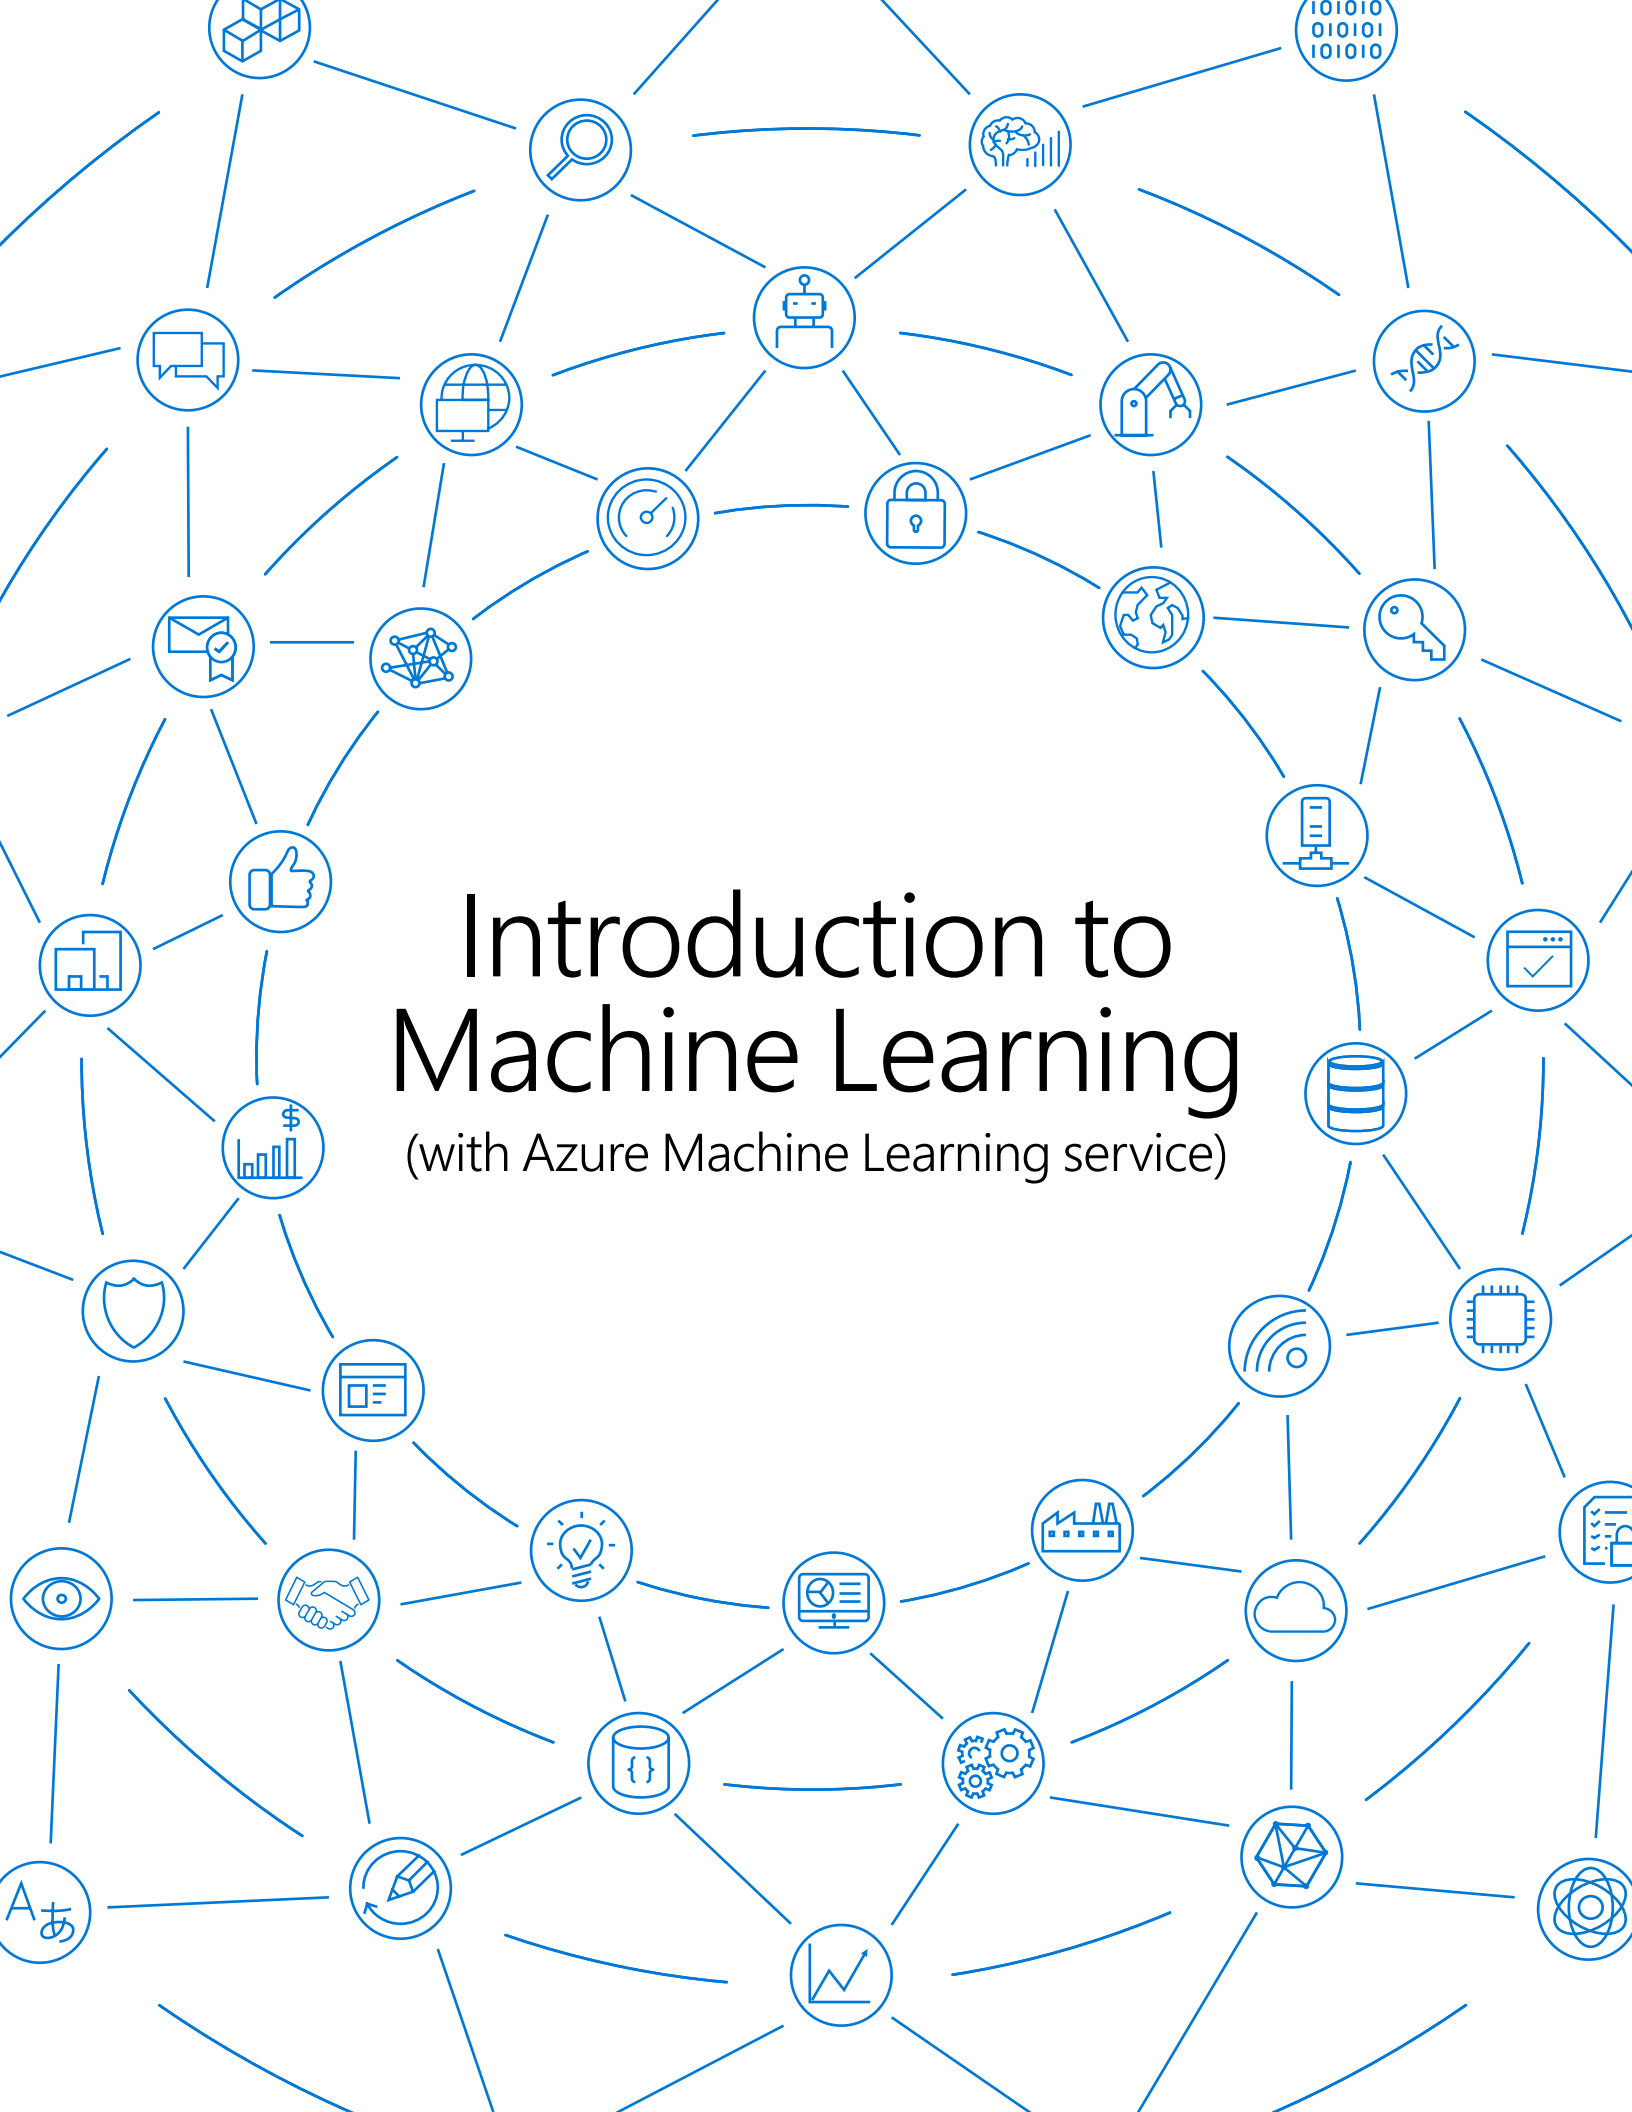

Introduction to Machine Learning

(with Azure Machine Learning service)

Contents

Who is this e-book for?

This e-book is designed to help business decision makers gain a high-level understanding of machine learning and assist IT decision makers in evangelizing the subject to stakeholders within their organization.

What is **machine learning?**

Machine learning (ML) is a rapidly developing field with a huge value proposition for organizations across industries.

It evolved from **artificial intelligence (AI)** decades ago and with new technological innovations, we are beginning to realize machine learning's full potential.

As new technologies like machine learning become critical to organizations, business decision makers and employees who interact with them need to expand their knowledge bases to achieve what is required for digital transformation.

Let's start with a basic definition of machine learning.

What is machine learning?

Data science technique: Machine learning is used to mine predictions from statistical models. Deep data science practices can also be applied that include modeless techniques and automatic programming.

Behaviors: Machine learning can predict future behaviors based on past behaviors, such as surfacing suggestions based on previously identified patterns.

Machine learning is a **data science technique** that allows computers to use existing data to forecast future **behaviors**, **outcomes**, and **trends**.

Outcomes: Complex models are constantly being tested to predict anything from sports outcomes to poverty outcomes to health outcomes, and everything in between.

Trends: Market index prices, forecasting the spread of diseases,[1] and other trends can be identified by recognizing patterns and existing trends in inputs.

By using machine learning, computers learn without being explicitly programmed.[2] A machine learning system uses data from sources such as apps, sensors, networks, and devices to build its own logic to solve a problem or extract insight (observations).* There are different methods for using algorithms that find patterns and insights.

1. https://www.beckershospitalreview.com/healthcare-information-technology/machine-learning-and-ai-are-taking-healthcare-by-storm-are-you-ready-to-rumble.html
2. https://docs.microsoft.com/en-us/azure/machine-learning/service/overview-what-is-azure-ml
*https://www.gartner.com/binaries/content/assets/events/keywords/catalyst/catus8/preparing_and_architecting_for_machine_learning.pdf

Machine learning subdisciplines

Supervised learning is about having access to labeled data that is used as a guide to teach the algorithm what conclusions to arrive to.[3]

Supervised learning requires that the machine learning algorithm's possible outputs are already known, that target values are included in the data used for training and testing the algorithm.[4] The algorithm looks for patterns in the attributes of the data samples to try to predict target values correctly.[5] After the algorithm has discovered which features allow it to make accurate predictions for the training data, its performance is confirmed by testing whether it can also make accurate predictions for the held-out test data.[6] The trained algorithm can then be used to make predictions on new data for which target values are unknown.

Example:

Historical stock prices can be used to make predictions for future prices.[7]

There are several specific types of supervised learning that are represented within Azure Machine Learning: classification, regression.[8]

Supervised learning

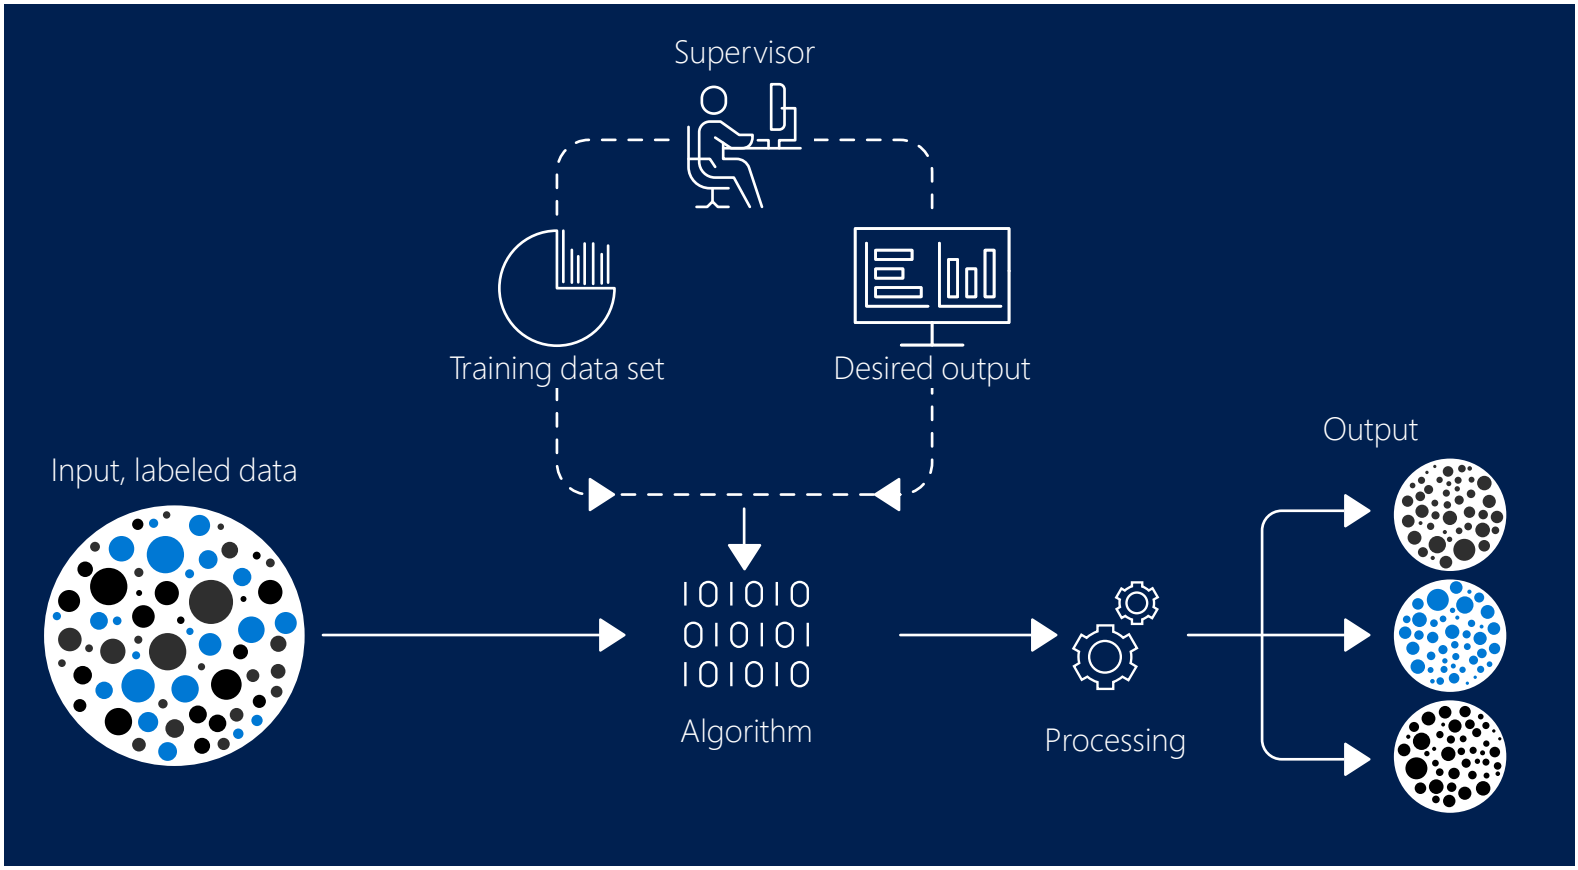

3-4. https://www.datascience.com/blog/supervised-and-unsupervised-machine-learning-algorithms
5-8. https://docs.microsoft.com/en-us/azure/machine-learning/studio/algorithm-choice

Machine learning subdisciplines

Unsupervised learning refers to the training of algorithms using information that is neither classified nor labeled and allowing the algorithm to act on that information without guidance.[9]

Here the task of the algorithm is to group data samples according to patterns in similarities and distances among them. Unlike supervised learning, no prior teaching is provided, that means no training will be given to the machine. Therefore, the machine is restricted to find the hidden structure in unlabeled data by itself.

Example:

Healthcare uses unsupervised learning to detect causality and to identify correlations humans may miss by inputting health data like blood pressure, heart rate, weight, prescriptions, etc.[10]

Unsupervised learning

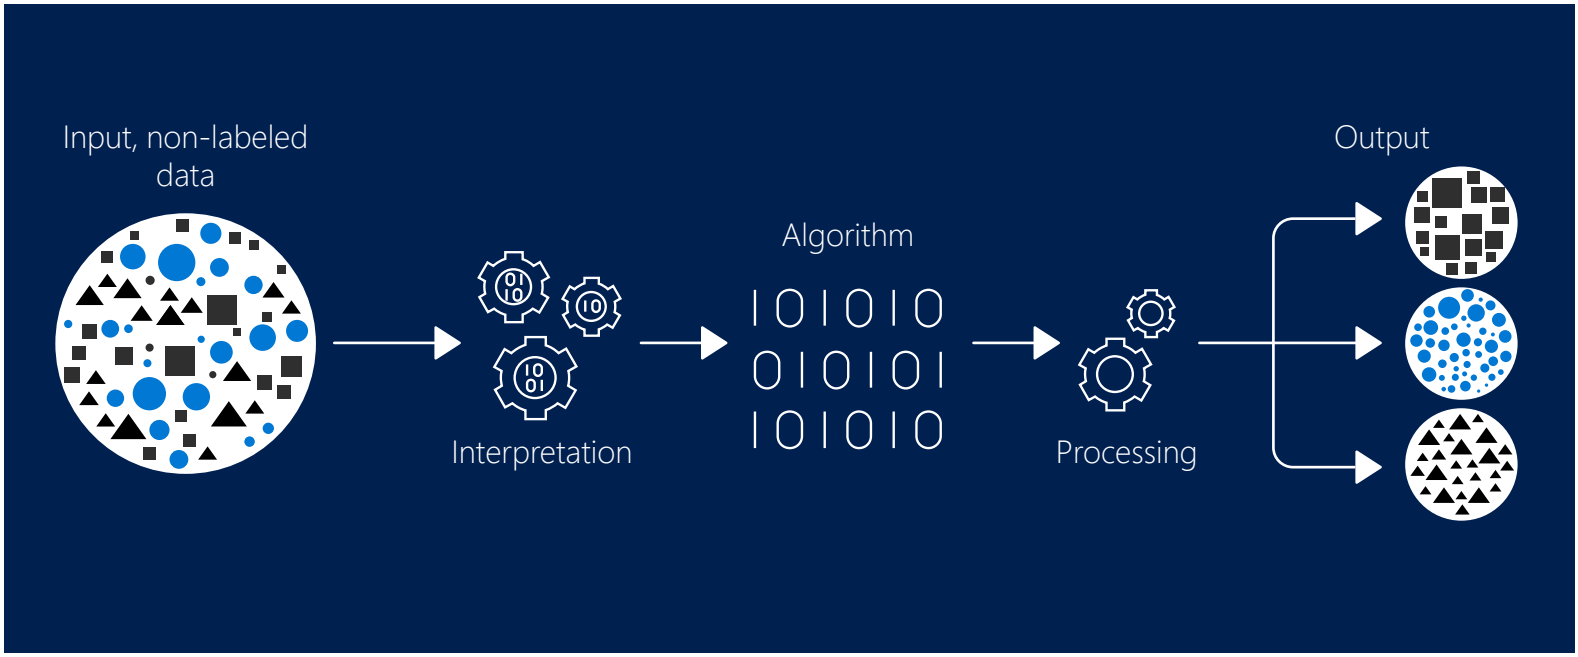

9. https://www.geeksforgeeks.org/supervised-unsupervised-learning/
10. http://people.csail.mit.edu/dsontag/courses/mlhc_summer18/day2/causal_inference.pdf

Machine learning subdisciplines

Reinforcement learning is when a machine learning algorithm interacts with its environment, either simulated or real, exploring different strategies to determine the optimal sequence of choices to maximize reward.

Reinforcement learning methods resemble how humans and animals learn: The algorithm tries various behavioral strategies, trying to maximize the reward it receives in response. Each choice the algorithm makes when it encounters a new data point determines which data point it will encounter next.[11] A sequence of good decisions is reinforced with rewards. With experience the algorithm learns to produce the optimal sequence of decisions within a given context.[12]

Example:

Enabling computers to learn to play games or drive vehicles.* Reinforcement learning is common in robotics, where the set of sensor readings at one point in time is a data point, and the algorithm must choose the robot's next action.[13]

Reinforcement learning

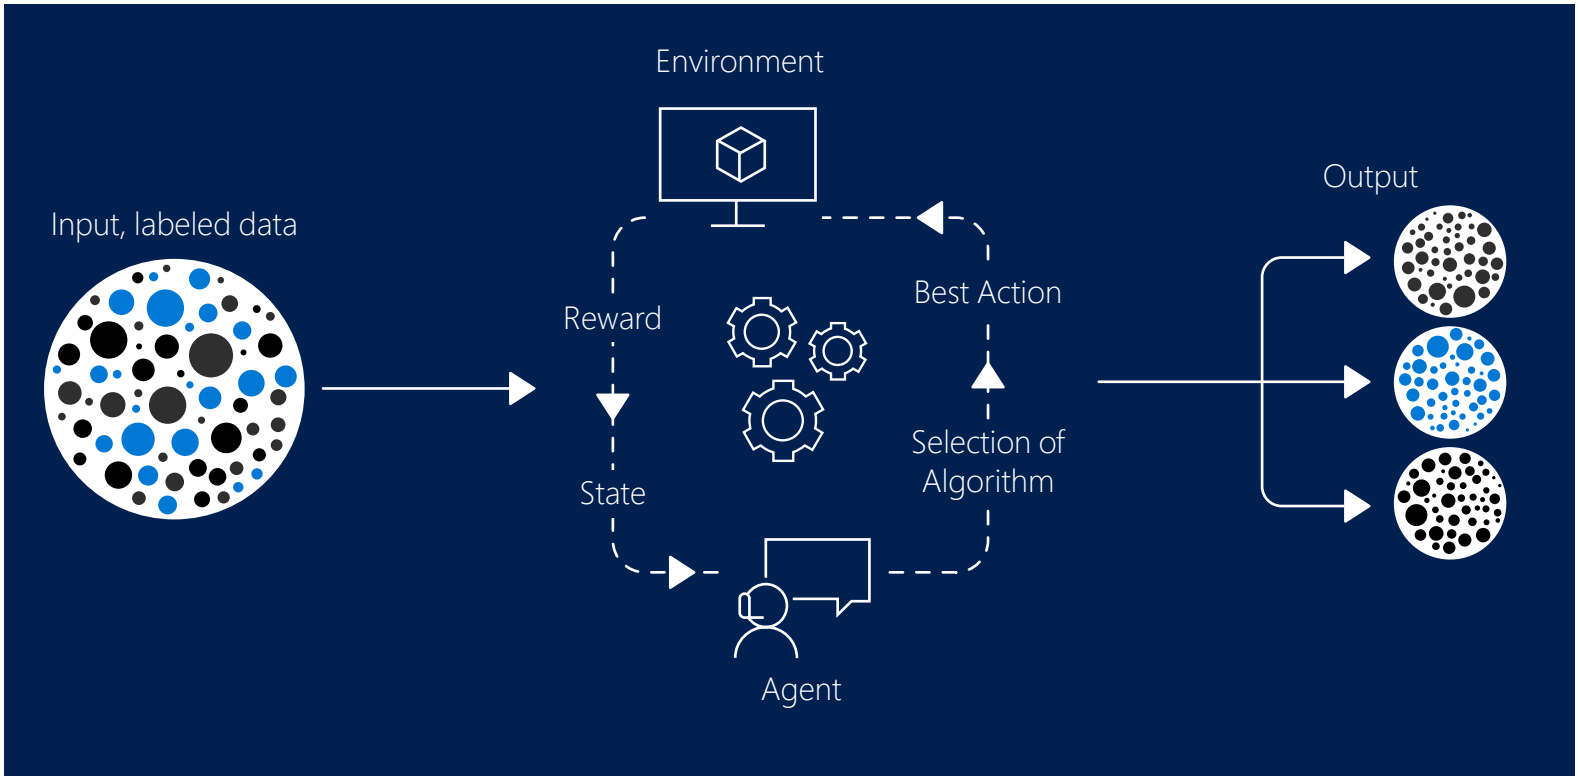

11-12. https://docs.microsoft.com/en-us/azure/machine-learning/studio/algorithm-choice
*https://www.gartner.com/binaries/content/assets/events/keywords/catalyst/catus8/preparing_and_architecting_for_machine_learning.pdf
13. https://docs.microsoft.com/en-us/azure/machine-learning/studio/algorithm-choice

Machine learning subdisciplines

Deep learning uses what's called a neural network architecture to do its learning. The neural network architecture mimics how a brain works rather than using traditional statistical frameworks. In this way, reinforcement learning could be placed under the deep learning umbrella.

The deep part of deep learning refers to the number of hidden layers in the neural network architecture, which are usually several and going up to one-hundred fifty layers or more. There are many different types of deep learning neural networks: convolutional, recurrent, multilayer perception and beyond.

Example:

Speech recognition, image recognition, and Natural Language Processing (NLP) are all typically achieved through the use of neural network architectures.

Artificial intelligence

VS

Machine learning

VS

Deep learning

Artificial intelligence vs machine learning vs deep learning

As a business decision maker, it can sometimes be difficult to understand the contrasts between artificial intelligence, machine learning, and deep learning. Here's a breakdown of what differentiates the three topics so that you can gain a better understanding of how they affect your business.

Artificial intelligence is the ability of a system to do tasks normally performed by humans. It is an umbrella term in which both machine learning and deep learning fall under.

Think of **machine learning** as a type of artificial intelligence, where its human task is to learn from data. It is a statistical method that uses features and attributes in order to describe an instance in a data set.

Deep learning draws inspiration from how the brain processes information. It is called deep learning, because these models are architectures with several stacked layers on top of each other. Just as in the human brain, deeper layers form more abstract interpretations of the input data. This approach is especially useful when working with data without structured attributes or features, leaving it up to the algorithm to come up with its own interpretation of what the input represents. In comparison to most conventional machine learning techniques, deep learning requires massive amounts of compute power, more training time, and very large datasets. Deep learning techniques have been around for many years, but only because of recent breakthroughs in the size of available datasets and computational resources has it become possible to apply them to hard real-world problems.

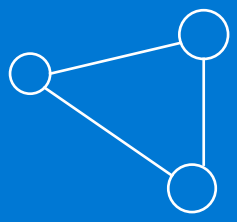

Machine Learning: The Value it Brings and case studies from 3 key industries

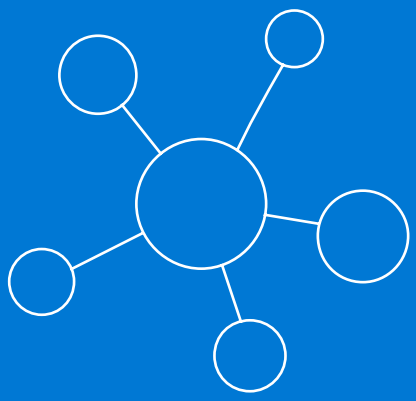

Machine learning:
The value it brings

One of the most intriguing aspects of machine learning's value proposition is that it doesn't require as much advanced programming as other methods to gain useful insights. This can help organizations cut down on time and costs, while gaining a better understanding.

The two key pain points for organizations and developers interested in deploying machine learning services are that they require both large amounts of raw data, as well as high computing power to perform the computations needed to learn.*

With powerful storage and data offerings, and world-class compute offerings, Azure is a great place to start your machine learning journey to meet your business objectives while alleviating some of the pain points that have been associated with machine learning in the past.

*https://www.gartner.com/binaries/content/assets/events/keywords/catalyst/catus8/preparing_and_architecting_for_machine_learning.pdf
14. https://www.forbes.com/sites/bobevans1/2018/06/04/microsoft-ceo-satya-nadella-on-the-extraordinary-potential-of-ai/#986818e162ff

Use cases and case studies from some key industries

To comprehend how these methods can be evangelized within your own company, it is worthwhile to see how they are used in real-world settings. This section will provide a few key industries that are already implementing artificial intelligence, machine learning, deep learning, and automated machine learning. We'll look at a few use cases that these industries benefit from, and then dive deeper into some companies within each industry that are utilizing Azure Machine Learning services to meet their own business objectives.

Healthcare and insurance

There are many use cases being tested in the healthcare and insurance industries. Below, we'll take a closer look at how leading life insurance company TAL (leading life insurance company in Australia) used AI to improve quality assurance and the customer experience.

DNA sequences	FAST-Q BAM SAM VCF Expression	Genomics and precision medicine Single cell sequencing Biomarker, genetic, variant, and population analytics ADAM and HAIL on Databricks	**Faster innovation for drug development**
Real-world analytics	HL7/CCD 837 Pharmacy Registry EMR	Clinical and claims data Claims data warehouse Readmission predictions Efficacy and comparative analytics Prescription adherence Market access analysis	**Improved outcomes and increased revenue**
Image processing	MRI X-Ray CT Ultrasound	GPU image processing Graphic intensive workloads Deep learning using Tensor Flow Pattern recognition	**Diagnostics leveraging machine learning**
Sensor data	Readings Time series Event data	IoT device analytics Aggregation of streaming events Predictive maintenance Anomaly detection	**Predictive analytics transforms quality of care**
Social data listening	Social media Adverse events Unstructured	Social analytics Real-time patient feedback via topic modeling Analytics across publication data	**Improved patient communications and feedback**

Healthcare and insurance

Traditionally, TAL's quality assurance team could only review a randomly selected 2-3 percent of cases.[15] In an effort to increase the number of cases they can review, TAL is embracing AI to improve quality assurance and the customer experience. Using Azure Machine Learning service, the team is now able to review 100 percent of cases.[16]

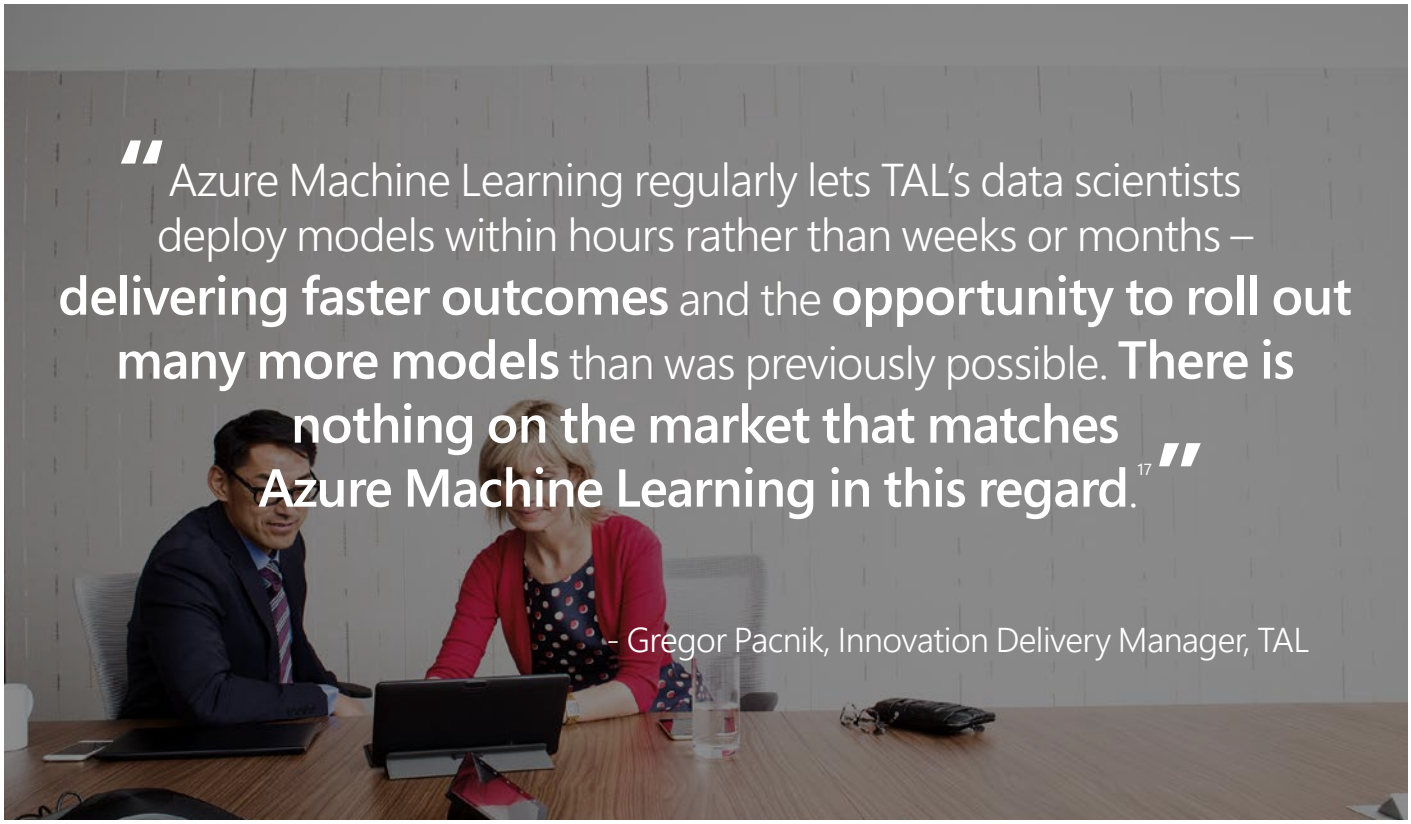

15-17. https://azure.microsoft.com/en-us/blog/azure-machine-learning-service-a-look-under-the-hood/

Oil, gas, and energy

Machine learning and AI are revolutionizing the oil, gas, and energy industry. Below, you'll see how the global oil and gas company BP is utilizing AI and machine learning to analyze their raw data quickly and with increased accuracy.

Upstream optimization, maximize well life	Field data Asset data Demographics Production data	Digital oil field/ oil production Production optimization Integrate exploration and seismic data Minimize lease operating expenses Decline curve analysis	**Faster innovation for revenue growth**
Grid operations, asset inventory optimization	Sensor stream data UAV images Inventory data Production data	Industrial IoT Pipeline monitoring Preventive maintenance Smart grids and microgrids Grid operations, field service Asset performance as a service	**Improved outcomes and increased revenue**
Supply-chain optimization	Transaction data Demographics Purchasing history Trends	Supply-chain optimization Trade monitoring, optimization Retail mobile applications Vendor management - construction, transportation, truck, and delivery optimization	**Optimizing supply-chain with machine learning**
Risk optimization	Sensor stream data Transport Retail data Grid production data Refinery tuning parameters	Safety and security Real-time anomaly detection Predictive analytics Industrial safety Environment health and safety	**Predictive analytics transforms safety and security**
Recommendation engine	Clickstream data Products Services Market data Competitive data Demographics	Sales and marketing analytics Fast marketing and multi-channel engagement Develop new products and monitor acceptance of rates Predictive energy trading Deep customer insights	**Improved customer engagement with machine learning**

Oil, gas, and energy

BP is a global oil and gas company that operates in 70 countries. They needed to overcome uncertainty in the exploration and production of hydrocarbons. There are close to 200 qualitative and quantitative features that BP records to summarize any reservoir. Knowing which of these is most important when predicting recovery factor is a challenge. They often relied on human bias, or other reservoirs with similar geological or geographical settings on which to base comparison studies, but what if other features not considered here could bring more insights to the problem and improve their predictability?

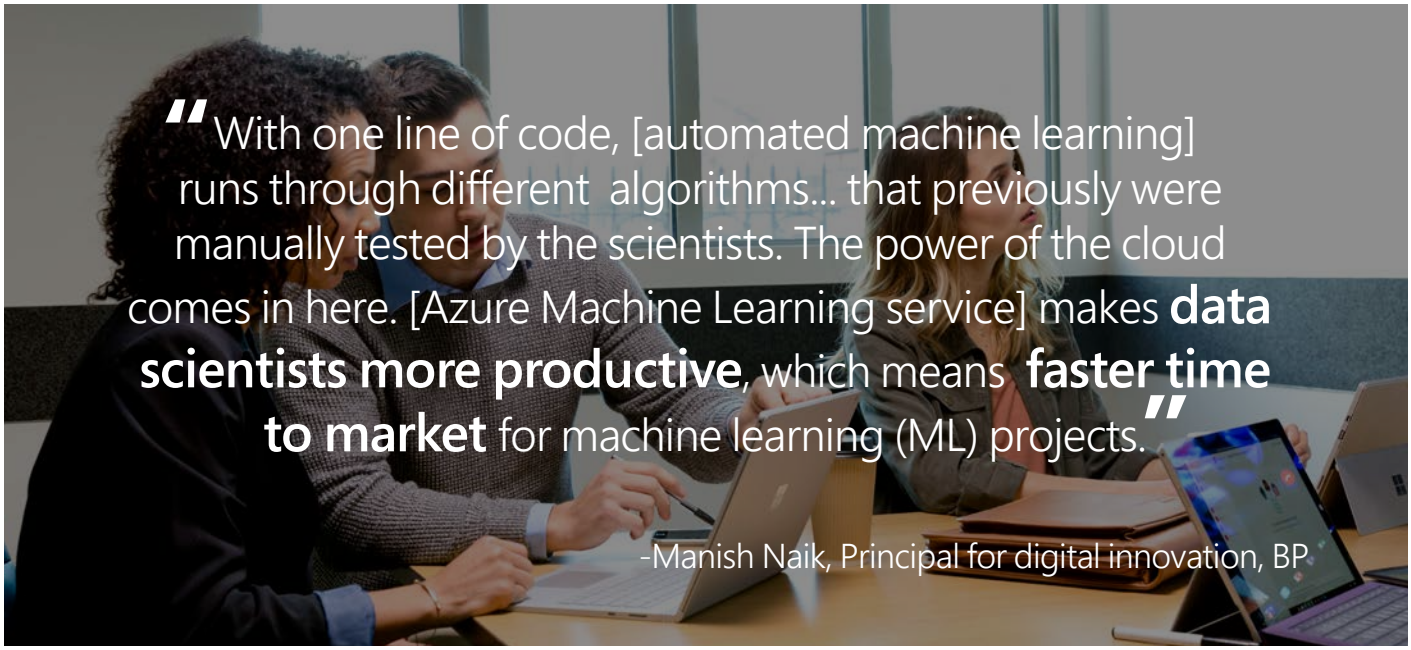

To address this problem, BP has worked with Microsoft to develop a machine learning model that can rapidly predict recovery factor. The model is now deployed and available globally for BP's internal community of reservoir engineers and geoscientists. They can use the tools to test and screen multiple assumptions and scenarios, comparing with parallel studies to help them make better-informed decisions faster.

Retail

The retail industry is finding some interesting ways to make use of machine learning. Find out how global retailer, ASOS, is improving their customer experience and innovating faster.

Recommendation engine	Customer profiles Shopping history Online activity Social network analysis	Next best and personalized offers Customer 360/consumer personalization Right product, promotion, at right time Multichannel promotion	**Faster innovation for a better customer experience**
Effective customer engagement	Shopping history Online activity Floor plans App data	Store design and ergonomics Path to purchase In-store experience Workforce optimization	**Improved consumer outcomes and increased revenue**
Inventory optimization	Demand plans Forecasts Sales history Trends Local events Weather patterns	Data-driven stock, inventory, ordering Predict inventory positions and distribution Fraud detection Market basket analysis	**Omni-channel shopping experience with machine learning**
Inventory allocation	Demographics Buyer perception Consumer research Market/ competitive analysis	Assortment optimization Economic modeling Optimization for foot traffic, online interactions Flat and declining categories	**Predictive analytics transforms growth**
Consumer engagement	Historical sales data Price scheduling Segment level price changes	Real-time pricing optimization Demand elasticity Personal pricing schemes Promotion events Multi-channel engagement	**Improved customer engagement with machine learning**

Retail

ASOS, a global fashion retailer, needs the ability to scale up and out, innovate quickly, and incorporate new customer engagement models to improve their customers' experiences. Azure and Azure Machine Learning service gave ASOS the ability to do just that.

18. https://customers.microsoft.com/en-us/story/asos

Democratizing AI with automated machine learning

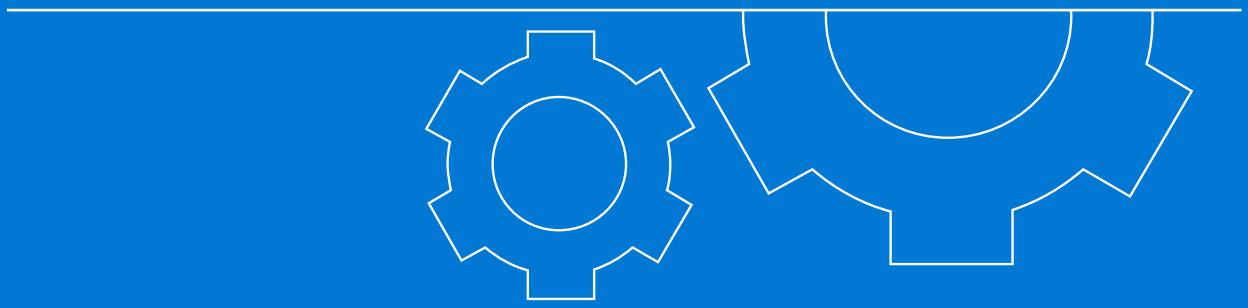

Democratizing AI with Automated Machine Learning

Automated machine learning makes machine learning more accessible across an organization.[19] It gives data scientists an automatic algorithm selection process and the ability to run automated, systematic processes on raw data. There are many repetitive, tedious and manual tasks associated with traditional machine learning. Automated machine learning automates many of these tasks so that the training aspect is not the only automated component. Automated machine learning can also help remove much of the risk of bias that manual model selection poses.

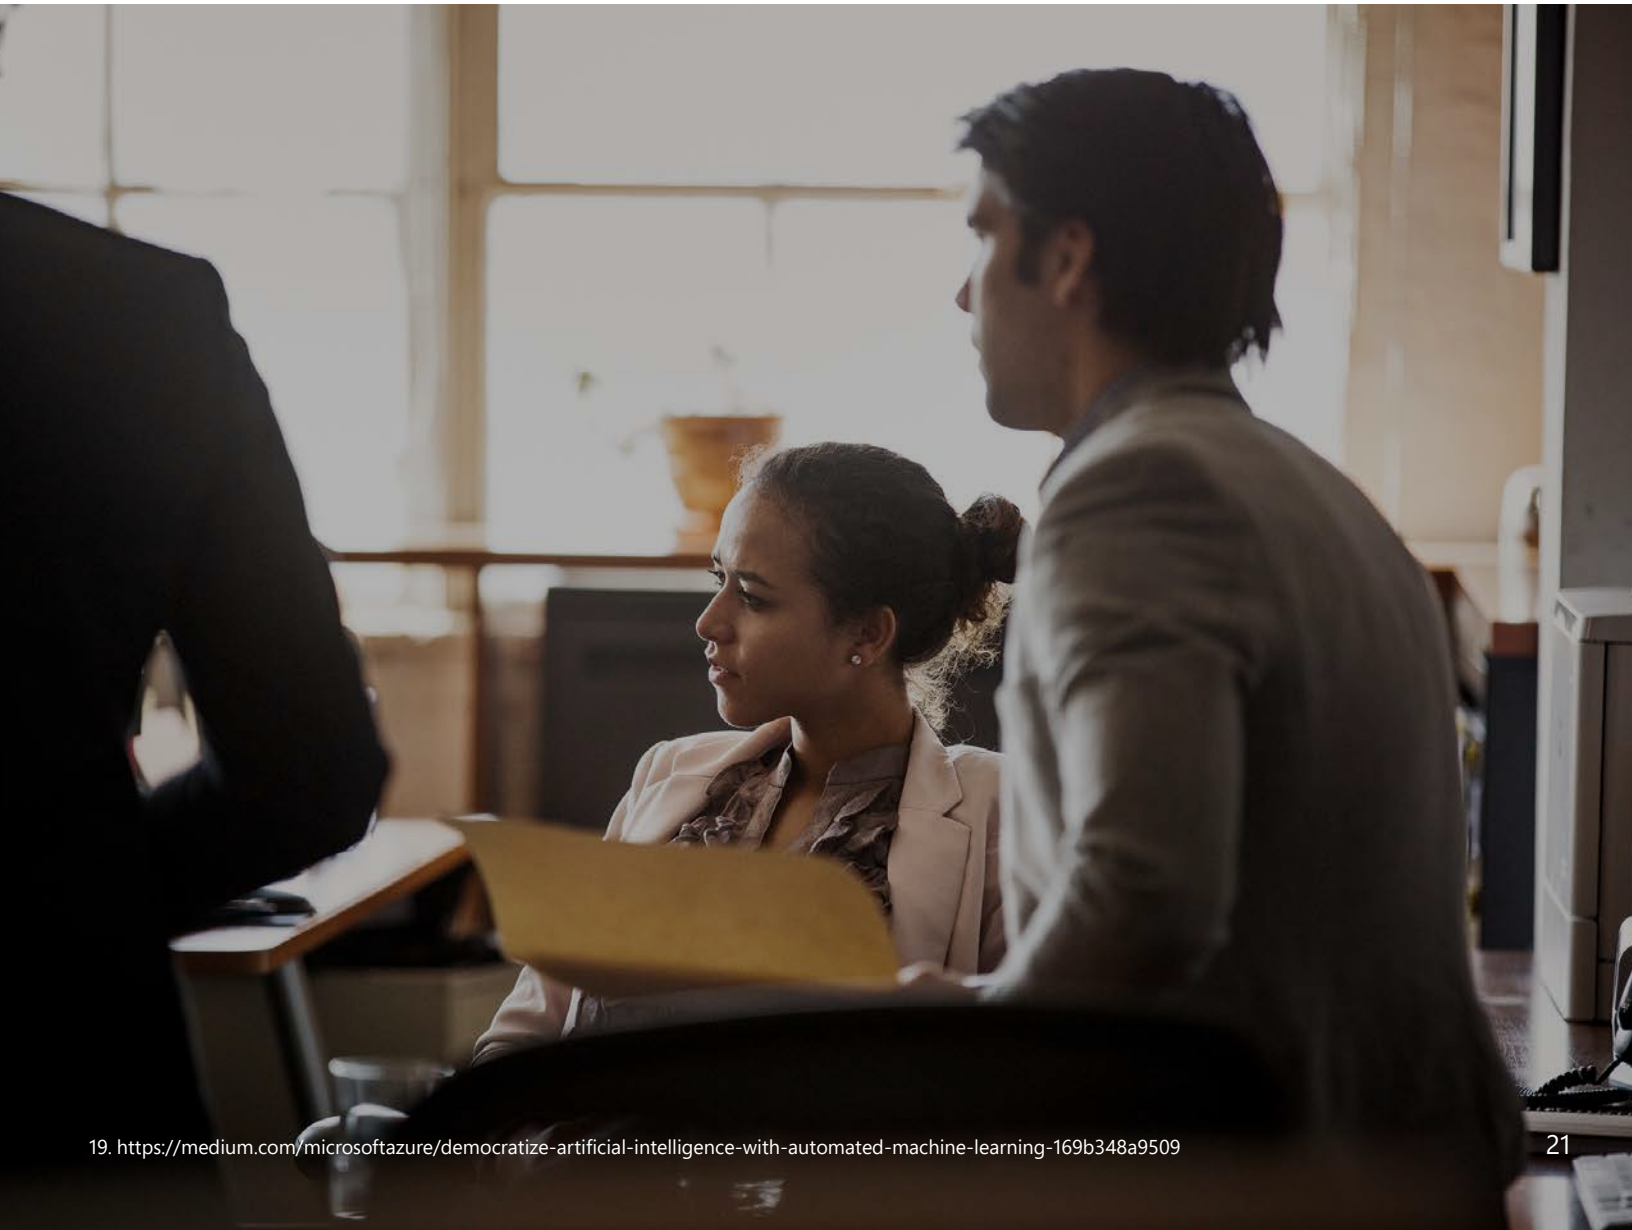

19. https://medium.com/microsoftazure/democratize-artificial-intelligence-with-automated-machine-learning-169b348a9509

Automated machine learning – top objectives

Simplify and accelerate building, training, and deployment

Slow tuning times, identifying suitable algorithms, and deploying to the cloud or edge are all challenges many developers face when deploying machine learning services. Automation allows developers to tune hyperparameters fast and identify the right algorithm to use and deploy to the cloud or edge with one click.[20]

Improve productivity and reduce costs

The typical machine learning life cycle consists of four parts: business understanding, data acquisition, modeling, and operationalization. Automated model selection and hyperparameter tuning reduces building time from weeks or months to days.

Automate analytical queries

Analytic queries are increasingly generated via search, natural language, or even generated automatically. Automated analytical queries allow data scientists to stop spending time on searches in order to focus more time on their business objectives.[21]

20. https://azure.microsoft.com/en-us/blog/new-automated-machine-learning-capabilities-in-azure-machine-learning-service/
21. https://searchenterpriseai.techtarget.com/feature/Automating-machine-learning-puts-analytical-models-on-autopilot

Automated machine learning – top objectives (continued)

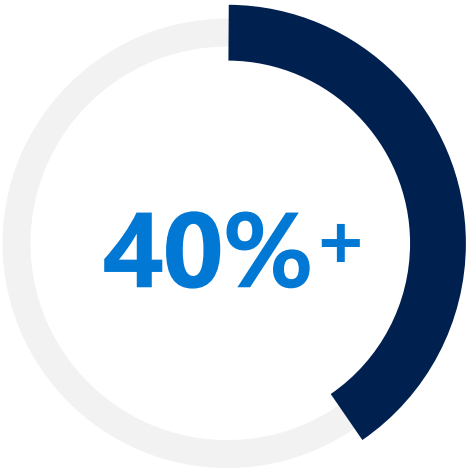

of data science tasks will be automated by 2020, resulting in **increased productivity** and **broader usage of data and analytics** by data scientists, according to Gartner, Inc.* There are many ways that data scientists can benefit from automated machine learning.

Reduce time-consuming tasks[22]

Many of the time-consuming machine learning tasks, such as hyperparameter tuning and model training, testing and selection, which typically take up as much as 80 percent of a data scientist's time, can increasingly be automated.

Explore more options[23]

It's virtually impossible for a data scientist to test every possible solution to a machine learning problem. Automated machine learning allows data scientists to rapidly test algorithms against a chosen data set and find the best one to use for that problem.[21]

Measure model performance[24]

With automated machine learning, it's possible to evaluate the importance of model features and measure the relationships between those features and a model's outputs automatically. This allows a data scientist to more quickly identify potential areas of improvement.

Achieve quicker time to value[25]

Automated machine learning tools can streamline the process of choosing algorithms, running them, and comparing the results, ultimately allowing a data scientist to put the right model into production much sooner.

*https://www.gartner.com/en/newsroom/press-releases/2017-01-16-gartner-says-more-than-40-percent-of-data-science-tasks-will-be-automated-by-2020

22-25. https://cdn2.hubspot.net/hubfs/532045/How%20to%20Use%20Automated%20Machine%20Learning%20.pdf?t=1500943259879

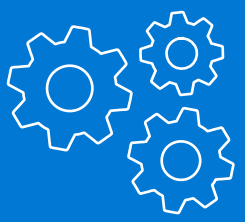

Azure Machine Learning service

Azure Machine Learning service

Scale easily on the cloud and to the edge

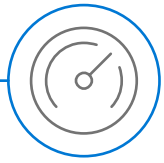

Brings the Azure approach to data science with enterprise-ready security, compliance and virtual network support. Users can use any data and deploy models anywhere with Azure Kubernetes Service (AKS) and lower costs with autoscaling cloud compute.[2]

Power productivity and ease of use

Users can access Azure Machine Learning service from any Python environment and use their favorite frameworks and tools. They can leverage multi-cloud interoperability with ONNX, rapidly build and deploy models across skill levels with automated machine learning, and quickly build Interactive dashboards with built-in Power BI integration.

Accelerate machine learning with MLOps

MLOps is DevOps for machine learning. With machine learning pipelines, Azure Machine Learning service brings collaboration and simplification across the end-to-end lifecycle, from data preparation to deployment and monitoring.

Azure Automated Machine Learning meets all the customer objectives listed above along with a few other considerations:

Automated Machine Learning

Data Professionals with all skill levels can easily use Automated ML to identify an end-to-end machine learning pipeline for any problem, achieving higher accuracy while spending less time. It enables a significantly larger number of experiments to be run, resulting in faster iteration towards production-ready intelligent experiences.

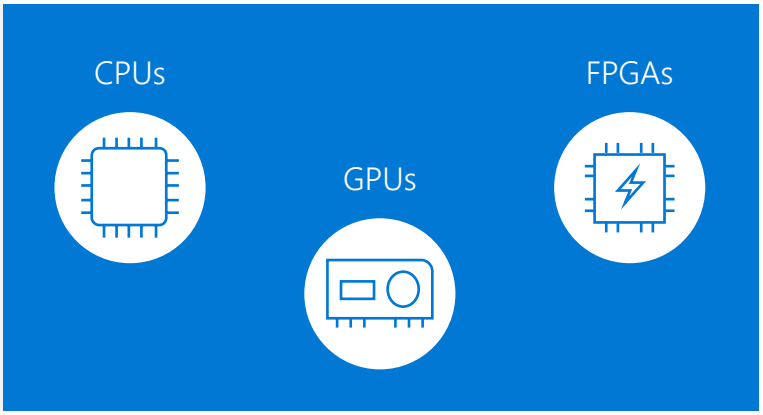

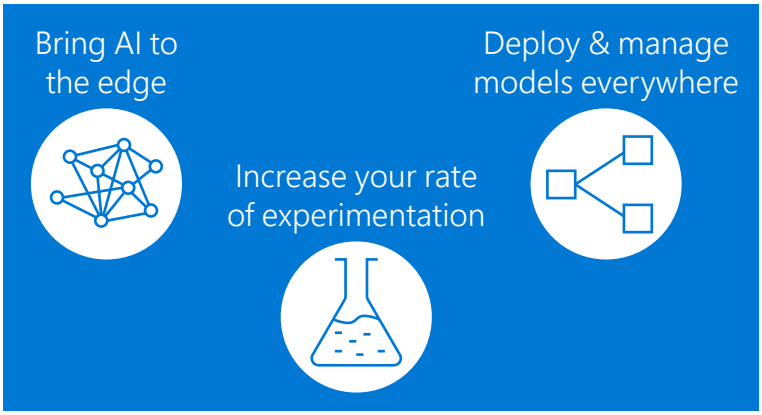

DevOps for machine learning

Increase productivity with experiment tracking, model management and monitoring, integrated CI/CD, and machine learning pipelines.[26] And with advanced machine learning pipelines, you can collaborate on all the steps of data preparation, model training and evaluation, and deployment.[27]

Powerful infrastructure

Optimized for flexibility and performance. Train models with ease and reduce costs by autoscaling powerful FPGAs, CPU, and GPU.[28] By using Azure Machine Learning service, you can start training on your local machine and then scale out to the cloud. You can train models faster and scale them up and out.

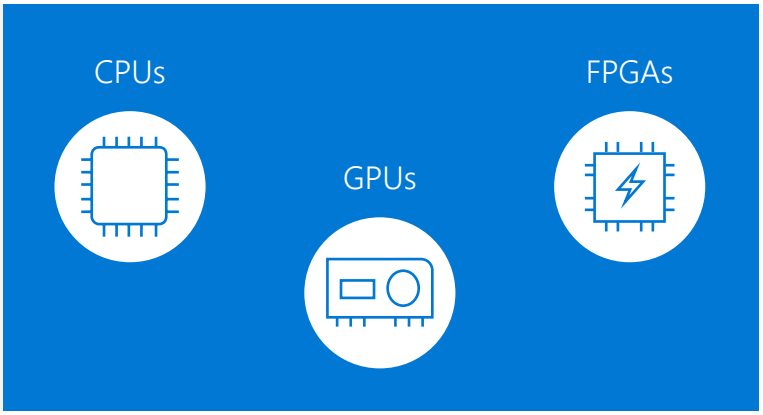

26. https://docs.microsoft.com/azure/machine-learning/service/overview-what-is-azure-ml
27-28. https://azure.microsoft.com/en-us/services/machine-learning-service/

Simple deployment

Deploy models on-premises, to the cloud, and at the edge with a few lines of code. When you have the right model, you can easily deploy it in a container such as Docker. So, it's simple to deploy to Azure Container Instances or Azure Kubernetes Service. Alternatively, you can use the container in your own deployments, either on-premises or in the cloud.

Users can deploy models on-premises, to the cloud, and at the edge, with a few lines of code.[29] When you have the right model, you can easily deploy it in a Docker container. Models can be deployed to compute targets including: Azure Kubernetes Service (AKS), Azure Machine Learning Compute, Azure Container Instances, Azure IoT Edge and Field-programmable gate array (FPGA).[30]

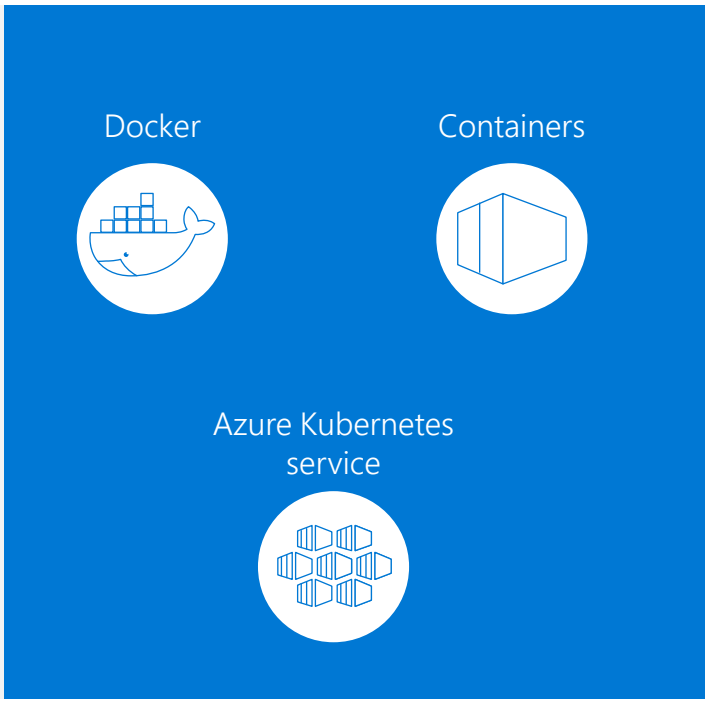

Tool-agnostic Python SDK

Azure Machine Learning service can be used in any popular Python development environment, including Visual Studio Code, PyCharm, and Jupyter and Azure Databricks notebooks.[31] Various rich tools are available for each of these environments, including a web portal for Azure, the Visual Studio code extension for Azure Machine Learning services, and widgets for Jupyter Notebooks. These tools make it easy to explore and transform data, to then develop and deploy models.[32]

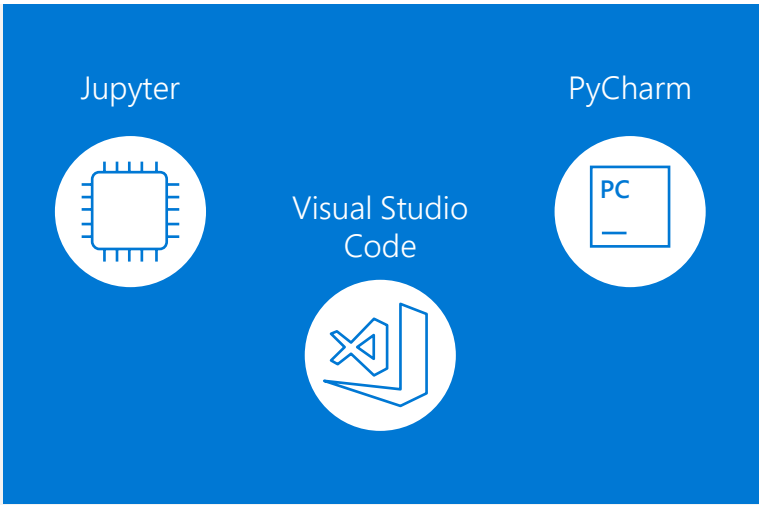

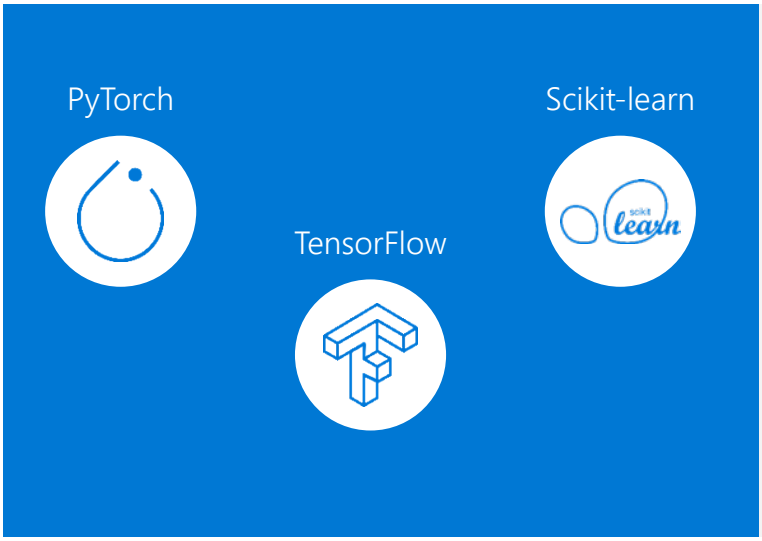

Support for open source frameworks

Use your favorite machine learning frameworks and tools, such as PyTorch, TensorFlow, and scikit-learn.[33] By using the Azure Machine Learning SDK for Python, along with open source Python packages, you can build and train highly accurate machine learning and deep-learning models yourself in an Azure Machine Learning service Workspace.[34]

29. https://azure.microsoft.com/en-us/services/machine-learning-service/
30. https://docs.microsoft.com/en-us/azure/machine-learning/service/overviewwhat-is-azure-ml
31. https://azure.microsoft.com/en-us/services/machine-learning-service/
32. https://docs.microsoft.com/en-us/azure/machine-learning/service/overview-more-machine-learning
33. https://azure.microsoft.com/en-us/services/machine-learning-service/
34. https://docs.microsoft.com/en-us/azure/machine-learning/service/overview-what-is-azure-ml

Azure supports experimentation by allowing users to leverage service-side capture of run metrics, output logs, and models. Users can manage training jobs locally then scale up or out. They can also use leaderboards, side-by-side run comparison, model selection, and conduct a hyperparameter search on ML or DNN.

To learn more about how to collect and prepare data at scale, check out Azure Data Factory, Azure Blob Storage, and Azure Databricks.

Machine learning with Azure Machine Learning service

Data Scientists can use Azure Machine Learning service to build custom machine earning models - train, deploy, and manage ML models at cloud scale.

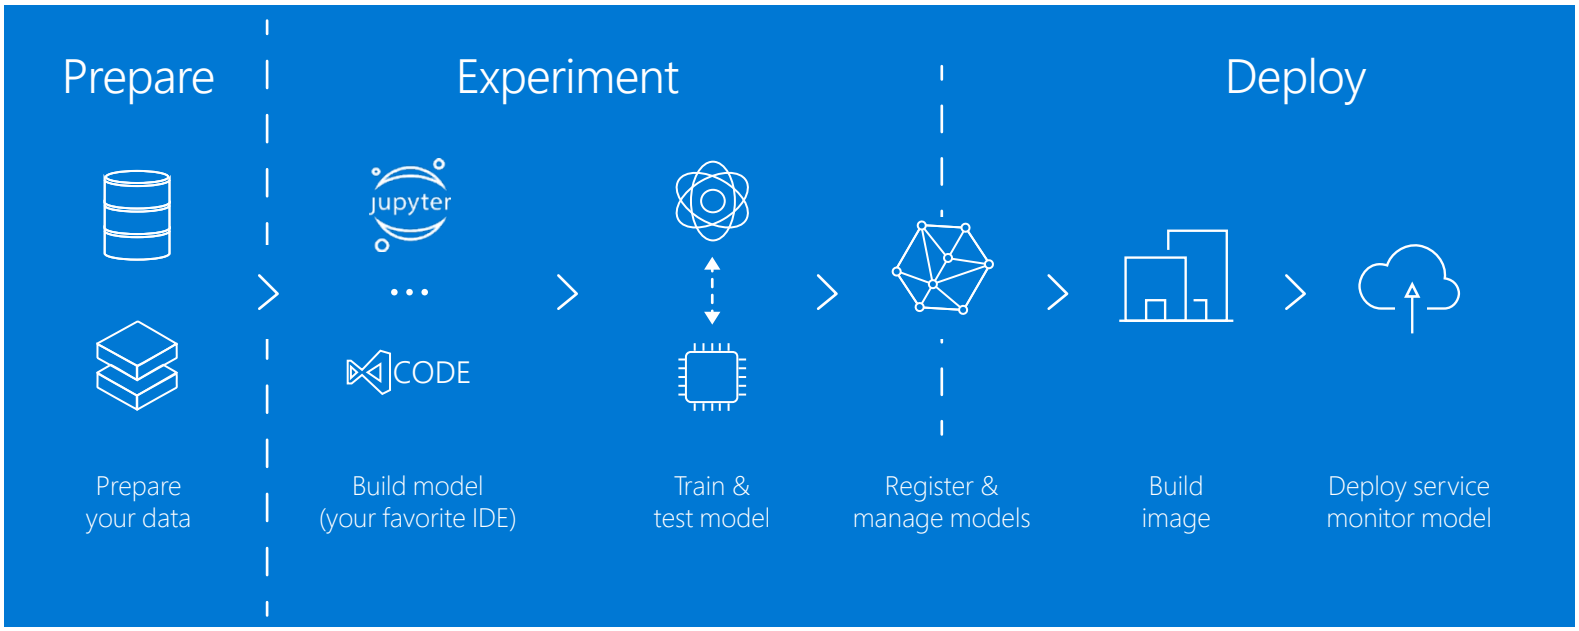

Developers can now use Azure Machine Learning service:

- Identify suitable algorithms and tune hyperparameters faster.[35]
- Improve productivity and reduce costs with autoscaling compute and DevOps for machine learning.[36]
- Seamlessly deploy to the cloud and the edge with one click.[37]
- Access all these capabilities from your favorite Python environment using the latest open source frameworks, such as PyTorch, TensorFlow, scikit-learn, and MXNet.[38]

35-38. https://azure.microsoft.com/en-us/services/machine-learning-service/

How to get started

How to get started

You probably downloaded this e-book because you're curious about machine learning and how to get started. Now, you understand what machine learning is; some of the customer objectives it meets; what Azure Machine Learning service delivers; some of the top, industry, machine learning use cases; and you've seen it in action through industry case studies. Here are a few options to start utilizing machine learning on your own, or for your business:

Get started with Azure Machine Learning service today

Build, train, and deploy machine learning models with an Azure free account

Sample Notebooks

Explore Azure Machine Learning service with sample notebooks

Quickstarts

Use a cloud-based notebook server to get started with Azure Machine Learning service
Use your own notebook server to get started with Azure Machine Learning service

Tutorials

Train models with Azure Machine Learning service
Deploy models in Azure Container Instances
Prepare data for regression modeling
Use automated machine learning to build your regression model


Microsoft