STRATEGIES FOR A SUCCESSFUL **MACHINE LEARNING JOURNEY**

LEVERAGING DATA FOR MORE ACCURATE PREDICTIONS AND BETTER BUSINESS OUTCOMES

argyle.

TATA CONSULTANCY SERVICES

The true value of data lies in the ability to extract information, derive knowledge from it, and efficiently and effectively share that information across the enterprise to make better business decisions.

Numerous machine learning tools and techniques are available to achieve these goals, but the effort from qualified people to extract the best from the techniques is still considerable. Machine learning has not yet reached a stage where it can extract information, knowledge, and wisdom without manual intervention.

While machines can help sort and analyze data, gathering raw numbers more efficiently through machine learning is just the beginning. The data must be presented in a way that makes it meaningful to stakeholders at all levels.

“Building a team – and then retraining them with the proper skills is critical for successful man machine engagement” stated Nidhi Srivastava, Vice President and Global Head of Enterprise Intelligent Automation, TCS.

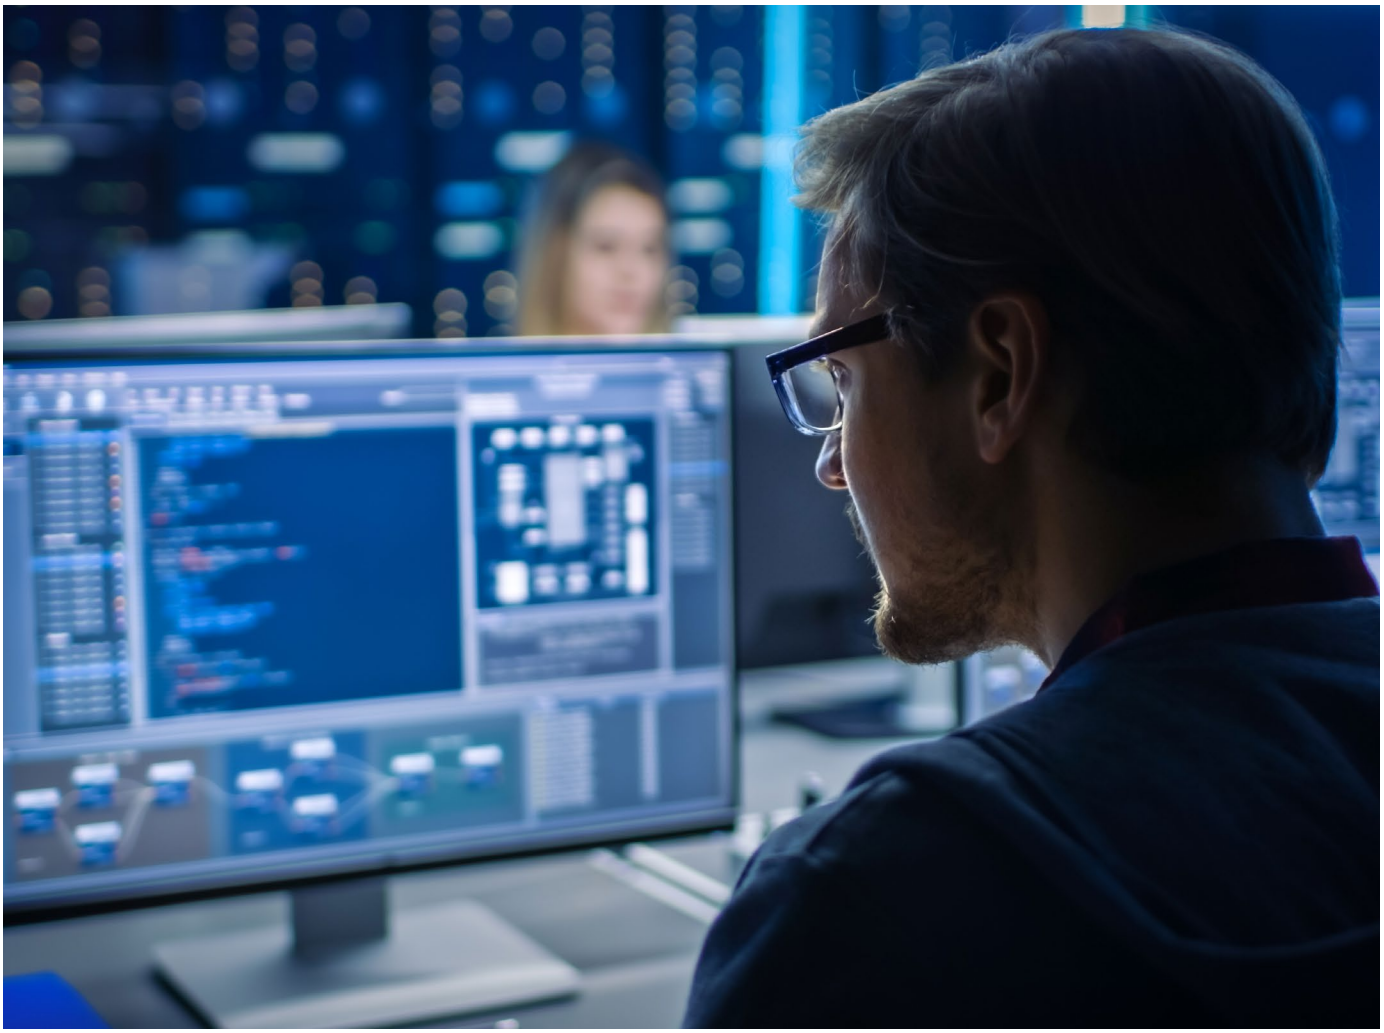

MACHINE LEARNING BASICS

Below is a quick overview of the various types of machine learning you may encounter as you explore ways that machine learning can impact your organization.

A subset of artificial intelligence, machine learning is a process or rule or model to predict future outcomes. Instead of hard coding the rules, machine learning uses data to develop a model.

There are a few ways that machine learning can use data. The first is **directed learning.** In this scenario, the machine learning engine is given examples of inputs and outputs and learns to create outputs from those examples. One example of directed learning in the financial industry is credit card authorizations. By feeding the system past examples of applications it learns to make decisions on future applications.

The second way that machine learning can leverage data is through **undirected learning,** in which the system aims to detect anomalies. A good example is a money laundering detection scheme. There are few examples of money laundering, but the machine learning engine can look at the data for irregularities.

Reinforcement learning is a combination directed and undirected learning, where models are refined along the way based on learned experiences and human input. An example in the finance world is fraud detection. When the system flags a transaction as potentially fraudulent, a human looks at it and determines if it is fraud. The system becomes more sophisticated over time using this feedback and gets more accurate at detecting suspect transactions. Treasury and account reconciliation can also benefit from reinforcement learning. When payments are being wired or going to a lockbox, for example, there is not necessarily an invoice number associated with the payment. ML can significantly increase the match rate as it learns over time.

Currently the most exciting branch of machine learning is the use of neural network models. Consumers are already seeing neural networks in action as they underpin technologies like the voice recognition of Google Assistant, Siri and Alexa, and the actor "de-aging" techniques seen in recent Hollywood and TV productions.

The advancement of computing power behind neural networks over the past five years, coupled with lower costs, has made it feasible for machine learning to take on many tasks that were not possible previously. Neural networks enable organizations to incorporate external data such as inflation, unemployment and other financial indicators to predict future outcomes.

Statistics, math, and probability are the key drivers of any machine learning application. You don't want to bring in inaccurate data or unformatted or data that is not normalized into the process. The objective is to learn patterns from data and use the patterns to influence the business processes.

Machine learning often relies on a combination of structured data, such as credit card transactions, and unstructured data, which includes voice, images, and unformatted text such as emails. Machine learning is a powerful computing platform which is now readily available at low cost from cloud providers.

The value of machine learning likely increases when data boundaries are blurred. Many organizations are siloed today, limiting the value of data to the respective silos.

As they have a responsibility to manage costs, finance leaders naturally want to know how much computing power will be required for machine learning. At times distributed machines will need to be put into service to process the data or build the models. Some options include:

- **Private cloud**
- **General processing, done through PPUs**
- **Specialized actualizers for AI**
- **Specialized vertical-specific hardware with in-house GPU**
- **Hybrid of private and cloud provider**
- **Third-party provider**

The issue comes down to generic purpose hardware versus accelerated hardware versus application-specific hardware. It could be in house, purchase on a cloud, or could be a hybrid approach. There are general cloud providers and specific AI cloud providers like DataRobot as well as our TCS ignio™ product. The many options available can cause confusion.

With this in mind, many organizations are taking a hybrid approach to meeting the computing needs of machine learning. Computing power is continuously evolving, and companies need flexibility. General mainstream public cloud providers such as Google and AWS offer inexpensive, powerful and flexible ML-optimized options to support complex ad-hoc ML tasks that won't be in long-term operation.

Due to the high computational requirements of the deep learning functions of machine learning, there are specialized hardware accelerators and optimized software offerings to exploit existing hardware, most of which are based on GPU hardware as they are optimized for the massive parallel calculations needed to draw realistic 3D game graphics, the same type of high-volume parallel calculations needed for training neural networks.

There are multiple ways to increase the throughput to meet the capacity required for machine learning with the help of mature, open source libraries. Software optimization is the first step and many software frameworks are being altered to perform better on heterogeneous hardware. Better utilization of the cores in a multi-core CPU is one such example. Other techniques include mixed precision computing to increase the throughput of GPUs.

Any machine learning solution requires a lot of data and it may not be easy to obtain. The data may reside in different silos across the organization. Organizations may not know the location of all their data or where it is being collected.

While many organizations face significant challenges when it comes to gathering data and building a data lake sufficient to accommodate machine learning, there are other hurdles to overcome. Machine learning involves collecting large training data sets that the model can over-fit to the training data. It becomes very good at identifying the correct classifications on the training data set, but it becomes so specific to that set that it does not generalize well when new data is introduced. This reduces the effectiveness at run-time when receiving new records that need to be classified.

It takes some experimentation, the expertise of data scientists and possibly third parties, to find the right amount of data to get relevant outcomes.

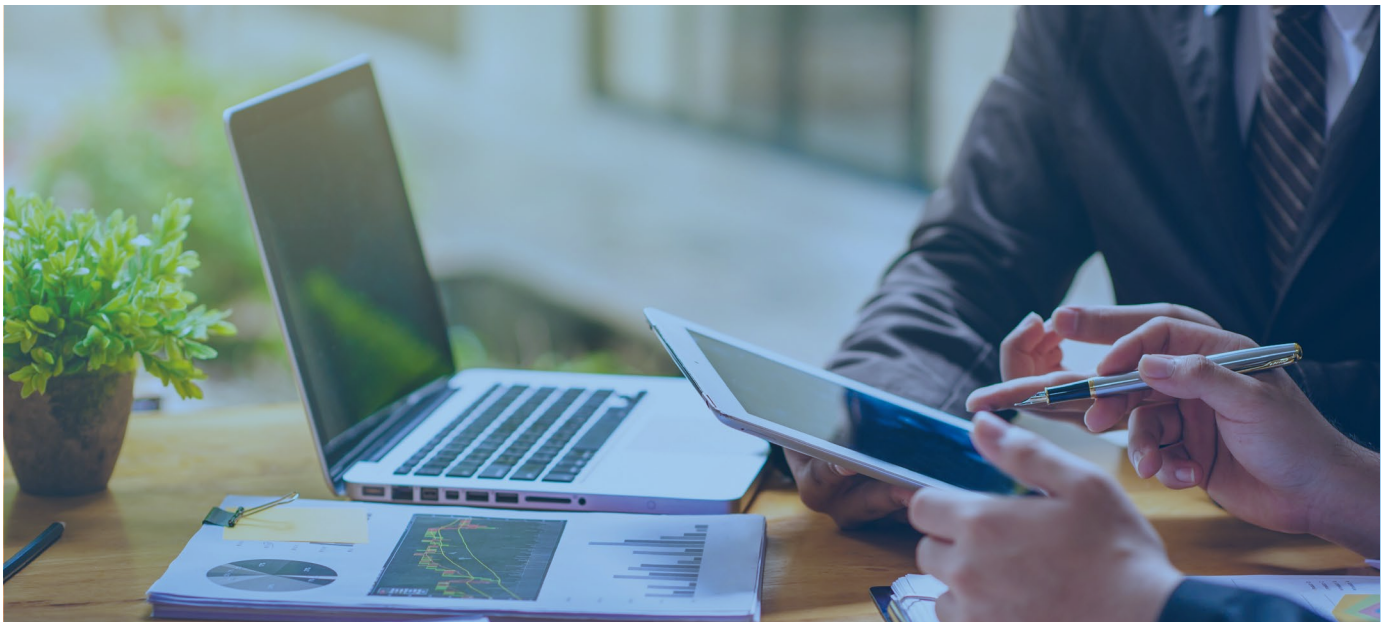

WHERE AND HOW TO APPLY ML

One of the main benefits organizations can gain from ML is the ability to learn patterns from data and apply them to influence the business processes. There are a number of scenarios where machine learning can provide a significant return on investment.

Some of the biggest gains can be seen when machine learning is used in areas that require high speed and accuracy that could only previously be achieved by adding more human operators. In finance, for example, the accounts payable team spends a good amount of time reconciling accounts and matching invoices with payments. Machine learning can handle many of those tasks and automate workflows so that the team can focus on enhancing vendor relationships and vendor consolidation.

Managing risk — something always on the mind of CFOs — is another area where machine learning impacts the bottom line by helping companies avoid bad decisions. Examples include flagging unusual volumes or types of orders involving a frequent vendor, identifying mismatches between a bill of lading and the items ordered, or detecting a fraudulent credit application based on the applicant's situation being similar to other incidents of fraud.

Machine learning can also mitigate risk in other ways, such as predicting events that can impact the business. One such example are utility assets that include existing transmission and distribution lines. Poorly managed vegetation could wreak havoc on a good portion of these assets. TCS has developed technology that combines unstructured imagery from satellites, helicopters, drones, land vehicles, and human inspections using a variety of camera and LIDAR imaging technologies. Heavy equipment like aero engines gather a lot of information from all the many sensors they have. Manufacturers can feed this data into machine learning algorithms to recommend maintenance and component replacement BEFORE the unit fails.

Other applications of interest to finance include, ML-based forecasting and scenario planning, Risk-based AP/AR decision making and approvals, and travel and entertainment fraud detection.

"Modern CFOs are also paying greater attention to the customer experience and learning more about how customers interact with the business is another area where machine learning can pay off," said Vikas Gopal, Global Managing Partner, Finance And Shared Services Transformation, TCS. For example, retailers can gather information about the shopping habits of their customers through sensors in the most popular displays as customers walk the aisles. They can leverage that information as they work with vendors on positioning their products in the store.

While some express concern that machine learning will replace humans, successful organizations don't deploy machine learning to replace workers. They invest in machine learning to supplement and enhance human knowledge and remove non-value-added effort from their people. The machine might make the recommendation and have a human approve it – especially for high-materiality items, or where the model's confidence level is below a set threshold. This saves anywhere from seconds to minutes on a transaction, and the manual approvals and rejections can be fed back to the model to improve its accuracy — the reinforcement learning technique described earlier in this article.

In addition, it takes human supervision and decision making to teach the machines over time. There's a growing concern that we are heading to a job loss scenario when decision automation happens. It might be more appropriate to think of it as job shift scenario because decision augmentation to decision automation will take time and great human effort resulting for the need for growing strategic roles.

There is also the issue of acceptance of machine learning among finance executives. Some ML algorithms, such as decision trees, can be reduced to a logical formula based on the inputs that are relatively easy to explain to non-ML experts. On the other end of the spectrum, neural networks do not produce easily justified formulas that can be laid out and explained to a layperson.

Building confidence that these "black-box" algorithms can be accurate and trustworthy requires focus on the quality of output. This requires a shift in mindset for many finance professionals used to auditable algebra as the basis for decisions and accounting.

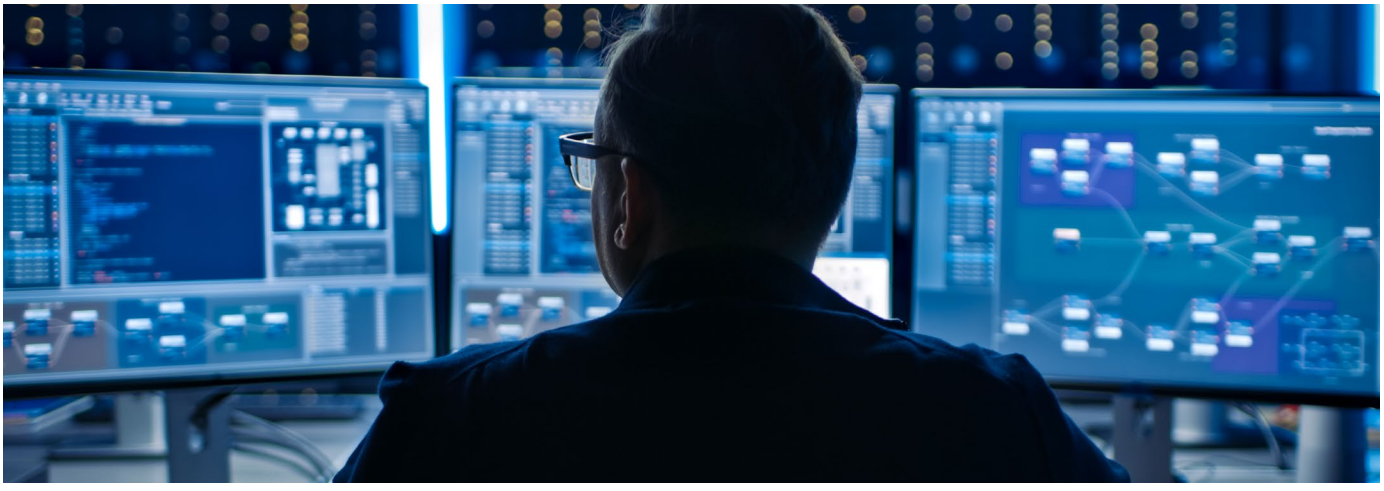

BEST PRACTICES FOR LEVERAGING ML

Enterprise architecture is critical to data governance, model governance and other data-centric functions. In addition to the CFO, the CIO, Chief Data Officer and stakeholders need to be aligned on enterprise architecture.

Short-term goals include increasing efficiency and reducing cost, but organizations need to take a longer view of machine learning as a value creator. While improved productivity can be achieved, companies need to have a vision for how those freed-up resources will work to improve the bottom line.

Some of the things that finance chiefs need to consider as they look to apply machine learning to their organizations include:

- **Do you have the right level of staffing? Are you able to leverage machine learning by upskilling your current team and adding new people with the necessary skill sets?**
- **How will a reduction of mundane work for employees improve the ability to focus on more value-added tasks, and what will those tasks encompass?**
- **How will machine learning increase the speed of decision-making? What will be the impact to customers?**

The ROI of machine learning is not always easily quantifiable, so you need to consider benefits beyond cost reduction. "You also need to consider the speed of decision-making and how that will impact your business," stated Vikas Gopal, Global Managing Partner, Finance and Shared Services Transformation, TCS.

CLOSING THOUGHTS

From a finance standpoint, machine learning is mapped to reducing headcount and minimizing cost but there is more strategic value it can provide. Machine learning predicts and detects errors, which has a huge impact on efficiency and accuracy. Machines also can detect patterns that human analysts would never find.

"Going forward CFOs will adopt machine learning for forecasting and planning across the organization as their finance departments expand the strategic value delivered in their roles. Machine learning algorithms can remove human bias from the numbers, and realizing their potential requires a cultural shift," stated James Byrne, Global Lead Partner - Finance Transformation, TCS, "The investment in machine learning to create forecast scenarios would be rendered worthless if it is routinely overridden by humans before submission."

Organizations are recognizing the value of machine learning goes beyond improving efficiency and handling routine tasks. They need to get on board now or they will lose out on the competition. Without investing in machine learning they not only risk becoming obsolete. They longer they wait, the more costly it will become to incorporate later.

ABOUT THE SPONSOR

Tata Consultancy Services is an IT services, consulting and business solutions organization that has been partnering with many of the world's largest businesses in their transformation journeys for the past fifty years. TCS offers a consulting-led, cognitive-powered, integrated portfolio of business, technology and engineering services and solutions. This is delivered through its unique Location Independent Agile delivery model, recognized as a benchmark of excellence in software development.

For more information, visit **www.tcs.com**.

TATA CONSULTANCY SERVICES