Leveraging Generative AI To Unlock Unstructured Customer and Market Data Value in Offline Retail Stores

Keqin Zhao, Qian Wang, Shincho Zu, James Tang, Ankur Manoj Maniar
Columbia Business School, April 21st 2025

TABLE OF CONTENT

I. PROBLEM ASSESSMENT

With an oversupply of both customer and market data, and without meaningful tools to quantify and understand this data for optimal decision-making, it is evident that retailers must assess new strategies to optimally harness and process customer information to discover actionable insights. The end goal is not necessarily to sell more products and contribute to mass overconsumption, as is the current state of the market, but rather to optimally pair customers with best-fit, well-made merchandise that improves their lives, create consumer retention and stickiness thereby boosting brand loyalty, and harness internal productivity and personalization to foster an enjoyable shopping atmosphere where customers feel cared for and inspired, without being aggressively sold to.

Currently, 80% of retail sales still happen offline, signifying the importance and relevance of in-store sales. If the study of retail data can be optimized, it can significantly evolve the in-store retail experience while creating a complementary digital environment, where both landscapes become integral facets that play off each other to improve the overall customer experience, rendering stronger sales across the board.

Currently, a lot of customer and market data value are lost by retailers:

1. Internal customer data is often lost due to being unstructured in language or multimodal formats; think verbal, video, customer text feedbacks etc. These data cannot be factored into current decisions because of their formats.
2. Additionally, external market data is also lost. Simply put, macroeconomic trends and external factors are difficult to predict or control, causing decisions to be made on forecasting and intuition, without structures or processes.

The loss in data gathering and usage leads to a host of problems, which specifically impact the in-store experience. We have identified two major pain points for customers engagement and retailer internal management respectively, though we also recognize there are other pain points due to data losses.

1. For customers engagement: Customers are treated uniformly, without personalized messaging or products that are optimally matched based on their current or future needs. The lack of individual customer data makes it hard to leverage personalized insights.
2. For retailer internal management: Inventory planning team and store managers struggle to create accurate forecasting decisions for in-store sales replenishment and product merchandising/buying, due to lack of algorithmic data when computing customer trends and tastes. Currently, retailers mostly rely on sales records when making such decisions, leading to a mere 50% best forecasting accuracy, signifying legitimate opportunities for major improvement - if the tons of verbal customer feedback and macro-data can be meaningfully integrated for more accurate predictive power.

We want to address these issues and help retailers make more effective decisions by optimizing the ways the unutilized data pool is harnessed. Here, we see the use of generative AI as our strongest opportunity for data-driven evolution: previously, unstructured data lacked analysis tools, but the onset of LLM models are now able to accurately process these different data formats and generate legitimate and meaningful insights. Specifically, we believe that there are two technical products that can be deployed to address the customer engagement- and firm management- side data loss issues, respectively:

1. To benefit the customer experiences and foster more optimized transactions, **an in-store digital "shop guide" solution** can be used to interact with customers in a more personalized manner, with an LLM model trained on proprietary individual customer data, feedback, and historic preferences.
2. To benefit inventory planning team and store managers, **a new replenishment forecasting algorithm solution** can be introduced adding on to the existing predictive AI tools already in use in the market, by utilizing the large unstructured customer data pool and public market trend data.

We would like to explore the status quo of these two use cases and picture solution blueprints, which could become end state strategies for retailers to directly deploy in the future.

II. USE CASE FOR CUSTOMERS – A PERSONALIZED SHOP GUIDE

1. Status Quo & Existing Solution Market Scan

As GenAI technologies mature, a growing number of companies are developing shopping assistant solutions, ranging from basic product finders to more intelligent, personalized agents. Most commercially available solutions remain in the pilot or early adoption stage, especially for offline retail settings. However, momentum is building.

1.1 Solution Landscape for GenAI Shop Guides

Table 1: GenAI Shop Guide Current Market Scan

Category	Company / Example	Channel Focus	Core Features	Maturity Stage
E-commerce Native	Amazon (Rufus)	Online	Natural language Q&A, personalized recs, product comparisons	▨ Mature

	Walmart AI Shopping Assistant	Online	Custom offers, product suggestions, AI coupon pushing	Mature
Retail In-store Trials	Macy's (On Call, via IBM Watson & Satisfi)	Offline	Mobile-based product location + basic Q&A	Pilot
	Uniqlo AI Robot	Offline	In-store robot, visual recs, inventory search	Early deployment
AI Solution Vendors	Kore.ai, Zowie, Yellow.ai	Online → Omni-channel	White-label GenAI chatbots, customizable intents	Mature online / Immature offline
Enterprise SaaS Players	Salesforce (Einstein GPT), SAP	Back-end	CRM + POS GenAI integration, data sync	Mid-stage backend tools

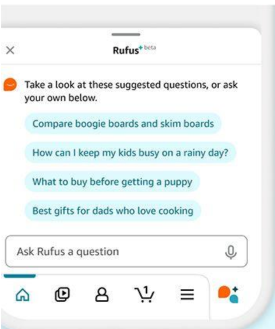

A few key themes emerge when scanning the current GenAI shop guide landscape. First, the development and deployment of GenAI-powered assistants have been significantly more advanced in online retail environments. The online AI shop guides mostly recommend products to consumers through chatbots based on the consumers' personalized shopping preferences and records. Amazon has introduced "the Rufus chatbot", its own AI shop guide, trained on extensive data from Amazon's product catalog, customer individual accounts and orders, and other sources such as social media trends. It also provides tailored recommendations and consolidates product information to customers. Walmart also has a similar GenAI shop guide to provide customized shopping recommendations, and personalized coupon pushes on its e-commerce channel. These online chatbots represent mature implementations that can handle complex queries and deliver tailored recommendations, leveraging vast digital datasets and customer histories. However, such tools have not yet been widely embedded into physical stores.

By contrast, the use of GenAI in offline retail remains largely experimental. Most existing in-store assistants—such as robotic guides or mobile Q&A tools—serve narrow functions like wayfinding or product lookup, and typically lack the depth of personalization, contextual understanding, or autonomous decision-making seen in online counterparts. Macy's has introduced an AI-powered mobile shopping assistant: shoppers can ask questions in everyday language about product locations, store services, and inventory. Similarly with an inventory locating function, Uniqlo's offline robotic shop guide also recommends products to customers by demonstrating pictures on its robotic screen. Customers on average use this robotic shop guide 1000 times per day, with a 10% conversion rate. These early pilots demonstrate the potential of in-store AI but fall short of realizing a fully interactive and intelligent shopping guide experience.

To bridge this gap, a convergence between online and offline systems is essential. Future GenAI shop guides must be designed as true omnichannel agents, capable of accessing CRM data, real-time inventory, and customer interaction histories across platforms. Only through such integration can they deliver consistent, personalized assistance that matches customers' expectations regardless of where they shop.

1.2 Expected Evolution

Table 2: Future GenAI Shop Guide Market Landscape Outlook

Stage	***Short Term***	***Long Term***
Online Chatbots	Wide deployment, already delivering measurable ROI	Enhanced personalization through integration of vision, voice, and behavioral cues
Offline GenAI Assistants	Pilot expansion and foundational in-store integrations	Agent-based assistants capable of taking real-time actions (e.g., reserve items)
Personalization	Primarily based on purchase history or stated preferences	Context-aware recommendations using sentiment, facial cues, and micro-behaviors
Agent Intelligence	Static responses and suggestions	Autonomous decision-making: reserving products, upselling, notifying staff, etc.

Looking forward, the evolution of GenAI shop guides is expected to follow a clear trajectory. While online chatbots are already well established, their in-store counterparts are likely to evolve from static, FAQ-style assistants to dynamic, autonomous agents.

These agents will not only recommend products but also take actions such as reserving items, suggesting outfit pairings, or notifying store associates. Moreover, advancements in multi-modal interfaces—including natural language processing, visual inputs, and contextual awareness—will elevate the GenAI assistant from a reactive chatbot to a proactive shopping companion. Retailers that can integrate customer data across channels and link GenAI agents with real-time inventory systems will be best positioned to lead in this space.

2. Solution Blueprint

While these existing trials are still preliminary, the future of GenAI in-store shop guide can already be imagined. Such an assistant would combine the functional pieces already explored by both online and offline channels, and better leverage customers' individual data to become a more personalized version.

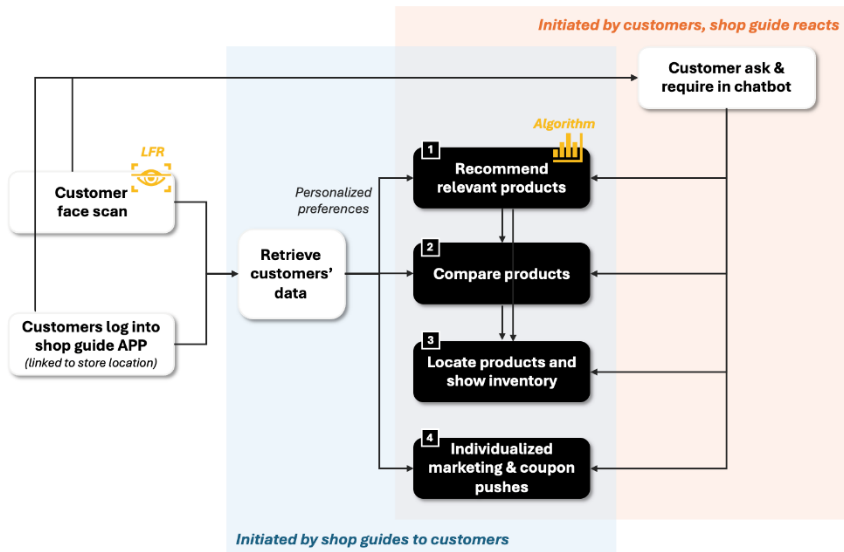

An GenAI in-store shop guide should have four main functional capabilities: (1) recommending products based on personalized preferences and requirements; (2) compare product dimensions with a well-rounded conclusion; (3) locate products in store and tell whether there are inventories; (4) push marketing ads or coupons to customers based on their tastes. The shop guide interaction process can be either initiated by customers or by the guide itself: whenever customer log into the shop guide solution and ask questions in the chatbot, the shop guide can react accordingly based on customer requirements, customers' personal preferences as in records, and other trends / macro-factor analyses; even if customers don't directly ask for help from the shop guide, the guide should be able to retrieve customer individual data whenever the customer logs in, and proactively recommend product portfolios to the customers.

The algorithm of product recommendation is quite important, which should be trained on a large amount of customer personal data and consider multiple perspectives: past customer sales records should be analyzed to see the customer's consumption pattern and product interests; customer persona should be embedded to predict customer reactions to chatbots and changing purchasing behaviors; product portfolio dimensions should be clustered into different types so that similar products or product bundles could be recommended together; trends should be considered so that certain products can be recommended during certain special periods; the customer's personal moods and life experiences can be considered (based on chatbot history dialogues) to make the recommendation and communication more effective. All these analyses in algorithm can make the shop guide a personalized one for every individual customer coming into the store.

With GenAI's continuous technological breakthroughs, chatbot has become a relatively mature tool, while AI agent is still under development. The solution blueprint is now imagined in a chatbot format, but it can be further elevated towards an AI agent format in the long-term. Compared to a chatbot, an AI agent should be able to take actions based on the chat results. For instance, when the customers have indicated their interests in products, the AI agent version of shop guide should know itself to directly take pre-orders for the customers to lock inventories, instead of just making recommendations without actions.

3. Value and Challenges

3.1 Value Added by GenAI Shop Guides

Table 3: GenAI Shop Guide Benefits

Value Area	*Description*	*Corresponding Challenge*
Sales Conversion	Personalized recommendations increase likelihood of purchase	Low engagement if customers do not initiate conversation with the assistant
Customer Experience	Faster, more relevant assistance throughout the shopping journey	Lack of trust or interest in AI interaction in physical environments

Online-to-Offline Integration	Unified view of customer behavior across channels enables better targeting	Data silos between online and in-store systems; fragmented infrastructure
Operational Efficiency	Reduces routine workload on store associates	Requires integration with store inventory, CRM, and real-time product data
Personalization Depth	Tailored suggestions based on inputs or preferences	Friction in obtaining user data; privacy concerns; limited data sharing

These assistants not only support conversion but also enhance experience quality and backend efficiency. For example, Uniqlo's in-store AI robot demonstrates a 10% conversion rate from roughly 1,000 daily interactions. As GenAI systems become more context-aware and multimodal, their role in bridging online–offline experiences will become increasingly central.

3.2 Challenges to Adoption

Despite these benefits, physical retail settings pose unique adoption hurdles. Most notably, many customers are not naturally inclined to engage with AI systems unless prompted. The lack of proactive interface design, perceived complexity, or unclear value may lead to low usage.

Moreover, personalization relies heavily on customer input, yet shoppers are often reluctant to share information unless the process is effortless, secure, and clearly rewarding. Without sufficient data, AI recommendations lose effectiveness, diminishing the perceived utility of the assistant.

Additionally, implementation carries significant cost and operational considerations:

Table 4: GenAI Total Cost of Ownership Analysis

GenAI Tools & Platform Access	*Licensing fees, API usage costs for foundational or fine-tuned models.*
Prompt Engineering	Crafting, testing, and iterating prompts for optimal performance across use cases.
Inference Costs	Ongoing compute cost incurred each time the model is queried (esp. large LLMs).
Fine-Tuning Costs	Additional model training to adapt to store-specific tone, product taxonomy, or brand needs.
Infrastructure	Cloud/GPU server infrastructure, edge computing devices in-store, bandwidth.
Data Management	Collection, cleaning, labeling, and integration of customer or inventory data.
Operations	Daily monitoring, bug fixing, model updates, incident response.
Compliance (AI Regulations)	Meeting legal standards like GDPR, data localization, model transparency.
Talent	Recruiting and retaining AI engineers, prompt designers, and product leads.

In addition, personalization relies heavily on customer input, which introduces friction. Customers may hesitate to provide information or engage in conversation with AI agents unless the process is perceived as effortless, secure, and rewarding. If the assistant receives minimal input, its recommendations become generic and less effective, weakening the user experience.

3.3 Incentivizing Customer Participation Through Gamification

To address data acquisition friction, retailers can embed gamification mechanics into the AI experience, turning interaction into an engaging and rewarding activity. Well-designed gamification increases both participation and the richness of input data, particularly among younger or more digitally native consumers.

Table 5: Gamification Strategies for Customer Data Contribution

Element & Mechanism	***Data Value***
***Quests and Missions**: Encourage task completion (e.g., "Find three eco-friendly items") via the AI assistant.*	Reveals category interests and in-store behavior patterns.
***Points and Loyalty**: Reward repeat interaction through a tiered points system.*	Captures session frequency and engagement depth.
***Surprise Rewards**: Offer randomized benefits after assistant use.*	Increases engagement across diverse visit contexts.
***Personality Quizzes**: Use preference or mood-based questions to recommend products.*	Generates soft persona and emotional-state data.

***Recognition Badges**: Acknowledge frequent or varied use with visible status.*	Encourages retention and broader interaction types.

By framing data sharing as part of a purposeful and enjoyable experience, gamification reduces psychological friction and builds trust. Over time, this enables the AI assistant to shift from a passive tool to a proactive, high-value shopping companion.

To fully unlock these benefits, retailers must invest in intuitive UX, seamless backend integration, and transparent communication around data use.

III. USE CASE FOR STORES – A NEW REPLNISHMENT FORECASTING TOOL

1. Status Quo & Existing Solution Market Scan

As retailers continue to modernize supply chain operations, many have begun exploring how GenAI can address limitations in traditional replenishment forecasting models, including predictive AI. While historical sales and inventory records have GenAI-enabled replenishment logic, such structured data alone often fails to capture the nuanced, rapidly changing nature of consumer demand. This is especially true in categories affected by external variables such as weather, social media activities, or local events. Recently, a number of leading retailers have adopted GenAI-enabled forecasting systems that integrate both structured and unstructured data sources, yielding measurable improvements in prediction accuracy and operational efficiency.

1.1 Solution Landscape for GenAI Forecasting Tools

Table 1: GenAI replenishment Forecasting current market scan

Category	Company / Example	Data Types Used	AI Application	Maturity Stage
Global Retail	Walmart	Sales, Weather, Social Media	Category-specific Forecasting	Mature
Grocery	Kroger	Voice Memos, Complaint Texts	Adaptive Replenishment	Early Adoption
Fashion	H&M	Social Media Images	Trend-based Stocking	Pilot
Convenience	7-Eleven Japan	Store Sales, Staff Notes	Store-specific Modeling	Mature
E-Commerce	Zalando	Reviews, Return Reasons	Restocking Decisions	Pilot

A few key insights emerge from the current GenAI forecasting landscape. Retailers are increasingly moving beyond traditional data inputs by using unstructured and external data sources to enable smarter, more adaptive replenishment strategies.

One leading example is Walmart. The retail giant has incorporated real-time external signals such as weather forecasts, social media conversations, and other data into its demand forecasting systems. By layering these insights on top of traditional sales and inventory records, Walmart has improved forecast accuracy by around 20%. This approach has been particularly effective for seasonal items and region-specific products. However, at present, Walmart mainly relies on basic weather data, like temperature, precipitation, and wind speed, and uses ZIP-code-level geographic granularity. This leaves room for improvement in how deeply and precisely weather factors are used in forecasting. There is significant potential in expanding this application to incorporate long-term weather patterns for broader regional demand forecasting. By doing so, retailers could not only improve inventory allocation at the individual store level but also reduce operational costs across entire geographic areas. Moreover, weather data can be strategically paired with dynamic pricing of selected items. For example, anticipating warmer-than-average temperatures over the next month could trigger price adjustments for seasonal products like beverages or air conditioners. This integration enables both sales maximization and inventory optimization, ultimately delivering greater business value.

In the grocery sector, Kroger takes a different but equally innovative approach. The company uses GenAI to capture and analyze operationally important but often overlooked inputs, such as voice memos and handwritten complaints from store staff. These unstructured inputs often reveal real-time issues, like empty shelves or sudden shifts in customer demand. The GenAI model transcribes, categorizes, and interprets this data to adjust replenishment schedules dynamically. As a result, Kroger has reduced out-of-stock incidents by around 15%, improving both customer satisfaction and operational flexibility. That said, the current system only processes complaints that are explicitly recorded in voice or written form. Feedback that is never expressed or reported still goes unnoticed, which limits the model's completeness. In the area of voice data utilization, combining GenAI models with in-store camera and sensor monitoring systems presents a powerful opportunity. This integration would allow the system to detect shelf conditions or unusual customer behavior, enabling retailers to address emerging issues before they lead to complaints. It could also help identify early signals of unmet or latent demand that would otherwise go unnoticed.

In the fashion industry, H&M operates in a fast-changing environment where trends shift quickly and product life cycles are short. To keep pace, H&M uses image-based GenAI tools that analyze millions of social media posts on platforms like Instagram. These tools identify rising fashion trends by recognizing patterns in colors, styles, cuts, and hashtags. This allows H&M to launch new collections more effectively and to restock trending items before demand peaks. However, not all products that go viral on social media lead to strong sales. Additional validation is needed to understand which social buzz actually converts into purchases. Social media activity, when combined with online purchase behavior, presents another promising avenue. By analyzing which social buzz actually converts into sales, companies can identify which content or product trends are most predictive of actual purchasing behavior. This insight can significantly improve the accuracy of replenishment planning and marketing strategies.

A particularly unique use case comes from the European e-commerce company Zalando. Zalando employs large language models (LLMs) to analyze qualitative customer feedback such as product reviews and return comments. This deeper understanding helps the company improve restocking decisions and avoid repeating past inventory mistakes. While this is an impressive application, it assumes that customer feedback is honest and accurate. The model also faces limitations with new products or items that have very low return volume, as there is not enough data for the LLMs to learn from effectively. LLMs can be supplemented with data from similar items. However, for this approach to be truly effective, product similarity must be defined not just by standard sales categories like gender or clothing type, but by design patterns, stylistic attributes, and functional characteristics. A more nuanced product mapping will enhance the model's ability to generalize insights across low-data SKUs. Similarly, in retail chains where store-specific GenAI models are deployed, there is an opportunity to reduce maintenance costs by grouping stores with similar characteristics into shared models. By periodically re-evaluating store classifications, companies can strike a balance between hyper-local optimization and scalable operational efficiency.

Finally, an example of a highly localized GenAI approach can be seen at 7-Eleven Japan. Each store operates its own individualized forecasting model, trained on local POS data, sales history, and regional delivery constraints. Recently, GenAI has been added to analyze staff-written notes and free-form customer comments, which often contain subtle insights into local preferences or shortages. While this allows for extremely precise forecasting, it also comes with high costs. Maintaining thousands of models, over 20,000 stores in Japan alone, requires significant operational resources. Since many of these models are still in the early stages, ongoing maintenance and updates are particularly expensive.

2. Solution Blueprint

Traditional models using past sales data can output accurate information and are currently relied on for inventory and replenishment forecasting by store managers. However, because they are based largely on lagging indicators, these models assume future demand will resemble historical patterns. As a result, they struggle to respond proactively to sudden shifts in consumer behavior or market conditions. GenAI introduces a fundamental shift in forecasting methodology, not just by expanding the amount of data used, but by unlocking entirely new types of data. Most importantly, GenAI enables forecasting models to integrate real-time, unstructured internal signals that were previously unusable. These include customer conversations, staff observations, behavioral patterns, and interaction logs; data points that often signal changes in demand before they are visible in sales. These leading indicators, such as the shop guide, customer search activity, and natural language feedback, can reveal rising interest in specific products, stock dissatisfaction, or unmet needs. Incorporating these signals into forecasting models allows retailers to move from reactive to predictive replenishment, anticipating demand shifts ahead of time and improving service levels while minimizing excess inventory.

External signals still play a meaningful role. For instance, weather data is often used for short-term forecasting, but there is untapped potential in using long-term weather patterns for regional demand planning. When paired with dynamic pricing, this can improve both inventory allocation and revenue optimization. Similarly, social media data, when cross-referenced with purchasing behavior, can help distinguish between viral interest and true demand. These external insights, however,

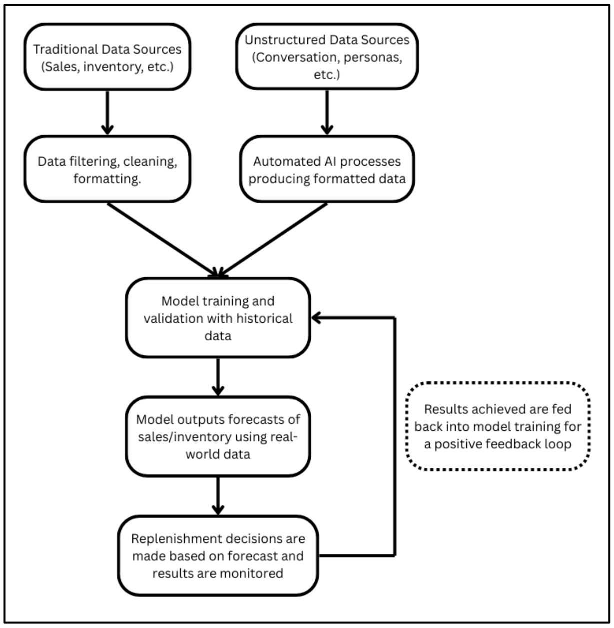

become exponentially more valuable when integrated with internal data that reflects what customers actually say, do, or experience at the point of sale.

As these GenAI-enabled systems are implemented, real-world performance, such as differences between forecasts and actual sales, can be fed back into the models to continuously refine and improve accuracy. Over time, the system learns from its own outputs, adjusting internal weights and feature importance to better reflect changing market dynamics. More importantly, by incorporating internal data such as customer interactions, shop guide queries, and staff observations, the model can detect shifts in demand before actual purchases occur, enabling not just reactive, but predictive replenishment. This forward-looking capability is expected to significantly enhance forecast precision and overall responsiveness to customer needs.

3.1 Internal Data Sources:

To enable next-generation forecasting, it is essential to go beyond structured data and tap into rich, real-time internal signals that reflect customer intent, behavior, and operational insights. These internal data sources provide a window into demand patterns before they are visible in sales figures. By capturing information from on-the-ground staff, in-store interactions, and customer feedback channels, retailers can identify emerging trends, gaps in inventory, or unmet needs much earlier. The list below outlines key types of internal data that can be collected and transformed into actionable inputs for AI forecasting models.

1. Customer Feedback & Interaction Logs
    1.1. Natural language input from in-store feedback kiosks or post-purchase surveys
    1.2. Complaint texts and voice memos stored in CRM systems
2. Persona Clusters and Behavioral Tags
    2.1. AI-generated customer segmentation based on past purchases, preferences, and interests
    2.2. Evolving interests tracked over time, tagged per persona group
3. Data from GenAI In-Store Shop Guides
    3.1. Query logs from product searches and chatbot interactions (e.g., “best dry shampoo for oily hair”)
    3.2. Frequency of product mentions or comparisons, especially if not currently stocked
    3.3. Customer reactions to recommendations, i.e., products clicked on, ignored, or rejected
    3.4. Drop-off points in customer journeys where products were not found or not purchased after recommendation
    3.5. Coupon usage and promotional response (helps predict demand surges)
4. SKU-Level Engagement Signals
    4.1. In-store dwell times or proximity (if combined with store camera/computer vision data)
5. Store Associate Notes and Observations
    5.1. Staff-input notes on recurring customer product questions or dissatisfaction
    5.2. Ad-hoc updates on unexpected inventory issues or popular emerging needs

3.2 Feature Engineering

Once the framework is in place, collection of data is trivial to the forecasting process. Perhaps the most important aspect that can make or break the implementation is feature engineering for the forecasting model. As one can see, most of the unstructured data is in text form, and to use it with regression models, it is crucial to extract useful information in a quantified feature. Here, I will provide examples for how each data source can be crafted into features; However, best combinations of features to include in the model is determined based on testing and cross-validating models (with different combinations and dependent variables) and cannot be determined at this stage.

1. Customer Feedback & Interaction Logs:
    1.1. Negative sentiment ratio: percentage of sentiment that is negative. Sentiment can be determined with computer language processing algorithms (both AI and non-AI options exist)
    1.2. Average sentiment score (past x days): The average sentiment score of the item for the past x number of days. Sentiment scores can also be generated by computer language processing similar to above.
2. Persona Clusters and Behavioral Tags:
    2.1. Number of profiles in category persona: This is an important metric that can be generated with gen-AI. Each SKU should be associated with a persona category that its style should most align with. Tracking the number of profiles fitting that persona category will be a strong predictor of sales.
3. Data from GenAI In-Store Shop Guides:
    3.1. Number of queries relating to theme/product type of SKU: From the queries asked to in-store AI shopping guides, named entity recognition (NER) can be used to extract useful features from conversations such as the theme demanded by the customer or the type of products they wanted.
    3.2. Growth rate of queries relating to theme/product type of SKU: The growth rate of the number of queries can be used to capture viral growth. However, it should be used in conjunction with number of queries since having an 100% growth rate could be a product with 1 query jumping to 2 queries, or 100,000 jumping to 200,000.
    3.3. Query to Purchase Ratio per SKU: An important factor to consider for queries is how many queries that recommended this SKU led to the purchase of the product. If a product is recommended many times but has a low percentage of purchase, the sales numbers should not rise.
4. SKU-Level Engagement Signals:
    4.1. Human minuets spent at category: Combining computer vision with AI algorithms, we can generate for each category the total amount of minuets customer spent at the location, which provides insight into in-person interest similar to how websites can measure traffic by category. When matched with the category of each SKU, it should allow the model to capture a more holistic sense of consumer interest.
5. Store Associate Notes and Observations:
    5.1. Staff flag for trending SKUs/Categories: a simple variable that is represents staff sentiment on particular SKUs and categories. It can be engineered in different ways, such as taking the most popular opinion, or taking an average. How it should be used is determined in testing.
    5.2. Staff Sales Predictions: Although not related to AI, allowing staff members to participate in a prediction market like scenario can aggregate information and inform the model.

3.3 Value and Benefits

To translate these technical capabilities into tangible business outcomes, each feature source must contribute a clear and measurable value to the overall forecasting and inventory management process. GenAI assistant logs, for example, enable early detection of product interest, which traditional sales data would only capture after-the-fact. Customer feedback provides a window into satisfaction or dissatisfaction drivers, enabling faster product lifecycle interventions. Persona-based analysis supports differentiated stocking strategies by location or demographic group, which can directly improve relevance and sell-through.

Furthermore, staff notes and frontline observations often surface localized product insights that are missed by aggregated sales trends; these become critical signals for capturing ground-level demand fluctuations. Historical sales and inventory data, while foundational, play a supporting role by anchoring newer AI-driven signals to longer-term demand patterns. Together, this multi-source framework ensures that forecasting decisions are not only data-rich but also contextually intelligent, responsive to real-world nuances, and aligned with both operational efficiency and customer experience goals.

Feature Source	Business Value	Impact Area

Gen-AI Assistant Logs	Early detection of interest and emerging demand	Improves forecast accuracy, increases sell-through rate
Customer Feedback	Identifies product satisfaction/dissatisfaction early	Reduces returns, supports product improvement and lifecycle management
Customer Persona Analysis	Personalized stocking per location or segment	Increases customer satisfaction and in-store relevance
Staff/Observational Notes	Surface real-world product requests missed by sales data	Captures ground-level demand signals and enables manual override
Historical Sales / Inventory Data	Baseline + reinforcement for demand learning	Anchors predictions in past trends while adjusting for new signals

3.4 Existing Ventures

As AI becomes more understood, more companies are creating custom solutions for themselves to reap the benefits of AI-powered analytics. Zara is one of the front runners in this space, using AI to take into account factors such as past sales, current purchasing patterns, and social media trends. It demonstrates how a model that combines internal and external sources can provide tangible value in terms of more efficient production and allocation due to better forecasting[1]. Stitch Fix is another company, with a focus on personalized styling, that has utilized AI for personalization services and inventory management. By analyzing customer data features such as style preferences, purchase history, and feedback, the company's algorithms can generate personalized style recommendations. Stitch Fix also employs a gamified AI tool "Style Shuffle" that asks customers to rate different outfits which can generate insight on future sales. It can also be used as a prototyping tool to test new products[2]. On the flip side of companies building customized AI solutions are companies that specialize in providing services to others. One such example is Autone. "Founded by Adil Bouhdadi and Harry Glucksmann-Cheslaw, who previously built Alexander McQueen's intelligence platform[3]" retailers using the Autone platform reported "a 55% reduction in inventory levels, a 25% increase in forecasting accuracy, a 30% increase in sales, and save 45 hours per week on manual tasks.[4]"

IV. OTHER AREAS OF DATA VALUE TO BE UNLOCKED

There are a host of additional generative AI use cases that can improve and evolve the customer experience. Of most interest include (1) dynamic pricing algorithms, (2) customer management, and (3) targeted advertising.

1. Dynamic Pricing. By incorporating unstructured customer feedback as well as live competitor pricing analyses, retailers can evolve their pricing algorithms and offer dynamic pricing for in-demand products. This is especially useful in industries like beauty, where social media virality plays a key role in determining which seasonal products will be the hottest and sell out the fastest. By responding in real-time to these popularity fluctuations, retailers can remain competitive and in sync with customer tastes.
2. CRM. Unstructured customer data can also be utilized to evolve the way customers are managed. Instead of grouping customer segments with similar characteristics, offers and services can become even more granular to appeal to niche consumers with specific needs - once those needs are discovered, analyzed, and met with specific messaging and products. Handling specific customer pain points becomes easier once unstructured data is processed and incorporated into the overarching customer-brand relationship strategy. This eases the burden of guessing what customers want.
3. Targeted Advertising. By understanding customer needs with more specificity, digital marketing pushes, print advertising, in-store visuals, and OOH campaigns can all be evolved to help attract, manage, and retain customers with higher degrees of success. Instead of generic catch-all communication, niche-specific communication will offer customers the attention and care they need in oversaturated markets.

---

[1] https://sites.lsa.umich.edu/mje/2025/04/04/ai-powered-fashion-how-tech-is-reshaping-the-future-of-zaras-fashion-empire/

[2] https://www.voguebusiness.com/story/events/how-stitch-fix-is-using-ai-to-predict-trends

[3] https://event.businessfrance.fr/nrf/autone/

[4] https://autone.io/series-a/

In all, the integrative use of generative AI tools across multiple pain points can provide a holistic strategy that evolves the customer experience, boosts sales, reduces returns, and offers consumer satisfaction and loyalty.