LogicMonitor

EBOOK

How Machine Learning Will Create a Brighter IT Future

Machine learning key concepts and how it's changing the game for IT organizations

Table of Contents

Chapter 1: What is Machine Learning? 2

Chapter 2: Understanding Machine Learning, Artificial Intelligence, and Algorithms 4

Chapter 3: What is the Future of Machine Learning for IT? 6

Chapter 4: The Rise of Quantum Computing 9

Chapter 5: The Trend Towards Hyper-Personalization ...12

Chapter 6: How Machine Learning Will Help IT Organizations 15

Chapter 7: Using Machine Learning to Monitor Your Infrastructure 18

Introduction

ITOps teams face a constant struggle of maintaining, monitoring, and managing IT infrastructure while also trying to develop strategic solutions to support business success across business units. Wouldn't it be nice to spend less time troubleshooting and more time innovating? This eBook examines the evolution of Machine Learning for ITOps and how it can benefit companies by offering context and meaningful alerting, discovering patterns, and enabling foresight and automation.

Chapter 1: What is Machine Learning?

Machine Learning: where did it start, and where is it going? The idea of algorithms has existed for a long time, but the modern concept of Machine Learning has its roots in 1950s computer science. With the strides made in data science and computer manufacturing—not to mention storage and data management—Machine Learning is becoming a great enabler of IT and business operations worldwide.

In this eBook, we're looking at how Machine Learning will ease the strain on ITOps teams in organizations across the globe.

This chapter:

- Outlines the history of Machine Learning.
- Provides examples of how Machine Learning is being used by businesses today.
- Shows how IT organizations, in particular, can implement Machine Learning concepts to improve operations.

LogicMonitor

The History of Machine Learning

Is Machine Learning just another algorithm? The answer is a definitive “yes and no.” While the concept of algorithms has been around for thousands of years (thanks to Persian mathematician Muhammad ibn Mūsā al-Khwārizmī), modern computational algorithms and the advent of Machine Learning (ML) has taken decades to evolve. Arthur Samuel is credited for coining the term “Machine Learning” while working at IBM, using the game of checkers in his research. Famed checkers master, Robert Nealey, lost a checkers game against an IBM 7094 computer in 1962. That famous match is still considered a major milestone in computer science.

Since then, the technological developments in storage and processing power have paved the way for modern products that we’re all familiar with, such as Netflix’s recommendation engine or personal assistants like Siri and Alexa. ML is typically considered to be a subset of artificial intelligence (AI) that enables computers to “learn” without being explicitly programmed with predefined rules. It focuses on the use of data and algorithms to imitate the way that humans learn, gradually improving its accuracy.

ML is an advanced algorithm (or model) that learns patterns in data and then calculates similar patterns in new data. So instead of setting up precise rules for an item you want to classify, for example, you can use hundreds of rules of an item. The computer finds the patterns in these descriptions and uses that pattern to identify and categorize the items.

ML concentrates on developing computer programs that can grow and change when exposed to new data. This reactive capacity, combined with our current ability to process massive amounts of data, means that ML can handle complex actions with efficiency and accuracy.

Machine Learning Gathers Steam with Big Data

The opportunities to use ML in myriad situations have grown exponentially in concert with computational power, device design, and data access and storage. From playing a simple game of checkers on a mainframe computer as large as your living room to shopping, banking, and video conferencing from a device that fits in the palm of your hand, ML has impacted business, education, and home life in equal measures. Today, ITOps organizations can use ML in various ways to support the business goals, particularly for companies that deal with volumes of data, media, documents, search needs, forecasting inventory and resources, and fraud detection.

For all these applications for Big Data, ITOps teams continue to be responsible for:

- Data backups and restorations, server management and configuration, optimizing the performance and resource allocation for effective delivery
- Management of overall tech stack, including IT infrastructure encompassing computation, networks, and hardware as well as the software applications plus configuration of any components
- Mitigation of disaster, including security implementation and management, along with disaster planning and management.

Large cloud computing companies—including Microsoft, AWS, and Google—as well as specialized tech firms like Databricks, MathWorks, and InData Labs, offer ML solutions to help optimize business operations, improve customer experience, and accelerate innovation. Data scientists and engineers use these tools to train and deploy models and manage ML operations.

How ITOps Use Machine Learning Today

While companies take advantage of ML to improve business opportunities and serve customers better through the scenarios explained in the previous section, ML can be especially helpful to ITOps organizations, no matter the industry. ITOps professionals know that when an incident occurs, and a network or service is impacted, it's usually the result of a sequence of incidents. Whether it be a security breach or a complete system outage, an issue with one service can spread to affect another service. Before you know it, the problem is compromising overall availability, performance, and customer experience.

Failure Detection

One obvious IT scenario where ML can help is the simple monitoring for and detection of network system failure. We've all read (or experienced) how missed signals can impact a company's stability and even the bottom line for weeks or months, from simple appliance failures to forgotten software updates. And the responsibility for these issues—large and small—remains with ITOps, a sometimes-under-supported organization in this age of "do more with less" budgets.

Performance Tuning

It's one thing to avoid disaster (failure detection), but ITOps teams can use ML to streamline processes and optimize the use of IT infrastructure in dramatic ways. For organizations like Managed Service Providers (MSPs) that manage networking and Wi-Fi, cloud services, IP telephony, hardware, software, and more, a little help can have a big impact. ITOps teams can use ML to gain visibility across entire environments, take advantage of data collection sets, set activity thresholds, and automate and manage alerts.

Conclusion

As Gartner notes in its Machine Learning Playbook for Data and Analytics Professionals,[1] "Technology is often an enabler of business operations, and ML should not be considered any different." This is particularly true for ITOps organizations, but it's important to acknowledge the complexity of such projects and understand the overlap and differentiators among the related data sciences. In our next chapter, we will clarify the distinction between Machine Learning, Artificial Intelligence, and Algorithmic Intelligence.

1 *Machine Learning Playbook for Data and Analytics Professionals,* Gartner 2020

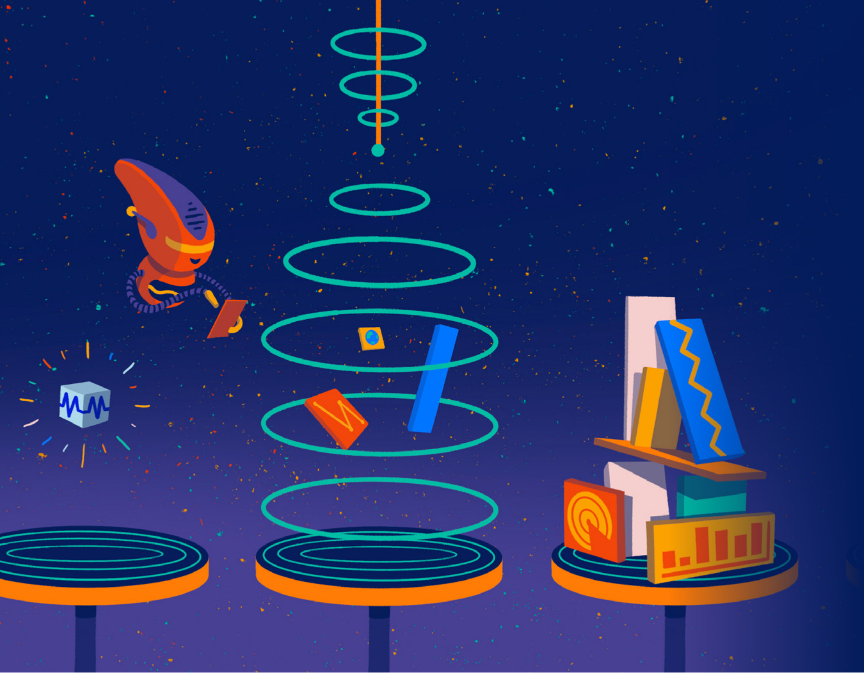

Chapter 2: Understanding Machine Learning, Artificial Intelligence, and Algorithms

Where do algorithms end and Machine Learning and Artificial Intelligence begin? In this chapter, we'll go beyond the buzz of trendy acronyms to understand the terms, the technology behind them, and how they fit together.

This chapter:

- Clarifies the difference between Artificial Intelligence, Machine Learning, and algorithms.
- Explains the nuances of supervised, unsupervised, semi-supervised, and reinforced learning models.
- Shows how these different learning-model types and algorithms can solve problems.

What's in a Name?

Machine Learning and Artificial Intelligence: the terms are thrown around, sometimes without much thought. After all, not just any algorithm qualifies for the label of "Machine Learning" or "Artificial Intelligence." We talked in the last chapter about how ML is an advanced algorithm (or model)—one that learns patterns in data and then calculates similar patterns in new data. But let's take a step back and define "algorithm."

Simply put, an algorithm is any step-by-step process with a defined output. So yes, adding water, coffee grounds, and pressing the button on your coffee maker counts as an algorithm. And if you set the coffee maker to start percolating at a certain time, well then, you've got an automated algorithm. Meanwhile, data science uses analytics to extract meaning from data. At its simplest, this starts as creating a business report and builds its way to Big Data Business Intelligence. Advanced analytics turns into ML and before you know it, we're into the broader AI field.

From Algorithms to Artificial Intelligence

At what point do we shift from simple algorithms to artificial intelligence? Some would say it's all in how the algorithm is programmed. As the complexity of the algorithm increases, it moves closer to the AI mark.

In between the simplest algorithm and the most intricate algorithm is what is sometimes referred to as "Algorithmic Intelligence." Algorithmic intelligence simply describes a process that can be defined as either basic or complex but with defined outputs.

Artificial Intelligence is the umbrella term covering the point at which computer processes can mimic human functions and outputs are not predefined. There are many subsets of AI, but perhaps the most well known include the following:

Machine Learning – ML refers to algorithms taking in data and performing calculations to find an answer. It includes software code that detects patterns in data. For example, a data scientist can use an algorithm to build (or train) a model that predicts outcomes. These models perform a range of different tasks on data. As noted in the previous chapter, ML is great for processes that search, detect, or forecast based on large amounts of data.

Speech – Another branch of AI applications includes language-related algorithms, including speech-to-text and text-to-speech applications. Common applications of this type of AI include voice assistants like Siri and Alexa, chatbots like the one in your banking app or e-commerce site, and translation engines such as Google Translate and Microsoft Translator.

Vision – Vision-focused algorithms process images. Computer vision and machine vision are often used in scenarios where large volumes of images need to be sorted or checked for errors, such as for goods packaging, braking systems for vehicles, and photo ID verification. Machine vision uses a camera to take a processable image and then partners with computer vision algorithms to process the image.

Deep Learning – Deep Learning (DL) can be considered a follow-on to ML. Traditional ML algorithms are linear, while DL algorithms are stacked in a hierarchy of increasing complexity and abstraction. Advanced solutions like driverless vehicles and medical diagnostic programs rely on DL.

Natural Language Processing – Natural language processing (NLP) is a more complex speech-related AI. NLP combines computational linguistics—rule-based modeling of human language—with statistical, machine learning, and deep learning models. These technologies work together to enable computers to process human language as text or voice data and to comprehend its full meaning, complete with sentiment and intent.

AI Model Supervision Types

Algorithms can be simple (defined input leads to a defined output) or complex (a system comes to a defined output by using a set of complex rules, calculations, or problem-solving operations). To reach the status of AI, outputs are not defined; instead, they are based on complex mapping of user data that is then multiplied with each output.

AI models also use the concept of supervision with categories including supervised, unsupervised, semi-supervised, or reinforced learning. Supervised learning relies on a full set of labeled data when training an algorithm. “Fully labeled” means that each item in the training dataset is tagged with the answer that the algorithm should deliver. For instance, a labeled dataset of cat images would indicate which photos were of Siamese, Bengal, and Persian cats. The model then compares any new image to the training examples to predict the correct label.

Supervised learning is best for situations that include many verified data points—very usable data. Unfortunately, that is sometimes difficult to provide. That’s where unsupervised learning enters. Unsupervised learning is a deep learning model that uses a dataset without explicit instructions on what to do with it. The training dataset is a collection of examples without a specific desired outcome or correct answer. The neural network then attempts to automatically find structure in the data by extracting useful features and analyzing its structure. There are a few ways to employ this type of learning:

Association – By looking at two or more key attributes of a data point, an unsupervised learning model can offer a prediction about the other associated attributes. For instance, if you are shopping online and look at mixing bowls, baking sheets, and wire whips, the site might suggest that you add a cake pan and parchment paper to your order.

Anomaly Detection – Unsupervised learning can flag unusual items in a dataset. This can be used in defect detection in manufacturing, fraud detection in financial transactions, or spotting potential risk or medical problems in health data.

Clustering – The most common application for unsupervised learning, clustering is a deep learning model that searches for training data that are similar to each other and groups them together. For instance, such a model could be used to sort photos of dogs roughly by breeds, relying on cues like the overall size, fur length and color, and head shape.

Autoencoders – Autoencoders compress input data into a code and then try to recreate the input data from that summarized code. One application of this type of learning is improving picture quality for video or medical scans. By training with both noisy and clean versions of an image, autoencoders can remove noise and gain clarity.

You can also use semi-supervised learning and reinforcement learning. As the name implies, semi-supervised learning tries to use the best of both unsupervised and supervised models. It uses a small proportion of label data and offers improved accuracy compared to a fully unsupervised model. Reinforcement learning is a model that learns through trial and error. An example of this is the Netflix recommendation engine—it learns by making suggestions, which you take or don’t (or start watching but don’t finish), and the model responds accordingly.

Conclusion

As we noted in the previous chapter, ML can be especially helpful to ITOps organizations. Using a variety of model types and construction, ITOps teams can use AI and ML to detect anomalies, predict usage, and suggest solutions to IT issues. In our next chapter, we’ll look at the future of ML for ITOps, including industry trends and use-case examples.

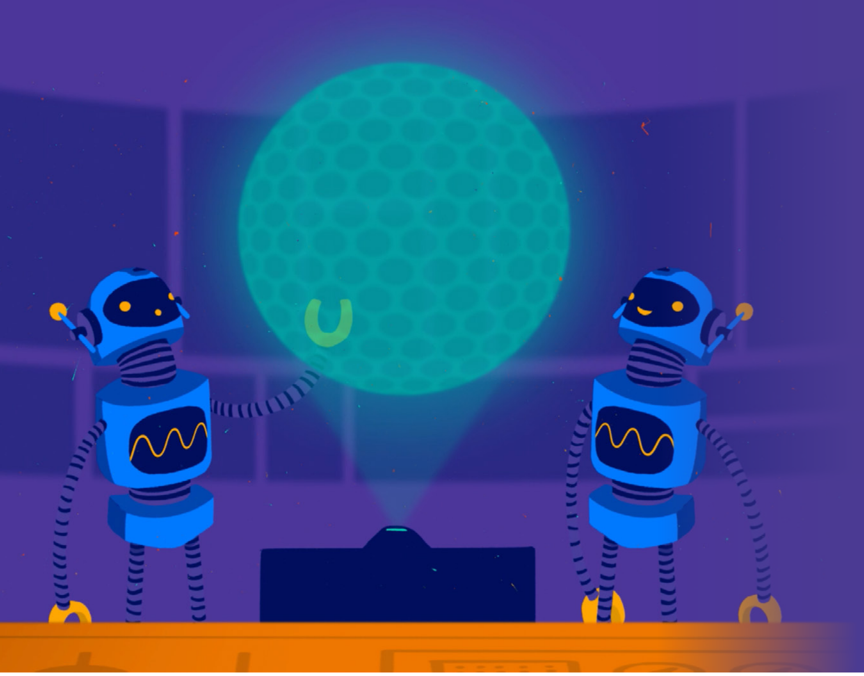

Chapter 3: What is the Future of Machine Learning for IT?

How will things change for IT organizations with the advancement of ML and AI? IT personnel will probably spend less time monitoring the network for issues and more time supporting innovative industry solutions. In this chapter, we'll find out how the workload will change for ITOps as ML matures.

This chapter:

- Highlights trends and advances in ML technology.
- Defines and explains multiple data algorithms, observability platforms, and federated ML.
- Presents use cases and developments in the healthcare and manufacturing industries.

Considering the Trends and Advances in Machine Learning

Many forces will impact how and when companies will use ML. Three particularly interesting areas include multiple data algorithms, observability platforms, and federated ML.

Multiple Data Algorithms

A big trend in AI for operations is applying capabilities from one data type to multiple data types. This began with different probabilistic methods such as AI, ML, and statistical analysis being applied to a single data type that was either metrics, logs, or transactions.

Soon, data scientists will design algorithms for multiple data sets together. The algorithm will look at the metric, log, and transaction data together, how they correlate, and what signals can be filtered out to make troubleshooting easier and faster.

These algorithms for multiple data types will help IT organizations save time by enhancing early warning systems and filtering signals more effectively.

Observability Platforms

Observability platforms are designed to look at metrics, traces, and logs, bringing them together to find the connection between the different data types. This data provides IT organizations with a broad view across the customer experience, employee productivity, and digital infrastructure to understand how the business is performing. In addition, incorporating ML and automation into these platforms reduces the time required to prevent system problems proactively.

Federated Machine Learning

Collaborative learning, known as federated learning, is an ML technique that trains an algorithm across multiple decentralized edge devices or servers holding local data samples without exchanging them. At the moment, it is being used to satisfy stringent global privacy regulations while making sure the business can still use the data. It adds a layer of data protection to ML technologies by decentralizing the ML model and distributing the algorithmic learning across more than one device.

This helps solve one of the key issues that arise from ML techniques today: aggregating large volumes of possibly sensitive data to train the model and keeping that data in one place. Gathering this data puts organizations at risk for malicious actors and privacy breaches. Federated learning models create a framework such that the devices do not exchange or share any data, and no centralized location is relied on to send information anywhere. Only the owner of the data has access to the information.

Still, federated learning is not entirely risk-free. But it's a new approach that has been gaining traction (since about 2017). Adaptive ML can be combined with federated ML to use model improvements in multiple locations and usage contexts. This approach can enable, for instance, autonomous systems, such as self-driving vehicles or smart robots. For ITOps organizations, this might lead to sophisticated, automated network monitoring for businesses or governmental agencies required to follow restrictive privacy regulations.

Driving Industry Innovations

Many industries started investing in ML and AI long before the global pandemic. Forced lockdowns helped companies in the healthcare and pharmaceutical industry and those in manufacturing to accelerate their digital transformation plans.

Machine Learning in Healthcare and Pharmaceuticals

Covid-19 was, in many ways, a catalyst for advancing technologies in the pharma, medicine, and health sector. As a result, IT organizations in this industry support clinicians and medical facilities that have a renewed focus on remote patient monitoring and telehealth, nursing, and patient care in general.

AI technology is rapidly expanding into other healthcare areas, including early detection of diseases, treatment, and research. And it's not as new as you might think. Three years ago, Japan was already addressing its doctor shortage with AI. The technology will continue to evolve and play a more prominent role, especially as the world manages Covid-19. From IoT to automation, ML and AI offer exciting opportunities to take healthcare to the next level.

Improving Patient Care with Connected Devices and IoT

Remote medicine today relies on connected devices and remote sensor technology. Devices include monitors for insulin levels and heart rate, which allow clinicians to manage chronic conditions without being in the same room with a patient. IoT devices can also produce the data needed to help determine a patient's healthcare needs in ways that can improve patient outcomes substantially.

Assessing Patients Using AI and Automation

AI is already being widely used as a diagnostic error prevention tool in radiology. Now AI is being used to help automate triage situations and help with prescription handling. This can ease the burden of clinicians and allow them to focus on more advanced diagnoses and ultimately see more patients. In addition, Chatbots—now ubiquitous across so many industries—are an easy, familiar tool for patients to engage with for care, and they reduce costs for healthcare companies.

Refining Drug Development and Efficacy

Device and drug efficacy can be improved through data analytics with AI engines. Comparative studies are made significantly easier and cheaper by using ML models to assess the effectiveness of new drugs to market or new medical devices. AI can also identify patterns between different drugs and diseases to help with the rapid prototyping of new drugs or new uses for existing drugs.

Machine Learning and AI in Manufacturing

Like many industries, manufacturers have been on a digital transformation journey for a couple of decades. The move towards automation is being accelerated by AI, robotics, and the public cloud. It shows up in nearly every aspect of the modern manufacturing business, even customer service and marketing. The evolution of IoT and sensor technologies alongside advanced digital modeling allows today's manufacturing businesses to drive innovation and change at an accelerated pace. Analyzing data from digital prototypes can accelerate the development of a product significantly. Let's explore the ML and AI trends in manufacturing to glimpse what the ITOps groups might soon be supporting.

Customer Service and Analytics – An effective AI implementation is built on data. Strategies for initiatives like predictive maintenance, digital customer success applications, and BPM rely on accurate environmental data. By leveraging environmental sensor data and AI, for instance, it is possible to predict changes in temperature, humidity, or other factors that could take a plant offline.

ChatBots – Companies in the manufacturing sector have successfully implemented ML and AI in several areas, but they are just starting to consider using chatbots, which are widely used in other industries. Implementing an effective chatbot could reduce customer service costs considerably and improve customer satisfaction.

Digital Twin – Originally pioneered by NASA, the Digital Twin has been one of the top technology trends for more than five years now. Mirroring physical environments in a digital framework is key to product innovation, improving processes, reducing downtime and waste, as well as improving customer experience. With the advancement of IoT and AI, Digital Twin technology is fast becoming the mainstream as manufacturers attempt to keep up with the market leaders in their field. Ensuring these models are populated with the right data is important—and that can be an easier task with ML.

Intelligent Automation and RPA – Automating repetitive processes with robotic process automation (RPA) can help beyond just the manufacturing process. It is also useful in other areas like order processing and logistical operations, where staff have to repeat multiple tasks that could be automated. The obvious contenders for RPA, like customer service, are now automated in most companies. RPA is increasingly being applied to other time-consuming processes within a business, resulting in better employee and customer satisfaction and productivity.

Industrial IoT – Industrial IoT (IIoT) offers manufacturers great benefits, including cost savings and better visibility of plant operations. It can also help IT organizations discover the root cause of major issues like outages.

Benefits of ML for ITOps across Industries

These industry examples showcase how ML and AI are changing healthcare and manufacturing for the better. However, these use cases also highlight the changing responsibilities for ITOps personnel. Fortunately, ITOps can apply ML to offload manual, repetitive tasks, such as monitoring networks, anomaly detection, and root cause analysis. This gives ITOps teams the ability to focus on supporting these innovative and strategic solutions.

Conclusion

In this chapter, we explored how the workload might change for IT organizations as ML matures, using industry examples. The future of an IT organization enhanced by ML will combine algorithms with advancements in personalization, resulting in IT personnel spending less time monitoring for IT issues and more time focusing on innovating. In our next chapter, we'll look at the rise of Quantum Computing and how it could benefit businesses across many industries.

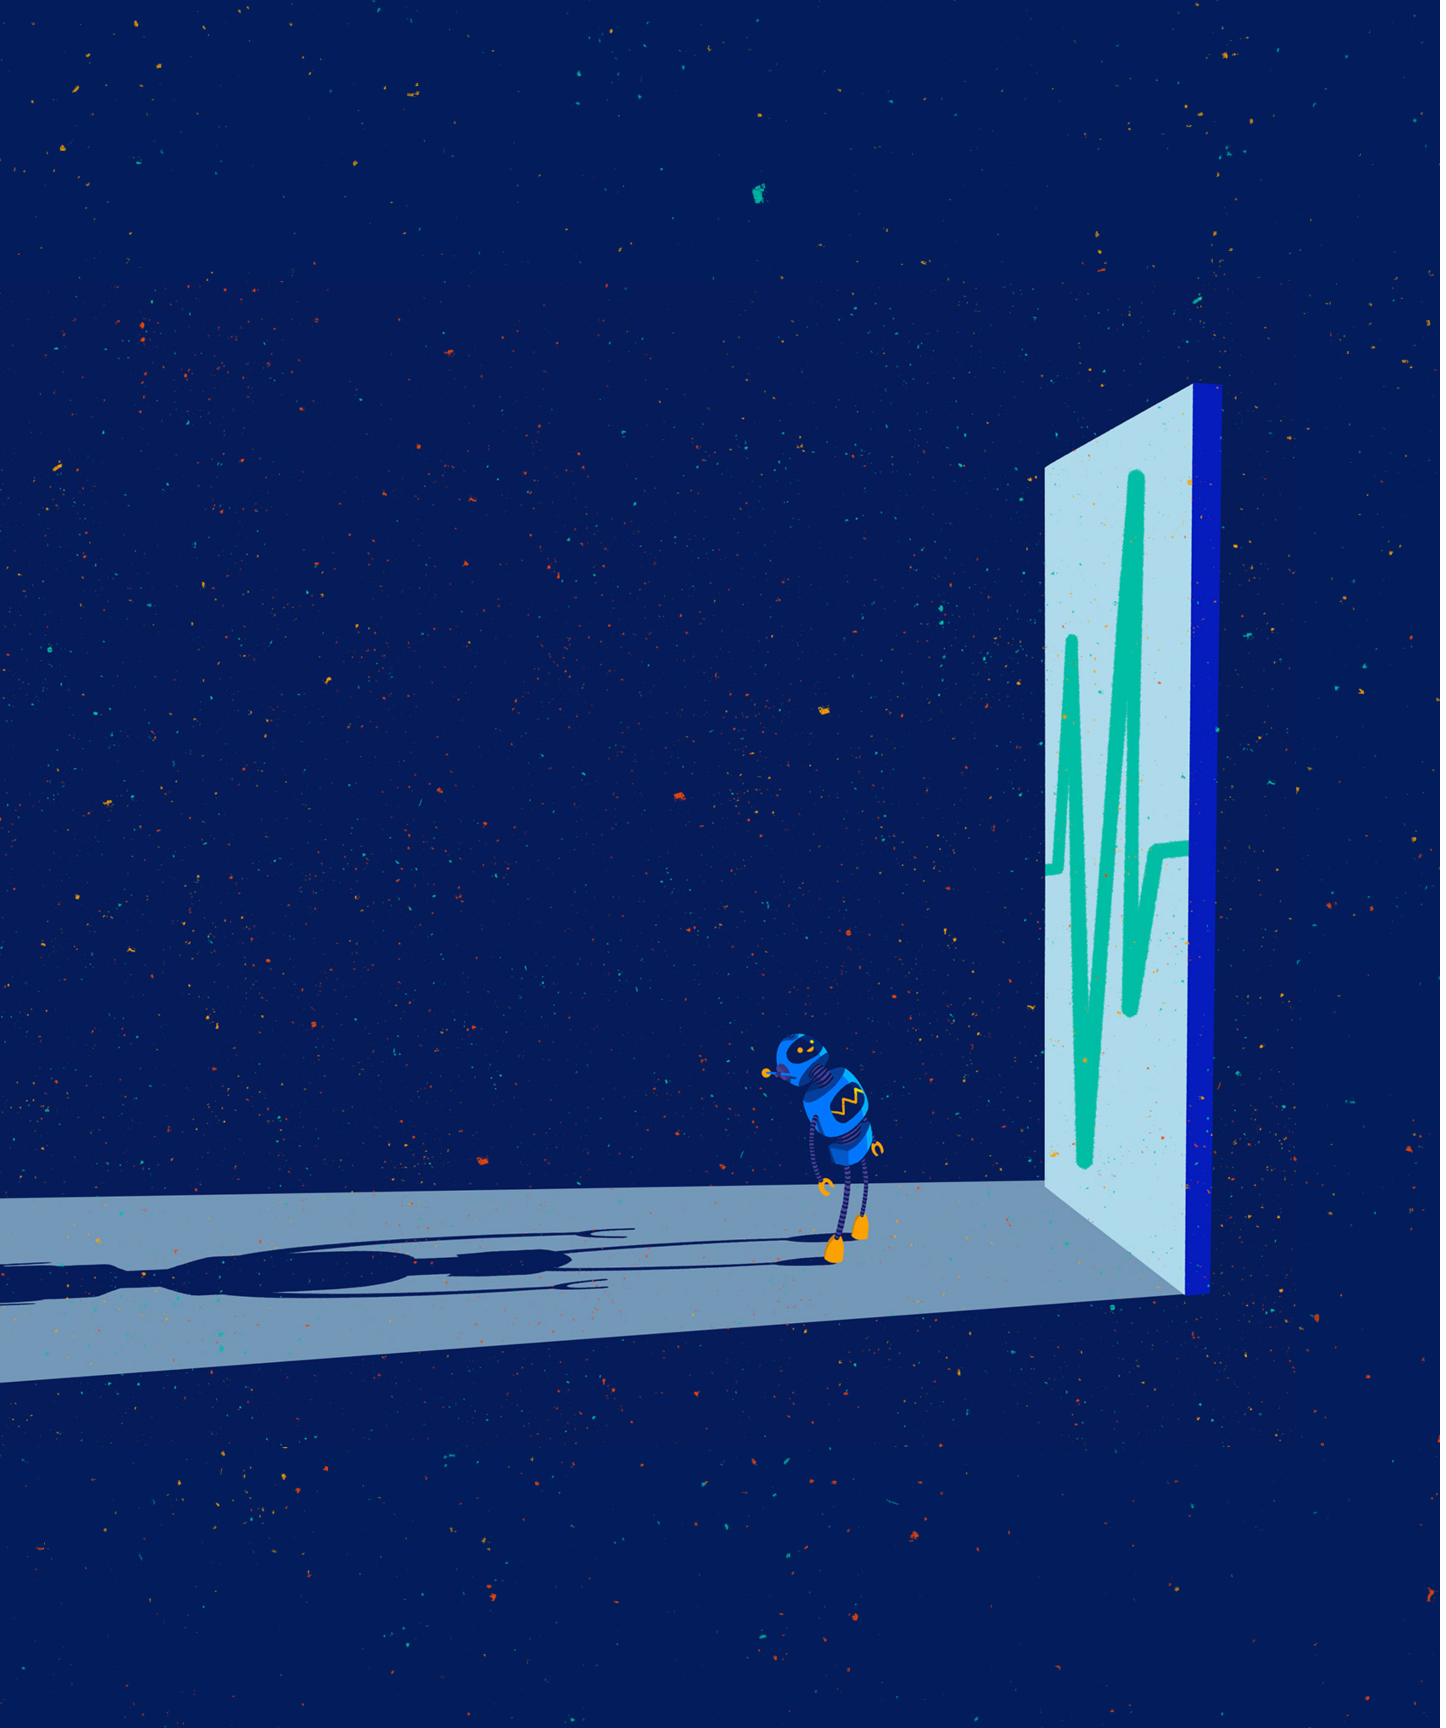

Chapter 4: The Rise of Quantum Computing

Machine Learning has always been aided and, at the same time, limited by computer power. As a result, research and development are laser-focused on quantum computing and how it can offer the computing power to do things that classical computing cannot. In this chapter, we'll contemplate how the rise of quantum computing has influenced and will continue to support Machine Learning innovations.

This chapter:

- Presents a brief history of quantum computing
- Describes the maturity level of quantum ML
- Offers use cases and developments in the healthcare and financial services industries.

How Quantum Computing Evolved

To this point, we've been talking about ML in the context of classical computing. With their thousands of classical CPU and GPU cores, the biggest classical computers, supercomputers can solve many complex problems. However, the truth is, even supercomputers can't manage to hold the number of combinations involved in real-world issues—supercomputers don't have enough memory. Since the 1980s, scientists like physicists Richard Feynman and David Deutsch have been exploring how quantum mechanics can improve computing speed and create solutions where none currently exist.

Quantum computing uses the collective properties of quantum states (superposition, interference, and entanglement) to execute computations. Quantum computers are being designed to solve particular computational problems, like integer factorization, significantly quicker than classical computers.

Large tech companies, including notably IBM, Microsoft, Google, and AWS spend a lot of R&D money developing analog quantum computers. IBM has built a network of partners in its IBM Quantum initiative. It includes a worldwide network of Fortune 500 companies, academic institutions, researchers, educators, and enthusiasts committed to driving innovation in this area. IBM, Microsoft, and AWS all offer cloud-based quantum computing options, primarily quantum computing simulators, so that interested parties can experiment with algorithms, test quantum hardware, design software, and explore viable applications for the technology as it matures.

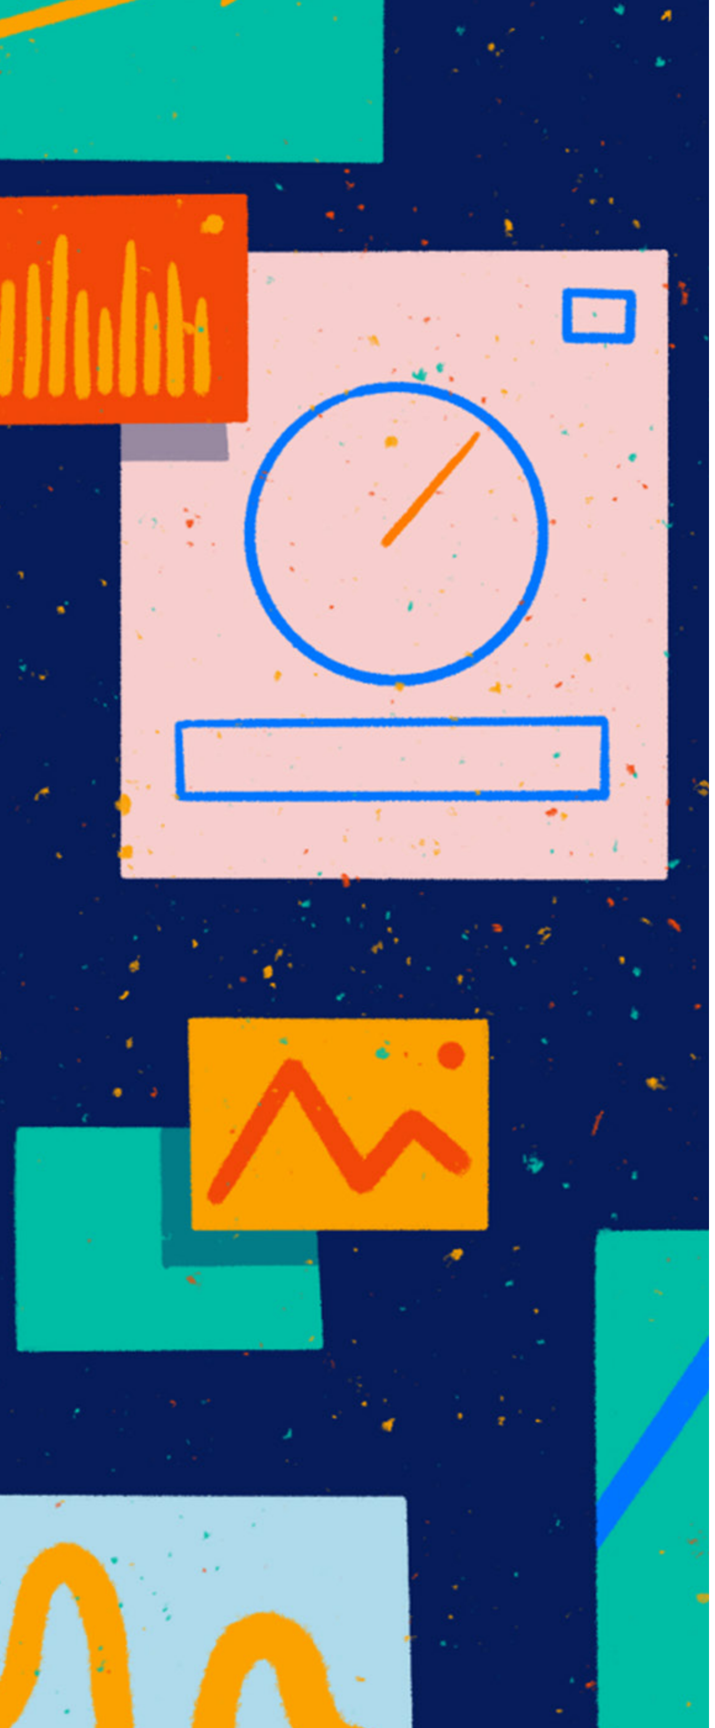

Quantum Machine Learning

As Gartner pointed out in its Hype Cycle for Data Science and Machine Learning, 2021,[2] quantum ML could enable a subset of ML algorithms to be run using the quantum computing paradigm. However, Gartner notes that only a limited number of quantum ML algorithms currently exist: “We have yet to see any evidence that ML could benefit from quantum computing over traditional alternatives. However, the parallel nature of some ML techniques could make quantum computing a viable path to explore. Increasing awareness of quantum ML capabilities plays a key role in determining its potential value.”

Industry Applications

As we mentioned earlier in this chapter, Google, IBM, Microsoft, and other major tech companies put substantial resources toward quantum computing research as they attempt to pioneer breakthroughs for industries like medicine, supply chains, financial services, and so on. From banking and financial services to healthcare, general science, and cybersecurity, quantum machine learning offers businesses opportunities to improve services and expand exponentially.

Financial Services

Banks and finance firms, for instance, are already using machine learning techniques such as reinforcement learning for algorithmic trading, and they use Natural Language Processing for risk assessment, financial forecasting and accounting and auditing. Some practical uses of quantum computing in the near future include the following:

Algorithmic Trading – Trade decisions can be sped up and complex models can be simplified by using quantum reinforcement learning methods in algorithmic trading.

Asset Pricing – Banks have been using Recurrent Neural Networks (RNNs) to run time series predictions. Institutions like JPMorgan are considering using them for asset pricing models but RNNs consume a lot of computing power. However, it can be advantageous to use parameterized quantum circuits and quantum Long Short Term Memory units that allow users to predict outcomes based on evolving processes from historical data.

Fraud Detection – Currently a wide area of concern that is being tackled by classical computing ML, fraud detection can benefit from quantum clustering algorithms when used to detect anomalies.

Volatility Prediction – Quantum methods can also be applied to tracking and predicting changes in a security’s price. A density matrix can be produced by deep quantum neural networks, and the implied volatility of an option can then be calculated using its corresponding element in that matrix.

Cybersecurity

Although quantum ML will improve anomaly detection, data protection is another cybersecurity area to benefit. Modern cryptography and its algorithmically encoded data have risks. Most notoriously, hackers can intercept the cryptography key to decipher the data or use powerful computers to predict the key. With enough compute power, a hacker will win.

Researchers are cautioning companies that more secure cryptography keys will be needed in the future. One solution is to make cryptography keys totally random and impossible to solve mathematically. Randomness is fundamental to quantum behavior, making quantum ML perfect for creating cryptography keys that are impossible to reverse-engineer.

Healthcare and Pharmaceutical Advancements

In previous chapters, we discussed how AI is being used in the early detection of diseases, treatment, and research. But there are also applications in the discovery of new drugs. The procedure currently relies on molecular simulation, which involves modeling how particles interact inside a molecule to configure molecules capable of treating a particular disease. Sounds pretty complex, right? And it requires a lot of computational power.

Quantum computers could speed up this process, whittling a typical 10-year cycle for bringing a new medicine to market down considerably. Pharmaceutical giant, Roche, is already working with Cambridge Quantum Computing to design and implement noisy-intermediate-scale-quantum algorithms for early-stage drug discovery and development.

2 *Hype Cycle for Data Science and Machine Learning,* Gartner 2021

Quantum Computing for ITOps

In the future, ITOps solutions may employ cloud quantum computing power to move beyond dependence on vendors and experts and design and implement self-learning and self-healing IT operations. The power of quantum computers combined with advanced AI models can help an IT system constantly learn about its environment and improve itself while adapting to changes.

ITOps teams may be able to let autonomic computing take on the complexities of network monitoring and maintenance and even have the system itself build a constantly expanding knowledge base that surpasses what humans can reasonably provide.

Limitations

Quantum computing is promising, but it's still fairly limited in actual application, including ITOps teams and the organizations they support. Generally, systems can only scale to tens of qubits (quantum bits). This means that algorithms executed on these systems are mainly exploratory. Although quantum computing can theoretically bring dramatic benefits to some classes of data, a big challenge is encoding. For quantum ML to scale successfully, great quantities of data must be encoded and loaded into the quantum system.

Likewise, new algorithms that can take advantage of capabilities offered by near-term noisy quantum systems are in the discovery stage. Research in this evolving field continues, and some scientists question the potential applicability of quantum computing in ML. And, of course, hardware and software must evolve as well.

Today, research and development tend to focus on creating different quantum algorithms for ML kernels. Companies like IBM have prototyped ML algorithms for select use cases, and IBM specifically has developed a roadmap to reach 1000+ qubits by the end of 2023.

Although quantum ML is still in an early stage, Gartner has noted that "potential applications of quantum computing in artificial intelligence and ML include quantum search, recommendation algorithms, quantum algorithms for game theory, and quantum algorithms for decisions and learning."[3]

Conclusion

In this chapter, we looked at the exciting area of quantum computing and how it may affect industries—and the ITOps that support them—in the future. IT organizations can start experimenting with cloud-based offerings like IBM Quantum, Azure Quantum, and Google Quantum AI. In our next chapter, we'll find out how the trend towards hyper-personalization built on data, analytics, and ML impacts IT organizations and benefits companies.

3 *Gartner Report,* ID G00747536, 2 August 2021

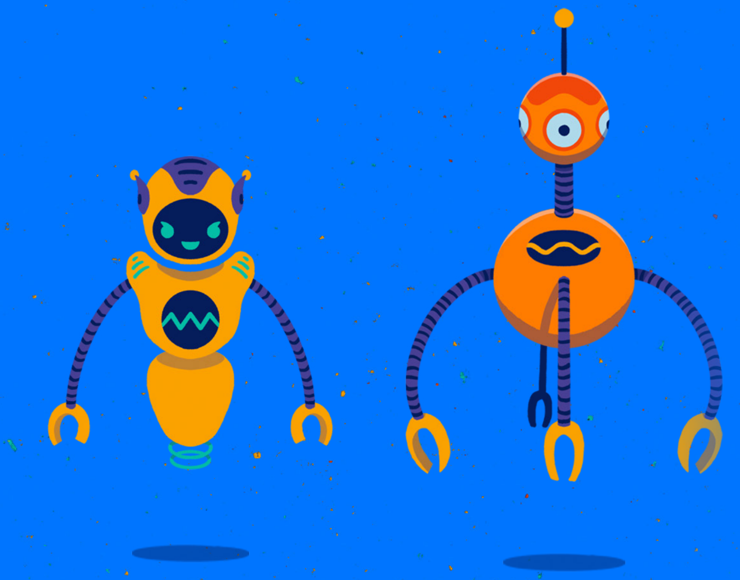

Chapter 5: The Trend Towards Hyper-Personalization

IT organizations that support e-commerce businesses will find that Machine Learning plays a big part in helping the shift towards hyper-personalization. In this chapter, we'll find out how this trend, built on data, analytics, and ML, impacts IT organizations and benefits companies.

This chapter:

- Defines hyper-personalization and why it matters.
- Describes how customer expectations across industries are evolving.
- Outlines how IT organizations will be supporting efforts to implement hyper-personalization across sectors.

What is Hyper-Personalization?

ITOps teams typically have focused on supporting the back-office business (think corporate email, intranet, HR systems, and so on) and customer-facing offerings, like e-commerce sites and online support services. Now more than ever, ITOps is involved in improving the customer experience (CX), and a huge part of that is personalization efforts driven by the CMO and the marketing organization. But just what is "hyper-personalization," anyway?

Modern marketing efforts began in the early 20th century with the concept of segmentation—defining who customers are and how to target customers using demographic and location data. It evolved to encompass customer details, including name, age, and purchase history once CRM technology arrived, which began the implementation of "personalized service" for the masses. As customers started leaving data crumbs all over the internet, companies began gathering that data to help design and deliver valuable offerings for those customers. For instance, Amazon's famous personalization engine offers product recommendations, personalized product re-ranking, and customized direct marketing. Likewise, Netflix knows just what you want to watch (or what they want you to watch).

Hyper personalization is a massive leap forward. It moves beyond demographics and purchase history to consider browsing activity, purchasing habits, and other behavioral data to help companies determine what the consumer wants or needs.

Modern CMOs see the value of pivoting to a digital-first mindset and are interested in using ML and AI to create deeper, more authentic interactions with customers while effectively capitalizing on data gathered to drive insights-driven results.

How Customer Expectations Are Changing

Companies like Amazon and Netflix have effectively trained consumers to expect companies of any category to understand what they want and make it easy to get. As a result, a variety of features and functionalities are making their way across industries:

Relevant, Custom Advertising – Customers want to see meaningful advertisements, showcasing goods, services, or experiences that relate directly to them.

Recommendation Engines – Content, service, and product suggestions should be tailored to individual needs and preferences.

Omnichannel Service – Customers want to shop and connect with businesses through online and offline channels.

Self-service Assistance – Chatbots are becoming more prevalent and useful to both customers and businesses, making it easier to answer typical questions in a more personal way without the need for office hours or a contact center.

Dynamic Pricing and Offers – It's easier to convert prospects when you can change the offer or price based on customer habits and preferences.

Next-generation Loyalty Programs – Re-engagement has previously centered on customer purchase history, but companies can now use sophisticated micro-segmentation and geospatial data to offer highly contextualized deals and suggestions.

Across industries, businesses are finding ways to bring their marketing efforts towards achieving hyper-personalization. We see it, of course, in retail and entertainment services, but it's emerging across all industries, particularly those that come with lots of data. In healthcare, for instance, clinicians and medical facilities focus on remote patient monitoring and telehealth, nursing, and patient care in general and early detection of diseases, treatment, and research.[4]

In manufacturing, the move towards automation is being accelerated by AI, robotics, and the public cloud. It shows up in nearly every aspect of the modern manufacturing business, even customer service and marketing, as companies implement customer service chatbots, order processing support, and logistical operations improvements. And in the financial services industry, ML algorithms are used in fraud detection, trading automation, and financial advisory services for investors.

Why Businesses Benefit from Hyper-Personalization

An often-cited study[5] conducted by the University of Texas indicates that the need to personalize comes from the desire to control and simplify decision-making. Interestingly, companies that invest in personalizing products and offerings often end up with customers at the center of their corporate decisions. Instead of making a product and convincing people to buy it, companies can know what products and services consumers want and plan accordingly. This is a win-win for customers and companies alike: customers get what they want or need, and companies can be more successful in sales and support efforts.

This might explain why so many companies spend time and money developing hyper-personalization marketing plans. As noted by Forrester in a 2019 study commissioned by IBM: "Even with immature personalization strategies, firms see an almost 6% increase in sales revenue, a 33% increase in customer loyalty and engagement, and an 11% decrease in marketing costs."[6]

4 *No longer science fiction, AI and robotics are transforming healthcare*, PWC, https://pwc.to/2weGo5v

5 *Consumer control and customization in online environments,* Laura Frances Bright, University of Texas, 2018

6 *Personalization Demystified: Enchant Your Customers By Going From Good To Great*, Forrester Consulting, December 2019

How IT Supports Hyper-Personalization

Hyper-personalization is an art and science that necessitates both a strategy and the right technology. The strategy might be typically driven at the C-suite level but with a close connection to frontline workers (sales and support) and the IT teams who support it all. There are several ways IT teams will be supporting hyper-personalization efforts:

CX Platform Support – For many businesses, the first step forward is selecting a CX platform that supports omnichannel journeys and connects with the CRM system. IT personnel will be an integral part of deploying and maintaining CX platforms, whether they be outsourced or developed in-house—or a combination of the two.

Big Data Management – IT organizations may find themselves adding more data scientists to the team to help direct and manage efforts to gather, store, and process the vast amounts of data required to support the hyper-personalization strategy. Compliance with data privacy regulations is also part of the equation that IT must support.

ML and AI Application – Part two of big data management is the use of ML and AI in CX and marketing solutions. Data scientists, ML and AI engineers, and product managers work together to create, maintain, and support personalized recommendation engines, chatbots, targeted ads, and so on.

Conclusion

Companies are increasingly personalizing along the customer journey, starting with product design and outreach and encompassing the end-to-end consumer experience. Data analytics, ML, and AI are indispensable tools for building a strategy to meet hyper-personalization goals. In our next chapter, we'll look at emerging and forecasted trends in ML and AI that will transform IT organizations.

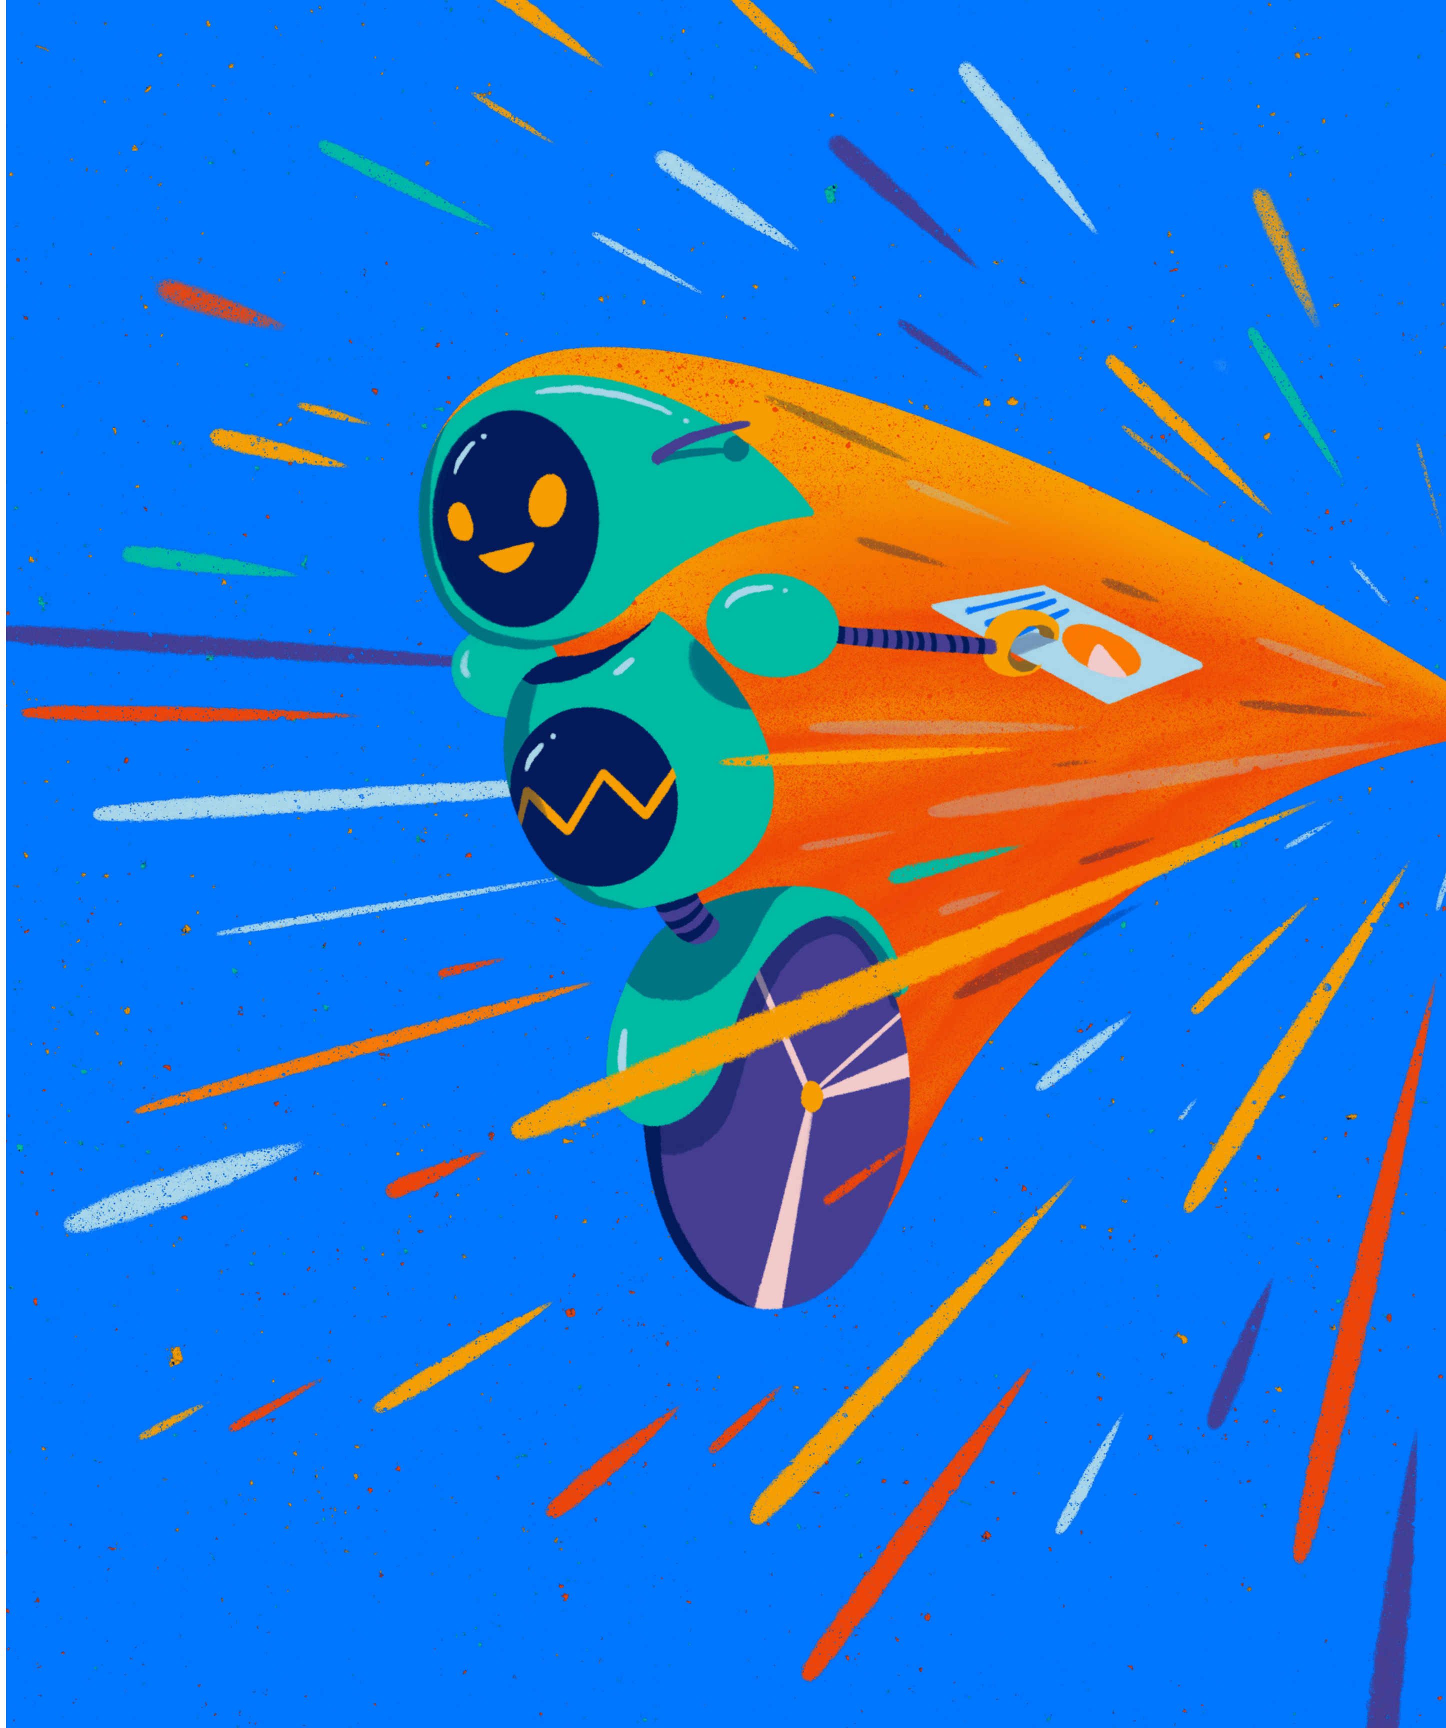

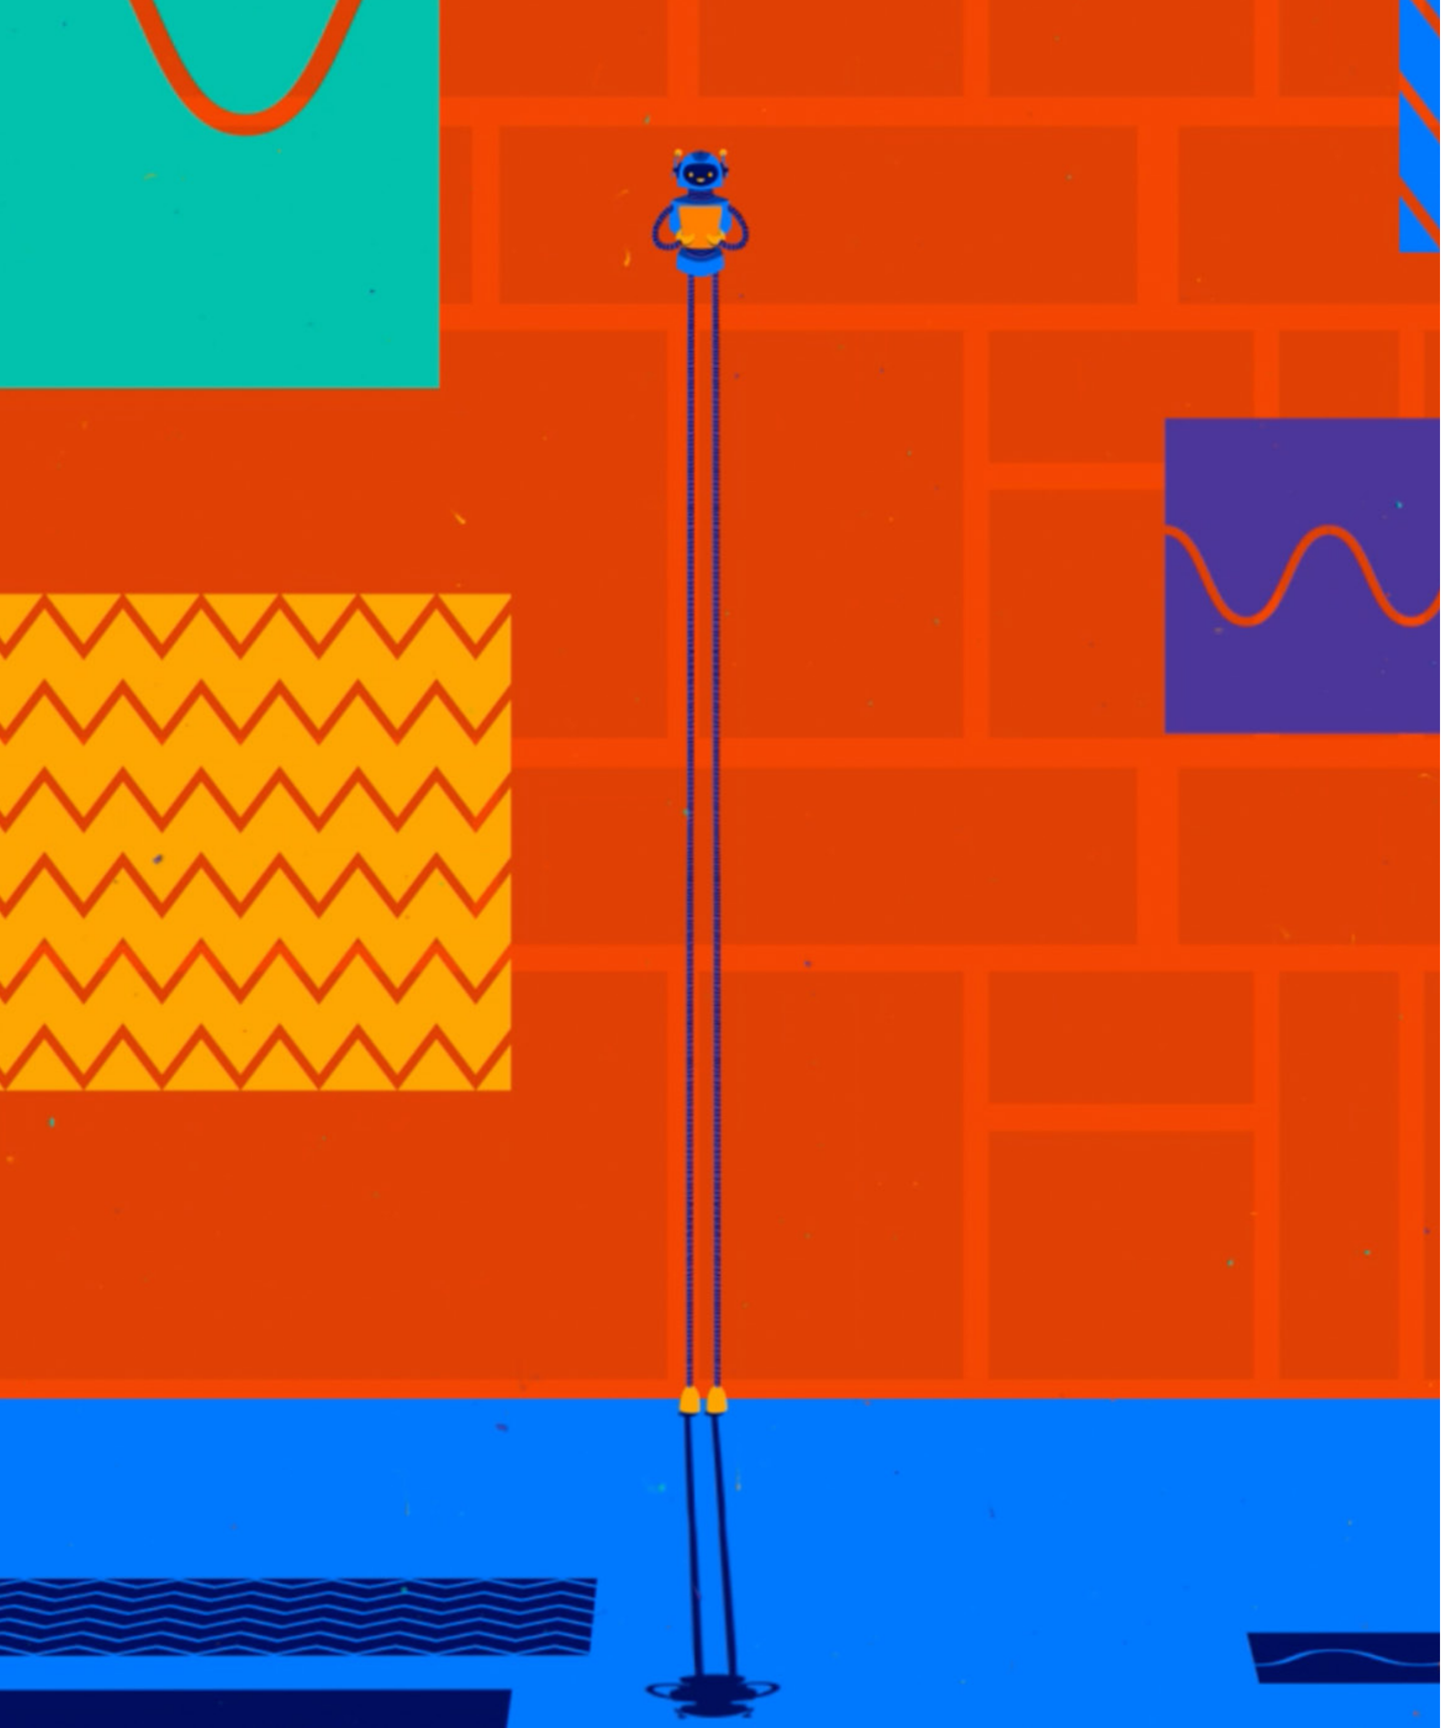

Chapter 6: How Machine Learning Will Help IT Organizations

How will AI and ML advancements impact IT organizations? As open-source tools like Python and TensorFlow mature, ITOps teams will find it easier to capture and export metrics, traces, and logs. In addition, solutions like OpenTelemetry, with its single set of standards and technology tools, will simplify the task of monitoring distributed cloud-born applications and lead to exciting advances like self-healing systems.

This chapter:

- Outlines the challenges of managing a distributed, cloud-based infrastructure.
- Describes the business and workplace changes that are leading to new opportunities and a transformation of how ITOps can support these endeavors.
- Highlights open-source ML tools and solutions that can help improve operations.

Workplace and Business Drivers

Many businesses have spent the past decade developing a digital transformation plan and slowly pushing that effort forward. The recent forced move to remote or hybrid working has accelerated the pace of this effort while adding the complexity of a further-distributed network. In addition, customer expectations across industries encourage companies to push the envelope with respect to customer experience, including hyper-personalization, 24/7 service, and custom offerings.

Meanwhile, cloud computing has advanced in a way that has made a ton of data available for capture and use. Companies can use that data in myriad ways: tackle cybercrimes, halt fraudulent transactions, minimize customer churn, recommend products, and more.

Trends in Machine Learning

As mentioned in previous chapters, ML is not new. However, recent advances make it easier than ever to exploit the value of ML, including the constantly increasing processing power, the rising sophistication of software to analyze data, and of course, the vast amounts of data available to train predictive models. In addition, IT organizations can use ML algorithms to create historical models to help clearly predict the future.

Marketing and sales organizations are likely to benefit from ML and AI almost immediately. According to Salesforce Research, these organizations have been picking up steam, using ML and AI to climb from 29% in 2018 to 84% in 2020.[7] In addition, McKinsey Global Institute estimates that ML and AI overall will bring $1.4 Trillion to $2.6 Trillion[8] in sales and marketing value over the next three years alone. These advances mean that IT organizations will be pivoting to support these efforts while also using ML for their IT tasks.

Will this highly specialized technology and the skills to apply it become available to all companies or just those with deep pockets and a technology focus? Fortunately, nearly all major tech companies are working hard to democratize ML, AI, and deep learning. Remember how difficult it was for non-technical people to sell products online before eBay and Etsy? Similarly, companies like Microsoft, AWS, and Google are working to ensure that companies will not need a massive team of data scientists and engineers to take advantage of ML.

In addition, organizations like the Open Machine Learning (OpenML) project provide a space for interested technologists to participate in an open, organized, online ecosystem for ML. OpenML builds open-source tools to find available data from any domain, easily draw them into just about any ML environment, and quickly build models with thousands of data scientists. It also helps members analyze results and offers advice on building better models.

Likewise, TensorFlow is a popular end-to-end open-source platform for ML. It has a comprehensive environment of tools, libraries, and resources that helps researchers drive advancements in ML while developers quickly build and deploy ML-powered applications. OpenTelemetry is perhaps even more helpful to IT teams because it can collect telemetry data from distributed systems. This data collection helps IT organizations tackle the troubleshooting, debugging, and managing of applications and host environments. It's a simple method for ITOps and DevOps teams to set up their code base for data collection and make adjustments that pave the way for growth.

Improving Operations with ML

Companies can use ML to drive business growth and secure operations for personnel and customers. But there are cautions to consider as businesses move forward with ML and AI-based solutions. Gartner highlights the extent to which ITOps, business units, and data scientists must work together to develop the right solutions.

"ML models often fail when data scientists initiate solutions without consulting the business unit owner or when the business team has unrealistic performance expectations from the model. Realizing the benefits of the ML model requires both teams to work together, while also owning several tasks throughout the process."[9]

IT organizations can employ ML and AI technologies in many ways to make it easier to gather and analyze data and help live up to digital transformation and security goals.

Infusing ML and AI Into Observability Platforms – More organizations are using observability platforms with the expectation that ML and AI will become more embedded soon. Observability is a progression of application performance monitoring (APM) data collection methods. It addresses the fast-growing, distributed nature of cloud-born applications and can enable better network monitoring and better APM. It can also help support the CX, improve employee productivity, and maintain the digital infrastructure. In addition, incorporating ML, AI, and automation into these platforms can decrease the time required to find and fix problems before they impact your business.

Automating Security – Securing machine and user data flowing through a company's infrastructure is a top priority of IT organizations. Algorithms can model the historical behavioral patterns and detect variances and anomalies compared to past patterns. This process could be further automated using ML and AI, ultimately enabling businesses to automatically block bad actors in near real-time.

Progressing Towards Self-Healing Systems – As ML and AI technology vendors work to improve and extend their platform capabilities, IT organizations will benefit, gaining increasingly more actionable insights and proactive functionality at a reasonable cost. This will set the foundation for integrated self-healing systems.

7 *State of Marketing, Sixth Edition,* Salesforce Research, 2020

8 *A technology blueprint for personalization at scale,* Sean Flavin and Jason Heller, McKinsey & Company, May 20, 2019

9 *Machine Learning Playbook for Data and Analytics Professionals,* Sumit Agarwal, Gartner, Inc. G00740149, August 2021

Use Case: Using ML in ITOps today

Big data and machine learning can combine to help automate IT operations using Artificial Intelligence for IT Operations, known as AIOps. For example, ITOps teams can use it for event correlation, anomaly detection, and causality determination. One of the advantages of this system is the ability to bring together data from disparate sources and apply AI algorithms to provide context and minimize the volume of irrelevant alerts. This is particularly helpful for companies that are spread out geographically. For example, a manufacturing firm based in the UK depends on the availability and performance of an IT and networking infrastructure located in data centers around the world.

The firm uses an IT infrastructure monitoring and observability platform with AIOps capabilities to consolidate views and gather analytics that help the company avoid duplicating monitoring data. This saves IT administrators time tracking down issues and providing the IT team with a single source of truth from a unified data set.

Likewise, support engineers use alert thresholds to identify and resolve issues before they require escalation. In one scenario, an engineer received a service call about an air conditioning outage at a data center. Because AIOps features were used to track the temperature of all the company's servers, an early alert prevented the A/C outage from turning into a data center outage, which would have affected hundreds of people across the organization.

Conclusion

Open-source tools like Python and TensorFlow can make it easier for IT organizations to capture and export metrics, traces, and logs. And solutions like OpenTelemetry will simplify the task of monitoring distributed cloud-born applications. In our next chapter, we'll show in detail how you can use ML to monitor your infrastructure.

Chapter 7: Using Machine Learning to Monitor Your Infrastructure

Machine Learning can make it easier to monitor all IT resources across your environment instead of using unique tools for each item in the stack. ML can help you develop an early warning system and even automate failure prevention.

This chapter:

- Describes how machine learning has automated and advanced monitoring
- Defines anomaly and root cause analysis
- Explains how IT organizations can start preparing to take advantage of ML for monitoring.

How Machine Learning Has Automated and Advanced Monitoring

Most monitoring products target one component of the IT stack—servers, databases, logs, cloud resources, and so on. These products were not constructed to monitor the entire network, combined contextually. As a result, changes to the environment (adding a cloud service, for instance, or a new application) frequently require an additional investment in a new monitoring tool. New tools, in turn, mean that IT staff must learn how to use them, or skilled people must be hired to use these products.

Surfacing Signals to Drive Rapid Action

An effective monitoring system will detect issues based on the signs and symptoms as they occur—patterns and anomalies in performance or log data, for instance—and warn users before major problems arise. Automatic anomaly detection based on ML or powerful algorithms can help detect issues and the root causes before extensive business impact occurs. This can be the first step to moving from a proactive to a predictive monitoring approach.

ML algorithms can support anomaly detection for a variety of data sets, including IT metric and log data, root cause analysis, automatic correlation, and dynamic thresholds. IT organizations can use ML to quickly examine vast volumes of monitored data, such as log events, surface the most important information, and then add context to take appropriate action.

In addition, ML can support advanced log analysis, correlating log data, for instance, with metrics and alerts that go beyond simple notifications and provide adequate information to determine why issues are occurring. An early warning system like this can help prevent incidents that can adversely impact a business by processing the signal through a rule-based engine tied into a robust automation framework. As a result, your IT team will be better informed and proactive and experience shorter mean times to recovery.

Preventing Down Time

Companies with frequent outages and brownouts can pay upwards of 16 times the cost[10] compared with organizations that experience less downtime. These businesses also need twice as many staff to troubleshoot problems and spend twice as much time resolving those issues.

10 *IT Outage Impact Study,* LogicMonitor 2019

The impact is not just on IT but the company and its brand. Consider associated lost revenue, poor customer experience, and harm to brand reputation. Downtime reduces the time and energy available to devote to strategic initiatives and drive business growth. IT teams that have to spend most of their time diagnosing root causes and resolving problems cannot be proactive and focus on major improvements.

An early warning system based on ML can help solve the problem of traditional monitoring, which focuses on static alerting and analysis configurations. Systems like this can help IT teams efficiently run dynamic, distributed environments that help support business goals and maintain a positive customer experience—without missing a beat because of downtime.

Minimizing Alert Fatigue

Broadcom reports that 47% of organizations polled receive a staggering 50,000+ alerts[11] each month. Unfortunately, this tsunami of notifications can overwhelm IT professionals, leading to ignored alerts, slow responses, and incident management failures. This is known as "alert fatigue."

A warning system that uses ML can ease alert fatigue by classifying alerts and workloads more effectively and surfacing the most relevant alerts from many data types. The system can detect the normal performance range for technical and business metrics and generate alerts based on anomalies by using dynamic thresholds. It can even establish alerts on historical performance and advanced algorithms, which helps IT organizations avoid exhaustion and surface anomalies sooner.

ML-based systems can make the deluge of data that your environment produces more manageable and actionable. And the IT team will have the capacity to support business goals and control costs.

Anomaly Detection and Root Cause Analysis

The root cause analysis (RCA) feature in a warning system is designed to identify the cause of an issue, giving ITOps staff the ability to focus on solving the issue quickly rather than spending precious time looking for the problem. When organizations mix RCA with an ML or AI platform's ability to monitor just about anything (containers, cloud environments, network, and so on), IT teams can reduce downtime even for highly complex hybrid infrastructures.

Identifying the root cause is the first step in taking a more proactive approach to keeping IT infrastructure healthy. A failure prevention system that uses ML also allows IT organizations to automate actions that remediate the root cause issue, essentially identifying and predicting anomalies and then automatically fixing and preventing them. This equates to reduced downtime plus more time that ITOps teams can spend on innovating and transforming the business.

Dynamic Thresholds Set the Stage for Automatic Remediation

Dynamic thresholds are built on ML-based algorithms focusing on anomaly detection based on the rate of change and seasonality, along with algorithms to contextualize issues. These algorithms automatically detect the normal performance range for any metric—whether it's a technical or business metric—and accurately send notifications based on values outside of this range that are considered anomalies.

Because dynamic thresholds and the resulting alerts are algorithmically determined based on the data point's history, they are well suited for data points where static thresholds are hard to identify, such as monitoring the number of connections, latency, and other criteria. Dynamic thresholds are also useful in situations where acceptable data point values aren't necessarily uniform across an environment.

Dynamic Thresholds are a Requirement for Proactive IT

While static thresholds are too complex and time-consuming for people to manage manually, dynamic thresholds function well in fast-changing environments. Unlike static thresholds, dynamic thresholds are calculated by anomaly detection ML algorithms and are continuously trained by a data point's recent historical values. By using dynamic thresholds, IT teams can calculate the thresholds to be set and continually adapt to environmental changes generating alerts only when unusual performance is detected.

11 *Reducing Alert Fatigue: How Your Automation COE Can Help*, Viki Paige, Broadcom, 2020

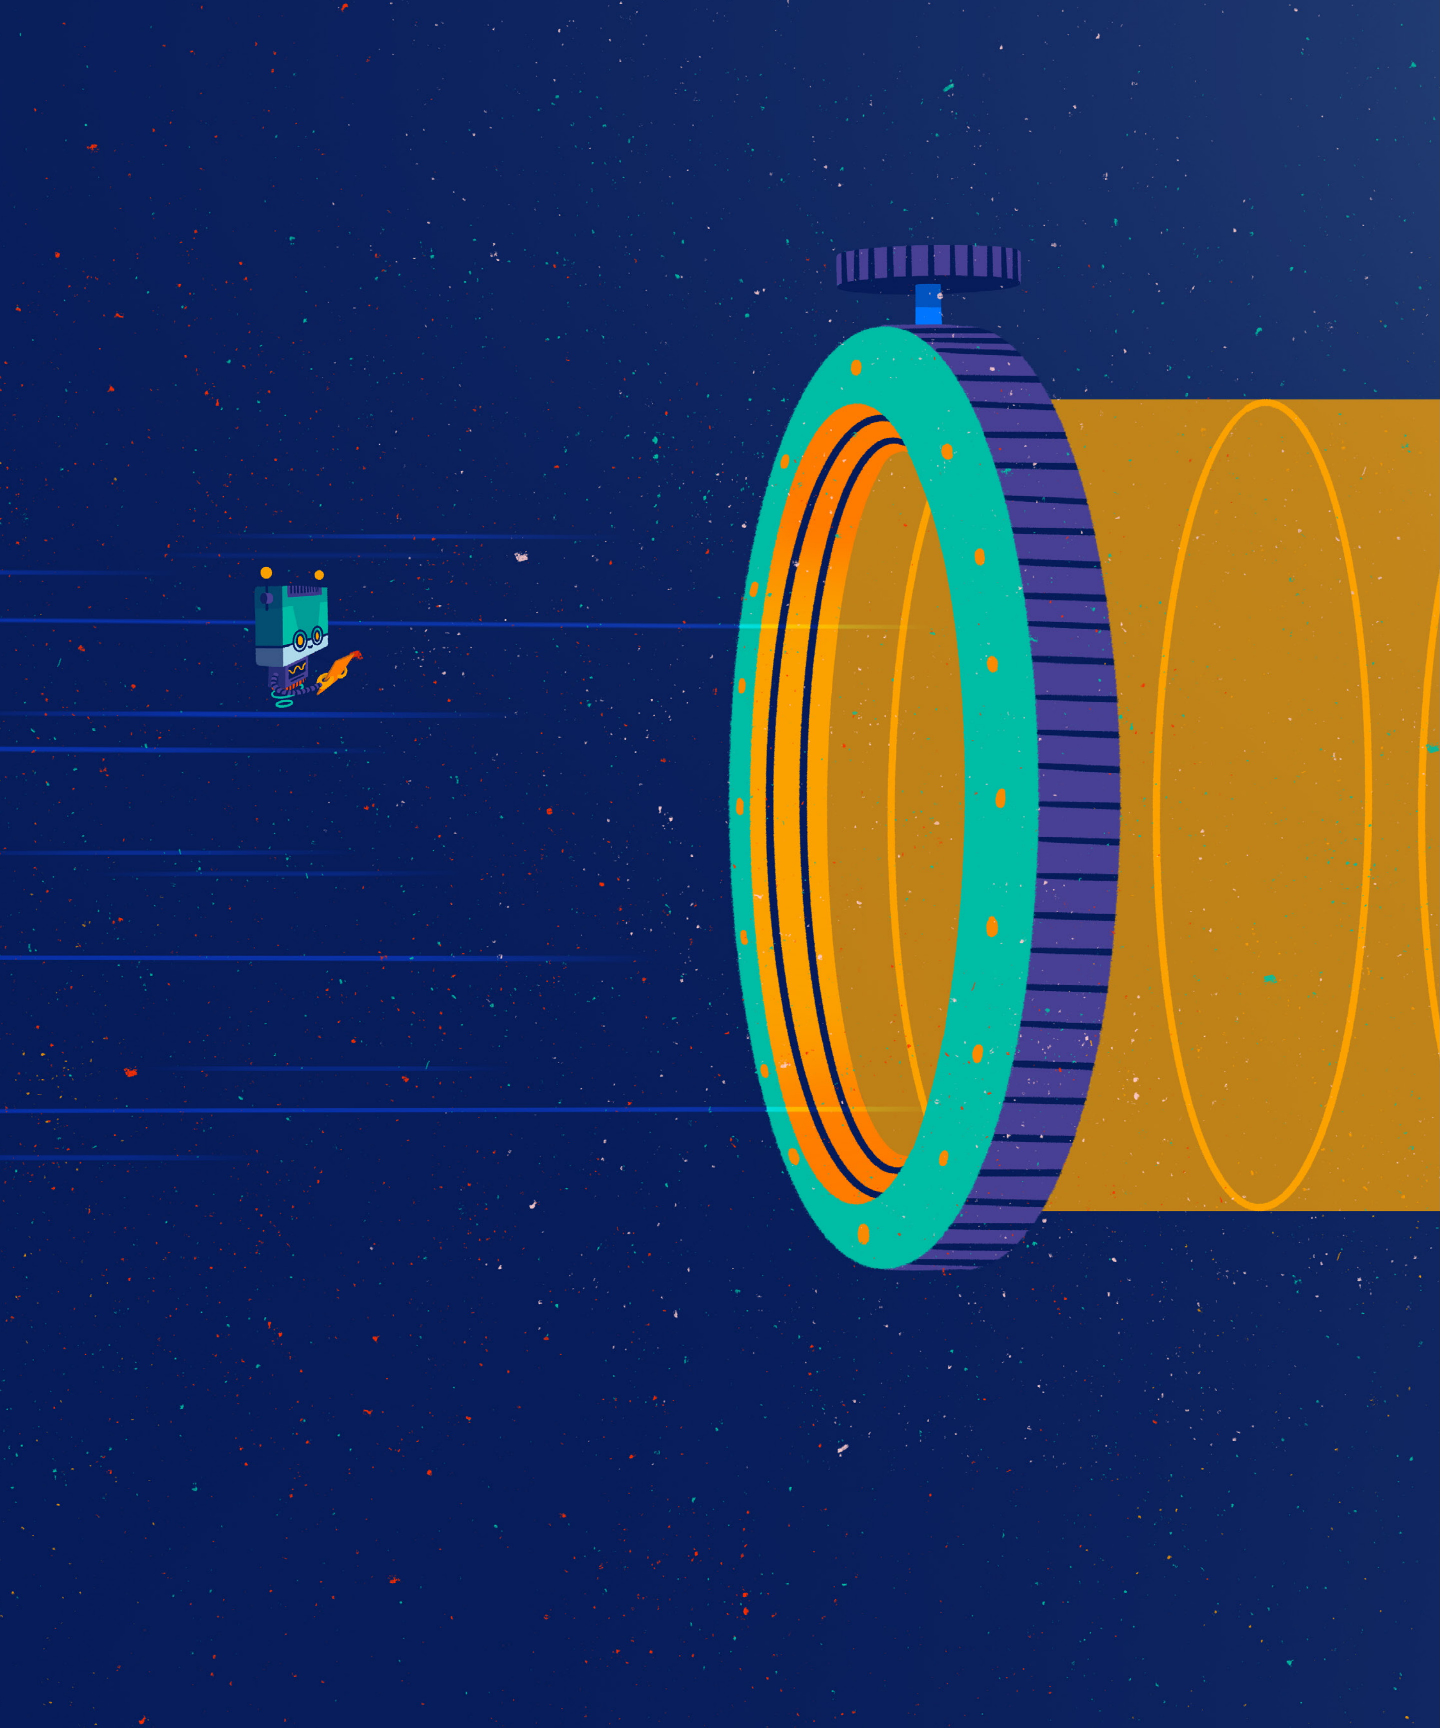

What IT Organizations Can Do to Set Up for the Future

ML provides strong advantages in rapid troubleshooting, but its capabilities go far beyond tracking and monitoring. It is also an effective way to predict future trends for a business's monitored infrastructure, using past performance as the basis.

ML can enable IT to address its increasingly strategic responsibilities and challenges, helping teams to discover and resolve issues more quickly and make proactive solutions. After the fundamental elements of an early warning system (anomaly detection, dynamic thresholds, root cause analysis, and forecasting) are set up, IT teams can deliver superior service quality and availability, collaborate and innovate better, and keep their technology aligned to the business outcomes they want most.

Driving Innovation with DevOps

Like ITOps teams, DevOps teams have to manage fast growth and complicated infrastructures. Traditional static thresholds cannot offer the context and agility needed to manage these environments, so modern DevOps teams rely on advanced ML algorithms.

ML algorithms can predict issues before they occur, averting severe problems that can impact business. To minimize excess noise, it also suppresses notifications on problems that don't require action. DevOps engineers can use the ML algorithms to troubleshoot issues as they occur, helping them to determine if a perceived issue is normal and understand if it was caused by or connected to a change in the environment.

Striving for Continuous Optimization

ML can help IT professionals gain insight and control that would not be possible using traditional monitoring and management approaches. However, it can help IT teams do more than simply respond to new problems. ML can help ITOps to:

- Define actions, such as the execution of a script.
- Set a predefined action in response to an alert.
- Automate those predefined actions in response to specific notifications using a rules-based engine.

Conclusion

As ML capabilities are more broadly used, organizations can collect additional data and apply these learnings to further enhance and automate monitoring and resolution. As a result, IT organizations can analyze a vast array of big data, examine patterns, make predictions, and drive business.

About LogicMonitor®

Monitoring unlocks new pathways to growth. At LogicMonitor®, we expand what's possible for businesses by advancing the technology behind them. LogicMonitor seamlessly monitors infrastructures, empowering companies to focus less on problem-solving and more on evolution. We help customers turn on a complete view in minutes, turn the dial from optimization to innovation and turn the corner from sight to vision. Join us in shaping the information revolution by visiting LogicMonitor.com.

LogicMonitor