NOVEMBER 2016

CONSUMER PACKAGED GOODS/RETAIL

The secret to smarter fresh-food replenishment? Machine learning

With machine-learning technology, retailers can address the common—and costly—problem of having too much or too little fresh food in stock.

Christoph Glatzel, Matt Hopkins, Tim Lange, and Uwe Weiss

Fresh food, already a fiercely competitive arena in grocery retail, is becoming an even more crowded battleground. Discounters, convenience-store chains, and online players are recognizing the power of fresh-food categories to drive store visits, basket size, and customer loyalty. With fresh products accounting for up to 40 percent of grocers' revenue and one-third of cost of goods sold, getting fresh-food retailing right is more important than ever.[1]

It's also more complex than ever. Fresh food is perishable, demand is highly variable, and lead times are often uncertain. Furthermore, many retailers now carry broader fresh-food assortments that include exotic and hard-to-find items, as well as "ultrafresh" items with a shelf life of no more than one or two days. Retailers are constantly having to make difficult trade-offs when placing orders with fresh-food suppliers: order too much, and the food goes to waste; order too little, and you lose sales and erode customer loyalty. But with demand fluctuating daily, how can retailers know the right amount to order?

Most traditional supply-chain planning systems take a fixed, rule-based approach to forecasting and replenishment. Such an approach works well enough for stable and predictable product categories, but fresh food is much more complicated. Because local demand and conditions vary from day to day, planners have to manually enter different types of data—price changes or promotions, for instance—into their replenishment systems. These daily manual processes are time consuming, error prone, and heavily reliant on individual planners' experience and gut instincts.

There's a better way. A number of leading retailers have found a solution that revolutionizes their supply-chain planning: machine learning.[2] Based on algorithms that allow computers to "learn" from data even without rules-based programming, machine learning allows retailers to automate formerly manual processes and dramatically improve the accuracy of forecasts and orders. Retailers that use machine-learning technology for replenishment have seen its impact in many ways—for instance, reductions of up to 80 percent in out-of-stock rates, declines of more than 10 percent in write-offs and days of inventory on hand, and gross-margin increases of up to 9 percent.

The advantages of machine learning in replenishment

Unlike standard supply-chain software systems, machine-learning solutions can collect, analyze, and adjust large data sets from a wide range of sources, without high investments in personnel. A machine-learning algorithm can make demand forecasts based not just on historical sales data but also on other influencing parameters: internal factors such as advertising campaigns and store-opening times, and external factors such as local weather and public holidays. Advanced algorithms currently used by leading retailers already analyze more than 50 parameters. And the calculations are done at a much more granular level than standard systems are able to do: retailers can determine the effect of each

parameter on each SKU in each store (and in each distribution center, where relevant) on a daily basis.

A machine-learning system can also take into account supply-chain constraints such as supplier delivery times and minimum or maximum order quantities. Based on all these considerations, it then generates order proposals for the entire product range every 24 hours. Each order proposal optimizes for product availability while minimizing the risk of waste and markdowns, as we discuss in further detail below. Central planners can spot-check these proposals, but we've found that with best-in-class machine-learning solutions, human intervention is rarely warranted. The manual data-entry and administrative work that staff needs to do in individual stores and at headquarters is therefore significantly reduced.

Machine-learning solutions are available as cloud-based, software-as-a-service (SaaS) applications. Unlike other replenishment-optimization solutions, which typically require upgrades to a retailer's IT systems, SaaS applications work as an intelligent overlay to existing enterprise-resource-planning systems, making them more flexible and faster to implement. They also don't require large investments in hiring new personnel—rather, a retailer can build its own staff's capabilities, eliminate a significant amount of planners' data-entry tasks, and reallocate valuable staff time toward more value-added activities.

By using predictive applications powered by machine learning, an international supermarket chain with more than 1,000 stores automated most of the central planning and decision making for daily orders in one of its largest fresh-food departments. And because the retailer operates several food-processing plants, it was also able to integrate warehouse and manufacturing processes—for instance, through just-in-time production—to reduce stock in the entire supply chain, increase in-store product availability, and get fresher products on store shelves.

Best-in-class demand predictions

For decades, retailers have extrapolated demand by looking at historical sales data—an obviously imperfect methodology that skews demand forecasts downward, since it doesn't measure unmet demand. Advanced machine-learning algorithms overcome this problem. The algorithms build demand-probability curves using sales and inventory data, making cost-benefit calculations that evaluate the risk of waste against the risk of out-of-stocks.
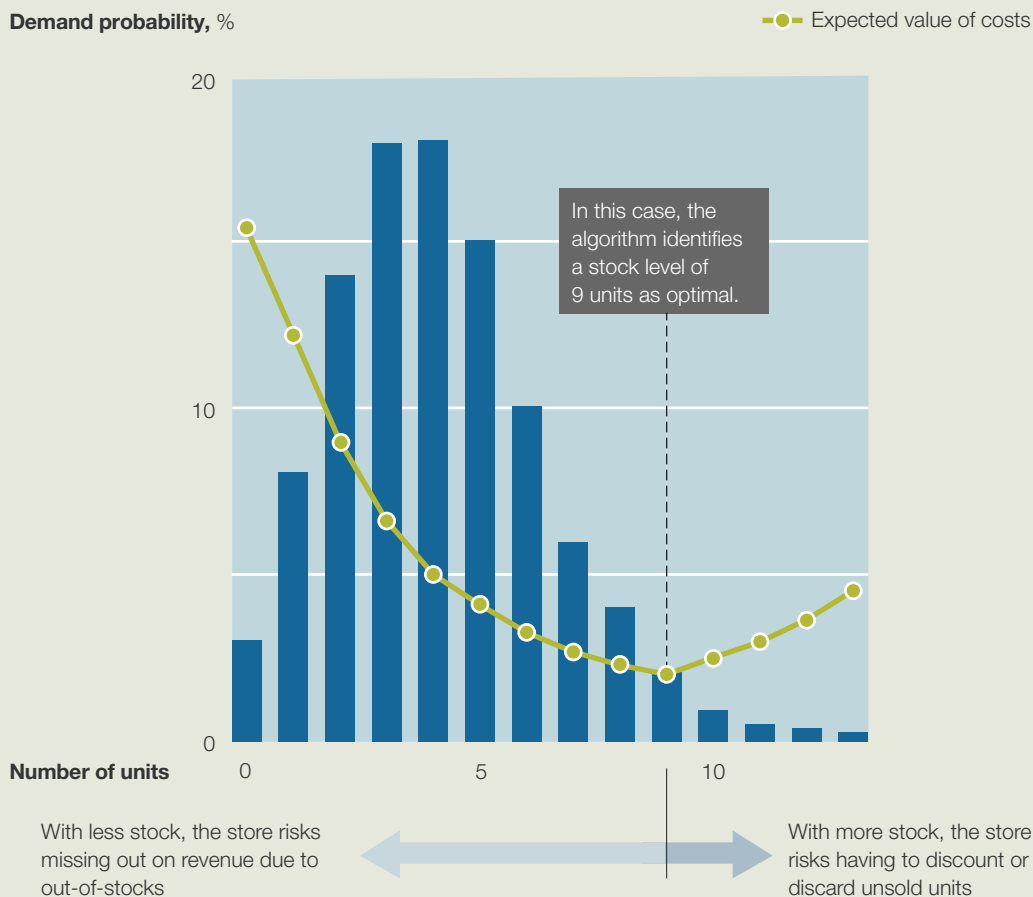
To illustrate: on the exhibit, the histogram shows the demand probability for a specific SKU-store-date combination. Let's use pineapples in Store #123 on June 10 as an example. The vertical bars show that stocking four pineapples in that store on that day will probably be enough to meet demand; the store will likely sell most or all of them, so the risk of having rotten pineapples in the store is small. But what if a customer wants to buy a fifth or sixth pineapple that day? The store would lose out on revenue because pineapples would be out of stock. The green curve on the exhibit represents the expected value of costs for each stock level, taking into account potential loss of revenue due to out-of-stocks, as well as potential markdowns and waste. In this case, the algorithm identifies a stock level of 9 units as optimal.

The system can align individual ordering decisions with the retailer's strategic goals and key performance indicators (KPIs). For instance, if the retailer is more concerned about margins than revenues, the algorithm will adjust decisions accordingly. It can also work toward improving several KPIs at the same time.

Exhibit

Machine-learning algorithms help retailers determine optimal stock levels, taking into account both waste and lost sales.

Demand probability, %

-●- Expected value of costs

In this case, the algorithm identifies a stock level of 9 units as optimal.

Number of units

With less stock, the store risks missing out on revenue due to out-of-stocks

With more stock, the store risks having to discount or discard unsold units

Newly developed algorithms can simultaneously optimize pricing and replenishment, leading to even greater profit increases in fresh categories. The algorithms simulate how changes in price will affect demand. For instance, if a small price reduction would lead to a significant increase in sales volume, the system would recommend larger order quantities. On the other hand, if demand for a product wouldn't change much even if it were sold at a deep discount, the system would recommend smaller order quantities so as to minimize losses due to markdowns. Because price elasticities change constantly—sometimes on a daily basis— machine learning's ability to continually integrate data on pricing and replenishment can make a tremendous difference in a retailer's profitability.

■   ■   ■

To capture the full value that machine learning offers, retailers shouldn't just stop at software implementation. Instead, they should take the opportunity to concurrently refine their business strategy for their fresh-food departments, and translate that strategy into detailed assortment rules—for example, defining which subcategories and SKUs should always be in stock at which hours of the day. Retailers should also review their end-to-end supply-chain planning processes, including store-level ordering and inventory management. Delivery frequencies may need to change; certain fresh-food products may require twice-daily deliveries, for instance, to ensure that the store won't run out of the product before the end of

the day. In addition, retailers should make adjustments to capacity planning and labor scheduling in distribution centers.

People processes, too, should change. Demand planners, both at the store level and at headquarters, will be spending their time differently and will need new performance-management metrics and incentives. And finally, retailers ought to reevaluate how they work with suppliers, renegotiating some supplier contracts and perhaps even helping suppliers adjust their own forecasting and ordering processes. By taking these actions, retailers will be better able to build on—and sustain—the vast improvements that machine-learning technology can bring. ■

[1] Raphael Buck and Arnaud Minvielle, "A fresh take on food retailing," *Perspectives on retail and consumer goods,* Winter 2013/14, McKinsey.com.
[2] Dorian Pyle and Cristina San Jose, "An executive's guide to machine learning," *McKinsey Quarterly,* June 2015, McKinsey.com.

Christoph Glatzel is a senior partner in McKinsey's Cologne office, where **Tim Lange** is a consultant. **Matt Hopkins** is a London-based principal at Blue Yonder, a retail-analytics firm with which McKinsey has a partnership. **Uwe Weiss** is the CEO of Blue Yonder.