



DEEP LEARNING

NVIDIA

NOV19

AI: 5 STEPS TO GET STARTED





GETTING STARTED IN AI

The adoption of artificial intelligence (AI) across every industry has outpaced predictions for growth. In fact, in a [2019 survey of CIOs](#), Gartner revealed that enterprise use of AI has grown 270 percent over the past four years. And the transformative impact is clear: AI is creating innovative new products, boosting revenues, cutting costs, and dramatically improving efficiencies. To get ahead of this widespread adoption, enterprises need to anticipate their future AI needs and transition current workloads.

[PARTNER] and NVIDIA are uniquely positioned to provide insight into the AI deployments that can drive these changes. As a cutting-edge leader in AI technology, we not only develop and implement AI within our organization, we help enterprises, organizations, healthcare firms, government agencies, and more adopt, implement, and scale their own AI initiatives.

Whether you're starting your first AI project, transitioning a team into AI workloads, or looking at infrastructure blueprints and expansions, these five steps will help set your AI projects up for success.

01

IDENTIFY A USE CASE

Transforming your business with AI doesn't happen overnight. Instead of trying to map a larger transition, start with one use case, ideally one that relates to a business-impacting problem that can garner executive support without budget challenges. These low-hanging fruit can be found in every industry. While selecting an AI pilot project, there are certain factors to consider:

- Can this problem be easily tackled with available data that can be mined for insights using AI algorithms? Tangibility and starting small is key.
- Does this use case mirror other problems your organization is trying to solve? Pick one that will have an immediate impact and will allow you to replicate success easily.
- Even though you're starting small, don't ignore return on investment (ROI). Make sure your investment is going to pay off.

Starting small and replicating success will create momentum that can help your team secure support and an increased budget for bigger and more complex projects. Once you've identified the project, the following tools for success are comparable to other initiatives: Set a schedule with attainable goals, track the project, be flexible, and adapt quickly when things aren't working.

02

EVALUATE YOUR DATA

AI projects aren't successful without data. But it's more complex than that. Successful projects require the right data—lots of it—and data needs to be prepped with a purpose. All of this can be time consuming. Before spending hours having a team prepare data, make sure you've set a strategy for your data.

Work with your business analyst to understand where the data currently sits and if it holds the potential for feeding AI models that can be trained for your problem. Data scientists can experiment with and engineer features from existing data to build models with the greatest predictive accuracy and business outcomes. A data engineer can build a data pipeline that supports the development workflow from initial concept and experimentation to production training and inference at scale.

Speed up and repeat: Wherever you are in the process—whether you're gathering data → preparing data → selecting an algorithm → training your algorithm → or testing (i.e., inferencing)—leveraging GPUs can accelerate it from start to finish and help you do more iterations for better, faster results and improved ROI.

Keep in mind, the data you choose to train with directly affects the quality of the resulting model. Selecting your datasets carefully, perhaps experimenting with a smaller subset, and spending time preparing the data properly will accelerate your time to usable insights.

03

BUILD THE RIGHT TEAM AND IDENTIFY ROADBLOCKS

Getting started in AI can involve many steps, from building the right team, to eliminating cross-workflow silos, to providing the right resources. Along the way, keep an eye out for roadblocks, which can appear in many different forms.

Create a dream team. Think outside the box when building your team to support this new initiative, which doesn't just mean "hire data scientists." Many are very well versed in the concept of hiring specialists to get the job done. But more often than not, if you have business subject matter experts, training them on AI—in other words, creating "citizen data scientists"—is often more cost- and resource-effective than seeking out talent that may not be as familiar with your business and objectives. In smaller organizations, a single person might be responsible for the entire AI project, but ideally you'd like your team to include a business analyst, a data scientist, a data engineer, a DevOps engineer, and an application developer who can launch your models into production.

Support the right workloads. In addition to AI applications, you'll want to make sure your organization supports implementation considerations for workloads like machine learning, deep learning training, and deep learning inference. All stakeholders must be on board with providing appropriate AI infrastructure to support modern AI workloads.

It's all about the data. Data gravity, which is the tendency for applications and data processing activities to gravitate toward where the data resides—data center, cloud, and edge—is also an important factor to consider. It's better to deploy intense processing such as AI training closer to where the data is located, as moving data can reduce efficiency, drive up cost, and introduce more errors.



Once you understand where your data lives, consider its characteristics. How large is the data you have available today? How much does it grow each year? Will you need to move it if it exceeds a certain size? Start developing processes and infrastructure today that can scale as your data inevitably grows.

Build a modern AI infrastructure. There are many options available to help inform your decisions about when to build, buy, or outsource your AI infrastructure. For example, there are hybrid cloud options, AI at the edge solutions, and colocation data center services. At NVIDIA, we specialize in delivering optimized, tested, and utilized GPU-accelerated computing infrastructure, including a robust, end-to-end software solution stack. And many resources are available to help you start your AI journey, including [NVIDIA Deep Learning Institute courses](#) and webinars like “[Deep Learning Demystified.](#)”

04 AVOID SHADOW AI

Following the steps above without IT leadership and a prescribed infrastructure to support your efforts creates shadow AI. That contributes to the sprawl of development silos that often prevent well-meaning enterprises from truly transforming with AI.

Early on, it's important to work collaboratively with the IT team to make sure resources and infrastructure are in place to support end-to-end AI development. This includes platforms and tools that are optimized for the unique demands of AI workloads in the data center. Many data centers weren't designed to support computing, power, and cooling of modern AI workloads. Forward-leaning enterprises are building AI centers of excellence that combine people, process, and technology to centralize AI development and achieve faster time to AI insights. You can plan for similar success with a robust and powerful platform that's scalable for your infrastructure to accelerate data science, machine learning, and deep learning.

Check out some of the resources around IT reference architectures that [NVIDIA has developed](#) across various use cases and industries.

05 CELEBRATE THE WINS

Along the way, make sure you recognize successes. These stories spark more green lights from stakeholders and guide future initiatives, creating a flywheel effect where each successful initiative adds to the momentum.

Companies that highlight the value derived from AI are more likely to secure support for scaling this technology across their organizations. That's why it's important to spotlight the wins gained from AI, such as revenue generation and cost savings. Higher levels of executive support will facilitate access to more resources and budget to tackle enterprise-wide AI initiatives.



THE BIG PICTURE

As you continue your AI journey, keep in mind the key components that will accelerate and sustain your AI efforts:

- **Purpose:** Reflect on what worked well to help guide you to other areas where AI can make an impact. This awareness can also help the organization to align on future strategic goals of AI.
- **People:** Invest in the professional growth of teams that perform the training, modeling, and analysis of machine learning and deep learning applications. Seek out networking events and AI conferences that will help individuals develop interdisciplinary and problem-solving skills.
- **Process:** Changing an organization's receptivity to AI can take time. Continue testing and gathering data to see the gaps in workflows. Initially, it's important to keep your AI initiative small and simple to achieve the quick wins that can inspire leadership confidence.

From defining your goals to getting the right people on board and optimizing your AI processes, you're poised to unlock AI's vast potential in your organization.

Learn more about NVIDIA's deep learning resources at:
www.nvidia.com/en-us/deep-learning-ai/

Partner
Logo

Lorem Ipsum Dolor
CTA LINK