



Improving efficacy of AI models during times of business disruption

Over the past few years, enterprises have increasingly deployed artificial intelligence (AI) in both core and non-core businesses. These AI algorithms range from basic predictive analytics to determine a course of action, facilitating interactions between a customer service representative and customer, to automated generation of business insights that can help guide business strategy in pursuit of increased revenue and improved margins.

While most AI in use today can be classified as early-stage advanced analytics, some enterprises have built large data science teams to apply machine learning and deep learning algorithms to business processes. Many of these enterprises have built the necessary support infrastructure to train these algorithms on large data sets, deploying the resulting AI models to production to generate business insights.

However, many consumer and business consumption patterns changed dramatically in 2020, causing these advanced AI models to fail or behave erratically. Many of these models that have been trained to make predictions based on historical data have not been able to deal with the data anomalies created by disruptive business conditions and changing preferences. Companies using AI for insights had a hard time making use of existing models in production. Many companies have reverted to using the standard analytics dashboards and tools that were in place before AI adoption.

The impact of macroeconomic conditions on AI models has been greatest in the customer service area, where next-best action and other recommendations are based on predictions from historical data. With data about changing economic patterns, sentiment analysis models predict negative customer expectations.

What this demonstrates is that machine learning techniques, such as supervised and unsupervised learning that depend on massive amounts of data (**Figure 1**), do not work when economic patterns change dramatically because the existing data is no longer relevant and useful for predicting meaningful business insights.

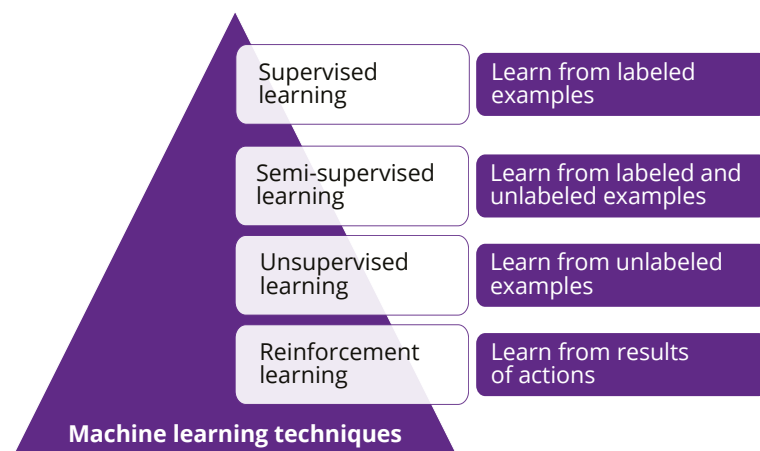


Figure 1. Machine learning techniques and their comparative advantages

Reinforcement learning is the process of a program ... adjust[ing] its logic, effectively learning to solve a problem over a period of time

As machine and deep learning techniques evolve, two promising approaches are reinforcement learning and deep reinforcement learning. These learning techniques are less dependent on data than traditional machine learning and offer greater insights when conditions are changing.

What is reinforcement learning?

Reinforcement learning is the process of a program learning as it accounts for the consequences (both good and bad) of actions it undertakes, such as predicting an insight. Based on these consequences, the program adjusts its logic, effectively learning to solve a problem over a period of time rather than by trying to produce insights from current data sources. Instead, the process is generating the data it needs from the actions it takes. Using this data and knowledge of the consequences of previous actions (collectively called state space), the program performs new actions to generate a beneficial consequence, which in turn produces more actionable insights.

Reinforcement learning is not dependent on input data sets as is the case with other machine learning algorithms. The actions and their consequences determine the generation of training data that helps the algorithm learn.

How to generate useful machine learning-based insights

To take advantage of machine learning in order to generate useful insights in a rapidly changed macroeconomic environment, enterprises will need to make some significant modifications to their approach. One way forward is to combine retraining current models with the adoption of reinforcement learning models and commit greater resources to employing experienced data science teams.

To start with, AI models currently in use need to be redeployed with new versions that should be trained with current business data. The models should be versioned along with the data sets used to train a version of the model. Versioning should include parameters used to tune the model.

Next, enterprises could use the opportunity to train these models on once in 100- to 500-year business disruption events. Thanks to the fact that data providers have been collecting data on business disruptions during unexpected events, a plethora of data is now available to train existing models. At the same time, the parameters used to tune the model with this additional information should account for anomalies in the data.

Trained data sciences team — integral to the process

Integral to the whole process are data science teams. First, enterprises need to ensure they have an experienced team in place, and these data scientists need to be focused on correcting any anomalies in the insights that the models are predicting.

This means data scientists should not be pushed aside in favor of “citizen data scientists,” users who are not employed as statisticians or analysts but use tools that automate some aspects of data science. While there is certainly an important role for the citizen data scientist, these individuals are not trained to identify anomalies, and this skill set is vital to addressing the problem with using AI models in unpredictable environments.

In addition, enterprises should increase the use of external data to train the models since having a more diverse set of data will help to normalize the dependence of

insights from internal data. However, which data and how much of this data will be relevant to the given situation is something only data science experts can identify.

The final step is to make use of reinforcement learning in parallel with current models to get another set of data points that will help in making decisions.

Pricing optimization — an insurance industry example

The benefits of reinforcement learning in generating useful insights is well illustrated in the insurance industry (**Figure 2**). Insurers are increasingly using predictive algorithms to price life insurance policies — both new and renewals — as part of a broader model that supplements the traditional role of actuaries. Actuaries, whose knowledge and expertise in statistics is integral to projecting risk and pricing in the industry, are now working with data scientists, who are increasingly being brought in to collaborate on initiatives that will enable the business to optimize premiums.

Evolution of insurance premium pricing



Figure 2. Evolution of insurance premium pricing based on reinforcement learning-driven insights

The current price optimization technique involves using data sets that include competitor prices, individual customer information — their buying patterns and socio-demographic profiles — market prices for similar insurance products and data about macrotrends in the economy.

This data needs to be prepared — from data cleansing, to moving the data into the AI model and then managing it — so machine learning algorithms can process it. To better prepare the data, data scientists use various algorithms and techniques — from linear models to more sophisticated neural network models — to understand

changes in premiums and their impact on sales and other business metrics such as combined ratio. Based on predictions from these models and after adhering to local regulatory requirements, they are able to price premiums for both new business and renewals, which is increasingly resulting in personalized premiums for individual customers.

However, changing macroeconomic conditions and lower demand have made these models more erratic. To continue to gather valuable insights for price optimization, machine learning algorithms that depend on data sets to make predictions need to be tuned based on newly identified parameters, and then redeployed. This is where reinforcement learning excels, helping insurers make sense of the data in the context of changing economic circumstances.

The way a reinforcement learning algorithm learns in this context is by setting a baseline strategy of increasing premiums where there is new business and retaining a defined threshold level in the case of renewals.

The algorithm would start with a set of customers who face an increase in premium (based on business strategy and multiple financial metrics) and learn from the consequence — either increased or decreased sales or retention at a threshold level or below. This consequence results in the algorithm learning the best action for the next set of customers. Over time, it learns what works and what does not work in terms of meeting business strategy and related financial metrics.

The result of the algorithm could be used by the insurer in conjunction with prediction from the current machine learning models to make a final decision on setting premium prices for new policies and renewals. Once that decision is known, the insurer would feed the result back to the algorithm for it to improve its logic for future decision making.

Conclusion

To adjust to changing economic circumstances in AI modeling, enterprises should expand their analytics and decision-making capabilities beyond traditional predictive analytics that rely on large, historical data sets. Adding reinforced learning algorithms will help models learn over time with much less input data. Combined with current analytics and machine learning capabilities, reinforced learning provides enterprises with valuable input, helping them to gather relevant insights and make good business decisions even during sustained unexpected business disruption.

Learn more at
dxc.com/insurance

Get the insights
that matter.

dxc.com/optin

