



The Complete Guide to
**Machine Learning
in Retail Demand
Forecasting**



Johanna Småros
Co-founder
Ph.D. (Tech)

Henri Kaleva
Head of RELEX Labs
M.Sc. (Tech)

Contents

Retail is detail at large scale	3
1. What is machine learning, and why should retailers adopt it now?	4
2. Machine learning tackles retail's demand forecasting challenges	6
2.1 Weekdays, seasonality, and other recurring demand patterns	8
2.2 Price changes, promotions, and other business decisions impacting demand	9
2.3 Weather, local events, and other external factors impacting sales	12
2.4 Unknown factors impacting demand	14
3. Make machine learning work for your retail demand planning	15
3.1 Working with retail's long-tail products	16
3.2 Leverage human expertise	17
3.3 Demand forecasting is only one part of retail planning and optimization	19





Retail is detail at large scale

The old adage is common but true: “Retail is detail at large scale.”

To ensure smooth operations and high margins, large retailers must stay on top of tens of millions of goods flows every day. At the center of this storm of planning activity stands the demand forecast.

A highly accurate demand forecast is the only way retailers can predict which goods are needed for each store location and channel on any given day—which in turn is the only way to ensure high availability for customers while maintaining minimal stock risk. A reliable forecast leveraged across retail operations can also support capacity management, ensure the right amount of staff in stores and distribution centers, or help buyers manage the complexities of long lead-time purchasing.

Generating an accurate forecast is actually quite simple under stable conditions, but we all know too well that retail is inherently dynamic, with hundreds of factors impacting demand on a continuous basis. Every day, retail demand planners struggle to consider an immense number of variables, including:

- Recurring variations in baseline demand, such as weekday-related and seasonal variations.
- Internal business decisions designed to capture consumer attention and provide a competitive edge, such as promotions, price adjustments, or changes to in-store displays.
- External factors, such as local events, changes in a store’s neighborhood or competitive situation, or even the weather.

With this much data, no human planner could take the full range of potential factors into consideration. However, machine learning makes it possible to consider their impact at a detailed level, by individual store or fulfillment channel. It’s not surprising, then, that so many retailers today are transitioning their technology strategies toward machine learning-based demand forecasting.

1. What is machine learning, and why should retailers adopt it now?



Machine learning gives a system the ability to learn automatically and improve its recommendations using data alone, with no additional programming needed. Because retailers generate enormous amounts of data, machine learning technology quickly proves its value. When a machine learning system is fed data—the more, the better—it searches for patterns. Going forward, it can use the patterns it identifies within the data to make better decisions.

Machine learning makes it possible to incorporate the wide range of factors and relationships that impact demand on a daily basis into your retail forecasts. This is enormously valuable, as just weather data alone can consist of hundreds of different factors that can potentially impact demand. Machine learning algorithms automatically generate continuously improving models using only the data you provide them, whether from your business or from external data streams. The primary benefit is that such a system can process retail-scale data sets from a variety of sources, all without human labor.




Figure 1: Machine learning enables retailers to capture the impact of recurring patterns, their own internal business decisions, and external factors for more accurate, granular, and automated short- and long-term demand forecasts.

Of course, machine learning algorithms are not new—they’ve been around for decades. But never before have they been able to access as much data or data-processing power as is available today. Though retailers may have struggled to update their forecasts quickly in the past, [large-scale data processing and in-memory technology](#) now enable millions of forecast calculations within the space of a single minute.

2. Machine learning tackles retail's demand forecasting challenges



Machine learning is an extremely powerful tool in the data-rich retail environment. It should be leveraged in any context where data can be used to anticipate or explain changes in demand. In some instances, it can even fill in the gaps where the data is lacking.

The types of demand impact machine learning can capture	
Recurring demand patterns	Recurring variation in demand caused by, for example, weekdays, holidays, and seasons.
Internal business decisions	The impact of promotions (including cannibalization and halo effects), price changes, and changes in how products are displayed.
External factors	The impact of factors not controlled by the retailer, such as weather, local events, and local consumer footfall.
Unknown factors	Changes in demand for which the impacting factor has not been recorded, such as a competing store opening next door or roadwork disrupting customer footfall.

Table 1: Machine learning addresses all of retail’s typical demand forecasting requirements.

2.1 Weekdays, seasonality, and other recurring demand patterns

Time-series modeling is a tried and true approach that can deliver good forecasts for recurring patterns, such as weekday-related or seasonal changes in demand. In our experience, though, machine learning-based demand forecasting consistently delivers a level of accuracy at least on par with and usually even higher than time-series modeling. Whereas time-series models simply apply past patterns to future demand, machine learning goes a step further by trying to define the actual relationship between variables (such as weekdays) and their associated demand patterns.




Figure 2: Retailers need accurate day-level forecasts for effective replenishment of fresh products as well as for managing capacity in all parts of their supply chains.

Machine learning also streamlines and simplifies retail demand forecasting. When using time-series models, retailers must manipulate the resulting baseline sales forecast to accommodate the impact of, for example, upcoming promotions or price changes. Machine learning, on the other hand, automatically takes all these factors into consideration.

In addition to taking an abundance of factors into account, machine learning also makes it possible to capture the impact when multiple factors interact—for example, weather and day of the week. Warm, sunny weather can drive a much bigger demand increase for barbecue products when it coincides with a weekend.

2.2 Price changes, promotions, and other business decisions impacting demand

Your own business decisions as a retailer are also an important source of demand variation, from promotions and price changes to adjustments in how products are displayed throughout your stores. Yet, despite the fact that retailers typically plan and control these changes themselves, many in the industry are unable to accurately predict their impact.

Machine learning allows retailers to accurately model a product's **price elasticity**, i.e., how strongly a price change will affect that product's demand. This capability is highly valuable as part of promotion forecasting, as well as when **optimizing markdown prices** to clear out stock before an assortment change or the end of a season. Furthermore, retailers must regularly adjust consumer prices to reflect supplier prices and other changes in their cost base.

Price elasticity alone, however, does not capture the full impact of price changes. A product's pricing in relation to alternate products within the same category often has a large impact as well. In many categories, the product with the lowest price captures a disproportionately large share of demand. Machine learning-based demand forecasting makes it quite straightforward to consider a product's price position, as shown in Figure 3 below.

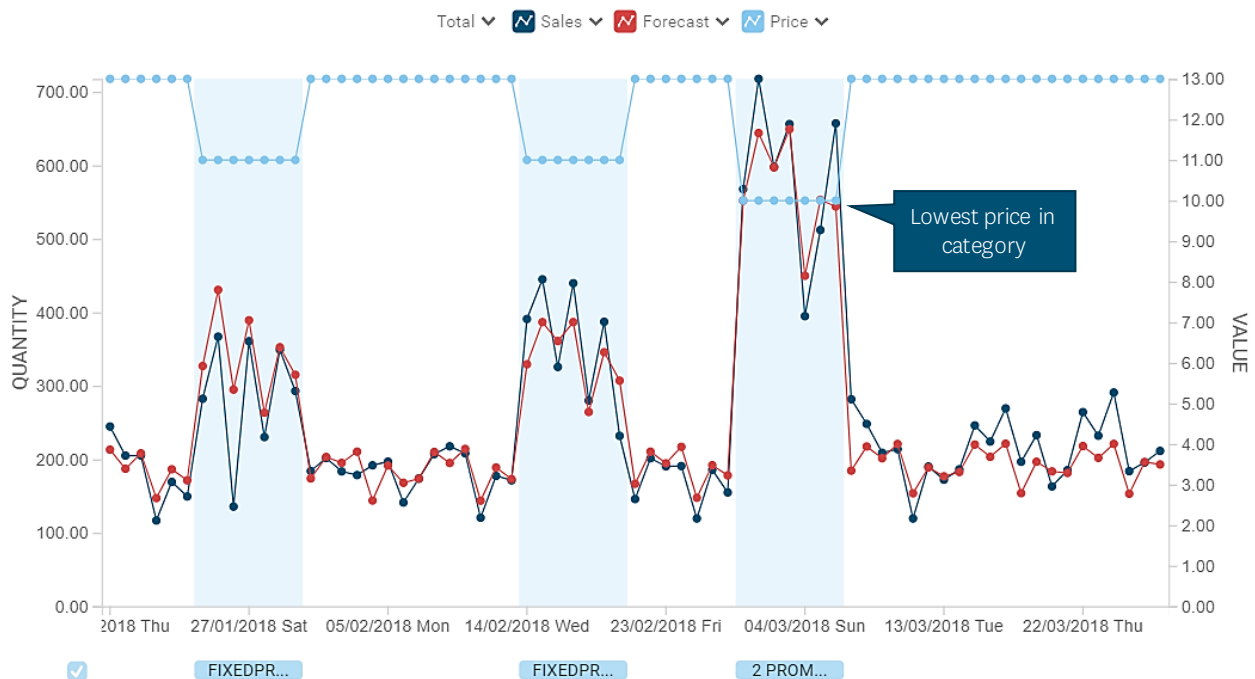


Figure 3: Demand for this product increases when its price drops, but the increase is bigger when the product's price drops to be the lowest in its category.

In a 2020 [study of North American grocers](#), 70% of respondents indicated that they could not take all the relevant aspects of a promotion—such as price, promotion type, or in-store display—into consideration when forecasting promotional uplifts. But they wish they could.

Here, too, machine learning can help. Using machine learning-based demand prediction, retailers are able to accurately predict the impact of **promotions** by taking into consideration factors including, but by no means limited to:

- Promotion type, such as price reduction or multi-buy.
- Marketing activities, such as circular ads or in-store signage.
- Products' price reductions.
- In-store display, such as presenting the promoted product in an endcap or on a table.




Figure 4: For this product, an endcap display with no price change results in a notable sales uplift, but the uplift is modest compared to the effect of a 50% price reduction.

Sales cannibalization, the phenomenon in which one product's promotional uplift causes a reduction in sales for other products within that category, is quite common and must also be accounted for in forecasts, especially for fresh products. For example, if a supermarket carries two brands of lean organic ground beef—HappyCow and GreenBeef—they should expect that a promotion on the HappyCow product will cause more people to buy it. As a result, though, some of the demand for the GreenBeef product will shift to HappyCow. If the demand forecast for the GreenBeef product is not accurately lowered, the retailer is at high risk of overstocking, which will ultimately drive waste.




Figure 5: Machine learning algorithms can accurately model cannibalization effects in relation to promotions or price changes based on historical sales data.

Manually adjusting the forecasts for all potentially cannibalized items is just not feasible in most retail contexts because the number of products to adjust is simply too high. The patterns are also typically quite specific to individual stores’ assortments and shopping patterns. This is where machine learning algorithms’ ability to automatically identify patterns and adjust forecasts accordingly adds enormous value.

On the other hand, a promotion for the HappyCow product will likely increase sales for some related products outside of the “ground beef” class in what’s known as the **halo effect**. Hamburger buns, for example, have an obvious and predictable correlation with ground beef. Unfortunately, the impact can be so diffused across the assortment that identifying every impacted product becomes more or less impossible, even with machine learning: think onions, potato chips, beer, watermelon, taco meal kits, salad fixings, oyster crackers, corn on the cob, Worcestershire sauce, soy sauce, or any number of other items shoppers might associate with ground beef-based dishes. But even if forecasting systems can’t identify all possible halo relationships, they should still make it easy for planners to adjust forecasts for the relationships they know to exist.



2.3 Weather, local events, and other external factors impacting sales

External factors such as the weather, local concerts and games, and competitor price changes can have a significant impact on demand but are difficult to consider in forecasts without a system that automates a large portion of the work. At a high level, the impact can be quite intuitive. On a warm day, you'll likely see increased ice cream sales, whereas the rainy season will see demand increase for umbrellas, and so on.

When looking at a retailer's entire assortment, though, the challenge gets more complicated. How can you effectively identify all products that react to the weather? Can you account for the full range of variables that comprise a "weather forecast"—temperature, sunshine, rainfall, and more? Will the weather-related impact of sunshine be stronger in summer than in winter? Or stronger on weekends than on workdays?

The use of weather data in demand forecasts is a prime example of the power of machine learning. Machine learning algorithms can automatically detect relationships between local weather variables and local sales. They can map these relationships on a more granular, localized level than any human endeavor could accomplish — and are also able to identify and act on less obvious relationships that human intuition or "common sense" might overlook.

When demand planners or store staff are asked to manually check weather forecasts to influence ordering decisions, they focus on securing supply for anticipated demand increases—pushing ice cream to stores during a heat wave, for example. Rarely, though, does anyone have time to adjust ice cream forecasts slightly downwards during rainy weeks or cold snaps in the summer. A planning team using machine learning doesn't have to worry about adjustments like that, as their system can suggest them automatically.

In our experience, automatically considering weather effects in demand forecasts reduces forecast errors by between 5% and 15% on the product level for weather-sensitive products and by up to 40% on the product group and store levels.




Figure 6: Machine learning makes it possible for retailers to incorporate the impact of weather and other external factors on sales in their retail demand forecasting.

But weather data is by no means the only external data that could or should be incorporated in your retail demand forecasting. Any number of external data sources, such as past and future local events (e.g., football games or concerts), data on competitor prices, and human mobility data can be used to improve outcomes in the same way.

As an example, RELEX used machine learning to help [WHSmith](#) improve their understanding of how flight schedules impacted demand patterns at their airport locations. By feeding external data from airlines into the system, WHSmith improved their forecasts and were able to significantly reduce their fresh spoilage rates while improving availability.

2.4 Unknown factors impacting demand

Thus far, we've explored contexts in which the factors impacting demand—weekly and seasonal patterns, business decisions, and external factors—are readily identifiable. But machine learning can help adjust forecasts even in situations where the influencing factors, whether internal or external, are unknown.

In brick-and-mortar retail, local circumstances—such as a direct competitor opening or closing a nearby store—may cause a change in demand. Unfortunately, data on the factor causing this change may not be recorded in any system. Sometimes, retailers' internal decisions also go unrecorded, such as adding a product to a special off-shelf display area in a store.

Fortunately, machine learning can help in these situations. Machine learning algorithms can tentatively place a “change point” in the forecasting model, then track subsequent data to either disprove or validate the hypothesis. This allows forecasts to adapt quickly and automatically to new demand levels.

Consider the example in Figure 7 below, in which a table display has been created in addition to the regular shelf space for a product. Though this change was not recorded in the master data, the system was easily able to track the demand impact as a factor of how the product was displayed in the store.

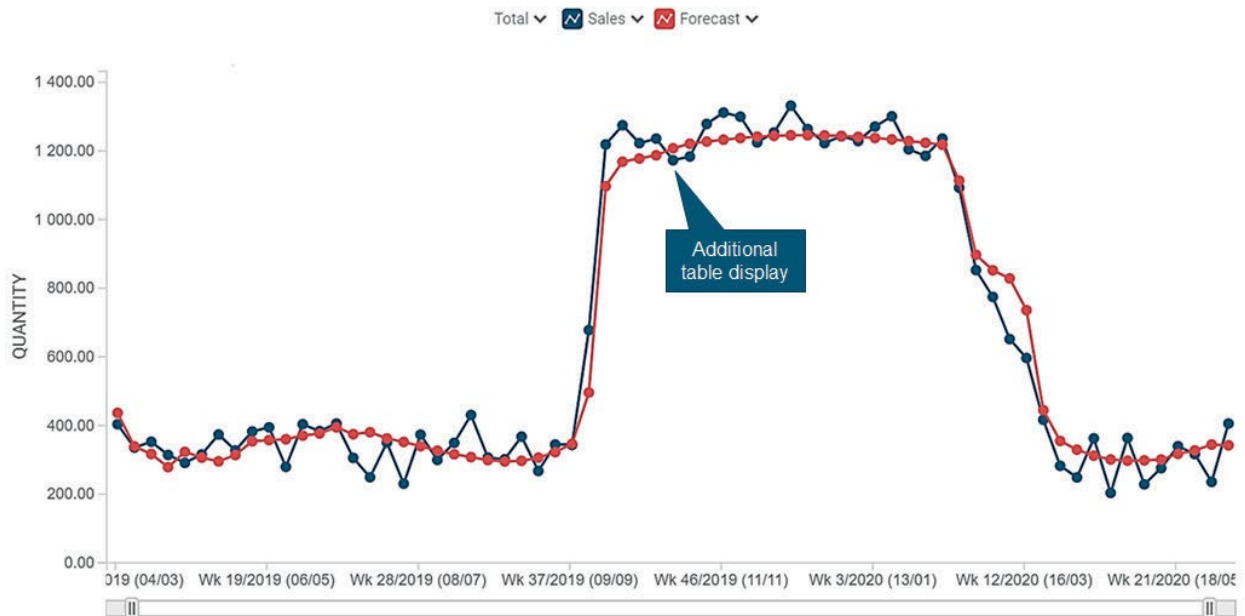


Figure 7: The step changes in demand for this product—first up and then back down—are a result of the product getting an additional off-shelf table display in the store. This additional display was not recorded in the store's master data, but the machine learning algorithm was still able to consider the step-change with very little delay.

3. Make machine learning work for your retail demand planning



Although machine learning is becoming increasingly mainstream, retailers should still keep some considerations in mind when determining how to utilize it in their business. Some considerations are specific to the retail context, whereas others—level of transparency, for example—are generic enough to apply to any situation that calls for computer-human teamwork.

3.1 Working with retail’s long-tail products

One retail-specific challenge is that despite the large amount of data available to retailers, the amount of data available per product, store/channel, and demand-influencing factor is sometimes quite small. Even if your annual sales are in the billions, that total volume is distributed among tens of millions of inventory flows and across hundreds of days.

The sales of so-called “long-tail products”—those that sell only a few units per day or week—often contain a lot of random variation, and it can be difficult to reliably identify relationship patterns within that noise. With few data points available—tens or hundreds, rather than thousands—differentiating the impact of demand-influencing factors like weather, price changes, display changes, or competitor activities from random variation becomes quite challenging.

When low-sales volume items introduce a significant amount of random variation, there is a risk of “overfitting,” in which the algorithm becomes too complex or contains too many variables. In overfitting situations, the algorithm can end up “memorizing the noise” instead of finding the true underlying demand signal. This overfit model would ultimately end up making predictions based on the noise. It may perform exceptionally well using its training data but extremely poorly when asked to incorporate new, unseen data. Typically, overfitting results in occasional “off the charts” forecasts or “nervous” forecasts, where the forecast reacts too heavily to minor changes in the data.

Furthermore, it might be impossible to detect a seasonal pattern at the product-store level for slow movers, but analysis of total chain-level sales for that product may easily identify a clear pattern. This pattern must be considered in sourcing and distribution center replenishment.

Due to low volumes and sparse data at the product-store/channel level in retail, it is very important that:

1. The machine learning algorithms used are robust enough not to deliver outlier results based on scant data points.
2. The machine learning algorithms avoid overfitting by minimizing or pruning out factors that have little-to-no demand impact.
3. You can apply the machine learning algorithms not only on a product-store/channel level but also at different levels of aggregation (e.g., product-region or product-chain) and with flexible groupings.




Figure 8: To avoid overfitting when applying machine learning to sparse retail sales data, the system must be able to 1) trend toward simplicity by pruning out factors that have little demand impact and 2) allow for expert input to more effectively find relevant relationships.

3.2 Leverage human expertise

The COVID-19 crisis has demonstrated that automated forecasting and replenishment is extremely useful when retailers face large-scale disturbances, as automation frees up a lot of planner time. However, planners are still needed to guide the system when dealing with highly impactful, novel events. In such situations, decisions should be about more than just trying to make good predictions—retailers must also judge the business risk of upside and downside scenarios. To create effective human-computer interaction, whether in exceptional scenarios like COVID-19 or during more normal demand periods, retailers need actionable analytics.

Because forecasts are never perfect, there will always be situations in which planners need to dissect a forecast. When planners can easily access which factors have been used to produce the forecast and how, they are more likely to trust the system to manage “business-as-usual” situations so they can focus on the exceptional ones that actually need their attention. However, “black box” systems with low transparency make it impossible to understand why automated recommendations are being made. They quickly erode user trust, often driving low system adoption rates.

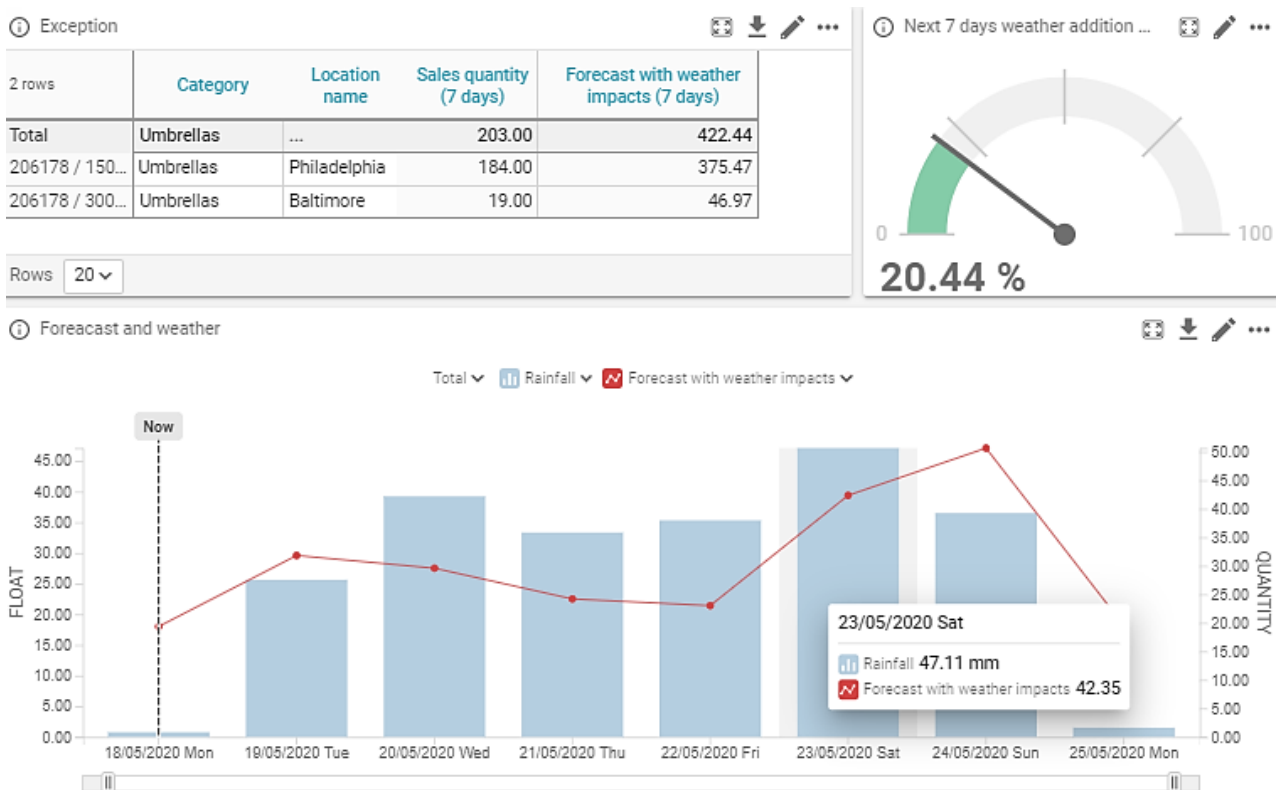


Figure 9: AI cannot be a “black box.” Planners need to understand how the forecast has been produced and be able to tune the calculations using their expertise.

A transparent solution also gives planners valuable insights for further improvements—be it better data, the need for additional product classification, or testing new combinations of factors (such as adding a “lowest price” variable in our earlier example).

3.3 Demand forecasting is only one part of retail planning and optimization

Finally, we must keep in mind that although retail demand forecasting is essential, even great forecasts amount to nothing if they're not used intelligently to guide business decisions. When managing slow movers, for example, forecast accuracy is much less important to profitability than replenishment and space optimization, which will drive balanced, low-touch goods flows throughout the supply chain. Retail forerunners are applying [AI across all their core planning processes](#)—demand, operations, and merchandising—for improved profitability and sustainability.

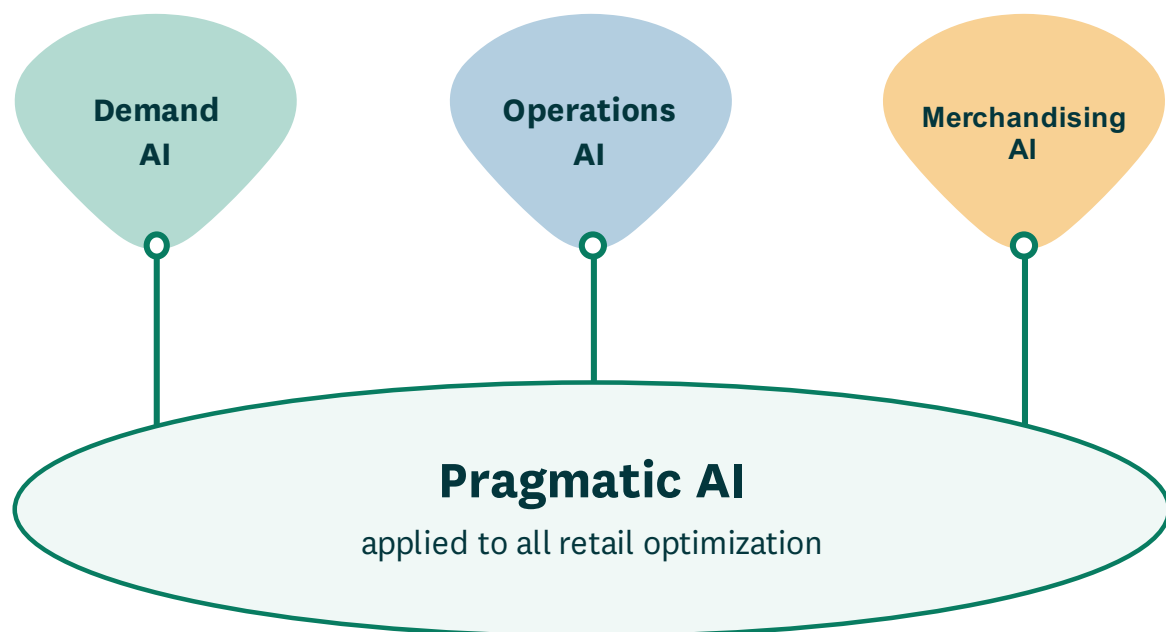


Figure 10: Demand forecasting is just one application area of AI in retail. Forerunners in retail also apply AI to merchandising and operations for improved profitability and sustainability.

While implementing machine learning-based demand forecasting provides a solid foundation for getting started with applied AI, your business's journey should not stop there. AI has already proven its value in addressing a wide array of retail's typical planning challenges: from [workforce optimization](#) to [more effective goods handling in stores](#) and [more automated and impactful markdown optimization](#). There is an abundant reservoir of surprisingly easy, quick wins to be earned by applying pragmatic AI throughout retail's core processes.

About RELEX Solutions

RELEX Solutions is a leading provider of cutting-edge retail optimization software that's built for the age of Living Retail, where change is the only constant. We help retailers adapt to every future, faster.

Our cloud-native Living Retail Platform delivers pragmatic AI across all retail functions and at retail scale, removing siloes, rigidities and inefficiencies along the way. We offer a fast lane to value that builds from a foundation of radically improved demand forecasting and supply chain optimization. Our customers leverage this enhanced supply chain visibility into exponential benefits — optimizing their space, allocation, workforce, pricing and promotion strategies, all within our unified platform.

Today, RELEX is a hyper-growth company with 250+ customers who love us — ask any of them for a frank and independent assessment of our team and solutions. RELEX is trusted by leading brands including AutoZone, Franprix, Morrisons, PetSmart and Rossmann, and has offices across North America, Europe and the Asia Pacific region.

www.relexsolutions.com