

Best Practices for Developing an AI Strategy

By: Priyank Mohan

Today, Artificial Intelligence (AI) is more focused on performing a single task very smartly rather than providing a comprehensive solution covering many areas that require intelligence. Research firm Gartner predicts that by 2020, AI will be pervasive in almost every new software product and service and that it will be a top five investment priority for more than 30 percent of CIOs. Today's AI in the communications service provider (CSP) space is primarily focused on machine learning—a branch of artificial intelligence that focuses on the development of intelligent computer programs with the capacity to predict. These programs can learn autonomously by training themselves using historical data and can improve when exposed to new data.



AI has been around for decades, so why the big push by CSPs into AI today? There are three key reasons. The first is Big Data. We have more data than ever, which itself is allowing machine learning to improve and provide more relevant insights. For telecom and cable service providers, today's digital world provides a plethora of data that can be put to good use and analyzed using AI technologies.

The second reason is reduced processing costs. Until recently, the cost to set up infrastructure and build a specialized team was high. AI also required huge R&D budgets and investment in made-to-order algorithms. Today, AI is going through exponential adoption because of four main factors: the rise of ubiquitous computing, low-cost cloud services, inexpensive storage, and new algorithms. Cloud computing and advances in Graphical Processing Units (GPUs) have provided the necessary computational power, while AI algorithms and architectures have progressed rapidly, often enabled by open source software. In fact, today there are many open source options and cloud solutions from Google, Amazon, IBM and more to address infrastructure costs.

The third reason AI is surging in CSP organizations today has to do with breakthroughs in deep learning technology. A subset of machine learning, deep learning has structures loosely inspired by the neural connections in the human brain. Most of these big deep learning technology breakthroughs happened after 2010, but deep learning (neural nets) have already demonstrated the ability to solve highly complex problems that are well beyond the capabilities of a human programmer using if-then statements and decision trees.

There are four key requirements that every machine learning project should meet:

1. Large volume of historical data with clear success/fail criteria
2. Well-defined information model, meaning the data can be understood and content can be parsed
3. Clear area of value, such as low levels of process automation, high levels of order fallout, poor customer experience, and/or unpredictable performance of network or processes
4. Ease of developing a proof of concept, so value can be demonstrated quickly

AI Best Practices

Like everything else in business, AI vision and strategy work best when they come from the highest level in the organization—which means every business leader should be aware of what AI makes possible. Simply hiring machine learning engineers or data scientists reactively is not a strategy. At the same time, operational employees should know the possibilities AI offers, too. For best results,

AI needs to be democratized and socialized throughout all levels in the organization. Customized AI training should be offered for all levels in the organization so that employees can make more informed business decisions. A lack of training may create uncertainty and a fear of lost jobs due to the advent of AI technology.

There are many categories of AI: computer vision, image recognition, deep learning/machine learning (applications and platforms), natural language processing, gesture control, personalized recommendations, smart robots, speech recognition, video analysis, content recognition, speech-to-speech translation, and virtual assistants, among many more. From a maturity standpoint, AI is still fairly new in many of these areas, and solutions are very specific and tailored to solve one problem. There are many AI solutions and products in the marketplace, but one size does not fit all.

While agility in machine learning deployment models will not be exactly like the agile software methodologies that many CSPs are familiar with, there are many common elements. The main difference is that machine learning agile methodologies are data-driven. To close that gap, good agile project management practices need to be customized and applied in a machine learning context for the organization. Often there is a disconnect between what the business needs versus what is able to be produced by machine learning engineers using the data and time available. How do you deploy models when the business needs them in two weeks, yet fine-tuning your model will take three to six months?

When you are dealing with millions of customers and vast amounts of data, as is common for many CSPs, the challenge becomes greater. Machine learning in production is becoming less about algorithms and more about the data workflows surrounding them—how to train machine learning models in the lab, deploy them into production, monitor and evaluate their performance, and improve them. If data flows are long, expensive or manual, they pose a big problem. In these cases, the strategy needs to be rethought and alternative solutions need to be considered to simplify the data flows.

Many CSPs question whether they should have one centralized AI division for their organization or let each group use their own strengths. The simple truth is that most large organizations are not ready to have one artificial intelligence division to address all internal and external organizational needs on day one. It is still too early to have a single AI foundation. To illustrate this point, let's use the example of one major U.S.-based telecom company. It has three AI divisions: one focused on operations and customer care; another focused on global supply chain strategy, which makes sure that products reach customers and that sourcing and procurement are effective; and a third focused on big data and artificial intelligence systems that create new data products. While these divisions remain separate, they do have a common organization that addresses data management, data governance, data warehousing, and data lakes, and common analytical and AI technologies. The goal of this common organization is to facilitate cross-functional, cross-organizational projects in a large enterprise. However, for smaller service providers, a centralized AI function may be a viable option.

AI in the Back Office

Business Support Systems (BSS) are foundational systems in any telecom company. They assist with taking orders, addressing payment issues, and tracking revenues, among other task and support processes such as product management, order management, revenue management, and customer management.

To improve the customer experience, most service providers are implementing new omni-channel and self-service capabilities. In these new revenue areas, service providers require support from their BSS platforms. The digital world is also changing the way service providers manage, sell, and support their core services. More service providers are adopting a digital-first paradigm, which puts the emphasis on automation.

Many traditional telco transactions like provisioning can occur without human intervention and in real or near-real time. Telecommunication companies are also aiming for more personalized interactions with their customers by using big data and analytics to tailor marketing and upselling

efforts to specific customers and segments.

When planned and executed properly, AI has tremendous potential for service provider organizations. To create success, it's best to stick to a tried-and-true methodology:

1. Look at the big picture
2. Define a use case for AI (Choose a use case which has high business impact and relatively low complexity)
3. Obtain the data to support the use case
4. Discover and visualize the data to gain insights
5. Prepare the data for machine learning
6. Select a model and train it
7. Fine-tune your model
8. Present your solution
9. Launch, monitor and maintain your system.

Simply put, success starts by focusing more on the use case, the user experience, and getting the right set of data than worrying about the latest and greatest AI algorithms.