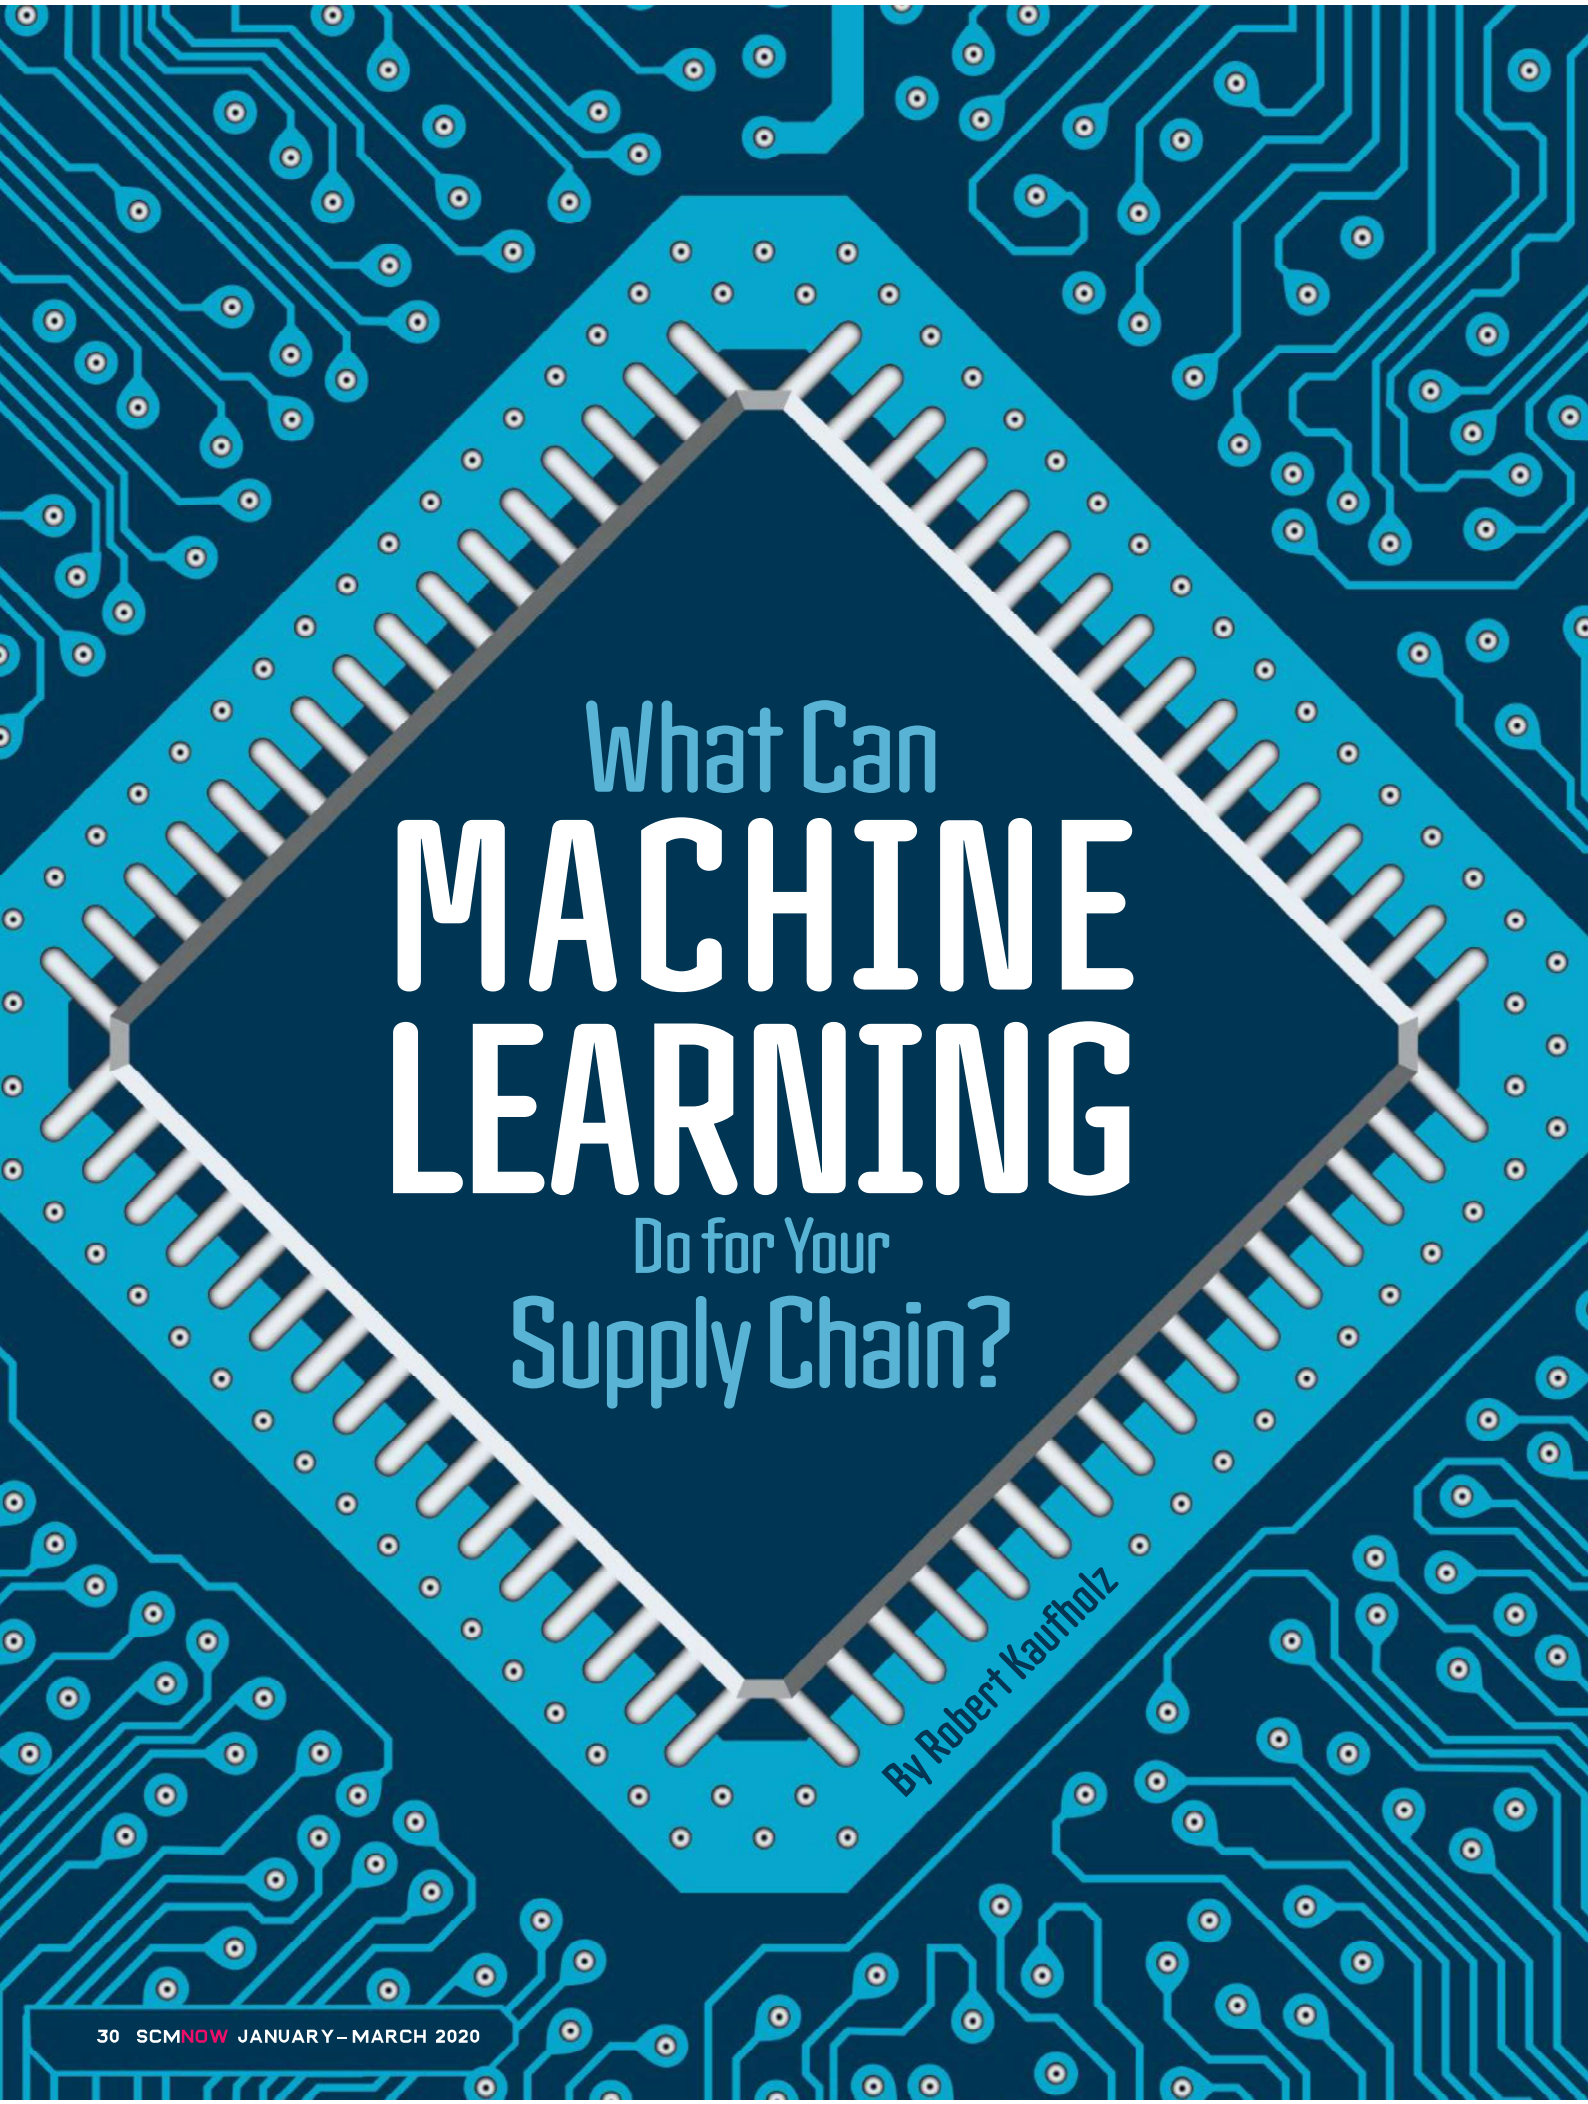

What Can MACHINE LEARNING Do for Your Supply Chain?

By Robert Kaufholz

WITH ALL THE FANFARE SURROUNDING MACHINE LEARNING, people are often surprised to learn that it has been around since 1959. This was when Arthur Samuel, a gaming and artificial intelligence (AI) engineer working for IBM, coined the term to describe the concept behind his revolutionary invention: a self-learning, checkers-playing program. Machine learning is a type of AI that enables computers to become smarter by performing tasks without any explicit programming.

The machine learning golden era is here, with affordable processing power, in-memory storage and an abundance of data driving mainstream adoption. The most important factor behind this machine learning revolution may be the unassailable business need. In the world of supply chain, escalating supply, distribution, sales channels and demand complexity make it increasingly difficult to plan accurately. Meanwhile, industry professionals are grappling with relatively new phenomena that create variability, including fast fashion and other types of hyper-seasonality, shorter product introduction cycles and the impact of social media. And wildcards such as extreme weather, stock market fluctuations and socioeconomic volatility add to the difficulties.

Machine learning augmentation is ideal for alleviating the headaches associated with many key processes.

What Can MACHINE LEARNING Do for Your Supply Chain?

SIX STRATEGIES FOR SUCCESS

Following are six ways supply chain professionals can ensure machine learning drives desired business benefits.

1. Set clear objectives at the outset. When you learn about other companies' success stories (see sidebar), it's tempting to want to just dive in. However, the scientific and iterative nature of machine learning calls for it to be implemented methodically. Create a clear charter of what you want to accomplish and why. This should include establishing baseline metrics so you can track how your machine learning application is performing.

Gather the necessary data to accurately compare previous results to those enhanced by machine learning. This is the best way to gain confidence in your initiative. And because machine learning systems get smarter over time, having a consistent measurement method is essential to tracking return on investment and outcomes. Consider making this a part of your sales and operations planning (S&OP) effort.

2. Embrace probability forecasting. After you've established project objectives, it's time to build a solid foundation for a successful and sustainable initiative. Probability forecasting is most compatible with machine learning because it enables planners to forecast at the most granular level and on different time horizons. It also helps people understand the range of possibilities of demand in the forecast.

This method is ideal for supply chains that include a high number of long-tail items and face demand variability and uncertainty. Using this approach, you still get one number that's associated with the most probable outcome. However, banded around this number is a range of other possible outcomes, each with a different probability.

3. Walk before you run. A phased approach helps ensure a sustainable machine learning solution that meets business objectives today and as needs change. This starts with establishing an adaptive, probability-based model for demand forecasting using existing historical data, then layering in more sophisticated machine learning using external data sources. A reliable demand forecast is critical to success with advanced machine learning and yields significant benefits on its own. Following is a typical progression of the phases a company goes through in modeling and applying data:

- Phase one involves the organization building a self-adaptive model for probability forecasting, taking into account trends, seasonality, calendars and daily sales patterns. "Small data" is used to create a baseline self-learning model for probability forecasting using historical demand history.
- Next, the impact of marketing activities — trade promotions, media events, product introductions and so on — is factored in. This improves the baseline probability forecast by applying machine learning technology to existing historical data.
- In this external demand-sensing step, the business looks for additional correlations by introducing external and big data sources, such as weather, point-of-sale information, social media and machine-generated data.
- In the final stage, specialized business knowledge is introduced from experts across sales, marketing and operations. Again, this information often emerges through the S&OP process. This is also the time to introduce external market intelligence into the model.

It can be tempting to go straight to the external demand-sensing in phase three. However, an iterative approach is the best way to determine which data is meaningful and to what degree. It also gives you the opportunity to change course based on what you learn along the way, drawing out correlations with each step.

4. Remember the four dimensions of data:

- Data volume. Big data plays to the strengths of machine learning. It is essential to have the right amount of data to draw upon, from which the model can derive statistical significance. Don't overlook "small data" related to historic demand.
- Data granularity. Unlike approaches where data is aggregated to weed out noise, machine learning thrives on granularity, examining that noise and using it to find correlations that train the model and make it more powerful.
- Data quality. As the saying goes: Garbage in, garbage out. Unfortunately, while machine learning can determine if a specific data stream has a correlative significance, it cannot tell whether that data is reliable. Therefore, machine learning projects usually include governance programs to clean, filter and maintain information quality through the data life cycle.
- Data variety. The more different types of data sources you factor in, the more robust and accurate the planning outcomes can be.

5. Operationalize your machine learning tool. Often, businesses will build a machine learning solution to tackle a one-off business challenge without considering its long-term worth. For sustainable business value, operationalize your results.

To achieve the stability and adaptability required for operational use, use models that are self-adaptive and do not require continuous tuning by experts. Otherwise, changing business environments will make them unreliable. This is common with traditional demand planning processes that use multiple forecasting algorithms

THREE REAL-WORLD MACHINE LEARNING APPLICATIONS

Heating and cooling systems provider Lennox Industries faced the challenge of managing an ambitious North American distribution network enlargement while transitioning to a hub-and-spoke model with 55 shipping and 161 selling locations. Company leaders wanted to improve service levels and optimize inventories in order to reallocate working capital and balance inventory in the changing network. But the supply chain environment was daunting, with a multi-echelon distribution network about to grow by 250%, 450,000 stock-keeping unit (SKU) locations, many slow movers and new product introductions.

Lennox implemented a transformational supply chain planning solution to dynamically rationalize the inventory mix and design an operational plan that sets inventory stocking targets and balances service levels with inventory cost. Machine learning enabled the business to reliably model highly variable seasonal demand patterns. It sifts through hundreds of thousands of SKU locations to identify clusters of similar seasonality profiles.

This has substantially increased peak period forecast accuracy. In addition, service levels have improved by 16%, inventory turns have increased by 26%, and Lennox is now able to support a significant growth in both sales and market share.

Dairy industry supply chain planning is particularly difficult because products have short shelf lives and demand varies greatly in response to promotions. The Italian dairy producer Granarolo runs thousands of promotions annually, producing 34,000 item-promotion forecasting combinations and causing demand peaks of up to 30 times baseline sales. Its complex supply chain environment is comprised of eight production plants, six logistics technology platforms, 35 transit depots holding inventory, a large fleet of refrigerated vehicles and about 750 merchandisers servicing daily sales. A network of 100 wholesale distributors covers other local markets.

Granarolo uses advanced planning software with machine learning to optimize the demand plan for perishable products. The software also identifies exceptions and optimizes inventory and logistics to meet demand and minimize waste. Average forecast reliability has grown from 80% to 85%, peaking at 95% for fresh milk and cream and 88% for yogurt and dessert products. Inventory levels and delivery times have been halved, resulting in fresher products and less waste. And Granarolo is enjoying significantly better customer service levels and lower transportation costs.

New demands posed by a British luxury car manufacturer's client base prompted its board to raise targets for first time availability (FTA) by 2% without increasing inventory. For the first time, the board also wanted to achieve FTA parity across all three of its car categories: heritage (pre-1997), recent production (mid 90s forward, but no longer in production), and current (today's models).

The company uses advanced machine learning to analyze historical data on consumer behavior to better anticipate customer needs. The tool dips into the company's vast store of historical data and finds eight completely new categories of behavior. The software then uses these new categories to generate a more accurate forecast. Each day, the machine learning engine tunes the safety stock for 80,000 SKUs, automatically reducing inventory before creating a replenishment plan to deliver the demanding new target service levels.

After only two months, the car manufacturer reduced the inventory value of its safety stock on the clustered items by 18% and raised FTA service levels to 97.1%. Outcomes continue trending toward significant further improvements in both service levels and reduced inventory.

assigned to each item or location according to the demand behavior. The forecast generated by these algorithms degrades as the demand patterns evolve over time. Discrete selection and tuning of algorithms require human skills that most businesses cannot afford.

One-off science projects create black boxes that only the developer can understand and support. Users remain skeptical, and, if the developer leaves the company, these models are shelved or discarded altogether. Isolated machine learning projects also require continual manual work to refresh the model when business needs change. The better method uses a self-adaptive model as part of a fully integrated solution, with frequent models updated automatically to react to changes.

6. Get the right people on board. Machine learning frees up planners to do more value-added, strategic work. As your business changes over time, you'll have new questions to answer and will need to adjust your existing models so they remain accurate and useful. It's critical to understand the needed skills and resources before kicking off the project.

HUMAN-MACHINE SYMBIOSIS

Part of a supply chain professional's strategic work is applying domain proficiency. Machine learning can only do so much; business knowledge and process expertise are required to fine-tune models and maximize results. The system will get smarter over time as it factors in human input, and humans will get smarter by learning from the success rate of the probability forecasts. This enables employees to focus on service, work on strategic projects and add their business insights to the system.

Robert Kaufholz is the director of solution design at ToolsGroup. He may be contacted at rkaufholz@toolsgroup.com.