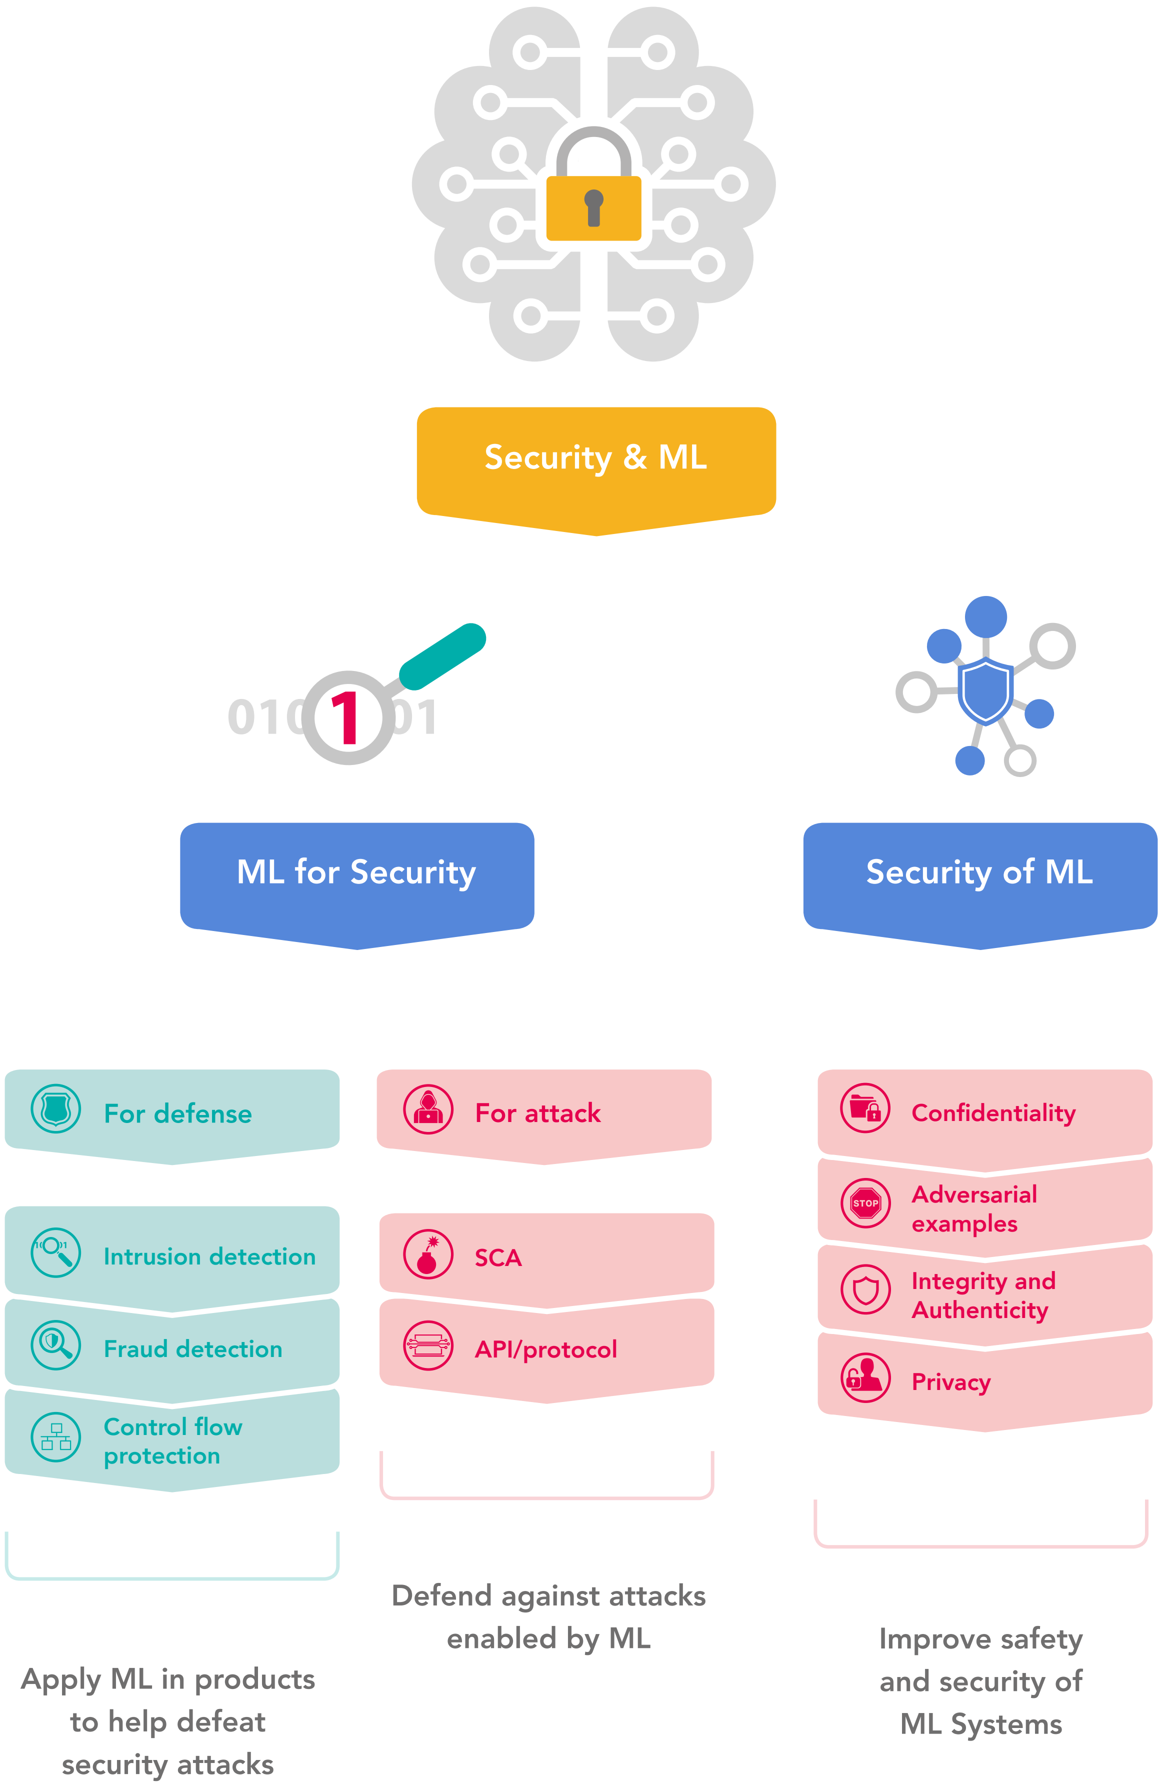

Security & ML

ML for Security

Security of ML

For defense

- Intrusion detection
- Fraud detection
- Control flow protection

Apply ML in products to help defeat security attacks

For attack

- SCA
- API/protocol

Defend against attacks enabled by ML

- Confidentiality
- Adversarial examples
- Integrity and Authenticity
- Privacy

Improve safety and security of ML Systems

PRIVACY-PRESERVING MACHINE LEARNING

It's easy to identify applications in which the data providers for AI, either in the training phase or in the inference phase do not want to provide their data unprotected. With the new EU General Data Protection Regulation (GDPR) in effect since May 25, 2018, privacy protection is mandatory for any business dealing with the data of EU citizens, and non-compliance can result in heavy fines.

Privacy protection is mandatory for any business dealing with the data of EU citizens

Non-compliance can result in heavy fines

In medical and financial applications, for example, businesses are held accountable for the privacy of users contributing to the dataset. A typical use case is the training of a diagnosis model from a patient's medical records. A related threat comes when the machine learning model is made publicly available, e.g., when hospitals perform diagnoses in the previous use-case. A malicious user having access to the model might be able to analyze its parameters and to recover some of the data used to train the model.

There are also applications in industrial environments where data privacy is crucial to system providers. For example, in predictive maintenance machine data is used to determine the condition of in-service equipment to precisely predict when maintenance should be performed. This approach achieves substantial cost savings over routine or time-based preventive maintenance because tasks are performed only when required and hopefully in advance of system failure. Machine owners participating in the service have a clear intent to benefit from the generated data, however, they also have a strong interest in not sharing their data with competitors use the same machines. This puts the maintenance service provider in a dilemma.

The key question is: How can businesses continue to respect privacy concerns while still permitting the use of big data to drive business value?

This has led to the general research area of homomorphic encryption, which is referred to as a privacy enhancing technology that encrypts data into computable cipher text. Any data being used in the computation remains in the encrypted form, and only becomes visible to the intended user. The result of the computation – once decrypted – matches the result of the same computation applied to the plain text. In a machine learning context, companies looking to feed data into an externally provided, cloud-based machine learning model can use homomorphic encryption to avoid giving access to unencrypted data while still allowing complex computations to be applied to their own data. NXP continues to do pioneering work in this arena.

Attribute-based cryptography is another privacy preserving technique that enables machine learning programs to run in compliance with strict data protection and privacy regulation. Attribute-based authentication, as used by NXP, is based on the Identity Mixer protocol developed at IBM® Research and allows for strong authentication and privacy at the same time. It relies on a combination of flexible public keys (pseudonyms) and flexible credentials that allows a user to share only the information required for a certain transaction, without revealing any other attributes. In addition, this enables external parties to create a profile of the user based only on his pseudonym.

The advantages are obvious:

> “The Internet is like the lunar surface — it never forgets a footprint. With Identity Mixer, we can turn it into a sandy beach that regularly washes everything away,” says **Jan Camenisch**, cryptographer and co-inventor of Identity Mixer.

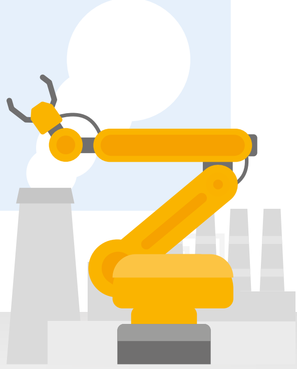

ATTACKING MACHINE LEARNING: ATTACKS AT TRAINING TIME

What happens when an attacker goes after the security of machine learning itself?

Before presenting potential threats in this arena, let's briefly recap how machine learning works. All machine learning starts with training data. The output is a set of parameters, essentially a model. In a second phase (inference), when given a new sample, the model infers the corresponding output. For instance, if the machine learning algorithm is an image classifier, one inputs a new image, the model returns its category (for example, that the image represents a cat). All the steps of this process, from training to inference, may be subject to attacks.

Attacks can occur even when the training data are collected and being fed into the ML model. While stealing data can be one objective of an attacker, changing the data or manipulating the outcome of the ML model may be another. For an AI model to make predictions that are in accordance with the physical reality, it is of utmost importance that the training data can be trusted. This property is sometimes difficult to achieve. A typical application is an anomaly detection tool trained from data sent by users. If a user "poisons" the training data by purposely sending incorrect inputs, this may result in inferior performance or even failure of the machine learning model at inference time.

ATTACKS AT INFERENCE TIME: ADVERSARIAL EXAMPLES

The user's privacy must also be guaranteed during the inference phase.
This is especially relevant when inference is done on private or sensitive data.

The user can also be the attacker. As a means of attack, the user may employ so called adversarial examples. An adversarial example is a valid input data that will cause the machine learning model to misinterpret it. This seemingly benign attack can create catastrophic consequences, for instance, if we think of road sign classification in safety critical situations.

By attaching a specially crafted sticker on a stop sign, researchers have shown that they can trick the image classifier into misinterpreting or not recognizing the sign at all. While the sign appears as regular stop signs to the human eye, the machine learning model is unable to see it as such. The concept of adversarial examples is not new. What is new is the severity of their consequences, like crashing an autonomous car in the stop sign adversarial example.

IP PROTECTION

The value of machine learning models mostly resides in the associated data-sets. Training data can be very expensive to collect or difficult to obtain. When machine learning is offered as a service, the user has access only to the inputs and the outputs of the model. For example, in the case of an image classifier, a user submits an image and gets in return its category. It might be tempting for a user to make a copy of the model itself to avoid paying for future usage. A possible attack is to query the service on chosen input data, obtain the corresponding output, and train the so-obtained data-set to get a functionally equivalent model.

The list of the presented attacks is certainly not exhaustive. Attacks can be combined to create even more damage. For example, once a model has been stolen, it can be used to try to recover training data or to craft adversarial examples. To ready ourselves against these evolving threats, the implications of AI must become an integral part of IoT system security. If training and inference of AI machine learning models are not to be become a wide-open gate for future adversaries, security by design and privacy by design principles must be considered from the beginning. Fortunately, it's not too late for the AI realm to apply the lessons already learned from IoT security.

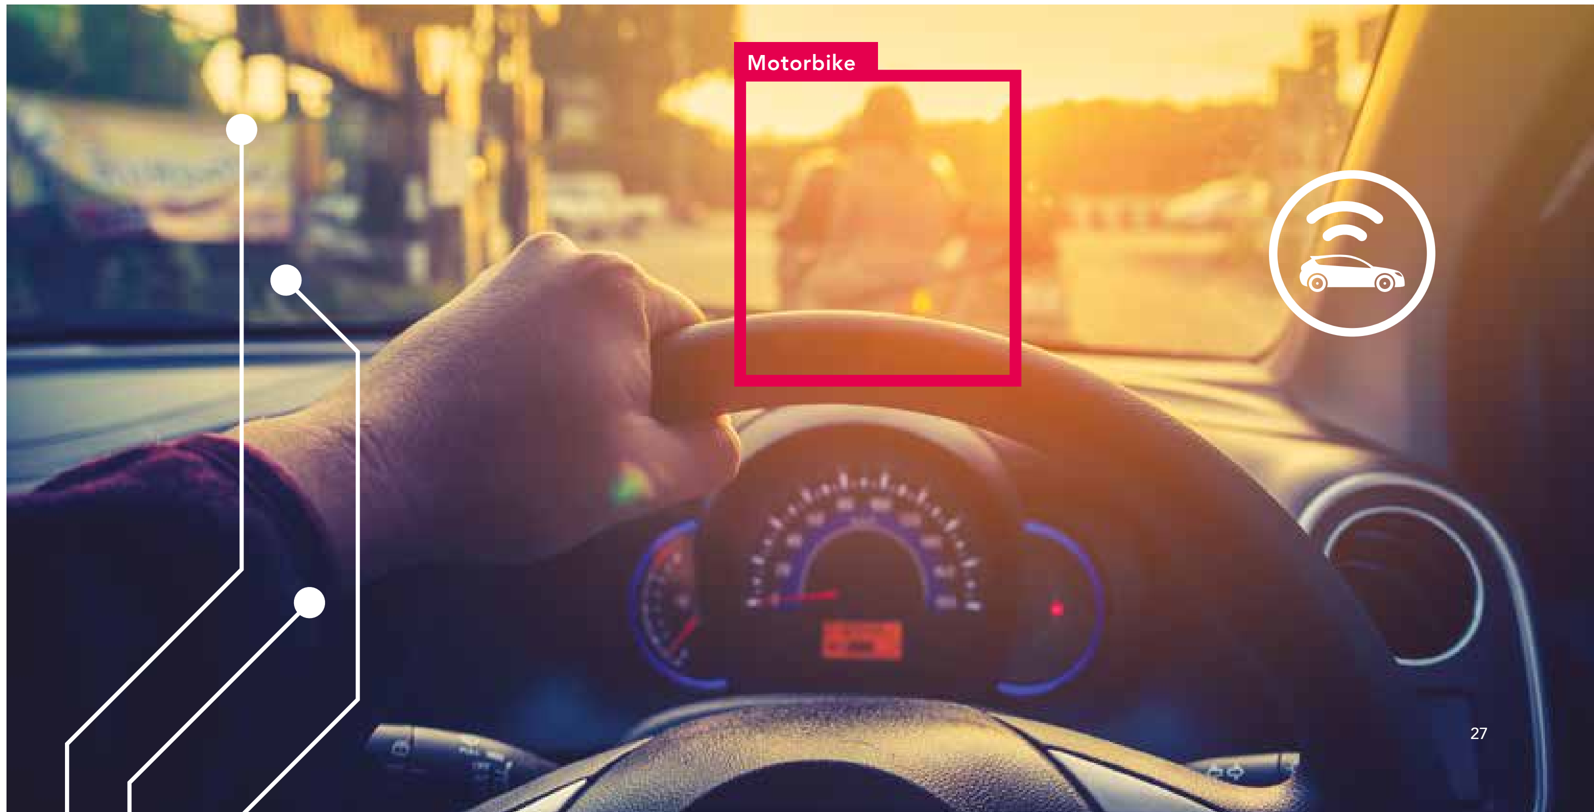

A GLIMPSE INTO THE FUTURE: ARTIFICIAL INTELLIGENCE OF THINGS

By designing things with smart properties and connecting them into the Internet of Things, we have created a global web of assets that have enhanced our lives and made them easier and more secure. The IoT gave us eyes and ears, and even hands, to reach out from the edge of the network into the physical reality where we gather raw information, which we stream to the cloud, where it's processed into something of superior value: applicable knowledge.

By adding high-performance processing, we've started to process and analyze information less often in data centers and the cloud, and more now at the edge, where we see the magic occur. We witness that magic in infrastructure, industry, personal devices and more making our lives more colourful: where all the action takes place that makes our lives colorful.

Smart traffic infrastructure

Smart supply chain factories

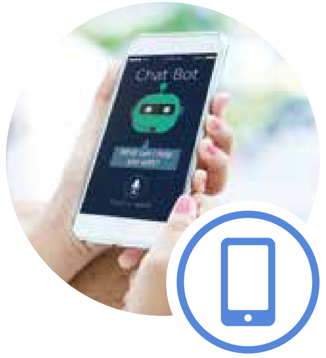

Mobile devices

Front-end stores

In real time

The IoT in its present shape has equipped us with unprecedented opportunities to enrich our lives. Yet, it is only a stopover on the way to something even bigger and more impactful. I'm talking about the artificial intelligence of things.

Today's smart objects, even though they stream data, learn our preferences and can be controlled via apps, they are not AI devices. They 'talk' to each other, yet they don't play together. A smart container that monitors the cold chain of a supply of vaccines is not an AI system unless it does 'something', such as making a prediction about the temperature development in the container and automatically adjusts the cooling.

An autonomous car or a search-and-rescue drone that autonomously navigates off-shore is in fact an AI system. If it drives or flies on behalf of you, you can trust that some serious AI capabilities are involved. Reading, speaking or translating language, predicting the mass and speed of an object, buying stock on your behalf, recognizing faces or diagnosing breast cancer, are all artificially intelligent characteristics when done by an algorithm.

Now, imagine a world in which the entirety of AI things was connected. Expanding the edge of the IoT with cognitive functions such as learning, problem-solving and decision making would turn today's smart things from mere practical tools into true extensions of ourselves, multiplying our possibilities to interact with the physical world.

As an integral part of the IoT, artificial intelligence is the foundation for entirely new use cases and services. Siemens®, for example, is using AI to improve the operation of gas turbines. By learning from operating data, the system can significantly reduce the emission of toxic nitrogen oxides while increasing the performance and service life of the turbine. Siemens is also using AI systems to autonomously adjust the blade angle of downstream wind turbines to increase the plant's yield.

GE has developed a drone and robot-based industrial inspection service, and it is using AI to automate both navigation of inspection devices and identification of defects from the data captured by them. In medical care, Thomas Jefferson University Hospital in Philadelphia seeks to improve patient experience with natural language processing that will enable patients to control room environment and request various information with voice commands. And Rolls-Royce® is developing an IoT-enabled airplane engine maintenance service that uses machine learning to help it spot patterns and identify operational insights that will be sold to airlines. On the consumer side, Google's Duplex offers an idea of what the future holds: A virtual assistant that can carry out "real world" tasks over the phone, performing functions like scheduling a dentist appointment or making a dinner reservation. These are tasks that typically require human interaction on both ends, but not anymore. Duplex's AI voice sounds so natural that the person taking the call could be unaware that they're chatting with a machine.

Now, what does this mean for the future? The truth is, even with a broad range of nascent AIoT applications emerging, we can't even fathom what else is coming. One thing is for sure – Today's digital age society is undergoing a fundamental change. The paradigm shift that comes with the convergence of AI and the IoT, will be even greater than the one we have witnessed with the introduction of the personal computer or the mobile phone. NXP is driving this transformation with secure, connected processing solutions at the edge, enabling a boundless multitude of applications in the future AIoT.

Effective security, based on the guiding principles of security and privacy by design, will be crucial to mitigate against the risks that come with it. If we keep this in mind when designing the infrastructure and devices of the future, the AIoT holds the power to transform our lives. And it's upon us, to turn the black box of the future into a bright one.

www.nxp.com