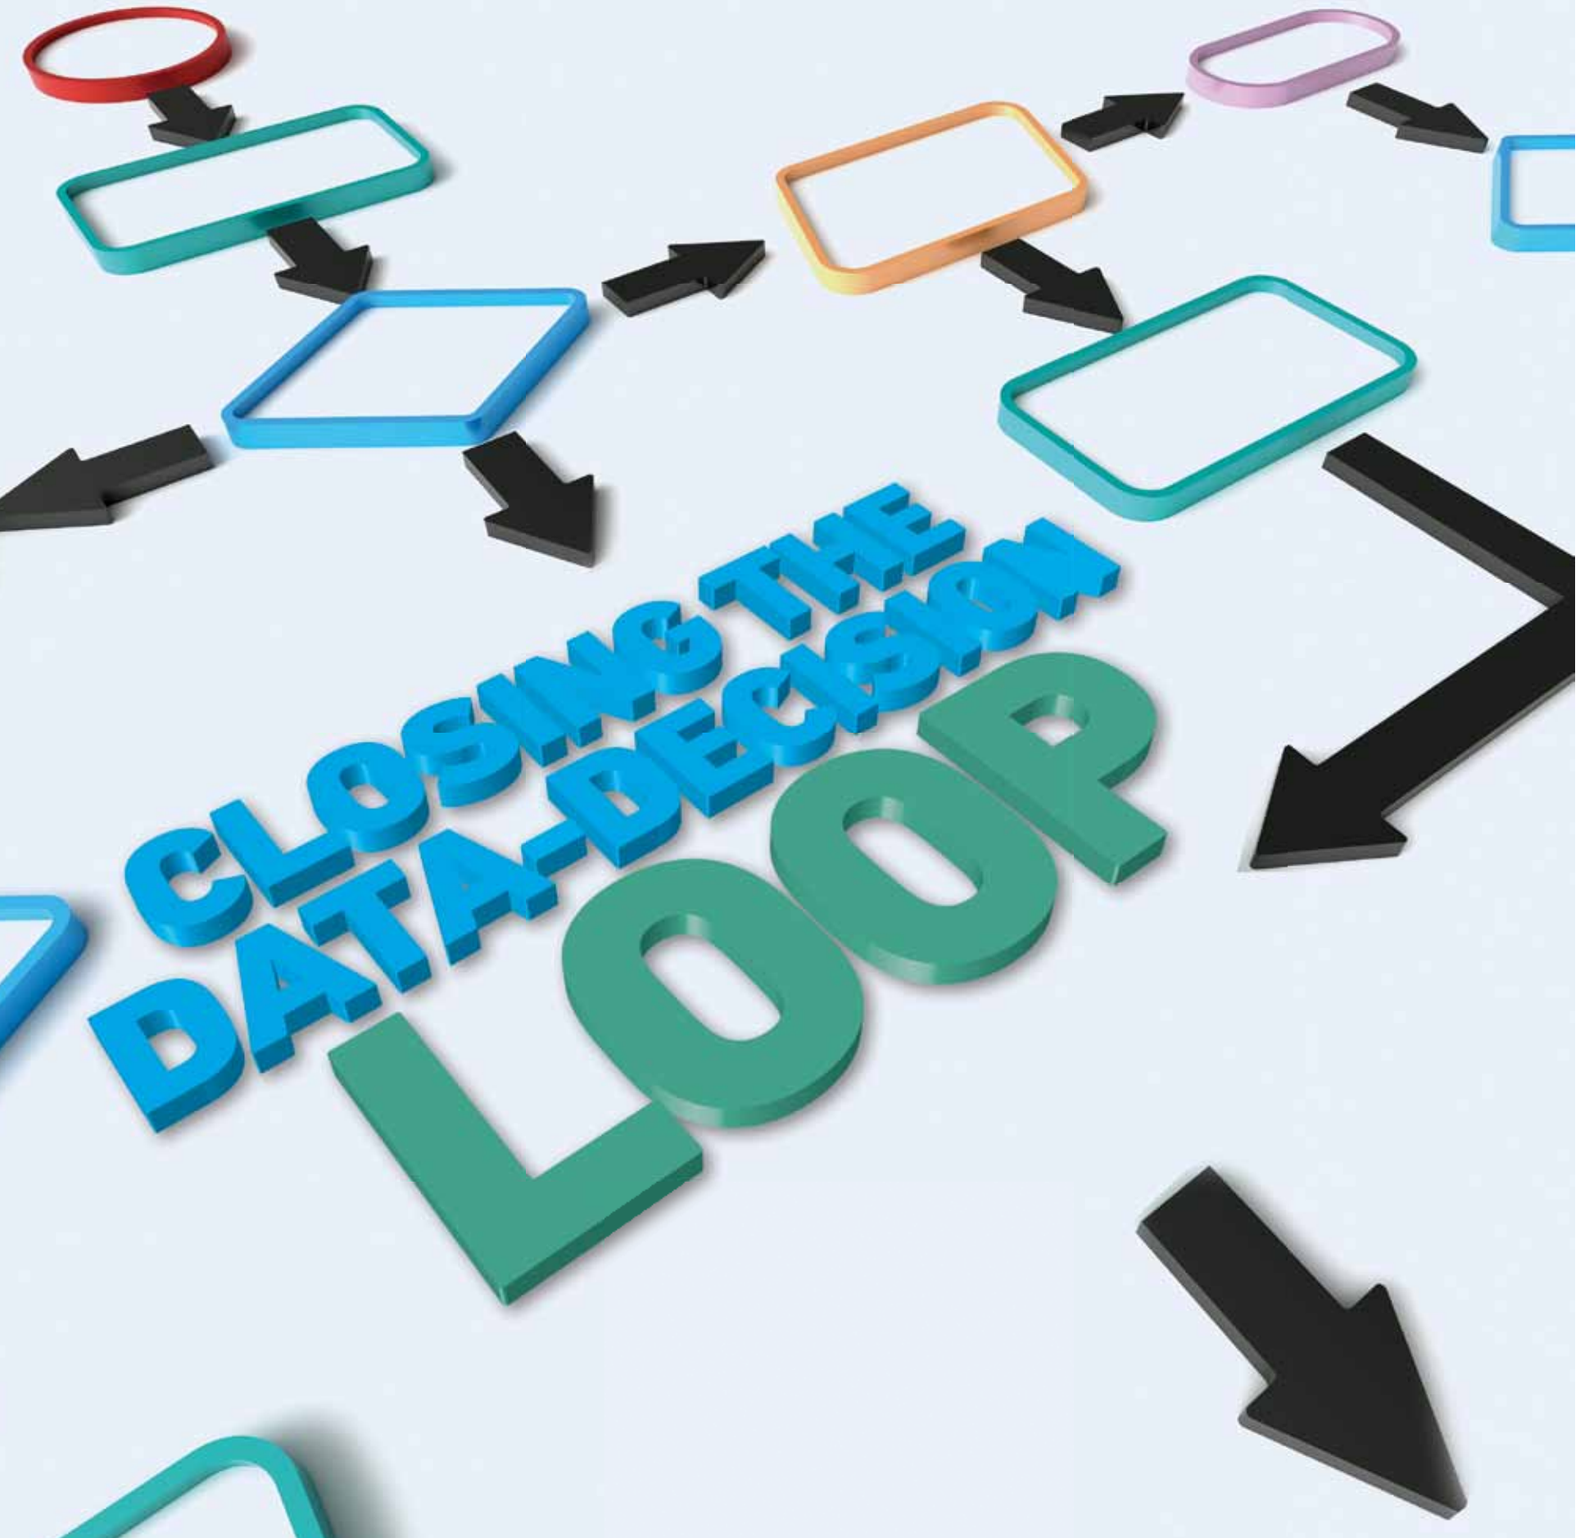

Deploying artificial intelligence for dynamic resource management.

By Pradeep Reddy Varakantham

For businesses everywhere, this is the epoch of artificial intelligence (AI). Major advances in deep learning and machine learning mean that machines can now dissect, evaluate, and interpret data in ways that were inconceivable just a decade ago. As a result, companies across all industries are harnessing AI to become more productive, to predict better and to market better. While we have been using computers for over six decades to help with business decision-making, today we find that AI systems are demonstrably delivering state-of-the-art performance for enabling business decisions, notably in resource planning, deployment and management.

The challenge of resource management

Every business is faced with the most basic, and also the most complex, decision of optimally allocating its limited resources, both human and non-human, to meet multiple demands in a way that brings maximum benefit to stakeholders. On the factory floor, this takes the form of scheduling labour and allocating raw materials to achieve the target output in the most efficient manner. In the warehouse, this calls for arrangements that optimise space and time while reducing errors and improving flexibility. Similarly, the transportation and logistics sphere requires optimal allocation of vehicles to ensure the fastest and cheapest ways to get products to market. At the last mile, sales and service delivery personnel also need to be deployed in a similar way, regardless of whether they are working for a brick-and-mortar store or an online business.

In the real world, this optimisation of resource allocation takes place in a dynamic setting, where the volume of resources (i.e. supply of labour, machines, raw materials, warehouse space, trucks, and service delivery personnel) and the demand for the goods and services they produce is constantly changing. Such dynamic deployment requires a configuration that can anticipate these changes and continuously adapt to them. Even this optimisation sequence that takes into account uncertainty in supply and demand is based on *ceteris paribus*, and can be disrupted by the emergence of new products, markets, competitors, business models, as well as new environment, health and governance mandates. Such adversarial influences can throw off even the best-laid plan, hence they require an immediate response to that particular incident. Such disruptions range from the uberisation of the taxi industry to the recent outbreak of the coronavirus.

While the challenge of resource management is not new to business, its scale and response time have significantly altered in recent times with the advent of e-commerce players, e-marketplaces, and social media. The scale at which companies like Amazon and Alibaba operate at entails planning, scheduling, and deployment of tens of thousands of people, along with their supporting physical assets. The number of cars managed by Uber in the U.S. today far exceeds that of any taxi company. The move towards omni-channel retailing further adds to the complexities of seamless delivery in the online and offline space, and this is even more so for services that combine both, where the online order is collected in store or store shopping is delivered.

Furthermore, time horizons are getting shorter as businesses compete to delight customers who are demanding ever higher levels of service. Amazon Prime is now delivering groceries within two hours, with many other companies scurrying to compete. One can only imagine the level of planning that goes into deploying transport and manpower (for shopping, packing, and delivery) to get a full range of fresh and frozen items delivered to one's doorstep in such a short period of time. Meanwhile, any news, especially bad news, goes viral within minutes on social media, even as stakeholders demand that the company acknowledge and produce a response plan in real time. Hence, businesses find themselves planning not only for the reporting month/quarter–but also for the next hour.

While we have considered the example of a merchandise supply chain, the same principles of constrained optimisation amid uncertainty apply when taxis are allocated to passengers; hospital staff and rooms to patients; ground staff and facilities to airport immigration, baggage and other passenger services; and in the deployment of emergency vehicles, security patrols and law enforcement personnel. Hence, the need for efficient resource deployment in business in the face of resource constraints, uncertainty and imperfect information, and unforeseen incidences, is pervasive. Businesses that are built to last are quick enough to respond and flexible enough to adapt to such ambiguous situations.

Where AI fits in

What do AI systems do better? While computers that run optimisation programmes have been around for a while, and are being used by businesses across the board, AI-enabled systems are able to: (1) deal with scale, process millions of data points within seconds, and provide answers in real time; and (2) learn from past data to continuously improve decision-making. Enhanced intelligence enables simultaneous processing and deployment of tens of thousands of resource segments under hundreds of scenarios–all in real time. They not only help to dynamically optimise the allocation of resources, but allow for short-term and medium-term planning to be done simultaneously.

CLOSING THE LOOP BETWEEN DATA AND DECISIONS

In a dynamic environment, demand, supply, and adversary patterns are constantly changing. One moment, you think you understand the universe and plan for it, and the next moment, the market environment can change dramatically. Business today is all about incident response and incident prevention.

While the use of enterprise software to collect operational data and build robust databases is commonplace today, very few businesses are able to use or leverage that data for better operational performance. AI-enabled tools help to close the loop between data and decisions, that is, they help to make meaningful extrapolations and predictions based on the data, which in turn help in decision-making. Once a problem is identified, a model constructed with algorithms that best represents the available data is used to obtain optimal strategies for matching supply with demand. These strategies are then tested out on a fresh set of real-world data (or on simulated data if real-world data is not available). Over time, as the experience set gets larger, the machine learns to further refine decision-making by taking into account a wider set of data. Decisions are made through an iterative adaptive process based on the observed data.

AI-enabled tools help make meaningful extrapolations and predictions based on data.

BIKE-SHARING IN SINGAPORE

Early bike-sharing systems suffered from two major challenges: first, when the fixed docking stations for bikes were full, riders were forced to drop off bikes at stations far from their destinations; and second, in certain locations, demand often exceeded supply due to the movement of bikes by riders. The critical issue, then, was how to make bikes available to customers at the 'right' locations and at the 'right' times.

Singapore sought to address these issues when it introduced bike-sharing in 2017. The three bike-sharing companies in Singapore—Ofo, oBike, and Mobike—did not use docking stations and allowed for bikes to be dropped off or picked up anywhere using GPS tracking. This addressed the fixed-capacity issue that traditional docking stations could not address. Unfortunately, it gave rise to an even more complex issue—of bikes being left at locations where they posed a hazard, impeded pedestrian traffic, or were not easily accessible. Given the thin profit margins, tackling this inefficiency was key to the sustainability of a privately-owned bike-sharing system.

The first way to address such inefficiencies was to employ trucks that take the bikes to places where customers need them the most, at different times of the day. All the bike-sharing systems typically store data on bike usage by customers, and algorithms help to identify customer bike-usage patterns. Truck routes were thus identified based on these usage patterns, so that bikes could be collected from areas where there was a surplus, and deposited in areas where there was a shortage. Locations with a similar level of demand were grouped together into regions, and the relocation of bikes was carried out at the level of these regions. While the trucks were useful in relocating bikes, they increased the company's carbon footprint, potentially offsetting the environmental gains that were supposed to come with bike-sharing systems.

Therefore, the second way was conceived, where the onus was placed on users, instead of intermediaries, to position bikes in the right locations. An incentive-based system in which customers were offered a reward to leave the bikes at the desired locations—meaning locations where there was likely to be greater demand for bikes—could work. The key challenge in providing such monetary incentives was balancing the trade-off between being attractive to customers (leaving the bike at a location desired by the bike-sharing company) and being financially feasible for the bike-sharing company.

While data-based prediction methods helped predict consumer demand, they were not sufficient to compel customers to move bikes to desired locations. The AI system—a data-driven decision-making system—could figure out the right financial incentives to use, given the multiple scenarios based on the future demand of rational consumers operating with a fixed budget. By combining data and decisions within one closed loop for bike sharing-systems, performance was significantly improved by up to 40 percent, compared to methods using multiple real data sets.

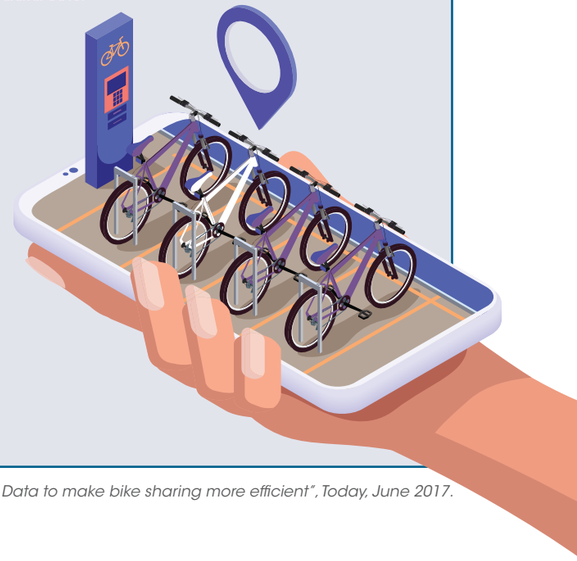

Source: Pradeep Varakantham, "Using Big Data to make bike sharing more efficient", Today, June 2017.

AI-enabled technologies can support complex decision-making processes for allocating/deploying resources to meet service delivery requirements through the following tools and processes:

Data-driven decision optimisation

By considering multiple scenarios of potential future demand (sampled from past data) and exploiting key structural properties, large-scale resource management problems (with tens of thousands of resources) can be optimised in near real time.

Reinforcement learning

This is a technique typically employed to learn from reinforcements (or reward signals) obtained from the environment. A key advantage of this method is that the technology can be trained offline for different potential scenarios that it has never encountered, and the strategy obtained can be executed in real time. By adapting reinforced learning to deal with large combinatorial action spaces, we can obtain better quality solutions than data-driven decision optimisation within near real-time. The only assumption that typically limits applicability for these methods is the requirement for an underlying simulator for the environment.

Adversarial AI

These methods, at the intersection of machine learning and computational game theory, are crucial for addressing large-scale resource management problems in adversarial (e.g., physical or cybersecurity problems) or competitive (e.g., multiple companies vying for the same customer demand) domains.

The potential of these technologies is immense, and we are already seeing their widespread application in many areas of business. In finance, AI software is helping to predict stock values, identify important events, and calibrate investment strategies. Acting like a crystal ball, AI software is able, for example, to forecast what events are going to cause a major slump in which stocks. In the transportation and logistics sector, AI systems are helping to predict demand and optimise supply to match demand.

Most importantly, by enabling resource management decision-making based on data or data-based insights, these methods help close the loop between data and decisions. In other words, this helps people to make decisions that are not only based on past data but also potential future data that will be generated due to the decisions taken.

In finance, AI software is helping to predict stock values, identify important events, and calibrate investment strategies.

Beyond resource allocation

As the rules of the game change, AI-enabled technologies are indeed a game changer for businesses looking to succeed. Beyond common applications like face recognition and language translation, AI applications can be used in key business processes to conduct the following:

Forecasting

AI systems are proving to be excellent at predicting demand, determining the success of a new product, and identifying the best way to source materials. This helps businesses make informed, astute decisions.

Optimising

With AI systems, fully-automated smart warehouses are now a possibility, where they can ensure smooth, optimal operations. AI systems are also helping in preventive maintenance, so that machineries and equipment are maintained in top-notch condition.

Customer assistance

AI systems are being used increasingly as virtual customer service agents. They allow stores to automate checkout and companies to handle round-the-clock customer enquiries. On the home front, digital voice assistants like Siri and Alexa are helping homeowners to do everything from ordering food, turning on lights, and adjusting the air conditioner thermostat to checking the weather.

Marketing

AI systems are able to simulate and model consumer behaviour. With their ability to gather customer data in real time and analyse trends in customer behaviour, AI systems can help companies roll out effective marketing campaigns. AI systems also allow for the creation of personalised promotions and individualised website displays for highly targeted marketing.

Indeed, AI-enabled technologies have a wide range of applications. In agriculture, AI systems are allowing farmers to identify issues in crop growth, harvest period and so on. And in law, AI systems are enabling lawyers to uncover gaps in legal documents. A case in point is an AI system that lawyers in Singapore have built to check if non-disclosure agreements are foolproof. While a lawyer would typically take a day to go through 20 agreements, this system sorts through the lot in just a few minutes.

In the medical industry, AI systems are enhancing operations at hospitals, diagnosing conditions and recommending treatments, among other revolutionary developments. Tan Tock Seng Hospital in Singapore, for instance, now draws up operating theatre schedules based on predicted demand. AI systems are also being used to process X-rays and MRI scans to determine whether a patient has cancer or is at risk. They are also being used for treatment. Based on how a patient is reacting to certain kinds of medicine, AI systems can recommend the next step in the treatment of that individual. A team from the Weizmann Institute of Science, Israel, developed an algorithm that can reliably pick out (in the early stages of pregnancy, or even before pregnancy) women at high risk of developing gestational diabetes. Also, a research team from the National University of Singapore used big data and predictive analysis of over 30,000 patient-derived biopsies to develop a new personalised tool to detect cancer, predict patient survivability, and forecast how well a cancer patient would respond to immunotherapy.

An intriguing future filled with possibilities

It is exciting to see how AI for business use is unfolding. With the technologies still in their infancy, businesses will continue to explore the way forward. What is clear, though, is that we are at the threshold of a new era of AI-enabled intelligent systems that can contribute towards better business performance in many ways. The key lies in closing the loop—learning from the data by observing what happened, identifying the gap between predictions and observations, and then using that knowledge to improve the model so as to narrow the gap, thereby improving predictions and allocations to determine the optimal matching of demand and supply in a dynamic, uncertain future.

Pradeep Varakantham
is Associate Professor at the School of Information Systems, Singapore Management University

Reference

[1] David Roe, "Is Augmented Artificial Intelligence Already Disrupting Artificial Intelligence?", CMSWire.com, August 27, 2019.