



# GET STARTED ON YOUR AI JOURNEY WITH LENOVO

from discovering  
your use case to  
implementing  
your solution

The CEO of eBay in 2017 put into words what many other industry leaders have lately been suspecting about Artificial Intelligence (AI). "If you don't have an AI strategy, you are going to die in the world that's coming," said eBay chief Devin Wenig<sup>1</sup>. On the other hand, the head of one of the world's largest and most technologically sophisticated tech firms may not quite appreciate how daunting the prospect of developing an AI strategy is. To begin with, what exactly is an AI strategy anyway?

This brief document will guide you through the process of developing AI projects for your company to tackle. If you have a rough sense of the kinds of AI projects your company faces, but would like to be smart and strategic about the focus of your projects, read section II below.

If you do not have a rough sense, do not worry. Look first to section III, where sample AI workloads in four representative industries are briefly described.

## Developing an AI project

Building out an AI use case in your organization will take time. But, other than step 3 - Pilot or Proof of Concept (whose time scale can vary greatly depending on the amount and kinds of data), the recipe described in this document from steps 1 to 5 should not take more than 4-6 weeks to do.

<sup>1</sup> Zia Daniell Widger, "6 Key Takeaways from Shoptalk 2017: An Insider's Perspective," Forbes (online) Mar. 28, 2017

## 1 Developing a use case.

The most important first step to take in developing your AI project is to look for the low-hanging fruit — the easiest win, the simplest implementable project. Because looking for a complex project that might (if successful) yield a big and conspicuous win, is equivalent to attempting to run before you walk.

You must prove out the business case first — and the case for AI. You might find, if you tackle an excessively complex initial project, a second problem known as “scope creep” can occur. Meaning the original goals you set for the project expand while in progress. Scope creep is a pernicious additional problem to address. You may begin with small adjustments to your goals or your means to achieve those goals, but you might also find you’ve set yourself up for missed deadlines or even projects doomed to ultimate failure.

Avoid all of this. Keep your initial use case modest. And then build on that success to scale up beyond your initial win.

## 2 Collect the data.

AI is nothing without good data. Especially if you’re trying to start quickly, it’s simplest to start with data that’s already been collected (this could in fact be a key factor in determining which is the lowest-hanging fruit in step 1). If you already have a big data system in place, this could be a good use case to get started with.

If, on the other hand, you have a great use case but haven’t yet collected data, you’ll need to pause and determine out how to collect that data — and how you’re going to build out the full data ingest pipeline.

Crucially, the best data to start with is clean data. Companies, that is, rarely lack for data files of some kind or another. But just because something takes up a lot of disk space and lives somewhere on your company servers doesn’t make it useful or meaningful for your AI project. Rather, data that’s already being used successfully for analytics or other data analysis is worth more than data that’s un-scrubbed and unanalyzed.

For instance, sales or financial data are often already relatively clean, coming from programs like Salesforce or Microsoft Dynamics or other CRM applications. This is especially true if someone on your sales or marketing team is already going through these files to make sure fields are labeled correctly, with the right deals in the right categories, you might already have a clean data source ready to be put to work.

## 3 Pilot or Proof of Concept (PoC).

As with the first step, developing a small-scale proof of concept implementation first will pay off with more robust and comprehensive results in the end. For example, in one of the many industries for which Lenovo develops AI solutions, Lenovo solution engineers have worked with a top car racing team to build out AI and HPC applications and use cases. During the changing conditions of the race, the decision to ‘pit’ a car to make repairs or change tires has to be made in seconds. Previously, the team was manually clicking through streaming pictures from the race track, trying to get good pictures of their car in order to make the decision to pit. Often by the time they found the right shot, the opportunity had passed and they lost critical time running a slow lap. With Lenovo’s help, we created a deep learning model that quickly and accurately identified their car from the stream of pictures, and the race team could then spend those critical seconds making the right pit call instead of manually clicking through pictures.

So the team wisely recognized, in this case, their real-time analytics may not provide the best and simplest source of clean data for an early AI PoC. Instead they recognized that in any given race they collect upwards of ten thousand images from their pit crew, from the track, from their car and from various places in the grandstands. Most of the images are meaningless to them. However, a small percentage contain crucial data points — tell-tale smoke plumes from the car, an unexpected tread of tire being rolled out on competing cars, a new passing strategy indicating a shift in the race.

The PoC implementation may not have been the most apparent initial use case for the team, but it leveraged a known AI strength and put it to work in a limited but meaningful way for the racers.

## 4 Implementation.

Now that you've run your pilot in a controlled setting, it's time to begin rolling your AI project out into a less controlled production environment. There will be challenges in this phase.

For starters, you will have less control over the data that goes into the system. In this phase you'll be taking the working algorithm that you've proved in your PoC into production. You will likely have to add substantial checks on the data before feeding it into your model, both that the new data is valid (such as excluding blank images) and also that characteristics (such as mean, variance) of the data itself are not changing over time.

Take for example an algorithm that scans liver cells for tumors. You may have proved this algorithm on a small subset of samples. But now you have to go implement it into an application. You must now put together a process about how you're getting the CT scans and how you're loading them into the system. If there are extra steps that your new AI workflow requires, your users are probably not going to use it - so user experience is a key consideration.

These are some of the challenges you face in the implementation to make sure it's used, efficient and effective. Again, remember the end goal is not to have an AI algorithm that stands out as an AI algorithm. Rather, it's for your AI algorithm to fade into the background as a tool that's useful to its users but also doesn't require a lot of extraneous care and feeding either.

## 5 Operationalize

Finally, comes the step that concerns, as it were, future proofing your AI process. Any AI system will be continuously updating itself and retraining on incoming data — and ideally improving itself over time. You will not be abandoning this AI project but will be returning to this from time to time. Operationalizing, then, involves querying the implementation you've rolled out — and ensuring as much as possible the acquisition of new data and the retraining on that new data is automated. This might require additional conversations around what is expected of the data scientist team in terms of acquiring and cleaning the data for your AI system's continued operations.

## What can I use AI for?

### Manufacturing

- Reducing supply chain forecasting errors and lost sales.
- Preventative maintenance.
- Improving manufacturing yields.
- Automating quality testing.
- Asset management & inventory management.

### Healthcare

- Voice-to-text transcription for doctors and nurses.
- Robotics-assisted surgery.
- Virtual nursing assistants.
- Dosage error reduction.
- AI-driven image analysis (e.g. radiology).
- Cybersecurity workflows safeguarding patients' health data.
- Administrative workflow assistants.

### Oil and Gas

- Automated hydrocarbon exploration & production.
- Online chat platform assisting navigation of product databases (e.g. Shell Virtual Assistant).
- Refining drilling and well completion.
- Modeling oil refining strategies.

### Life Sciences

- AI-powered drug discovery (esp. expensive early stages).
- Predict future behavior of compounds in clinical trials.
- Surfacing of promising new drug pipelines from public data sets (e.g. PubMed, Clinical Trials. gov and patent filings).
- Discovering gene target links for diseases.
- Scanning EEGs, ECGs, retinal scans and genetic data for warning signs.



## How Lenovo can help you on your AI journey

At Lenovo, AI stands for “Augmented Intelligence” — meaning using traditional AI solutions to augment human capabilities. A customer’s journey to Augmented Intelligence involves discovering, developing and then deploying. This brief paper has, ideally, provided ideas and incentive toward accelerating this journey.

To assist you in the Discovery stage, Lenovo has opened AI Innovation Centers in Morrisville, NC, Stuttgart Germany, and Beijing China. These Innovation Centers allow customers to test their AI workloads on a variety of systems, with access to experts from both Lenovo and partners like [Intel](#)<sup>®</sup>, [SUSE](#) and others. In the Develop stage, Lenovo has also announced Lenovo Intelligent Computing Orchestration ([LiCO](#)) — an optimized orchestrator that allows the open source AI frameworks to better take advantage of the training libraries running on our Intel<sup>®</sup> Xeon<sup>®</sup> powered ThinkSystem servers.

Because there’s an industry expert for your market at Lenovo, you can be sure your bespoke AI solution is exactly that: designed with focus and understanding of your goals. To round out a complete solution the vast amounts of data you’re using can easily be managed by Lenovo’s market-leading storage solutions, increasing your capacity for analysis. Utilizing the Lenovo AI development pipeline ensures your time is spent discovering and achieving; not finding IT workarounds or utilizing un-optimized AI platforms that delay analysis.



**DISCOVER LENOVO AI SOLUTIONS TODAY**

