

anexinet

 EBOOK

How to get Started with Machine Learning

www.anexinet.com



How to Get Started with Machine Learning

The field of Machine Learning is still young, incredibly young, especially when compared to the advanced state of the digital analytics industry. While we understand the problems inherent in creating a “get started” guide for a complex technology still in its infancy, recent advances in GPU/TPU technology, and the development of new frameworks and platforms from Amazon and Microsoft has enabled an explosion in Machine Learning and placed pre-built frameworks in the grasp of developers. This guide aims to provide a valuable introduction to the very latest best practice approaches to Machine Learning and enterprise adoption.



Table of Contents:

Chapter 1	<i>Machine Learning - Hype vs. Reality</i>	2
Chapter 2	<i>Top Inhibitors of ML in the Enterprise</i>	5
Chapter 3	<i>5 Tips for getting started quickly with Machine Learning</i>	7

Machine Learning - Hype vs. Reality

The promise of world-changing innovation from Machine Learning (ML) has expanded from the enterprise to everyday consumer life. Some claims even present a utopian world made possible by advancements in Machine Learning. But this technology has also polarized experts: some claim Machine Learning is all hype, while others stand resolutely behind idealist predictions. So let's take a look at what's hype and what's reality when it comes this technology some herald as "the new electricity."



Hype

For the time being, set aside any preconceived notions that Machine Learning is a panacea. Machine Learning will not perform complex operations that connect through every area of your business. What Machine Learning excels at is the rapid performance of simple, monotonous tasks. Its automation capabilities have the potential to relieve workers of tedious, mundane work (e.g. compiling, organizing information), but it won't replace human decision-making anytime soon.



Reality

Most enterprises have not begun to benefit from Machine Learning. One reason enterprises are hesitant to incorporate the new technology is due to the dissonance between current hype (that AI & ML promise long-term disruption) and present impact.

But don't be discouraged. Even present-day Machine Learning creates tremendous value for organizations. Here are some current-day use cases for Machine Learning in the enterprise:



Digital Assistant

For companies with low cloud maturity, the hybrid cloud is an excellent toe-dipping option. An IT department can experiment by migrating non-critical workloads before moving more sensitive data and applications to the new private environment. Additionally, a multi-cloud strategy mitigates the risk of vendor lock-in; even if one service provider goes out of business (or experiences other issues), your business's operation should be unaffected.



Anomaly Detection

New Machine Learning algorithms in the healthcare industry crawl the web and social media to absorb a large group of data inputs. From these inputs, the algorithms determine where anomalies exist and whether the drugs associated with this data are having an adverse effect.



Facial/Fingerprint Recognition

You need only pick up an iPhone or enable biometric security on your laptop to see this one in action.



Signature Validation

In the past, banks would have to verify signatures manually. Machine Learning allows banks to process a signature-image automatically and automatically cross-reference it with the original.

In addition to today's successful use cases, executive leadership across organizations are committed to the integration of Machine Learning. A recent MIT/BCG survey revealed these sentiments:

- 84% feel AI is critical to obtain or sustain competitive advantage
- 75% believe ML provides an opportunity to enter new businesses
- 75% think AI will be the basis for new entrants into their industry
- 69% expect their competition to leverage AI



Top Inhibitors of ML in the Enterprise

So, if organizations are finding success and achieving executive buy-in, why are they failing to adopt Machine Learning? Many internal roadblocks inhibit industry-wide adoption of Machine Learning, including an organization's data, skills, and ability to overcome bias. Let's take a closer look:



Data

Machine Learning requires a significant amount of data. This data must be classified, formatted, scaled, and vectorized. Organizations who try to adopt Machine Learning without first curating their data properly are destined to fail. Enterprises often face more complex issues as they encounter problems that stem from siloed data. Disconnected data sets lead to latency issues, integration issues, accessibility issues, etc. These issues impede the implementation of Machine Learning projects and can cost organizations massive amounts of time and money.



Skills

Most companies lack the proper skills to implement Machine Learning. But hiring a team of data scientists won't necessarily bridge this gap. They must assess the specific skill-gaps relevant to their organizational needs. This presents a chicken and egg issue: how can an organization accurately assess its skill gaps in an area without any experts in said area? There is no generic solution. A viable first step is to seek help from specialists in Machine Learning skill-assessment (e.g. Anexinet) to illuminate the path forward.



Ability to Overcome Bias

Data-driven decision making seems, at the surface, to be unbiased. However (as you may know), data itself often consists of bias. For example, Amazon developed an AI-based recruiting program to find the best hires for any given position. But as it turned out, using past data to determine future success presented a significant issue: their company had historically been male-dominated. This data skewed the Machine Learning algorithms to believe being a male was a determinant of success. Unable to patch the issue and convince the algorithms of the value of the female gender, Amazon scrapped the tool entirely.

Data bias continues to be a major issue across organizations. An organization's ability to vet data properly plays a pivotal role in their ability to benefit from Machine Learning. Left unvetted, flawed data creates major lapses in decision-making, the possibility of scandal, and inoperable programs that drain resources.

Awareness is the key to overcoming these obstacles. Audit your organization's data readiness and assess its skill gaps. Heeding these precautions creates a strong foundation for building Machine Learning models. So now, let's talk about how we can achieve this ML foundation as quickly as possible.

5 Tips for getting started quickly with Machine Learning

The worst reason to get started with Machine Learning is because you think you should. At Anexinet, we have a saying: “no data-science hammers.” Often, when empowered by new technology (hammer), we seek out an endless list of applications to apply it to (nails.) But if you’re truly interested in getting started with Machine Learning, these five steps will provide a framework.

1 Identify Redundant Processes

Redundant processes and Machine Learning make a perfect pair. As an organization grows, certain operational costs balloon, leading to a lower return on investment (ROI). Machine Learning automates these tasks, reducing the cost of scaling.

Answer the following questions to identify your organization’s redundant processes:

- *Is the process repeatable and straightforward?*
- *Are the costs of the processes scaling?*
- *Is ROI rising accordingly?*
- *Would your organization benefit from automating this process?*

If you answered yes to all of the above questions, congratulations! You’ve successfully identified a valid use case for Machine Learning.

2 Brainstorm and Prioritize Use Cases

After you've identified your organization's redundant processes, the next step is to hold brainstorming sessions that focus on the Machine Learning use cases that correspond to these processes.

Reference external use cases. If the use cases aren't industry-specific, break them down by type (e.g. anomaly-detection). Capitalize on the successes and failures of early adopters.

Ask yourself the following questions

- *Would automation improve the accuracy or speed of specific operations?*
- *Should we mine certain bodies of data for latent information?*
- *Which tasks would be easier to complete with AI/ML?*

3 Brainstorm and Prioritize Use Cases

Earlier, we discussed the complexity of assessing data readiness. Not every organization can simply run a classifier. There are a host of concerns. So, to complete this step, be sure you've answered the following questions:

- *Is your data freely available between the sectors of your organization?*
- *Are there any corporate silos to overcome?*
- *Is your data warehouse reliable?*
- *Do you provide data concurrency?*
- *Is your data governance solid?*
- *Will you have future availability of your data?*
- *Do you have access to historical data for trending and analysis?*

Thoroughness at this step will help you avoid problems in the future. Once you're sure you've identified the right data, the next step is to ensure it's appropriately formatted for scaling, normalization, and vectorization.

4 Demystify ML Platforms and Models, and Start Building

Once you've assessed that you have the right data, and that it's properly formatted, conduct your due diligence on Machine Learning platforms and models. Perform a cost-benefit analysis of the available pre-canned solutions and custom alternatives.

Here are some examples of pre-canned AWS enterprise data sets:

I need to know how my customers feel.

Blazing text sentiment analysis

I want to forecast sales over time.

DeepAR forecasting

What's my projected click-through rate?

Factorization machines

Who am I looking at?

Image classification

Where should I target my ads?

K-means clustering

Should I extend credit to this person?

K-nearest neighbors classification

I need to group emails/tweets by topic automatically.

Neural topic model

I need to monitor and analyze video feeds automatically.

Object detection

I need to translate and/or convert speech to text.

Seq2Seq

If none of the pre-canned examples align with your situation, developing a custom ML model is a viable alternative. Building, training, and tuning your own model will your organization time and resources through process-automation.

5 Acquire the Skills to Sustain ML Efforts

Sustaining a Machine Learning model requires a diverse (and sought-after) skill set that varies depending on the character and scope of the project. Assess the skills and capabilities of your team and evaluate the requirements of your Machine Learning endeavors. If your organization has sufficient resources and skills to form a team, do so. However, realize that your organization may need to hire new talent to meet the skill requirements of effective Machine Learning. For enterprises just starting out on their Machine Learning journey, the latter is more often the case.

Every organization is interested in leveraging the benefits of Machine Learning, but not everybody understands what it takes to do so. This is only natural since Machine Learning is still in its infancy. As we've discussed, organizations face issues with data readiness, skills alignment, and data bias. While this guide alone can't sufficiently guide organizations through all the complexities of Machine Learning adoption, we hope it helps you on the road to asking the right questions.

Lastly, if you're interested in a more hands-on experience, please check out our Machine Learning kickstart. In just two weeks, we'll arm your organization with a cohesive strategy for introducing Machine Learning into the enterprise. We flesh out your ideas and priorities, enabling you to overcome the top inhibitors and reap the benefits of this new and exciting technology. Don't wait! Get started now. After all, there's a high price to doing nothing.

Reach out to Anexinet today to schedule your **Machine Learning Strategy Kickstart!**

[Get Started](#)

