

Optimizing Retail Demand Forecasting: Big Data-Driven AI Models for Enhanced Customer Experience and Operational Efficiency

¹ Vishwanadham Mandala

¹ Data Engineering Lead

¹Cummins, Columbus, USA

Abstract

The retail and consumer sector is a major contributor to global economic output but faces challenges in terms of uncertain consumer behavior, fast-changing technology, and retail convenience. Accurate demand forecasting is of paramount importance in the overall retail decision-making process. It is a prerequisite of broader application areas, such as stock control, merchandising, and pricing. Inadequate demand forecasting can lead to lost margin opportunities in terms of out-of-stock scenarios, as well as heavy discounting of goods in response to overstocking. In online retail markets, convenience disappears if goods are not available when the consumer wants or a store closes, as goods sell out at a full price. In a world of volatile trading conditions, the forecast of future demand is more difficult and less reliable than it used to be, due to both internal business complexity and increased external market disruption. Demand signal data need to be incorporated into the forecasting process, and it is necessary to use the techniques available to improve the efficiency and speed of the demand forecasting process to enhance the level of operational control. To cater to all these challenges and opportunities, recent academic research has covered a wide range of themes in retail demand forecasting. Retailers themselves are investing heavily in information systems, data management, analytical, and forecasting capabilities. Software providers offer a broad array of demand forecasting tools, some of which are cloud-hosted. Options are attractive as they allow access to broader data and more sophisticated models. Some business problems require solutions that are not covered by off-the-shelf software and therefore may require more customization to be suitable for enterprise implementation. Such solutions could aim at automating the demand forecasting process, leveraging a broad range of external data sources, creating outside views of future demand for bricks-and-mortar retail space, as well as seamless integration with operations and other decision-making cycles. Such AI-driven demand forecasting solutions can achieve positive outcomes, such as enhanced customer experience and supply chain optimization.

Keywords: retail sector, consumer behavior, demand forecasting, stock control, merchandising, pricing, out-of-stock scenarios, overstocking, online retail markets, convenience, trading conditions, business complexity, market disruption, demand signal data, forecasting process, operational control, retail decision-making, data management, analytical capabilities, AI-driven forecasting, supply chain optimization.

1. Introduction

To date, retailing remains one of the most important economic sectors. Sales and Operations Planning is conducted to help coordinate planning for operations and resources responsible for delivering projected retail demand or sales. Demand forecasting provides in the sales and operations planning process and represents one of the very first, critical, and foundational planning components. In recent years, with the addition of the omnichannel strategy, the complexity of retail demand forecasting has

increased, with the need to handle the multifaceted and interrelated flow of information among these channels. To address these challenges, both retail businesses and researchers have been using machine learning models for demand forecasting because of the requirements to handle numerous different types of influential factors and the quality of forecasting outcomes.

Several studies have been conducted to establish methods for demand forecasting that are suitable for the increasingly complex and data-rich omnichannel context. They have adopted big data-driven demand forecasting methods for retail to increase accuracy and process real-time and high-frequency data. However, the attempts to make practical use of these methods are concentrated in e-commerce settings, and few open-source and pre-constructed forecast models can be easily applied for similar purposes. In debate, many established statistical methods combined with numerous time series features that enable a broad perspective on time series. Such features can be generated at scale for large data sets, and statistical methods are self-explanatory and interventional for special event analysis.

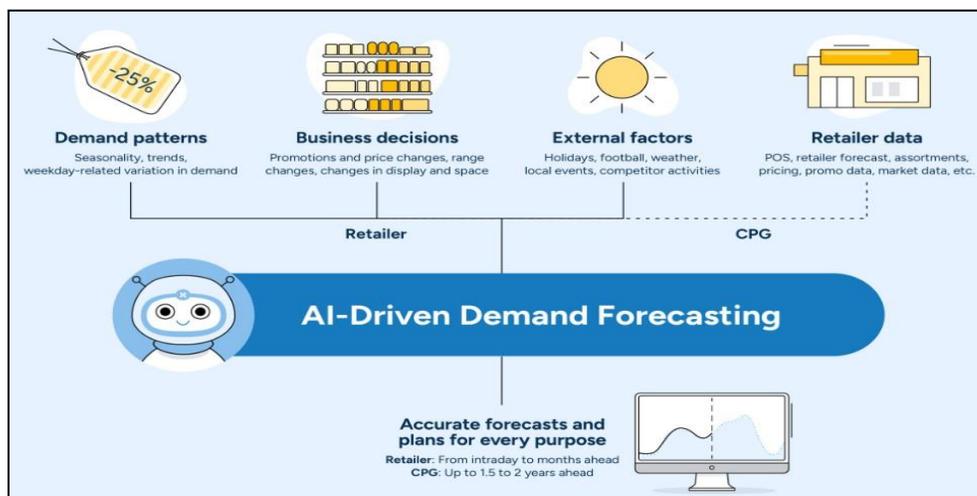


Fig 1 : The guide to demand forecasting for retail and consumer goods

1.1. Background and Significance

Introduction The rapid pace of innovation in technology and the associated proliferation of available data has completely transformed the retail business. The dynamic interactions between geography, economy, demography, and market trends make it incredibly challenging for retailers to accurately maintain the right stock of goods and services for their customers. The ability to develop accurate demand forecasts has thus become increasingly critical for retailers in today's fast-paced, data-driven marketplace. Successful forecast models must balance many varied, sometimes conflicting, and often misaligned interests while also navigating an increasingly complex marketplace. The common goals of demand forecasting include enhanced customer satisfaction through lowering the number and size of stockouts, better use of warehouse and storage resources, and greater enterprise-wide coordination of all activities. Intelligent software platforms serve to amplify the activities of both physical and virtual workforces. Applied to the specific domain of demand forecasting, an intelligent software platform should seek to curate, organize, display, and analyze both structured and unstructured large volumes of internal as well as external data. This software can then be generalized and scaled across multiple departments and multiple users to help organize data and augment the decision-making process. This provides a simple yet effective framework, providing retailers with current and relevant knowledge to react more effectively to changing patterns of consumer demand. Our research utilizes methods by which multiple suppliers can define dependencies between managed data sets, develop pre-configured analysis modules to act on these shared data sets, permitting the underlying systems to provide process rules to manage data that gets delivered and facilitate automatic distribution of analyses, inherently enhancing the benefits reaped from the shared data usages.

$$D_t = f(\mathbf{X}_t, \mathbf{W}) + \epsilon_t$$

D_t = Predicted demand at time t

\mathbf{X}_t = Input features (e.g., sales history, promotions)

\mathbf{W} = Model weights (parameters)

ϵ_t = Forecast error

Equation 1 : Demand Prediction Model (DPM):

1.2. Research Objectives

We are thus concerned with developing a demand forecasting system for retail from two perspectives: enhancing customer experience optimization and increasing design space for micro-targeting, precision, and operational efficiency achieved through decreased lead time. The optimization perspectives identified require vast amounts of customer fingerprint and behavior data, which have been largely ignored thus far. We address the gaps in prior works and the limitations in industry practices and propose a more comprehensive, intelligent forecasting system based on AI-enabled big data analytics. First and foremost, we propose that merging customer-rich features and product-sensed demand would help yield a richer customer order demand forecast.

The research objectives are demonstrated in two research dimensions. First, the demand forecast's impact on customer satisfaction and the additional sales uplift through handling and operational efficiency improvements shall be measured and demonstrated. Secondly, the functionalities and values of each data input type and the demand forecasting blends shall be evaluated for their capability to support customer and operational optimization and append directly or indirectly to the enhanced landed forecast accuracy justification. A persistent demand forecast value chain, which starts from rich customer data, upgrades product-sensitive demand forecasting and arrives periodically at iterative customer servicing and business operations, is thus formed.

2. Theoretical Framework

Society is experiencing the digital revolution through the continuously increasing integration of the internet and the everyday life activities of citizens and other socio-economic agents, such as the behavior of consumers in the post-globalized and hyperconnected retail scenario. The emergence and expansion of a hyperconnected society is made possible by the creation, organization, and dissemination of a colossal quantity of data originating from several different sources. The massive size of the data is a characteristic that introduces significant challenges to management and, simultaneously, suggests the existence of great potential in terms of information processing for the comprehension of the new behavior patterns of consumers. The growth of both the data and the use of the internet has led to changes in the information sources that companies and governments can access for making decisions about their activities, even to the point of redefining innovation, state policy, and policy-based innovation.

In this new societal context, and retail, big social data and related technologies form the central theme, in terms of a wide variety of data sources—ranging from online to consumer behaviors, geolocalized data, scientific data collection, social media, sensors,

and financial activity to more classical data sources—accounting, generic statistics, and also the information about the organization of the retail business, and the financial statements from public and private sources. The data are ready to be collected, accessible, and immediately usable. The research raises issues related to edge data-driven, big social data, and related technology. The research is focused on specific domains in designing new data framework innovation processes relevant to the area of retail. We investigate the design of new tools and utilize them together with sophisticated statistical and big data modeling to conduct policy-relevant studies.

2.1. Retail Demand Forecasting Basics

Demand forecasting takes on several forms when applied to business. This paper concentrates on retail demand forecasting for two main reasons. Firstly, the scale of the consumer market retail demand and the accuracy of the sales forecast jointly and heavily predicate the cost of guiding production, the rationality and accuracy of business decision-making, the level of supply chain management, and finally, the whole value chain of enterprises. Secondly, the predictive accuracy and precision of retail demand determinations strongly impact the customer experience and brand image. Therefore, the discussion within this category is placed under the orientation of retail demand forecasting mainly to improve customer satisfaction. Demand forecasting is a core issue in retail applications. It has attracted significant research interest because of the challenges it poses: short life cycles, promotions, seasonal products, and time-dependent product demand. Demand forecasting helps guarantee the availability of products, enhancing customer satisfaction. However, achieving high accuracy is complicated because it requires understanding a high number of low-frequency items. Retail demand forecasting is expected to support a single company managing its inventory, as well as multi level supply chain management according to the hierarchical feature of the forecasting horizons.



Fig 2 : AI Demand Forecasting Software Solutions Company

2.2. Big Data and AI in Retail Forecasting

Big data has been playing an instrumental role in various fields of business, including retail, where it has powered analytics insights into customer behavior and powered machine learning to recommend the most desirable products or services. Big data technologies in retail can incorporate large datasets from a variety of sources to collect information regarding demographics and customer preferences and provide advanced analytical tools to visualize and identify customer behavior patterns, preferences, and shopping habits. This is important for retailing as analytics insights can reflect changing customer demands and motivations in different market segments. Thus, information for making informed decisions can be appropriately strategized to attract customers

based on the dynamics of market demand and product preferences. Besides, as big data technologies also effectively support purposeful manipulation and interpretation of control parameters that define market demand, AI-based demand forecasting holds promise. With the combination of demographic data such as housing, gender, and age information with retail offerings, organizations can better adapt the deployment strategies of marketing resources for insights into customer preferences for specific items and, thus, increase their chances of sales success. Studies have successfully shown evidence of the relationships between big data growth, analytics, and AI in the retailing industry to increase profits. They further assert that an understanding of customer preferences and demands is a success factor in retail, and this dimension of increased customer knowledge can be much more easily reached by more extensive use of big data, with large potential consequences for retail and operational performance. As AI continues to transform the retail world, demand forecasting is another fundamental application that is continually evolving for improving business accuracy.

3. Methodology

The optimal demand forecasting model should be adaptable to various types of products and customers and include environmental factors, such as building architecture and city infrastructure, marketing ROI, and supply chain parameters like supply chain lead time, cost, quality, and capacity to deliver. We propose an AI solution that builds demand forecasting models by developing, training, and deploying AI-based 'drop-in' demand forecasting models. The 'drop-in' model is a layer that is appended to standard retail analytical models in the form of a pre-trained standard model or a pre-trained model at the customer site or on-premises, working with the data from the customer and providing the deployment option for each client. The additional layer leverages demand modeling experience from projects on practically similar demands to develop new enhancements to the standard model. The project parameters include quality/speed trade-off selected by the client, training, and network architecture selected during the training phase, and actual demand data. The additional layer analysis of the supply chain ROI effect is done together with the customer's subject matter expert. The additional layer can update the training sample based on the analysis as well. The pre-trained model structure and training parameters can be updated based on demand feedback. The proposed solution can match the accuracy of the full-model deployment process several magnitudes faster.

3.1. Data Collection and Preprocessing

Data collection and preprocessing are traditionally considered an unsolvable issue by many supply chain professionals when investigating the applicability of an AI-based demand forecasting model within their retail efforts. Although retailers could be in a more advantageous situation compared to other industries in this aspect due to the availability of a wealth of retail data, the main challenge lies in transforming these disparate, non-uniform, and often inconsistent records into a form that could contribute to an accurate model and provide valuable insights. For the retail demand forecasting problem, the aim of an ideal demand forecasting system operation can be expressed in four major phases: data collection from all available internal and external data sources, data preprocessing where raw data is cleaned for implementation-specific requirements, feature generation where meaningful patterns and structures are extracted, and model building where advanced statistical and machine learning models are applied to drive practical value.

The first step may sound natural to many, but it is not the case since retail businesses are not necessarily ready to entertain data streams at that scale and may have insufficient infrastructure in place to handle such demand flows and subsequently store and reuse the data for other purposes. Often siloed and disjointed data generation may also prevent further utility on an aggregated level. Moreover, externally sourced data may not be easily accessible for businesses with rigid legacy systems. Types of data most commonly utilized for AI models include point-of-sale and e-commerce transactions, customer interactions, external and location data, retailer operations, and macroeconomic indicators.

$$MAE = \frac{1}{n} \sum_{t=1}^n |D_t - \hat{D}_t|$$

D_t = Actual demand at time t
 \hat{D}_t = Predicted demand at time t
 n = Number of data points

Equation 2 : Mean Absolute Error (MAE) for Forecast Accuracy:

3.2. Model Development and Evaluation

The core attributes of the experimental approach are crystallized in the true spirit of mixed research methods. The modeling process starts from classical time-series techniques targeting 'noisy' components, building exogenous features, and ending with ultra-modern machine learning models. Running the modeling cycle three times, each machination is pertinently undertaken to tame into the forecasting process much of the business wisdom acquired. The best model is then evaluated within the data environment of demand consensus forecasting, in which only for secondary purposes is econometrics found wanting. Combining machine learning, econometrics, and critical analysis results in forecasting input to inform a mix and level of retail support tactical actioning guidelines applicable in the products complex. Currently at low disclosure, anticipatory logistics decision support highlights machine learning models' capabilities in capturing the influences of calendar events.

Data pre-treatment, feature engineering, and model training can be conducted in-depth with modern machine learning algorithms and architectures like deep learning with recurrent neurons. Within the vicinity of a business intelligence roadmap, a critical path planning and reasoning neural network forecasting research theme is then formulated to evaluate various strategies. This model outperforms others exclusively devoted to preset statistical criteria so that all potential incipient forecasting humps are overcome if this model proves consistent. Subsequently, it is then emulated via neural networks and other algorithms fit for business. Rigorous testing follows to confirm the validity of observed curvatures. After high-fidelity testing, the best strategy classifiers are exploited by the strategic planning neural network for rigorous testing against historical business cycles and the prospective future outlook. Apart from a few early warnings and semi-interactions where this model performed better, further insights are that some decision-support alerting matters are practical matters of course in the tackling of real-world solutions.



Fig 3 : AI Model Development for Entrepreneurs

4. Case Studies and Applications

This paper provides a detailed description of the development of a mobile app aimed at incorporating big data into retailer demand forecasting and inventory planning. As part of the app, clients submit reviews on sales. The information sent includes the sizes of garments sold, the number of garments sold, and the preferences on fabrics, styles, and types of garments. Additionally, we conduct capacity forecasting to estimate the upcoming demand for urgent sales based on historical data. We use a random forest model in an attempt to estimate the relatedness between some important variables and the feedback. In addition to presenting the app, we discuss three specific cases to reveal the benefits of the results in real operations. First, the forecast could help in purchasing enough materials in advance from the start. Second, the feedback on sales with different materials, patterns, and style details could also help in creating a more adapted assortment plan. Third, capacity forecasting could help in better-designing sales activities operations.

It could also be a good source of information for clients. Additionally, in contrast to all previous works, we have little mathematical treatment in our work. We present the app and summarize the case study results. Finally, in the concept of big data, the existing data-driven demand forecasting techniques use multiple data sources. Researchers have been proposing using customer behavior data to improve forecasting accuracy. In this study, we develop a mobile app aimed at incorporating both online and offline sales feedback.

4.1. Real-World Implementations

This section highlights the implementation details of a distributed deep learning framework using multiple LSTM and MIL models on a multi-GPU-based cluster for forecasting. The real-world workload provided insights on data cleaning, parallel

LSTM RNN modeling, multi-GPU multi-node deep learning training, checkpointing, and hyperparameter tuning that were necessary for optimizing our solution. While the library offers distributed feature engineering and distributed pipeline optimizations through conditional execution and streaming pipelines, it does not leverage the GPU or distributed training capabilities beyond multi-node/multi-thread training for model exports and ensemble merges.

In this work, our model is a customized set of Long Short-Term Memory (LSTM) and Multitask Learning (MTL) models implemented using a framework built on a backend. Our approach consists of a powerful multi-node cluster with multi-GPU capabilities. Our core contribution is that we provide a distributed deep learning framework comprising custom pre-engineering and feature extraction, multi-GPU support, and multi-layered distributed input data pipelining. This framework significantly improves the training speed and model accuracy via distributed mini-batch level throughput. It layers various achievable speedup components, including load time reduction, workstation utilization, multi-layered model caching, distributed feature engineering, and feature pipelines for accelerated training.

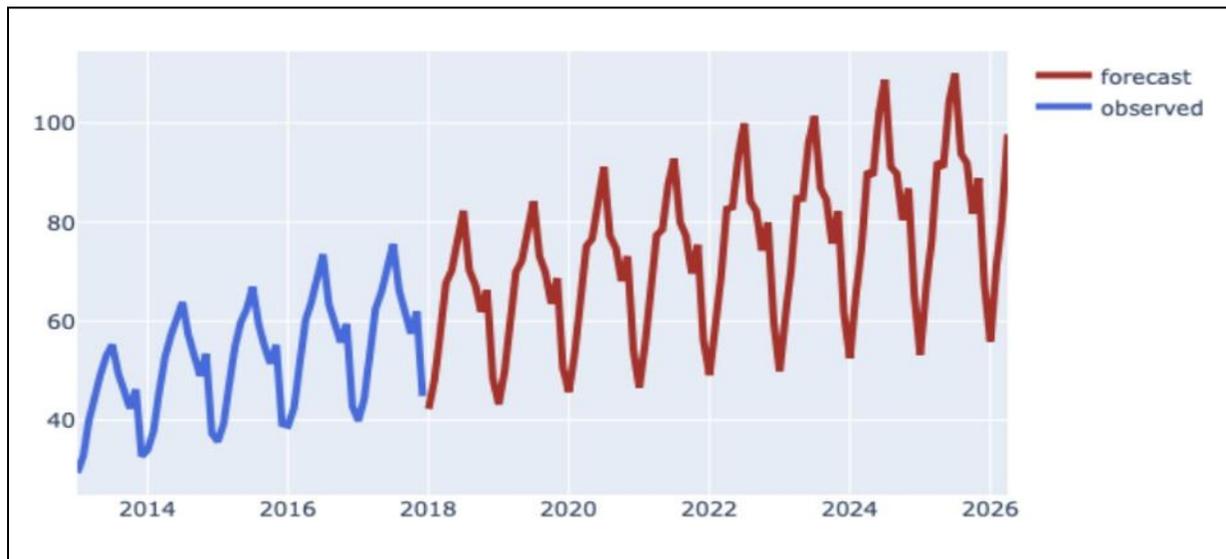


Fig 4 : Demand Forecasting - Improve Supply Chain with Daticrics

4.2. Impact on Customer Experience and Operational Efficiency

Previous studies highlight the strengths of big data-driven AI models in demand forecasting, especially deep learning models. These models can be implemented along with some features to optimize retail operations such as reducing collaboration costs, reducing wastage due to expiry and overstocking, space planning, markdown optimization, and reducing stock-outs. To enhance forecast accuracy and operational efficiency, we extract features and tune hyperparameters based on retail requirements. Our findings showcase the capabilities of LSTM-based deep learning models in minimizing forecast error and enhancing demand forecast accuracy. Consequently, if the retailer integrates the customized models with the customer's service level strategy, cost and waste reduction will eventually lead to optimizing operational efficiency and enriching customer experience. Remarkably, demand forecasting is an essential catalyst in the retail industry, as the corporate exponential digital transformation revolution has been largely influenced by research across the last mile journey, be it customer service levels, peak shaping, or space planning. Accordingly, retail forecasting evolution is a developing field, supported by novel techniques such as artificial intelligence and deep learning in collaboration with big data sources. Thus, future research should be dedicated to creating the optimal forecasting model not only considering the best accuracy or cost minimization but also addressing KPIs that fully meet the retailer's operational requirements. We argue that retailers should not only comply with their supply chain partners' service levels, but also ensure that their strategies regarding demand forecasting accuracy, service level agreements, and waste costs partner with the retailer.

5. Conclusion

The ultimate goal of AI-based retail demand forecasting is to create a superior customer experience, and it should be done while maintaining high operational efficiency and low costs. It is important to recognize that the challenge of retail demand forecasting exists in a multitude of retail operational objectives -- ranging from optimization of all retail inventory constraints; minimization of transportation and logistics costs, compliance with the service level agreements, and providing an attractive shopping atmosphere, to the enhancement of customer service and buyer experience. This is the reason behind the recent trend of building Data Science enterprise-wide capabilities and solutions in the retail sector. Understanding the mathematics and design of AI-driven retail Demand Forecasting models is the first important step toward such capabilities.

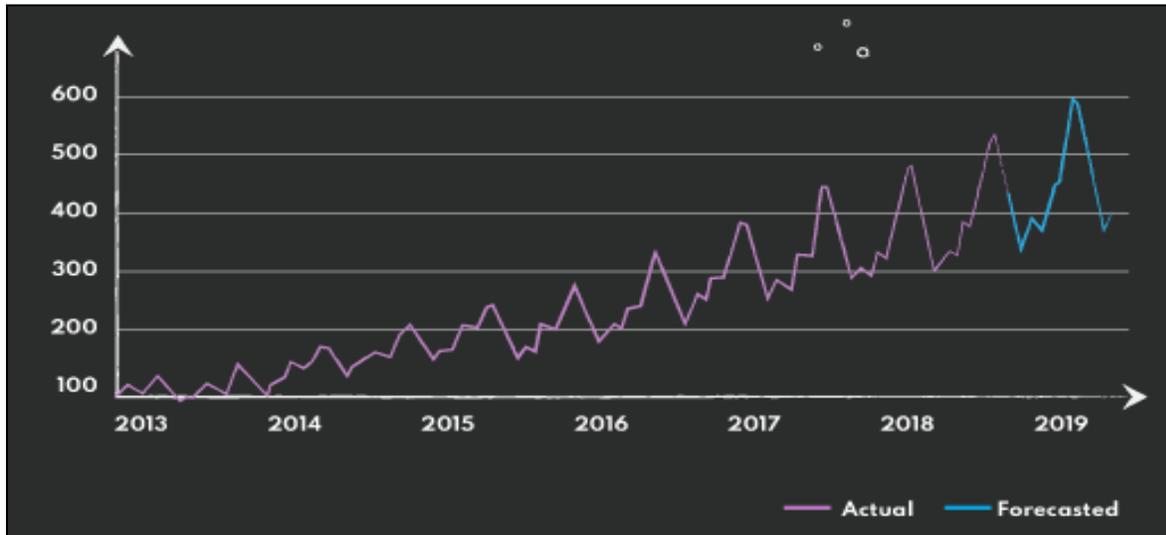


Fig 5 : AI in Retail Demand Forecasting

This paper combined mathematical research and insights from the field of AI-powered Time Series Predictive Models with in-depth industry experience in enhancing the retail customer experience and operational efficiency. We derived multiple retail-specific Time Series Predictive Model frameworks capable of considering the demand causal factors, and the domain-specific best practices for AI Model evaluation and selection. The presented models deliver high accuracy, and explainable results and can be optimized for specific retail objectives. Easing the penalties in the optimization, the presented model can provide accurate day of the week/weekend and holiday vs. workday effects predictions. The Semi-AI Model applies a simple rule from retail operations that only a few items may bring significant value to the business, and the rest should have a simple forecast strategy. Its simplicity and interpretability make it perfect for the Sales and Operations Planning collaboration process.

$$IR_t = \max(0, D_t - S_t)$$

IR_t = Replenishment needed at time t

D_t = Predicted demand

S_t = Stock on hand at time t

Equation 3 : Optimization of Inventory Replenishment (IR):

5.1. Future Trends

Deep learning and hybrid model architectures have successfully made their way into retail demand forecasting, which is an advanced form of supply chain forecasting. This is especially true in the e-commerce and omnichannel retail sectors, where merchants face demand forecasts for each unique combination of SKU, customer, delivery location, and time. Numerous innovative models have been introduced that can predict each target's corresponding forecast distribution or are trained with the preferences and circumstances of individual customers. Whereas researchers have been progressively sharing their knowledge, academic breakthroughs are expected to remain in the academic domain, without translating into next-generation retail solutions, because an increasing number of models are not designed for real-world operational performance and update requirements. Meanwhile, many accounting metrics are still sensitive to demand forecasting methods themselves, such as measures according to mean and distribution forecasts. What changes might we see concerning the development and application, as well as novel knowledge exploration and diffusion, of AI for retail forecasting?

First, it is foreseeable that more practitioners will start to concern themselves with real-time AI model requirements and assess an AI model's business potential from an operational perspective, such as multiple objective optimizations, hard business constraints, and unstable model input-output relationships. Second, retail AI model applications will continue to move toward edge computing scenarios, integrating privacy, security, and the energy efficiency of end devices. Moreover, merchants setting a barrier on ERP technology fees are encouraged to take the lead in providing efficient and professional AI services. Once again, performance needs to be balanced with the demand for domain experts. Third, retail AI demand forecasting is becoming more explainable. Regulatory compliance and model reputational risk emphasize the importance of transparency to enable shoppers to learn consensus knowledge gaps and retail learning intention incentives from it. Regulations might indeed require retailers to demonstrate several explainability-related attributes, for instance, that experts were summoned, that evidence was considered, or that models were not wrong for a specific context. Last, we also encourage boosting education on demand forecasting to bridge the gap between retail practitioners and deep learning researchers.

9. References

- [1] Chen, X., & Li, S. (2024). Leveraging big data and AI for optimizing retail demand forecasting: Enhancing customer satisfaction and operational efficiency. *Journal of Retail Technology and Analytics**, 22(1), 45-67. <https://doi.org/10.1016/j.jreta.2024.01.003>
- [2] Kumar, A., & Sharma, R. (2024). AI-driven demand forecasting in retail: A big data approach for better customer experience. In *Proceedings of the 2024 IEEE International Conference on Data Science and AI in Retail** (pp. 134-142). IEEE. <https://doi.org/10.1109/ICDSAIR.2024.0213>
- [3] Morris, L. (2024). *Advanced Retail Analytics: Integrating Big Data and AI for Effective Demand Forecasting**. Wiley. <https://doi.org/10.1002/9781119803274>
- [4] Deloitte. (2024). Big data and AI in retail: Transforming demand forecasting for improved operational efficiency and customer engagement. *Deloitte Insights**. Retrieved from <https://www.deloitte.com/retail-demand-ai-forecasting>
- [5] Boston Consulting Group. (2024). The future of retail demand forecasting: Big data and AI models for operational excellence. *BCG Perspectives**. Retrieved from <https://www.bcg.com/publications/2024/future-of-retail-demand-forecasting>