

Getting started with Artificial Intelligence (AI)

Lenovo Infrastructure Solutions
for The Data-Centered

intel.
XEON
PLATINUM

Lenovo

- **A Brief Summary**
- **AI, Machine Learning, and Deep Learning**
- **Making AI Autonomously Intelligent –
Training and Inferencing**
- **Getting started with AI**
- **Conclusion**

A Brief Summary

For the data-centered, insights are everything - improving efficiency, profitability and increasing revenue - ultimately giving your organization a competitive edge. Data is the key to unlocking insights driving automation. The challenge for many organizations is finding the resources that best enable their Analytics and AI strategy in the midst of all the noise that exists today.

Let's begin with what AI means and what it means to businesses. AI is the simulation of human intelligence by computer systems. And increasingly, in a wide range of industries, organizations are using AI to increase business value.

AI is implemented in two phases - Training or model development, which is the process whereby data scientists develop and optimize models with a curated dataset - and inferencing, which is applying a trained model to new data to derive insights and enable automation.

Here is where the other two buzzwords of Machine Learning and Deep Learning come into the picture. While they seem to be used alongside AI, and sometimes even synonymously, each of these terms has a specific meaning, a connotation, and a business case.

Real innovation happens at the intersection of the art-of-the-possible and the art-of-the-practical."

-Robert Daigle, Global AI Leader, Lenovo ISG

Machine Learning is the use and development of computer systems that are able to learn and adapt without following explicit instructions, by using algorithms and statistical models to analyze and draw inferences from patterns in data.

Deep Learning, on the other hand, a subset of Machine Learning, uses Neural Networks (NN) which are designed to mimic the human brain in order to discover and evaluate patterns in the data.

AI can increase value for a variety of business cases, in many industry verticals. Examples include identifying product defects, speeding response time to earthquakes by analyzing seismic data, and reducing inventory shrinkage in retail. From identifying product defects to sifting through tons of seismic data to reduce response times to earthquakes and tsunamis.

According to a 2021 PWC report, in today's global markets, 33% of businesses have started implementing AI, while 25% have fully enabled AI processes with widespread adoption. These statistics show how important AI is in today's marketplace. This trend is further supported by a McKinsey report estimating that the potential value created by AI could be anywhere between \$3.5 and \$5.8 trillion annually!



**AI can be applied in a number of ways to improve business outcomes- the important considerations are - what is the expected business outcome?
Will we build or buy the capabilities?
Do we have leadership buy-in to ensure we have resources to be successful?**



Artificial Intelligence

AI is essentially a broad umbrella definition of any automation system that learns on the job. Quite literally, it is a machine that performs any task associated with human intelligence. But that goalpost is always changing. Once, handwritten digit recognition or later, Google maps, would have been considered AI. Now, however, those examples are so common that they're no longer thought of as AI.

VS



Machine Learning

These days, when we talk about AI, we are usually talking about Machine Learning (ML), in which a program performs tasks without explicit programming. ML uses techniques like clustering and regression, among a host of other generalized techniques that crunch a lot of data to develop predictions.

With ML, you need to manually select “features” to train based on the problem (for example, color, shapes, and edges). Subsequently, you pick the appropriate algorithms to use and give the program examples of these features. From there, the program can analyze data and improve its understanding of those features.

ML, although pretty effective, is human effort-intensive when it comes to extracting the features from the data to begin with.

VS



Deep Learning

Deep learning improves upon Machine Learning capabilities as it uses neural networks to learn data representations and perform tasks without explicit programming. What we mean by explicit programming is essentially that the extensive human intervention required with Machine Learning is reduced significantly, if not eliminated altogether. This is because the neural networks Deep Learning utilizes simulate how the human brain works.

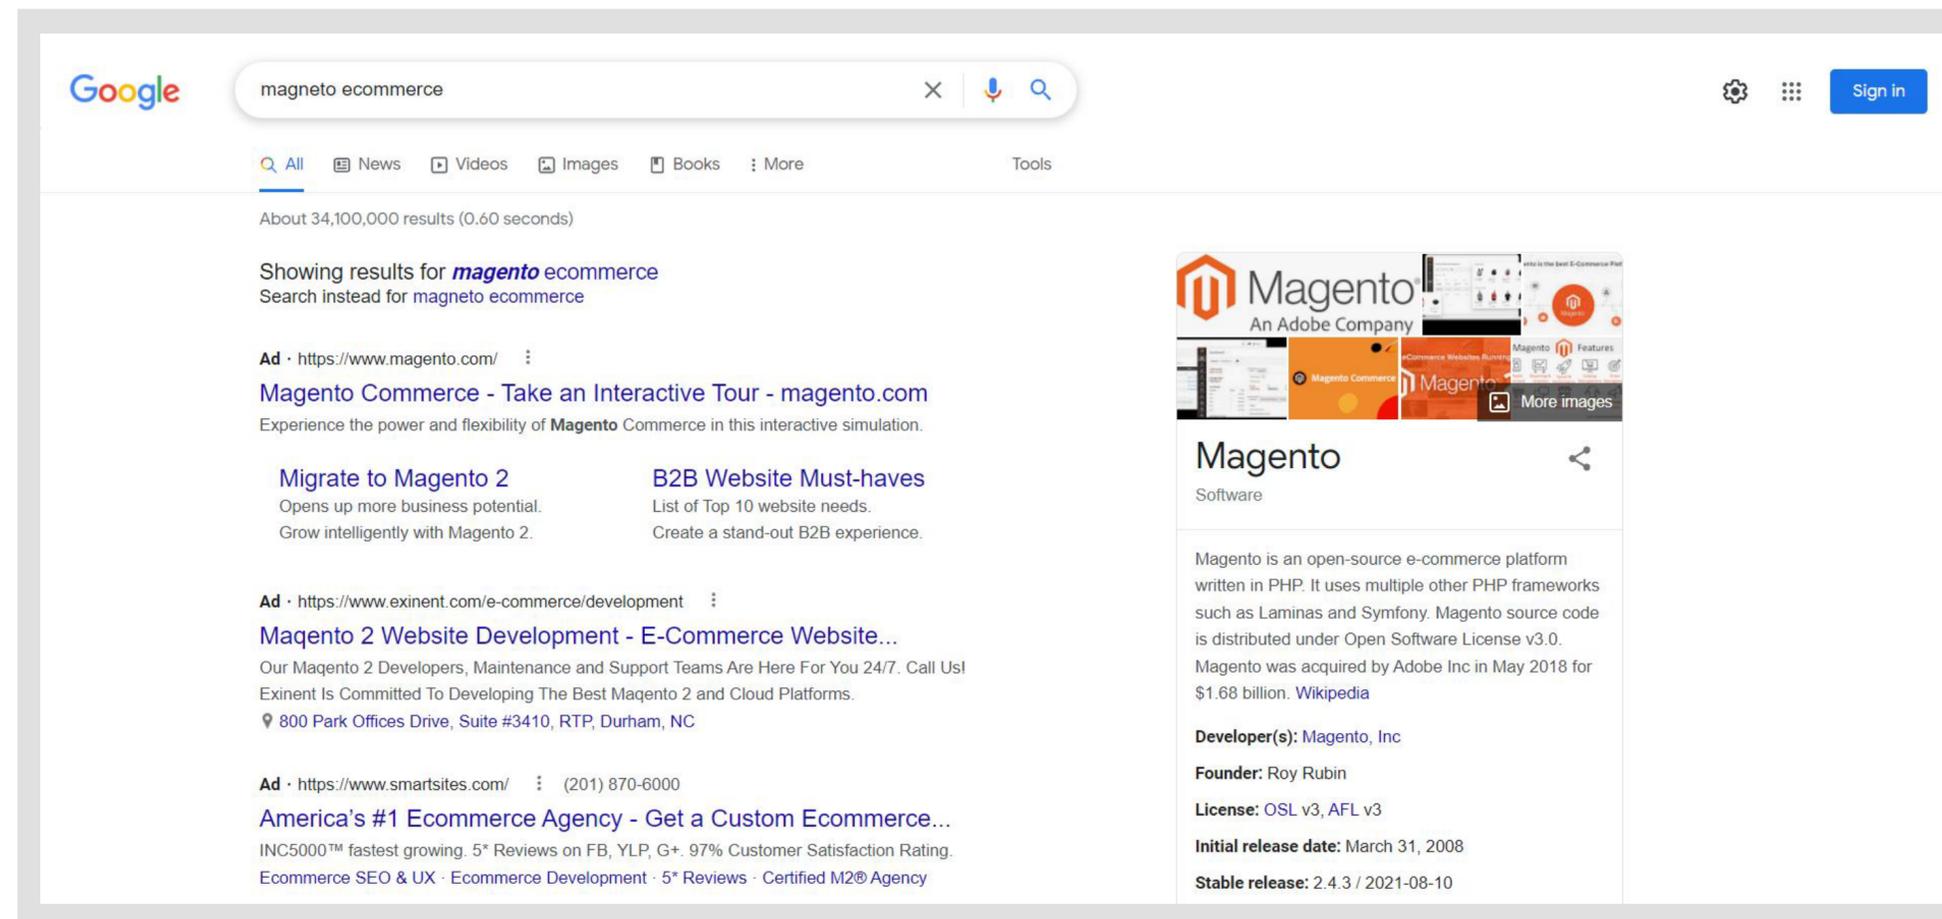
Thus, it is not necessary that you know the features you are looking for beforehand. The neural network trains itself on the data and discovers the features on its own.

Consequently, you may not be able to articulate why an algorithm derived a certain decision which is one of the drawbacks with deep learning. Fortunately, many new tools are being developed to address this 'black box' effect of deep learning.

There are a few subcategories of AI, Machine Learning and Deep Learning:

- **Natural Language Processing (NLP)/ Natural Language Understanding (NLU):** tries to understand human voice and text to recognize the intent. While traditional automation systems will take whatever input is received and process it, NLP analyzes the input and applies its machine or deep learning capabilities to recognize patterns and correct errors. A simple example would be when you enter a typo in your google search.

Note how Google's NLP engine identifies the error in the spelling of Magento because of its association with eCommerce, which is the probable intent of the person searching. Likewise, NLP can also recognize different accents and pronunciations of words to improve voice-command recognition for devices such as Alexa, Echo, Siri, and Cortana.





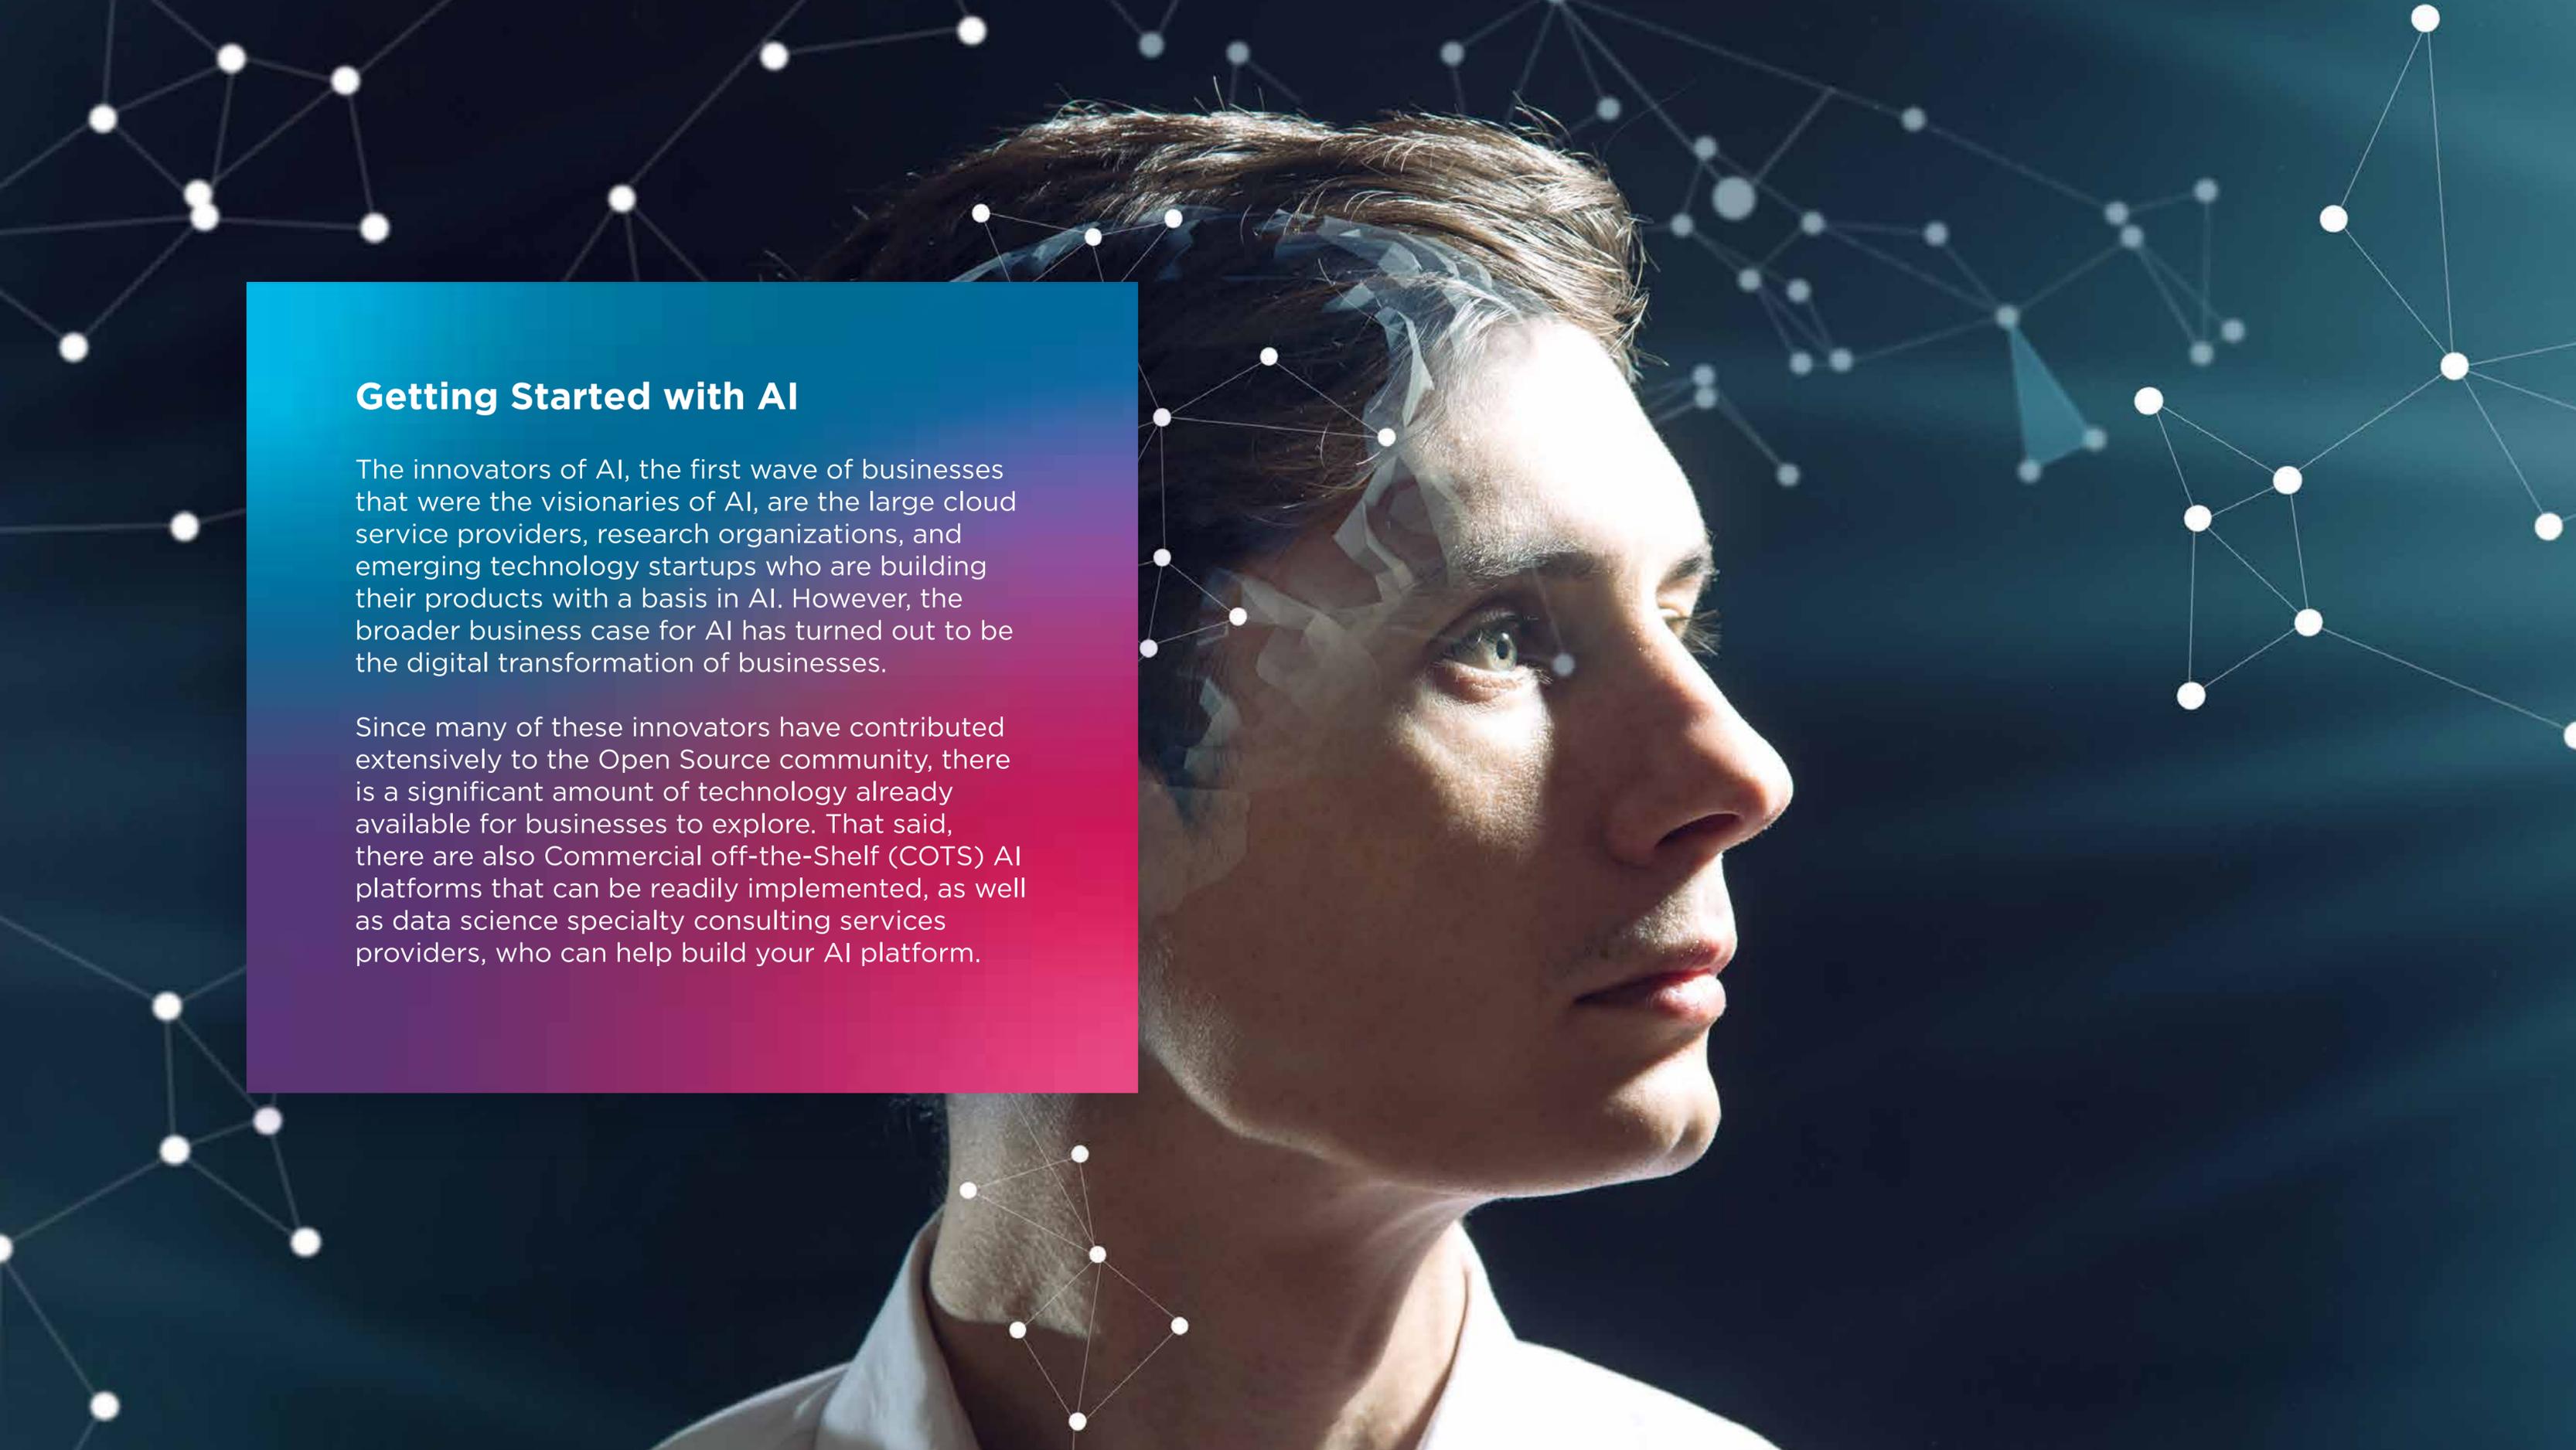
- **Computer vision** – Used interchangeably with machine vision aims to automate image recognition, analysis, interpretation, and recreation. The technology is applicable in a variety of business cases, including self-driving cars, automated analysis of medical imaging, reconstruction of 3-D models from a series of images, and pattern recognition for automating defects in manufacturing processes. Lenovo, for example, uses machine vision as a means to take images of devices produced and identify potential defects in them by analyzing the images against the deep-learned image data benchmarks.
- **Robotics** is another area where AI, ML, and DL can play a significant role. Smart manufacturing plants utilize AI to reduce errors, replace humans in hazardous conditions, or increase productivity in repetitive processes such as riveting. In healthcare, robots are being trained using AI to help nurses, practitioners, and surgeons improve their productivity and efficiency while also reducing the cost of service. In agriculture, robots using computer vision and Deep Learning techniques are increasing the yield and efficiency of harvesting by detecting produce ready for harvesting, as well as unwanted flora requiring removal. One of the best examples of AI and Robotics coming together to speed up and improve the quality of work is smart warehousing. From Amazon to Alibaba, to Ocado and Walgreens, retailers around the world are switching to entire armies of robots that use computer vision to navigate and recognize product aisles. This switch to robots has seen unprecedented improvements in everything from space utilization and error reduction to collection speed.



Making AI Autonomously Intelligent - Training and Inferencing

The foundational step in implementing a sustainable AI system is training the model. The crucial role of training lies in its end objective - reducing human intervention. Training typically takes a lot of data and can take hours to weeks, depending on the size of your model. Apart from volume of data, training also requires large computing power and is typically done at a data center. The approach to training the AI engine is directed at getting the results within a defined range. This approach allows the AI engine to recognize the tolerance for aberrations, thus improving the accuracy and consistency in performance. For example, training a robot to avoid obstacles would require calibrating the computer vision engine based on observations and mistakes. The famous videos of Boston Dynamics engineers kicking and whacking the robots are nothing but them training the AI engine to compensate for the loss of stability.

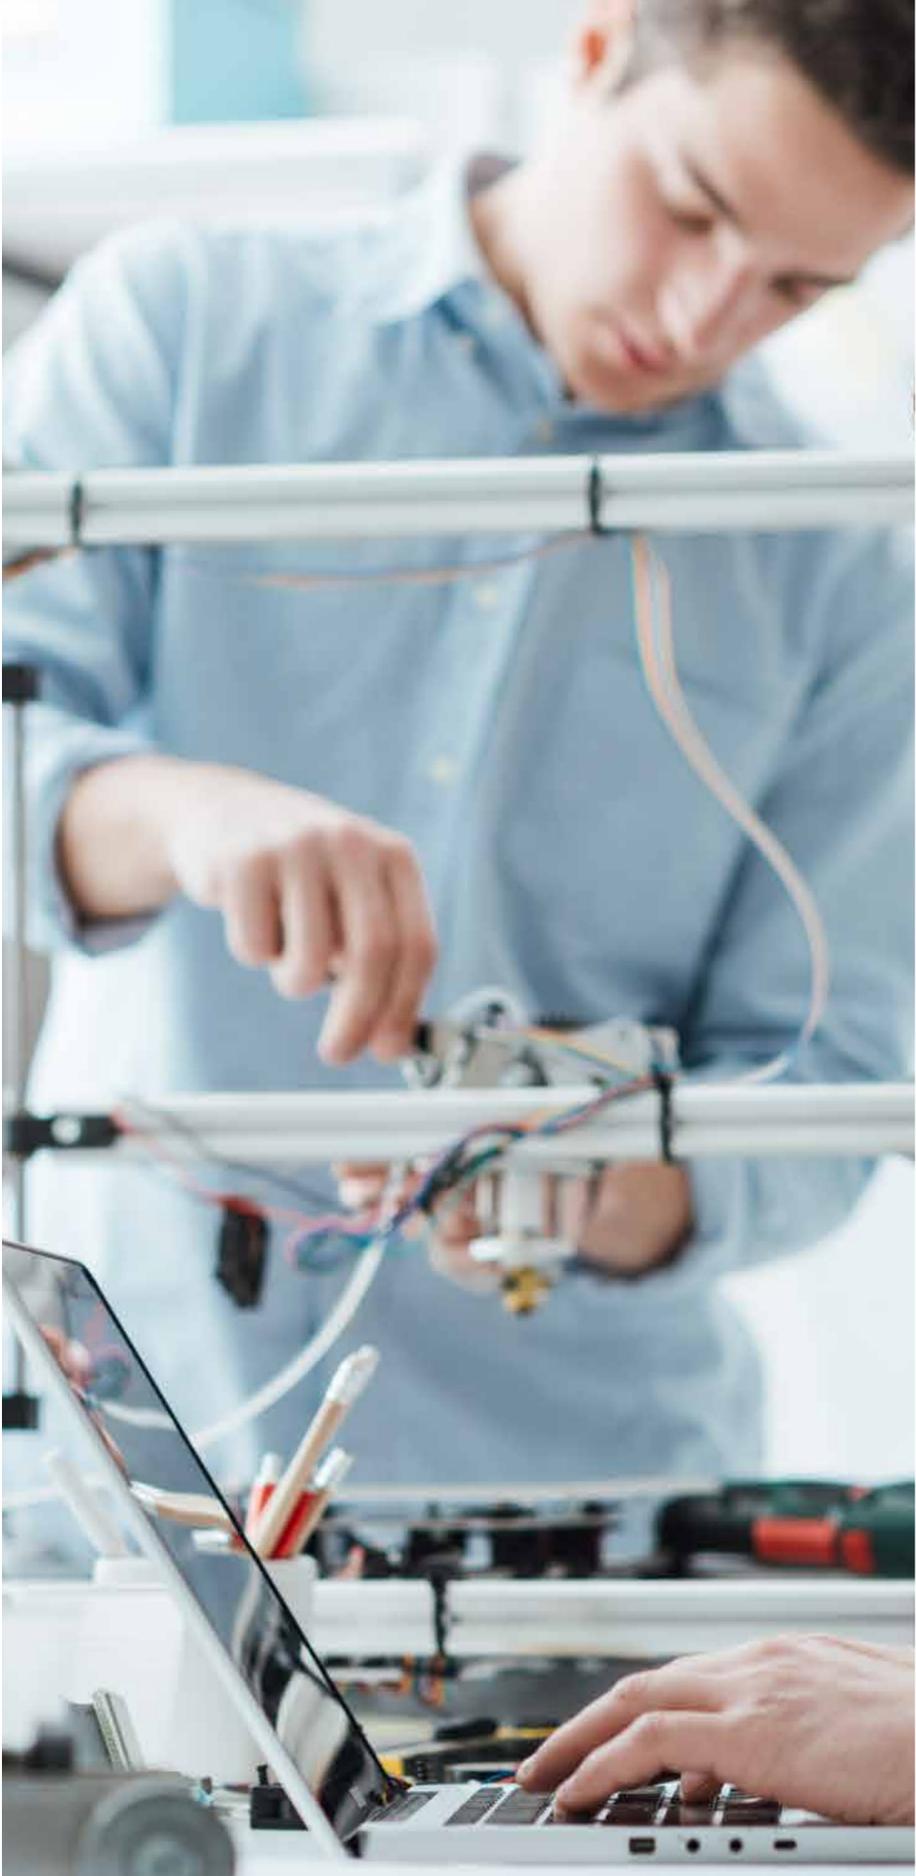
Once the AI is trained, you take the model and deploy it. Now comes the second part - using new data to have the model infer information from it (the inferencing stage). Inferencing allows us to gain new insights by using the trained model on new data, while also validating the model against that new data to make sure the training is complete, comprehensive, and that the model is within the acceptable tolerance for errors. Inferencing can be done with a lot less computing power than the training and can be performed at the location of deployment, perhaps at the edge, for real-time processing.



Getting Started with AI

The innovators of AI, the first wave of businesses that were the visionaries of AI, are the large cloud service providers, research organizations, and emerging technology startups who are building their products with a basis in AI. However, the broader business case for AI has turned out to be the digital transformation of businesses.

Since many of these innovators have contributed extensively to the Open Source community, there is a significant amount of technology already available for businesses to explore. That said, there are also Commercial off-the-Shelf (COTS) AI platforms that can be readily implemented, as well as data science specialty consulting services providers, who can help build your AI platform.



The principal considerations before getting started with AI, are:

✔ **What is your desired business outcome or expected result?**

The most important questions to ask yourself before you choose your approach to AI, are firstly, if you have a business case that has a readymade solution. Not all industries and use cases have COTS solutions available, so it may actually be helpful to take the DIY or Consulting Services approach.

✔ **Build or Buy?**

Secondly, you need to ask yourself if you're willing to take it slow and steady. If so, a DIY approach may be best, but there are pitfalls here as well, because building the right capabilities in-house can be pretty challenging. Ultimately it is the tradeoff between cost, time, and effort that will define your choice of approach with AI.

✔ **What is your budget?**

For those who wish to start small, DIY can be a convenient option. For those who would like to invest upfront, COTS and Consulting Services can help get there faster, and easier.

✔ **How do you want to go about setting up the infrastructure required for AI?**

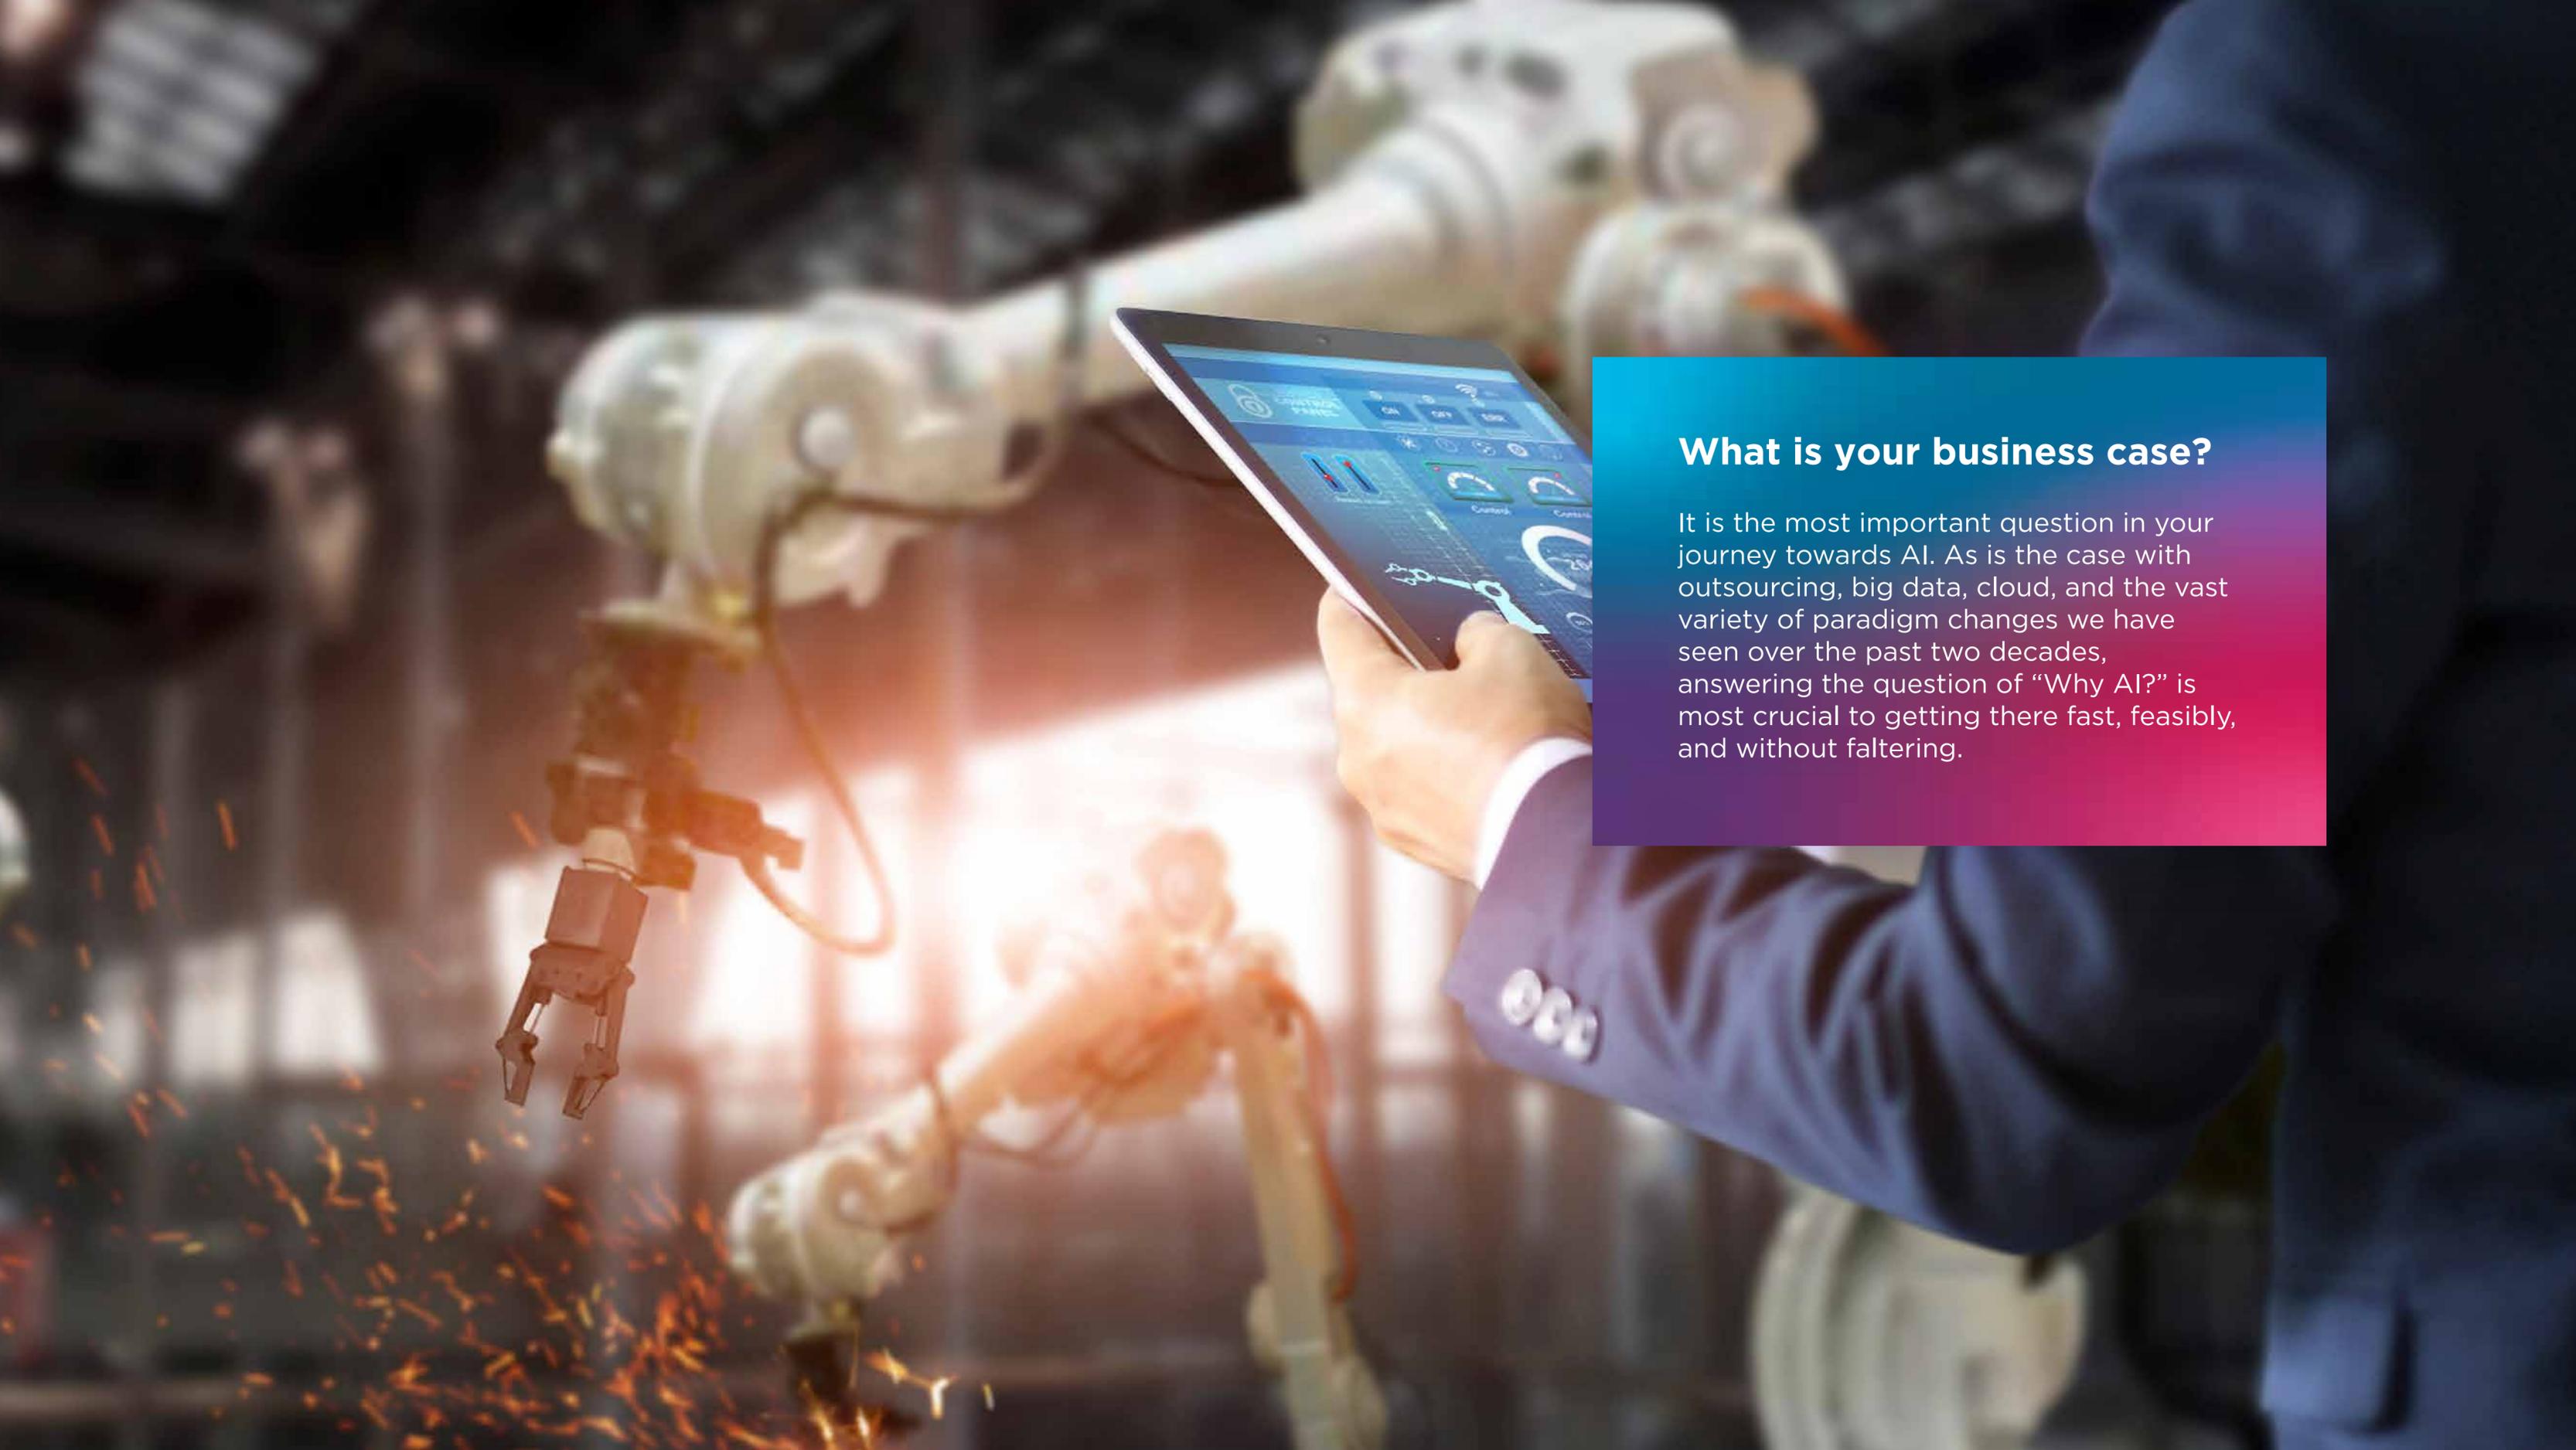
With the rise of AI as a Service, choosing to build your proof of concept on a pay-as-you-go model could be a good start for businesses that do not have the in-house capability. However, another question that needs to be asked is how easy, economical, and feasible it is to bring AI to production, once the PoC is complete. Architecture decisions are critical for training and inferencing. Training typically happens in the Core or cloud, inference can happen at the core, cloud or edge but more and more of inferencing is moving to the edge. Core is typically significantly lower cost than cloud but may need IT staff to upskill. Many inferencing workloads are moving to the edge due to latency and bandwidth requirements - it's not economical or feasible to run some inference workloads in the core or cloud.

According to Mckinsey's 'The state of AI in 2020' report, as many as 50% of businesses have not started on their AI journey. It is not too late to start now. Consequently, there really are no wrong answers when it comes to how you would like to approach AI, as long as you start before your industry is saturated with competitors that have strong AI capabilities. It eventually boils down to the tradeoffs you make. The four most important considerations as you choose your AI strategy are investment, skill, risk and infrastructure.

For those who are still on the sidelines waiting and watching, there are a few industries that are implementing AI-based solutions quite effectively. While the manufacturing sector is perhaps the most frequently cited example, stemming from smart factories and Industry 4.0, Retail, Healthcare, and Finance have also been extensively involved with implementing AI systems.

One of the best examples of AI is perhaps in the retail sector, where self-checkout is becoming the norm at brick and mortar stores. While the system has improved significantly, a consequence of the lack of staff oversight of the checkout process has been shrinkage. Whether due to mis-scans, missed items, or intentional evasion of scanning of the purchases, the retail industry is looking at a \$61 billion/year loss in the US alone. Using a combination of Computer Vision and video analytics, retailers can now identify mis-scans in real time and alert store attendants.

However, the scope for AI doesn't end there. From analyzing CCTV footage in real time to identifying choke points in aisles, to building customer profiles and demographic analytics using facial recognition, AI is improving the shopping experience at retailers tremendously. Computer vision can also help with sentiment analysis, thus giving retailers real-time feedback on the shopping experience, as well as product quality, based on facial expression recognition by AI systems.



What is your business case?

It is the most important question in your journey towards AI. As is the case with outsourcing, big data, cloud, and the vast variety of paradigm changes we have seen over the past two decades, answering the question of “Why AI?” is most crucial to getting there fast, feasibly, and without faltering.



Although AI is still in its relative infancy, examples of business cases abound. Healthcare providers have implemented computer vision based on AI for diagnosing patients with COVID-19. By training the AI engine to recognize the COVID-19 infected lung from x-ray images, diagnosis has become exponentially faster. Computer vision has also been used in conjunction with infrared cameras to detect people with fever, which is a significant improvement from manual recording of temperature. Additionally, computer vision can also analyze CCTV footage in real-time to enforce social distancing by measuring the distance between people, as well as automate checking of PPE.

It is important that you ask the right questions of whether you need AI, and if you do, what it is that you expect AI to do that is harder, time consuming, or expensive otherwise. Ultimately, like automation, AI is useful, but only when put to use on business cases where it offers improved efficiency and cost-effectiveness. Lenovo can help you explore our extensive solution partner ecosystem and provide efficient consulting services to help you decide on the business case you would like to approach.

What is your business case?

1



What is your business objective?

2

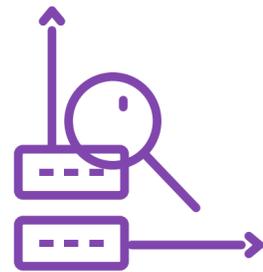
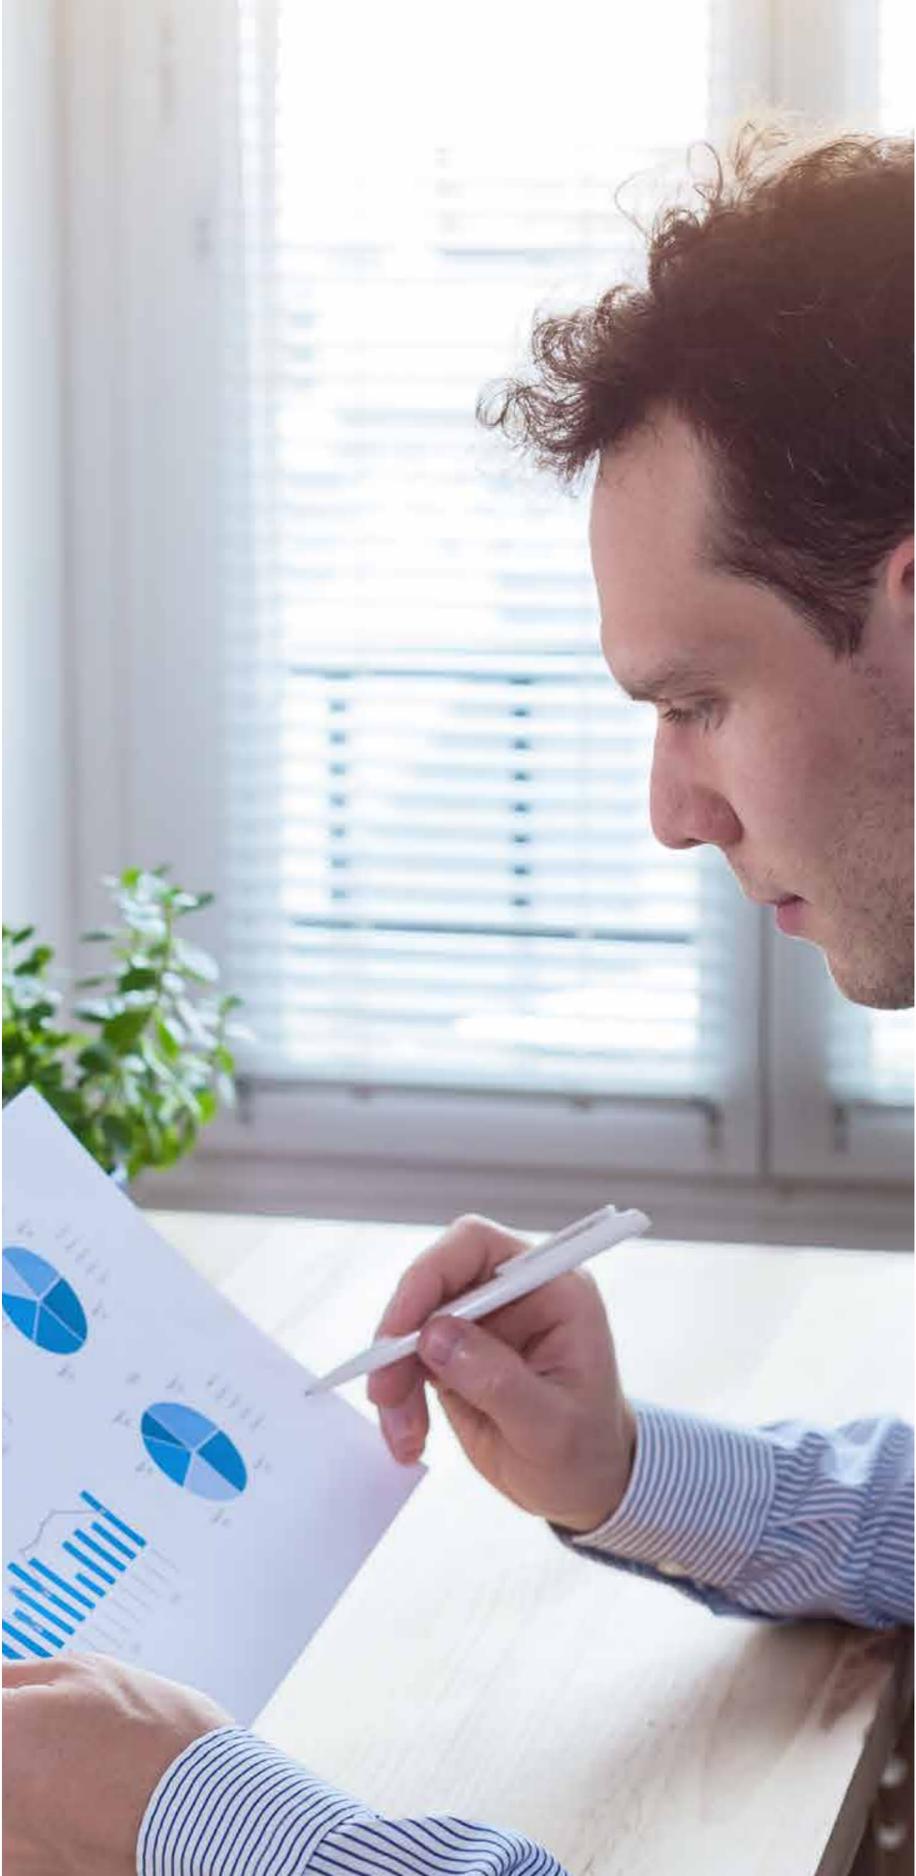


What kind of data do you have?

3



Build vs Buy



What is your business objective?

Understanding your data sources, types, and locations is key to building the right business case for AI. For example, if you intend to use computer vision, you may be able to train the AI system on a variety of scenarios. Exploring the possibilities within your data set requires understanding your data and how they can be mapped to different AI business cases. Lenovo can help you identify your business case, deliver proofs of concept, and also help build your AI landscape with solution guidance and professional services.



What kind of data do you have?

Depending on where you would deploy your AI--at the data center, in the cloud, or at the edge-- the corresponding platform for your AI system can change dramatically. Lenovo can guide you through infrastructure requirements, offering hardware, software, development tools, as well as plug-and-play platforms. With innovation centers around the world, Lenovo can help you can gain an in-depth understanding of the different AI platforms and how to distinguish between each, based on cost, complexity, and capability, as well as identify the platform best suited for your AI journey.



Build vs Buy?

Choosing your approach, whether it is DIY, COTS, or Consulting Services, is critical to your success. It is important to vet your vendors and identify the partner with the right skill-set, capability, and fit for your industry, business, and corporate culture. Lenovo can help here, offering guidance on the different vendors in the marketplace, as well as consulting services to help you make the right choice.

Lenovo can also help by introducing you to our vast AI partner ecosystem, based on validated and tested solutions from leading ISV partners, optimized for Lenovo Infrastructure. Through our extensive partnerships and validated solutions, we ensure we're providing the highest performance for your AI workloads.

Conclusion

Whether you're just getting started on your AI journey or you're moving PoCs into large scale deployments, now is the time to leverage AI within your organization to reach your business objectives. Lenovo solutions powered by Intel® Xeon® Platinum processors are here to help you at every stage of your AI journey, from planning and deployment to optimization. Let the innovation begin today, at the Lenovo AI Innovation Centers.

[Explore Analytics & AI Solutions](#)



Lenovo

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo.
Intel, the Intel logo and Xeon are trademarks of Intel Corporation or its subsidiaries.

© Lenovo 2021. All rights reserved.