Afresh

Putting AI in grocery aisles

Look at those avocados!

Creating a display of avocados with the perfect mix of ripe, unripe, and soon-to-be ripe produce is one of the most valuable puzzles to solve in food. By crafting beautiful, bountiful displays of a customer favorite, your stores can double their sales, gain deep customer loyalty, and represent the pinnacle of quality in fresh. And yet, most grocery shoppers in the United States have to choose from full displays of rock-hard avocados or piles of mushy, near-guacamole. Why? Because grocery stores aren't using the right technology.

The avocado conundrum presents a host of variables and challenges that simply don't exist in the cereal aisle, especially when it comes to freshness and quality. Avocados are scrutinized and chosen based on a variety of reasons. By contrast, every box of cereal on the shelf is identical and your customers will enjoy the same crunch whether it's 8 months from expiration or 6 months.

Keeping that cereal shelf full is already a frustrating task. Cereal management runs on safety stocks and perpetual inventory, and they're both near impossible to get right. Startups have cropped up around the country promising accurate inventory counts with flying drones, friendly robots, and constant webcam supervision. Amazon needs a staff of 8 and a remote labeling workforce to keep a 7/11-sized autonomous store's shelves full.

And the math behind cereal ordering is trivial. Legacy providers determine a safety stock level based on historical demand. When the inventory count dips below its safety stock, the mainframe reorders a set amount. The system does not recompute its safety stock daily. It doesn't need to. Cereal ordering is easy. But it still doesn't work. And it definitely doesn't work for your produce department.

Afresh

Good for cereal isn't good enough for blueberries

For too long, grocery retailers have attempted to digitize fresh operations. But the technology available has always been center-store-centric, and once built and shipped out, those systems are never updated again. Models for these solutions on the IBM mainframe don't get retrained, have no feedback loop for customer complaints, and don't have online monitoring that addresses issues when they arise. In short: legacy systems are too static for the dynamics of today's grocery store.

Additionally, these systems are built for non-perishable goods and using them in the perimeter requires perfect maintenance of upstream data and store-level inputs—an impossible task given fresh's complexities. Stores end up deploying massive amounts of resources in an attempt to adapt legacy solutions to the unique nature of perishable foods. Store teams are left wasting hours tracking scan outs and weigh outs, and completing tedious data entry in an attempt to maintain accurate perpetual inventory.

Despite the best efforts of IT and store teams alike, these systems have proven they don't work for the fresh paradigm. Inevitably, reliance on high-volume, low-skilled administrative processes and perpetual data inaccuracies cause loss of trust, low adherence, and, ultimately, abandonment of the system. Store teams are left relying on manual, inefficient processes that drive suboptimal fresh retail operations and lead to lost profits.

Static legacy systems, even those that do leverage artificial intelligence (AI), don't work with the inherent dynamics fresh food. Fresh requires an adaptable system that can expect the unexpected and can accept the inherent imperfections of perishable foods.

 |

Investigating the key failure points of legacy solutions

Perpetual inventory: The case of the non-organic mango

Perpetual inventory is typically maintained by taking shipments into the store, subtracting sales, and then subtracting scan outs. If shipments, sales, or scan outs are inaccurate, then the inventory estimate will be inaccurate as well.

Another common case of perpetual inventory breakdown occurs when shoppers at self checkout or untrained cashiers scan an organic mango as a non-organic mango. This seemingly minor mistake skews the system's understanding of both organic and conventional mango quantities—leaving it to mark organic stock as too high and conventional too low.

Demand forecasting: The case of hidden demand

Demand forecasting usually relies on models that incorporate past sales data, current pricing and promotions, timing of the order, and potential upcoming holidays. But in order to work properly, the incoming data must be correct. However, sales data doesn't always reflect demand and recommendations can quickly become flawed because of that.

Here's the core problem: When a store only has 50 apples in stock the most apples they can sell is 50, regardless of customer demand. So when items are out of stock or low, the forecast may not reflect true demand for the product.

This is called "demand censoring" because what's in stock conflicts with (or censors) the predicted demand, and systems that use this can seriously undermine future recommendations.

Order quantity: The case of the inaccurate schematic

Unlike non-perishable departments, which change schematics only once or twice per year, fresh departments change their schematics up to once per week. A fresh item can go from having 2 cases worth of product displayed in its “home” location, to adding 6 more cases of display volume on an endcap or center bin the next week, and then back to its 2-case home display size the following week. Think about Thanksgiving. Sweet potato bins grow significantly before the day, and quickly shrink when there are no more pies to make.

 |

Why fresh food requires purpose-built AI

Fresh food, by its very nature, is complex. Unlike a box of cereal, a basket of blueberries is impacted by a number of factors that affect supply and demand, and when it comes to fresh food every item varies when it comes to quality, quantity, and shelf life. AI is uniquely suited to address these challenges as long as it's built for fresh.

Challenge	Technique	Description
Handling Massive Data Sets	Deep Learning	Algorithms that can ingest and scale the 10-100 million data points generated by grocery chains, and extract more features from raw data.
Adapting to Rare Events	Multi-Task and Few-Shot Learning	Predictions that are accurately adjusted for the demands of rare events (like holidays which only happen once a year and can be considered "rare), ensuring high in-stock rates with minimal shrink
Measuring Uncertainty	Probabilistic and Bayesian Methods	Predicts an entire distribution over model demand rather than a point forecast. This allows for more accurate planning and confidence assessments.
Minimizing waste and stockouts	Planning Algorithms	Outputs inventory decisions that balance minimizing waste and limiting out of stocks.
Establishing trust with store teams	Human-in-the-loop Feedback	Ability to gather real-time information from end users to reconcile data gaps where we are not confident in recommendation accuracy, building trust and adherence.

Built for fresh: How it works and how it's different

To be truly designed for fresh, a system built for perishables should leverage AI as a decision-making engine and apply techniques jointly across inventory estimation, demand forecasts, and order recommendations.

This system must also prioritize the retailer's top goal: Keeping shelves stocked without creating excess waste. It's easy to keep stockouts low by ordering far more than you need, but doing so leads to excessive and unnecessary waste. And the reverse simply leads to empty shelves. The right AI-powered system will create balance through learning and optimization.

Let's look at the key components that make a fresh-first solution like Afresh perfect for ordering, inventory, merchandising, forecasting, and store operations.

Data validation: Ensuring no garbage in

As discussed, one of the main challenges presented by grocery retail is massive data sets and time series data for every item in every store, every day of the year. This is even more challenging taking into consideration fresh data is often imperfect and messy. If any inputs have errors, the system will generate bad orders.

Take, for example, a watermelon. While most grocery shoppers think a watermelon is just one item, there actually are multiple UPCs and PLUs used to track essentially the same item. This happens for many reasons, mostly due to the variety of suppliers a single grocery distributor may use. In order to gather accurate historical data, it is important to normalize that information.

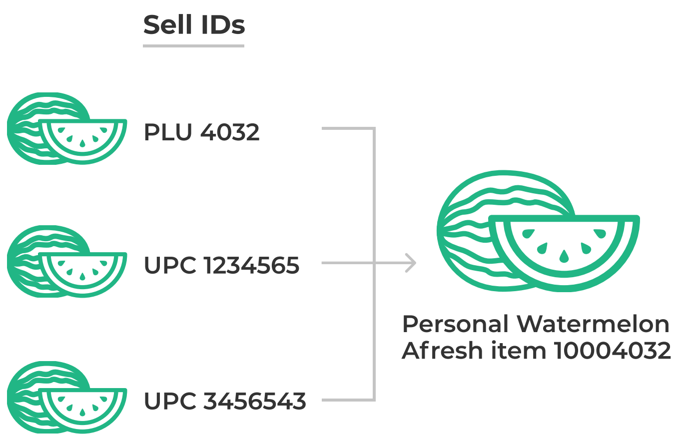

Data validation is absolutely necessary in fresh, where imperfect information and uncertainty are unavoidable. A data validation module assesses the health of all data inputs and its own estimation of inventory, demand forecast, and order recommendations. As a result the probability of any variables being incorrect is known, and the module passes its estimation of uncertainty to the AI engine.

 |

Afresh

Feature extraction

To increase the accuracy of predictions, Afresh's AI works to derive new features from raw data sent to us. Here are some of the types of features we generate to refine our recommendations:

Auto-regressive features

Examining historical demand for any and every item to identify trends like average demand for the current weekday that needs prediction or variance in demand over the last 30 days.

Example: Sales of limes tend to be higher on Fridays.

Time features

Considering the days of the week and time of month an order is placed, as well as proximity to holidays to understand and predict impact on sales.

Example: A week out from Thanksgiving, we expect potatoes and onion sales to increase.

Price and promotions features

Reviewing the current price, future price, and specific promotion associated with that to determine potential changes in demand.

Example: Departments need to order enough to prepare for a buy one, get one free sale.

Item features

Comparing the sales of an item in one store with the sales of that same item in another store to help highlight data discrepancies.

Example: Sales of avocados at store 415 generally follow the same pattern as sales of avocados in store 510.

Understanding the future: How to forecast for fresh

Looking back at the avocado conundrum: The ideal scenario would be to know exactly how much of each ripeness of avocado your local customers want to buy, and then to stock those exactly. But the factors that influence who buys avocados, how much they buy, and how ripe they want them are almost endless. We know that consumer demand for avocados is highly influenced by price, promotions, weather, holidays, the price of tomatoes, the price of organic avocados vs. conventional, and the current store location's season. We also know that large external factors like COVID-19 have dramatically disrupted the weekday/weekend shopping trend in many areas of the U.S., and that certain foods go in and out of fashion depending on recipe and diet trends.

Together, these factors beg for highly adaptable demand forecasting that's responsive to changes in consumer demand, rather than a "set it and forget it" system.

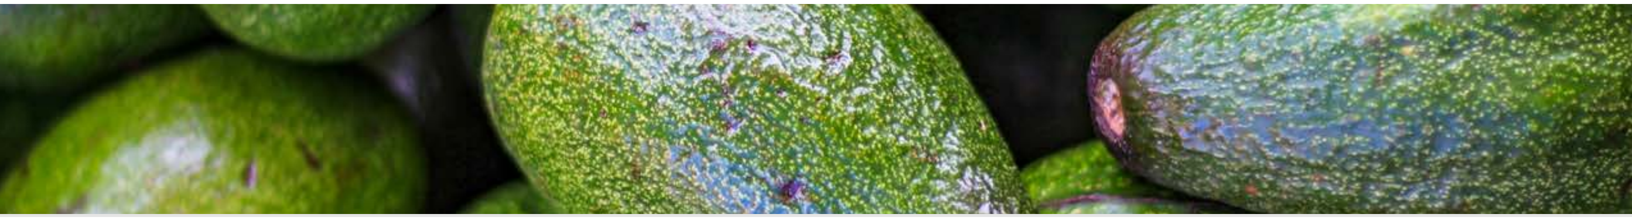

Afresh's AI solves this complex problem with state-of-the-art neural networks that take into account more than 40 different variables for each forecast. Our networks predict demand trajectories over long ordering horizons to forecast how demand will evolve over the next few orders. These networks are vastly superior to off-the-shelf AI solutions and are rigorously tested, validated, and monitored for forecasting quality.

No matter what we do, a perfect point demand forecast is impossible for a grocery store to achieve. Because all forecasting is built from historical data, there will always be factors that cannot be captured by AI. Maybe the local high school is having a football game tailgate event or the town newspaper published a new recipe for fruit salad. Afresh's models are trained to predict distributions of demand rather than relying on point forecasts. We know uncaptured events are inevitable, making predictions tricky to get right. But through rigorous validation of data and an approach focused on fresh, store teams can rest easy knowing our recommendation engine is always working to provide accurate suggestions.

Equipped with highly accurate probabilistic forecasts, grocery retailers can now tackle the real problem. How do we build that perfect shelf, day after day?

Simulation

The first thing we developed at Afresh is a high-fidelity simulation software library. This software builds a lifelike model of your store operating environment—shelf space, shelf lives, inventory accuracy, backroom space, hourly demand patterns, missing demand—and runs the Afresh system in this simulated world.

In this system, we are able to take any strategy that produces an order—say, always ordering 200 pounds of avocados—and evaluate its shrink, out-of-stock, and fullness percentage over a wide range of adverse conditions. For example, given a strategy, we can backtest the average impact to avocado fullness if a truck arrives late, identify weekdays or seasons where age distributions will skew young, and find and correct underperformance on promotions or price changes.

Before we launch to your store and before we launch any new modeling strategy, we know exactly what shrink performance to expect at your desired levels of shelf fullness. We'll never release a model that underperforms your current solution, and we'll never decrease performance by updating our system.

Optimal Ordering

Now that we can model consumer demand and evaluate strategies, we can get to the hard part. How do we get that perfect floor? If we order too much today, we'll skew the age distribution the day after tomorrow. If we order too little, we'll risk going out of stock and lose sales. Meanwhile, the demand distribution changes under our noses: an undersized reaction to a promotion today means that there will likely be an underreaction tomorrow.

How do we control this constantly changing system? We rely on approaches initially developed for state estimation in the Apollo mission and control algorithms for large chemical plants. Each order day, we plan out our next few order days with our current knowledge, so that we have full understanding into how an order today affects the order we'll make tomorrow. We make that first order, and, on the next order day, throw everything out and replan. This allows us to constantly react to changing demand and unexpected waste in order to always make the best order.

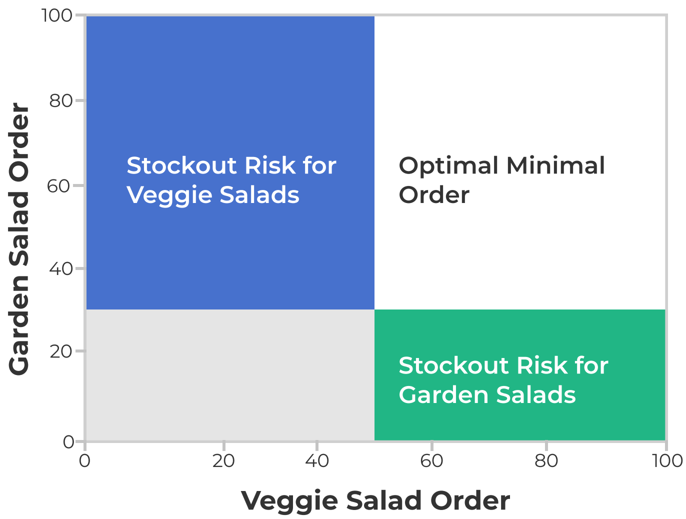

Afresh

It’s your store: Do what you wanna do

Afresh was created to help produce managers run their produce departments. We rely heavily on our expert partners to ensure store teams have a solution that works for them—consistently and accurately. In the pilot stage of any new launch, we develop close relationships with produce managers at stores in your chain to provide technical support and gather candid feedback about the quality of our order recommendations, inventory estimation, and recommendation coverage. We use this feedback to finetune our simulation and adjust our modeling parameters to fit your store’s needs, not what some research paper prescribes.

Once we launch in your stores, we continue to gather and react to produce manager feedback, as well as real store-level performance. We set a shrink reduction target with every customer for every quarter, driving results and improving forecasting, ordering, and product to achieve those goals.

Expertise in fresh and deep product knowledge can get you a long way and your fresh departments are technically running fine—but why settle for fine when fresh could be your competitive edge? Your customers certainly won’t.

What if your fresh departments could be running 10% leaner with 30% fresher product? What if your consumers could rely on you to give them avocado toast today, a fish taco tomorrow, and guacamole on gameday? With the right technology in hand, produce managers can finally run the departments of their dreams. That’s not AI magic beans, it’s rigorously built and maintained engineering systems that make each store the best version of itself. That’s Afresh.