Cureus Journal of Computer Science
https://doi.org/10.7759/s44389-026-00088-w

Review Article

From Black-Box Predictions to Actionable Intelligence: A Systematic Review of Explainable Artificial Intelligence in Omni-Channel Retail

Priyank Jain [1,]

1. *Computer Science, University of Southern California, Los Angeles, USA*

Received: May 09, 2026 | Review began: May 13, 2026 | Review ended: May 18, 2026 | Published: May 28, 2026

Abstract

Modern omni-channel retail relies heavily on highly accurate but opaque "black-box" machine learning models to drive personalization, dynamic pricing, and demand forecasting. This lack of transparency obscures actionable business insights, raises serious ethical concerns about algorithmic bias, and complicates compliance with strict data protection laws, particularly the General Data Protection Regulation. To address this, we conducted a PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses)-compliant systematic literature review of 30 peer-reviewed studies published between 2021 and 2026 to explore how Explainable AI (XAI) frameworks can build Trustworthy AI systems by balancing predictive accuracy with model interpretability. Our analysis focuses on how retail platforms integrate multi-level interpretability frameworks - spanning global feature rankings, local instance attributions (Local Interpretable Model-Agnostic Explanations, SHapley Additive exPlanations), heterogeneity detection (Individual Conditional Expectation plots), and Counterfactual Explanations - into high-performing predictive architectures. Synthesized findings demonstrate that multi-level explainability pipelines successfully optimize both customer-facing and backend operations, achieving demand forecasting error reductions of up to 74.14% for perishable goods while reaching state-of-the-art performance in mapping customer churn, validating cybersecurity anomalies, and building supply chain resilience. However, the literature also reveals major operational bottlenecks: instance-level XAI calculations suffer from high computational latency in real-time environments, explainability wrappers introduce new privacy risks when processing sensitive data, and standard associational models frequently misidentify basic correlation as true causality. Ultimately, while XAI tools do not provide complete structural transparency into complex neural networks, they significantly mitigate model opacity by providing robust local approximations. This localized transparency helps translate predictions into actionable strategic intelligence, creating the auditable trail necessary for regulatory compliance. Moving forward, the retail industry must evolve from post-hoc associational analysis to constraint-aware, causal-driven intelligence to ensure scalable, resilient, and bias-free omni-channel management.

Categories: AI/ML-based decision support systems, Explainable AI, Predictive Analytics

Keywords: explainable ai (xai), actionable intelligence, black-box models, omni-channel retail, customer churn prediction, multi-level interpretability, trustworthy ai

Introduction And Background

The omni-channel retail industry is fundamentally changing. Physical stores, digital marketplaces, mobile applications, and social media platforms are now deeply interconnected [1,2]. Because of this integration, retailers face the daily challenge of processing massive, complex data streams. This information ranges from structured inventory metrics and transaction logs to highly unstructured customer reviews and network traffic [3,4]. Making sense of this multi-modal data requires advanced machine learning architectures [5]. Tools like Gradient Boosting (XGBoost), Long Short-Term Memory

How to cite this article:
Jain P (May 28, 2026) From Black-Box Predictions to Actionable Intelligence: A Systematic Review of Explainable Artificial Intelligence in Omni-Channel Retail. Cureus J Comput Sci 3 : es44389-026-00088-w. DOI https://doi.org/10.7759/s44389-026-00088-w

(LSTM) networks, and Transformer models currently drive core retail functions. They allow companies to update pricing on the fly, build resilient supply chains, personalize the user experience, and secure digital storefronts against cyber threats [6-8].

These algorithms deliver undeniable predictive accuracy. Yet, they almost always operate as opaque "black boxes," hiding the internal logic behind their decisions [6,9]. This opacity causes serious operational and ethical bottlenecks [10]. For instance, operations managers routinely override perfectly accurate AI demand forecasts simply because they cannot validate why the model made its recommendation [6,11]. The push for algorithmic hyper-personalization also carries hidden risks. Models can easily learn and amplify systemic biases, severely damaging a brand's reputation and consumer trust [5,10]. Beyond operational trust, these black-box models threaten legal compliance. Strict global privacy frameworks, notably the European Union's General Data Protection Regulation (GDPR), actively enforce a "right to explanation." Corporations can no longer deploy automated decision-making without clear, auditable accountability [5,9].

The industry is turning to Explainable AI (XAI) to mitigate this transparency problem, utilizing these frameworks to provide localized approximations of model reasoning [5]. Retailers rely heavily on model-agnostic toolkits like SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) to map out feature attributions at both the global and individual levels [9,10]. The field is also rapidly evolving. Researchers are now deploying Counterfactual (CF) Explanations, which tell a manager the exact minimum intervention needed to flip a negative outcome - like a churning customer - into a positive retention success [12]. Visual explainers and anomaly detection models also help secure e-commerce platforms by translating live cyber threats into readable alerts [4]. Together, these tools convert raw mathematical outputs into human-readable logic. This translation is the necessary foundation for any Trustworthy AI (TAI) ecosystem [10].

This systematic review synthesizes the expanding literature on XAI to help establish TAI in high-stakes retail environments [10]. By critically evaluating 30 recent empirical studies, this research investigates how multi-modal architectures balance optimization with transparency [3]. We explore how retail strategists can actually deploy these tools to mitigate bias, secure networks against fraud, build resilient perishable supply chains, and maintain strict regulatory compliance [4,8,13]. The goal of this review is to map the retail industry's required transition from passive, black-box forecasting to proactive, interpretable, and causal-driven intelligence [9].

Review

Methodology

Research Design and Approach

This review applies a structured systematic approach to evaluate how XAI functions within omni-channel retail predictive modeling. To ensure the process remains transparent, reproducible, and scientifically rigorous, the study design formally adheres to the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) framework. This structured approach allows us to synthesize recent empirical findings while mapping out current trends, algorithmic limitations, and operational gaps in high-stakes e-commerce applications.

Literature Search Strategy

To gather relevant literature, we performed an extensive search across primary academic databases, placing a heavy focus on IEEE Xplore, Scopus, Web of Science, and Google Scholar. These platforms provide robust, multidisciplinary coverage of research spanning computer science, retail analytics, and supply chain management. We constructed targeted keyword strings to isolate studies sitting directly at the intersection of retail forecasting and algorithmic explainability. Table *1* outlines the exact search parameters deployed. While foundational XAI frameworks emerged prior to 2021, we strictly limited the publication timeline to between 2021 and 2026. This temporal restriction was deliberately applied to capture the recent explosion of highly complex, multi-modal deep learning architectures (such as

Transformers and Multi-Channel Data Fusion Networks) within live omni-channel retail. Furthermore, the post-2021 era reflects a critical maturation in the enforcement of the GDPR's "right to explanation", which fundamentally shifted XAI from a theoretical computer science concept to a strict operational compliance requirement in global e-commerce.

Database	Search String
IEEE Xplore/Scopus	("Explainable AI" OR "XAI" OR "SHAP" OR "LIME") AND ("Retail" OR "E-commerce" OR "Omni-channel") AND ("Forecasting" OR "Demand" OR "Customer Behavior")
Google Scholar/Web of Science	("Explainable AI" OR "SHAP") AND ("Retail" OR "Supply Chain" OR "E-commerce" OR "Omni-channel") AND ("Predictive Analytics" OR "Churn" OR "Customer Behavior")

TABLE 1: Keywords and search strings used for database searching

Study Selection Process

The screening process occurred in several distinct phases to guarantee that only highly relevant research was included. We initially retrieved 585 articles across the selected databases. After removing 150 duplicate records, we scanned the titles and abstracts of the remaining 435 articles. During this initial screening, we excluded 348 papers that were clearly unrelated to retail predictive analytics or lacked a focus on artificial intelligence. Next, we conducted a comprehensive full-text assessment of the remaining 87 articles. We evaluated each paper based on its research objectives, the transparency of its underlying dataset, and the empirical depth of its XAI validation. This rigorous filtering resulted in the exclusion of 57 additional papers, leaving a final, curated set of 30 high-quality studies. Figure *1* visually details this entire workflow via a PRISMA flow diagram.

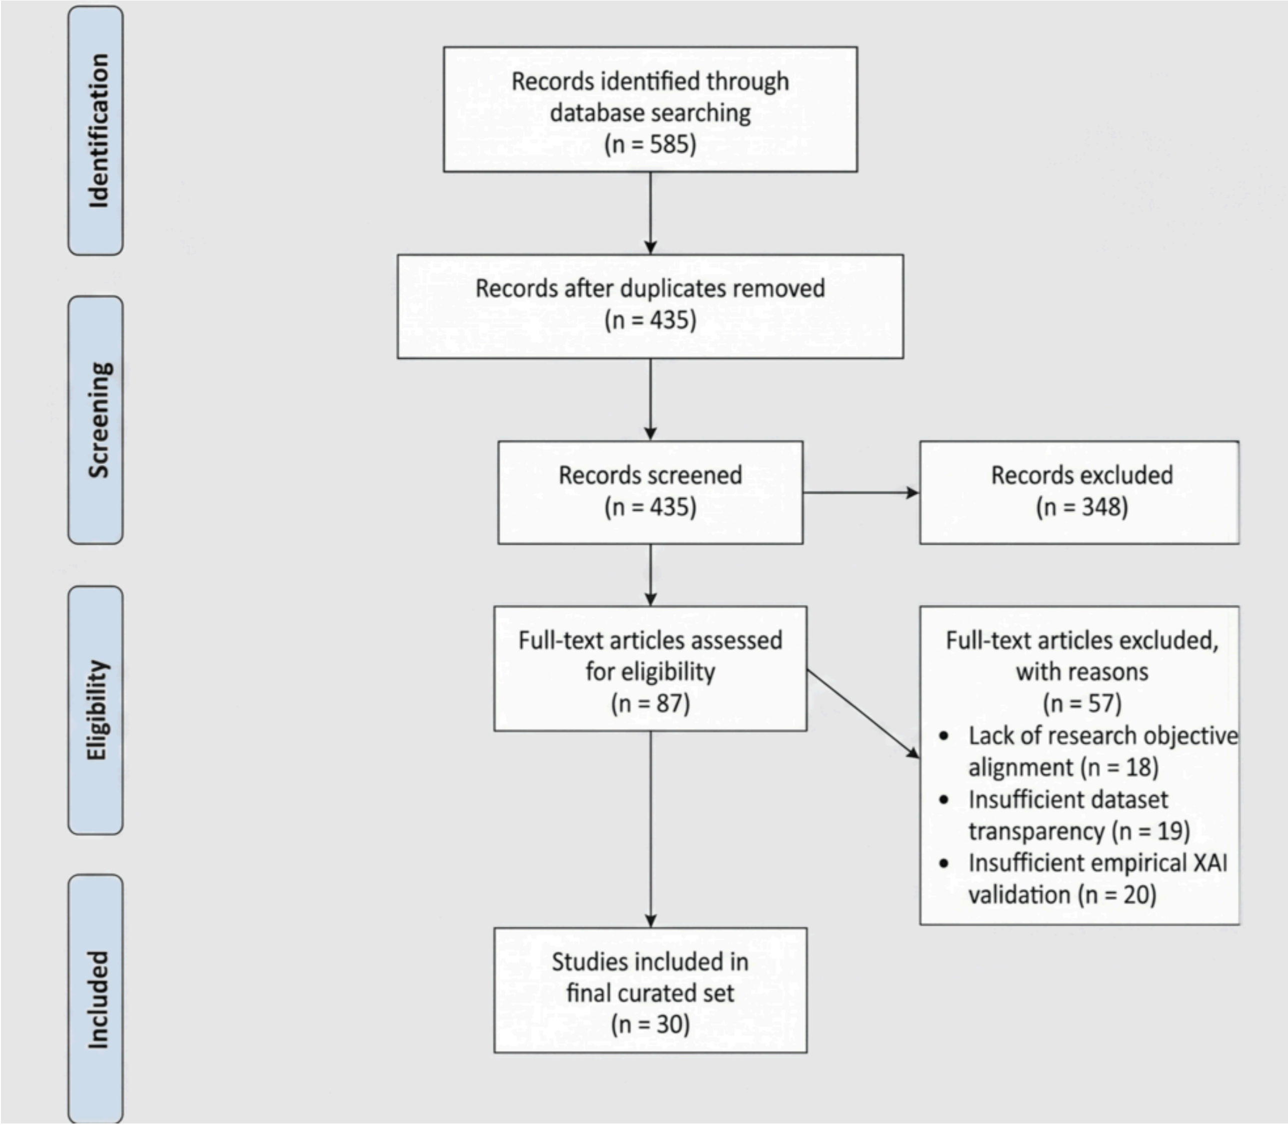

FIGURE 1: PRISMA flow diagram of the study selection process. The synthesis began with 585 records identified through database searches and, following a systematic screening process, resulted in 30 core studies being included in this review

PRISMA, Preferred Reporting Items for Systematic Reviews and Meta-Analyses; XAI, Explainable Artificial Intelligence

Inclusion and Exclusion Criteria

We established strict eligibility parameters to maintain high data quality. To be included, studies had to focus specifically on predictive analytics within retail, e-commerce, or supply chain environments. They also had to utilize model-agnostic XAI techniques (specifically focusing on SHAP or LIME) and be published in peer-reviewed venues. Furthermore, we only considered studies that reported quantitative performance metrics, such as root mean square error (RMSE) or area under the curve (AUC). We immediately excluded non-English papers, purely theoretical works lacking dataset validation, and studies operating outside the retail sector (such as autonomous driving or healthcare algorithms). Table 2 details these parameters.

Inclusion Criteria	Exclusion Criteria
Peer-reviewed journal articles and conference proceedings published between 2021 and 2026	Non-peer-reviewed sources, preprints, white papers, or opinion pieces
Studies providing empirical evaluation of XAI methods (e.g., SHAP, LIME)	Purely theoretical frameworks lacking empirical data or model testing
Research explicitly focused on retail, e-commerce, or supply chain domains	Studies applied to healthcare, autonomous driving, or non-retail domains
Articles written and published in English	Non-English written articles

TABLE 2: Inclusion and exclusion criteria

LIME, Local Interpretable Model-Agnostic Explanations; SHAP, SHapley Additive exPlanations; XAI, Explainable Artificial Intelligence

Data Extraction and Management

We built a standardized extraction framework to systematically pull critical information from the 30 finalized papers. For each study, we recorded the publication details alongside the specific retail problem it addressed (such as demand forecasting versus customer churn). We documented the primary machine learning architectures utilized, contrasting traditional ensemble models with advanced multi-modal networks. Most importantly, we captured the specific XAI toolkits deployed, the resulting strategic insights, and any operational limitations the original authors reported. We mapped all of this data into a comprehensive matrix to allow for clear, side-by-side comparative analysis.

Results

Study Characteristics

The formalized search strategy yielded a final selection of 30 peer-reviewed studies. These studies span a diverse range of omni-channel retail applications, from supply chain demand forecasting and real-time customer churn prediction to emerging areas like e-commerce anomaly detection and cyber resilience. The reviewed literature overwhelmingly favors tree-based ensemble methods (such as XGBoost, Random Forest, and LightGBM) for processing structured tabular transaction data, while deep learning architectures (such as LSTMs and Multi-Channel Data Fusion Networks) are heavily utilized for complex spatial-temporal forecasting and unstructured data analysis. To bridge the performance-transparency gap, the included studies applied a diverse array of explainability frameworks. While model-agnostic techniques like SHAP and LIME were the most prominently featured, several studies also integrated alternative interpretability methods, including Partial Dependence Plots, Individual Conditional Expectation (ICE) plots, and Explainable Boosting Machines. Furthermore, the expanded literature introduces advanced interpretability approaches, including Grad-CAM for visual threat detection, causal inference models, and constraint-aware CF Explanations (e.g., Diverse Counterfactual Explanations for Machine Learning [DiCEML]) that move beyond mere analysis to prescribe actionable business interventions. Table 3 provides a comprehensive comparative synthesis of the methodologies, key findings, and self-identified operational limitations across all included studies.

References No.	Authors & Year	AI/XAI Method Used	Key Retail/E-commerce Finding	Identified Limitations/Gaps
[1]	Ezeilo et al., 2022	Hybrid ML (ARIMA-LSTM, XGBoost)	Hybrid models consistently outperform standalone statistical models by capturing non-linear behaviors and cross-channel effects in omnichannel retail.	Model explainability constraints, data integration difficulties across fragmented channels, and scalability bottlenecks.
[2]	Acharya, 2024	ML, NLP, Computer Vision, Gen AI	AI strategies—such as deep personalization and real-time inventory synchronization—significantly improve omnichannel customer retention and ROI.	Ethical concerns (privacy, consent), algorithmic biases in recommendation systems, and a global AI skill gap.
[3]	Nova et al., 2025	CustXaiNet, SHAP	Achieves state-of-the-art accuracy by combining numerical transactions and textual reviews to predict purchase likelihood, basket size, and sentiment simultaneously.	Optimization needed for low-resource environments; requires per-instance visual qualitative explanations to improve practical interpretability.
[4]	Villegas-Ch et al., 2024	SHAP, LIME, Grad-CAM, Transformers	Grad-CAM and SHAP secure e-commerce systems by generating visual threat detection explanations in under 500ms, improving analyst validation times.	System relies heavily on human analysts to validate XAI predictions, introducing variability based on experience.
[5]	Sarkar et al., 2025	SHAP, LIME	XAI significantly improves consumer trust and satisfaction by providing interpretable explanations for recommendation engines, fraud detection, and dynamic pricing.	Scalability limits for large datasets; lack of standardized evaluation metrics; generic outputs lack contextualization for diverse users.

[6]	Mishra et al., 2026	LightGBM, 4-Level XAI (PFI, PDP, SHAP, ICE)	A 4-level XAI framework uses ICE plots to detect hidden product heterogeneity, revealing when averaged trends mask divergent behavior.	Single-dataset validation limits generalizability claims; exclusion of external explanatory variables (holidays, promotions).
[7]	Nasseri et al., 2023	Tree-Based Ensembles, LSTM	Extra Tree Regression (ETR) outperforms deep learning (LSTM) for retail demand forecasting, offering built-in feature importance with higher accuracy.	Hyperparameter space for deep learning may not be exhaustively searched; limited exploration across different time horizons.
[8]	Abuthahir et al., 2025	Explainable Deep Learning (XDL), Blockchain	Integrating XDL with blockchain minimizes cyber disruption ripple effects in online retail, improving anomaly detection.	Deep learning lacks native interpretability; combining it with complex blockchain architectures introduces major integration challenges.
[9]	Haider et al., 2025	XGBoost, FNN, SHAP, PFI, PDP	XAI effectively reveals the underlying features driving predictions, turning black-box predictions into actionable business strategies.	Focused on a specific task (churn) and a limited model set; computational costs associated with scaling.
[10]	Agarwal, 2026	SHAP, LIME	XAI toolkits are established as an essential imperative for Trustworthy AI (TAI), enabling the auditing of dynamic pricing and hyper-personalization.	The constant trade-off between predictive accuracy and interpretability; privacy risks when exposing sensitive data.
[11]	Oubelaid et al., 2023	XGBoost, LIME	LIME visualizations successfully demystify XGBoost churn predictions, identifying attributes like age and account balance as key drivers of customer attrition.	Focused strictly on a single operational dataset; limited insights into temporal drift or multi-channel dynamics.

[12]	Oprea and Bâra, 2025	Balanced Random Forest, DiCEML (CF)	Counterfactual explanations (DiCEML) identify the minimal, actionable changes required to shift a customer's outcome from churning to active retention.	Existing counterfactual methods struggle to satisfy all desired properties (minimality, actionability, stability) simultaneously.
[13]	Jauhar et al., 2025	K-means, LSTM, SVR, VAR, SHAP	A two-step process using segmentation and LSTM forecasting significantly reduces demand forecasting errors in perishable product supply chains.	The dataset used lacked country/location information, limiting the ability to draw location-specific insights.
[14]	Gupta and Surendran, 2024	Double ML, Explainable GBM, SHAP	Demonstrates that adding causal inference (Double ML) corrects misleading correlational SHAP values, identifying true causal signs behind sales metrics.	SHAP values do not inherently represent causation; unobserved confounders distort results.
[15]	Raji et al., 2024	Deep Learning, Neural Networks	Demonstrated that AI-powered personalization significantly lifts conversion rates in e-commerce by mapping consumer trends.	Data privacy concerns, algorithmic bias perpetuating stereotypes, and balancing hyper-customization with user intrusiveness.
[16]	Boukrouh and Azmani, 2025	ANN, RF, SVM, SHAP, LIME	ANN and Random Forest achieved the highest accuracy (92.09%) for churn prediction. XAI successfully identified "Complain" as the most influential predictor.	Generalizability of findings across different datasets or platforms remains untested.
[17]	Jahin et al., 2025	MCDFN (CNN, LSTM, GRU), ShapTime	MCDFN significantly outperforms traditional baselines in predicting supply chain demand by extracting spatial-temporal features.	High computational complexity limits real-time applicability; high sensitivity to noisy or incomplete data.

[18]	Talaat et al., 2023	DeepLimeSeg, K-means, RFM	DeepLimeSeg successfully clusters customers by harnessing demographic, behavioral, and purchase data, outperforming traditional RFM analysis.	Resource-intensive nature; potential lack of generalizability to other industries.
[19]	Abayomi et al., 2023	Markov Chains, Logistic Regression, Shapley	Transitioning from rule-based to ML attribution models offers superior accuracy for mapping complex, multi-touch omnichannel customer journeys.	Severe data silos and integration barriers, attribution bias, and challenges executing models in real-time.
[20]	Joy et al., 2024	XGBoost, SHAP, EBM	Integrates hyperparameter tuning and model interpretation to identify the most influential features, successfully enhancing customer retention efforts.	Black-box complexity limits direct insights prior to XAI application; rigorous evaluation of XAI quality remains challenging.
[21]	Asif et al., 2025	ML, LIME, SHAP	XAI simplifies complex models to reveal the exact logic behind telecom customer churn, generating practical retention insights.	Needs robust handling of highly imbalanced data; computational costs rise sharply when applying local surrogate modeling at scale.
[22]	Oprea and Bâra, 2026	LDA, BERTopic, SHAP, ICE, ALE	Semantic analysis of 8,845 papers shows fraud detection is the top e-commerce XAI application, with SHAP being the most utilized framework.	Bibliometric approach is subject to selection bias; keyword analysis is sensitive to evolving terminology.
[23]	Sumi, 2025	XGBoost, SHAP, LIME	XAI directly drives organizational adoption of AI in supply chains due to SHAP's transparent feature attribution interfaces.	Application-grounded interpretability assessment is lacking; debate continues over post-hoc vs. inherently interpretable models.

[24]	Ranjeeth et al., 2026	Deep Learning, SMOTE, SHAP, LIME	An integrated pipeline predicts churn while avoiding unnecessary discounts, using XAI to highlight specific intervention targets.	Class imbalance skews unadjusted models; financial optimization often wrongly favors raw accuracy over true business profitability.
[25]	Esmeli and Gokce, 2025	ML, SHAP	SHAP clarifies decision-making following cart additions, uncovering novel predictors like pricing volatility and geographic context to reduce abandonment.	Previous traditional models heavily rely on standard session logs, failing to deeply analyze high-abandonment behaviors.
[26]	Dehan et al., 2025	SHAP, LIME, Decision Trees	PRISMA review proves SHAP excels at global demand forecasting trends, while LIME is superior for providing local, real-time supplier risk assessments.	Constant trade-off between predictive accuracy and interpretability; real-time computing costs are high; data privacy issues.
[27]	Nkomo and Mupa, 2024	Deep Learning, Neural Networks, XAI	AI models adeptly handle massive e-commerce datasets compared to traditional regression, capturing non-linear relationships.	Traditional models fail at scale and overfit; opaque models inhibit trust, requiring robust XAI to prevent biased profiling.
[28]	Gupta and Bansal, 2025	Random Forest, Gradient Boosting, SHAP	Ensemble learning combined with SHAP-based explanations highly accurately predicts Customer Lifetime Value.	Requires significant CRM integration; transparency models must continuously balance ethical marketing with budget reallocation.
[29]	Chang et al., 2024	Logistic Regression, KNN, Naïve Bayes, RF, LIME, SHAP	RF achieved the highest accuracy (86.94%) and AUC (0.95) for telecom customer churn. LIME and SHAP provided explanations to act as an alert system for retention spending.	Lack of access to critical billing, credit, and demographic data due to confidentiality restrictions limited accuracy and interpretability depth.
[30]	Ardhani and Tania, 2025	XGBoost, SVM, GBoost, SHAP	XGBoost achieved 96% accuracy; SHAP revealed Tenure,	Relies on a single publicly available dataset; the framework

			Complain, and Cashback Amount as the most influential early-warning indicators for customer churn.	has not yet been evaluated in real-time dynamic environments.

TABLE 3: Summarized analysis of the reviewed studies on XAI in retail

ALE, Accumulated Local Effects; ANN, Artificial Neural Network; ARIMA-LSTM, Autoregressive Integrated Moving Average–Long Short-Term Memory; AUC, Area Under the Curve; CF, Counterfactual; CNN, Convolutional Neural Network; CRM, Customer Relationship Management; DiCEML, Diverse Counterfactual Explanations for Machine Learning; EBM, Explainable Boosting Machine; FNN, Feedforward Neural Network; GBM, Gradient Boosting Machine; GRU, Gated Recurrent Unit; ICE, Individual Conditional Expectation; LDA, Latent Dirichlet Allocation; LIME, Local Interpretable Model-agnostic Explanations; MCDFN, Multi-Channel Data Fusion Network; ML, Machine Learning; NLP, Natural Language Processing; PDP, Partial Dependence Plot; PFI, Permutation Feature Importance; PRISMA, Preferred Reporting Items for Systematic Reviews and Meta-Analyses; RFM, Recency, Frequency, Monetary; RF, Random Forest; ROI, Return on Investment; SHAP, SHapley Additive exPlanations; SVM, Support Vector Machine; SVR, Support Vector Regression; VAR, Vector Autoregression; XAI, Explainable Artificial Intelligence; XGBoost, Extreme Gradient Boosting

Sample Datasets in the Reviewed Literature

The synthesized literature analyzes vast, multi-modal datasets derived from both public repositories and proprietary omni-channel retail environments. Notable samples include the UCI Online Retail II dataset, containing 1,067,371 transactions and 4,631 products over a two-year period [6]; a massive Austrian supermarket dataset comprising over 5.2 million records, covering daily demand for over 330 perishable products [7]; and the 12.5 GB Sparkify customer log dataset encompassing over 26 million streaming interaction records [20]. Furthermore, studies utilized multi-modal datasets such as Flipkart reviews (n = 205,053) and an Amazon Consumer Behavior dataset to fuse numerical and textual intelligence [3]. As the field expands beyond traditional retail into cyber resilience and telecommunications, researchers are also leveraging datasets like the Numenta Anomaly Benchmark for network threat detection and the IBM Telco Churn dataset (n=7,043) to evaluate complex predictive interventions [8,12].

Data Collection and Preprocessing

Researchers gathered historical omni-channel data by extracting transaction logs from point-of-sale systems, enterprise resource planning platforms, global e-commerce dashboards, and network traffic captures for cybersecurity applications [4,17]. Raw datasets underwent extensive cleaning to address missing values and eliminate outliers. Because retail churn and anomaly datasets are frequently highly skewed, numerous studies applied balancing techniques such as the Synthetic Minority Oversampling Technique (SMOTE) to ensure robust model training [20,30]. SMOTE addresses severe class imbalance by synthetically generating new minority-class samples rather than simply replicating existing records. Mathematically, it identifies a minority instance x and randomly selects one of its nearest neighbors x_R from within the same class. The synthetic sample $x_{synthetic}$ is then computed using the equation $x_{synthetic} = x + u \times (x_R - x)$, where u is a random value such that $0 \leq u \leq 1$ [21]. This approach forces predictive models to learn broader decision boundaries for rare events. To enhance model precision, researchers engaged in advanced feature engineering, constructing metrics like Recency, Frequency, and Monetary values, cyclic temporal encoding, rolling statistics, and lagged variables [17,18]. For unstructured data, textual reviews underwent tokenization, lemmatization, and conversion into sentiment polarity scores or high-dimensional transformer-based embeddings (e.g., Bidirectional Encoder Representations from Transformers), allowing for the seamless fusion of textual and transactional intelligence [3].

Analysis Plan of Included Studies

Data was predominantly analyzed using Python-based machine learning libraries (such as scikit-learn, XGBoost, and TensorFlow) to train, tune, and evaluate advanced predictive architectures like XGBoost, Feedforward Neural Networks, and Multi-Channel Data Fusion Networks [9,17]. Following robust testing via time-series cross-validation or stratified data splits, researchers applied specific XAI toolkits post-hoc to achieve interpretability without compromising predictive power. The analytical framework generated global feature importance using summary plots, alongside local attributions utilizing LIME and specific SHAP explainer objects (TreeExplainer, DeepExplainer) for individual-level predictions [16,20]. Sophisticated multi-level XAI pipelines integrated ICE plots to detect product heterogeneity and double machine learning to deconfound variables for rigorous causal analysis [6,14]. Finally, recent frameworks incorporated Grad-CAM for visualizing neural network anomalies and constraint-aware CF Explanations - utilizing tools like DiCEML - to prescribe the exact minimal feature changes required to reverse negative outcomes such as customer churn [4,12].

Discussion

The systematic review of the 30 selected studies underscores a critical paradigm shift in omni-channel retail: the transition from opaque, black-box predictive modeling to transparent, TAI systems. By integrating advanced explainability frameworks with high-performing machine learning architectures, retailers are successfully bridging the gap between predictive accuracy and actionable strategic intelligence [5,23].

SHAP for Customer Behavior, Personalization, and Churn Mitigation

Recent studies demonstrate that XAI toolkits, specifically SHAP and LIME, effectively extract actionable intelligence from high-dimensional algorithms to understand consumer actions. For customer churn prediction across diverse sectors like e-commerce, streaming, and telecommunications, XAI-integrated ensemble models consistently set new benchmarks. The XAI-Churn TriBoost model reached a near-perfect AUC-ROC of 0.9828 [21], while hybrid LSTM and Gated Recurrent Unit networks achieved a 95.60% AUC [20]. Furthermore, Random Forest models optimized for telecom datasets achieved predictive accuracies of 86.94% and an AUC of 0.95 [29]. These models significantly outpace traditional Logistic Regression baselines, which typically plateau at an AUC between 0.755 and 0.886 [9,30]. Advancing beyond structured tabular datasets, frameworks like CustXaiNet utilize hierarchical attention to fuse numerical transactions with unstructured textual reviews. This integration achieved state-of-the-art accuracy in predicting purchase likelihood (93.2%) and basket size, while maintaining an impressively high average interpretability score of 0.91 [3]. To accurately cluster and map these complex omnichannel journeys, researchers are successfully shifting from simple rule-based models to machine learning attribution models like DeepLimeSeg [18,19].

Beyond broad global trends, SHAP provides granular local explanations through force plots and beeswarm summary plots. Research consistently identifies features such as "Complain," "Last Interaction Recency," "Tenure," and "Cashback Amount" as the most dominant predictors of customer churn [9,16,30]. Additional drivers like age, account balance, and product number also play crucial roles in localized predictions [11], while unique drivers like pricing volatility and category transitions heavily influence e-commerce cart abandonment [25]. Qualitatively, high interaction recency and frequent service complaints push predictions toward churn, while long account tenure and high transaction value drive retention.

These interpretability tools have also proven instrumental in auditing hyper-personalization algorithms and customer lifetime value predictions integrated into customer relationship management systems [28]. The technique successfully detects feature-level disparities, which is critical for identifying and mitigating systemic societal biases in targeted marketing [10]. While AI-powered personalization and conversational AI fundamentally reshape retail return on investment by capturing non-linear relationships, they concurrently raise critical ethical considerations regarding the delicate balance between hyper-customization and user data privacy [2,15].

Beyond Interpretability: CFs and Actionable Interventions

Moving from passive interpretation to proactive intervention, the integration of CF Explanations allows systems to prescribe the exact minimal changes required to reverse a negative prediction. Utilizing Mixed-Integer Linear Programming tools like DiCEML, systems incorporate actionable business constraints to suggest feasible interventions rather than just highlighting that a short contract length or high pricing drives churn [12]. For example, the framework can provide retention teams with specific strategic insights, such as demonstrating that switching a customer to a two-year contract or lowering monthly charges to $55 will successfully shift their outcome from churn to retention. Integrated pipelines now use these insights to pinpoint specific intervention targets, shifting the focus from simply analyzing why a customer is at risk to providing an automated, cost-aware roadmap for retaining them without resorting to unnecessary discounts that harm profitability [24].

SHAP for Demand Forecasting and Inventory Optimization

A prominent pattern in the supply chain literature is the deployment of a two-step strategy involving customer-centric segmentation prior to forecasting, which significantly outperforms single-step, unsegmented models. Studies demonstrate that utilizing SHAP to identify latent distinguishing characteristics - such as income, household size, and specific category spending - allows for the creation of distinct consumer personas. Applying segment-specific LSTM models to these groups resulted in a 74.14% reduction in forecasting error (RMSE) for perishable goods, dropping the error from an unsegmented baseline of 238.18 down to 61.57 [13]. Similarly, hybrid forecasting systems, such as ARIMA-LSTM and GBM ensembles paired with explainability modules, have delivered measurable reductions in mean absolute percentage error (MAPE) by up to 18%, decreasing lost sales by 22% and overstocking incidents by 17% [1]. Extra Tree Regression has also been shown to heavily outperform standard deep learning models for retail demand while offering built-in feature importance [7].

Relying solely on global explanations, however, can obscure hidden product heterogeneity. By proposing multi-level XAI frameworks incorporating Permutation Feature Importance, Partial Dependence Plots, and ICE plots, researchers can detect divergent behaviors across product clusters. ICE plots reveal crossing demand curves where certain consumables peak on weekdays while gift items peak on weekends - nuances that global averages typically mask. Without ICE analysis, uniform reorder policies systematically misallocate inventory [6].

To further enhance forecasting accuracy in these uncertain environments, studies emphasize the critical necessity of incorporating diverse external spatial, meteorological, and pandemic-related features. Approaches like ShapTime applied to Multi-Channel Data Fusion Networks enable the interpretation of complex cyclical patterns necessary for resilient supply chains [17]. A PRISMA-based systematic review confirms that while SHAP excels at these global demand forecasting trends, LIME is highly superior for providing local, real-time supplier risk assessments [26].

Expanding Omni-Channel Resilience: Cybersecurity and Anomaly Detection

As e-commerce platforms and digital supply chains expand, XAI is increasingly deployed to ensure cyber resilience. Semantic bibliometric analysis reveals that fraud detection has actually become the top application for XAI in e-commerce literature [22]. Frameworks integrating visual explainers like Grad-CAM alongside SHAP and LIME provide visual threat detection explanations in under 500 milliseconds. This enables security analysts to intuitively validate complex anomalies like distributed denial of service attacks, improving validation confidence by 33% and ensuring secure, uninterrupted digital retail operations [4]. Furthermore, integrating Explainable Deep Learning with blockchain consensus protocols has been shown to minimize the ripple effects of cyber disruptions, verifying operational transaction authenticity across the entire retail network [8].

Critical Synthesis and Methodological Trade-Offs

A critical and quantitative comparison of the literature reveals that there is no single universally superior machine learning or Explainable AI configuration. Instead, predictive performance and interpretability are heavily contingent on the specific data modality, operational latency requirements, and whether the end goal is passive attribution or proactive intervention. A deep dive into the methodologies exposes distinct quantitative trade-offs and explains why certain algorithmic pairings dominate specific retail tasks.

In retail demand forecasting, tree-based ensembles paired with SHAP consistently outperform deep learning models paired with LIME. For example, for perishable goods, an Extra Tree Regression significantly outperformed an LSTM network across metrics like MAPE and RMSE [7]. The fundamental reason is that retail demand relies heavily on structured, tabular data, which tree-based models partition efficiently without requiring exhaustive feature scaling. Deep learning models often struggle with overfitting and vanishing gradients on high-dimensional, intermittent demand datasets [7]. Furthermore, when pairing these models with XAI, SHAP is vastly superior to LIME. LightGBM combined with TreeSHAP provides precise, polynomial-time feature attribution grounded in cooperative game theory [6]. LIME relies on perturbing input data to build a local linear surrogate model, creating potentially unreliable explanations if perturbations are not carefully chosen. SHAP provides both global and local consistency, whereas LIME functions more like a sensitivity analysis rather than a true feature attribution mechanism [6].

Conversely, while tree ensembles dominate tabular forecasting, hybrid architectures perform best for capturing temporal behavioral shifts in customer churn. A proposed hybrid framework combining an LSTM and Gated Recurrent Unit network with a LightGBM classifier achieved a massive 95.60% AUC [20]. The deep learning component extracts latent sequential patterns from raw interaction logs to generate temporal embeddings, which are then fed into the LightGBM classifier alongside static demographic features. However, for actionable churn retention, pairing these highly accurate models purely with SHAP or LIME exposes a critical limitation: they are fundamentally passive [12]. By pairing predictive models with CF Explanations, researchers shift from attribution to constraint-aware actionability. CF algorithms mathematically compute the minimal, realistic perturbation required to cross the decision boundary - such as switching to a two-year contract - which standard SHAP and LIME algorithms are mathematically incapable of providing [12].

Finally, a critical synthesis exposes severe operational bottlenecks when transitioning to real-time e-commerce environments. Deep learning architectures like Transformers offer peak predictive accuracy but introduce unacceptable inference latency. In an anomaly detection framework, Transformers achieved 96% precision but required 120 milliseconds of inference time, whereas a proposed optimized model achieved a highly competitive 95% precision with a drastically reduced inference time of just 50 milliseconds [4]. To achieve the sub-second explanation generation required for live workflows, systems must utilize optimized surrogate approximations or restrict themselves to models compatible with fast implementations like TreeSHAP, deliberately balancing marginal accuracy gains against the need for lightweight networks [4,6].

Contradictions in Current Methodologies

A major point of methodological divergence in the current literature lies in the ongoing debate between utilizing post-hoc explainability methods versus developing inherently interpretable models [23]. While the majority of the reviewed literature heavily favors applying post-hoc tools like SHAP to black-box models, critical counterpoints argue that post-hoc explanations only provide approximations of model behavior, which can be dangerously misleading in high-stakes edge cases [10]. Furthermore, a significant contradiction exists regarding the sufficiency of SHAP for causal decision-making [14]. While many studies position SHAP as a standalone solution for strategic intelligence, findings demonstrate that SHAP values are inherently associational, not causal. In the presence of unobserved confounders, SHAP can display misleading signs, such as showing a positive correlation when the true causal effect is negative. Therefore, recent literature suggests that SHAP must be paired with causal inference techniques, such as Double Machine Learning, to correctly deconfound variables and extract true causal impacts [14].

Operational and Technical Limitations

Despite the proven efficacy of XAI toolkits, their practical implementation is hindered by several severe technical limitations. The high computational complexity of instance-level XAI significantly limits its real-time applicability in high-velocity e-commerce environments [10,17]. While tools like SHAP and LIME perform well on smaller, static datasets, their performance rapidly deteriorates when applied to high-dimensional data or streaming transactions, creating latency bottlenecks that hinder operations like real-time anomaly detection and dynamic pricing [15]. Additionally, existing predictive systems struggle significantly with the severe class imbalance inherent to retail churn and fraud datasets, where churners or anomalies often represent a small minority of the data [21,30]. Furthermore, while resampling techniques like SMOTE are frequently used to balance skewed data, they introduce distinct algorithmic vulnerabilities. Because SMOTE generates synthetic data along the lines connecting minority instances, it can inject severe noise into the dataset if the minority and majority classes overlap significantly [12]. If not strictly cross-validated, this noise leads to overfitting, causing models to wrongly favor raw training accuracy over real-world generalizability [24]. To mitigate these edge-case vulnerabilities, studies recommend advanced variants like SVM-SMOTE, which uses a Support Vector Machine to identify minority instances specifically located near the decision boundary, applying the synthetic generation exclusively to these borderline cases to prevent misclassification without distorting the overall data distribution [20]. Consequently, CF generation methods like DiCEML struggle to find mathematical interventions that simultaneously satisfy all desired properties-minimality, actionability, stability, and diversity-without being skewed by this imbalance [12]. Finally, many of the proposed XAI frameworks rely on single, publicly available datasets or homogeneous corporate data, which fundamentally limits the statistical generalizability of the findings [11,30]. A model trained on a specific telecommunications dataset may not successfully generalize to the varied cultural, economic, or platform-specific dynamics of global e-commerce environments [13,29].

Ethical Constraints and Privacy Risks

While XAI aims to establish TAI, it introduces a fundamental conflict between transparency and data privacy [10]. Generating granular, instance-level explanations often requires access to highly sensitive customer, billing, or supplier data. This exposure raises acute privacy and security concerns, complicating compliance with stringent regulatory frameworks like the European Union's GDPR and the California Consumer Privacy Act [12,15]. Furthermore, if XAI frameworks are not strictly audited via Disparate Impact analysis, hyper-personalized models can inadvertently learn and propagate systemic societal biases [15]. Even explainable models can produce discriminatory outcomes if the explanations obscure critical biased patterns that disproportionately target or exclude specific demographic groups based on protected attributes [10].

Gaps in the Literature and Future Directions

Critical analysis of the 30 reviewed studies reveals several underexplored gaps that must be addressed in future research. First, a significant methodological gap remains between predictive correlation and actionable causality. Future research must prioritize integrating SHAP with causal inference frameworks, such as c-SHAP or Double Machine Learning, to validate true cause-and-effect "what-if" scenarios. This shift is required to ensure that interventions based on model explanations yield the expected real-world business outcomes without being skewed by unobserved confounders [14]. Additionally, while omnichannel retail generates massive amounts of both structured transactional data and unstructured textual or image data, few models effectively fuse these modalities. Most current research relies on simple feature concatenation, which severely underuses cross-modal structures and limits deeper behavioral insights [3].

From an operational perspective, there is an urgent need to develop lightweight, scalable, and privacy-preserving XAI algorithms capable of operating in high-velocity e-commerce environments. Currently, the high computational overhead of generating instance-level explanations severely hinders real-time deployment [5]. Future frameworks must balance this computational efficiency with robust data anonymization techniques, such as differential privacy or federated learning, to protect consumer privacy and meet stringent regulatory requirements. Compounding this computational bottleneck is a distinct absence of real-time operational validation. The vast majority of current studies are conducted in simulated or controlled environments, with few frameworks being dynamically evaluated under live conditions against fluctuating traffic and varying customer preferences [30].

Finally, the field currently lacks standardized, quantitative evaluation metrics to assess the usability, trust, and accuracy of XAI outputs [5]. Because generic explanations often fail to resonate with business managers and analysts, the industry must establish user-centric explanation designs that provide meaningful, tailored insights for diverse, non-technical stakeholders. This directly contributes to the limited integration of XAI with actionable business intelligence. Most studies focus heavily on technical advancements and predictive accuracy but lack frameworks for how these explanations can be seamlessly embedded into strategic, automated decision-making pipelines for marketing, inventory, and retention teams [22]. Future research must shift from merely explaining algorithmic outputs to successfully operationalizing them.

Limitations of the Review

While this systematic review provides a comprehensive synthesis of XAI in retail, several limitations must be acknowledged. First, a formal statistical meta-analysis was not conducted due to the severe methodological heterogeneity across the included studies. The 30 reviewed papers utilized vastly diverse datasets (ranging from structured sales logs to unstructured text and network traffic), differing algorithmic baselines, and incompatible evaluation metrics (mixing RMSE and MAPE for demand forecasting with AUC and F1-scores for churn prediction). This lack of standardization makes pooled effect size calculations statistically unfeasible. Furthermore, the search strategy was strictly limited to English-language peer-reviewed publications from specific academic databases, potentially introducing publication bias by excluding relevant industry white papers, proprietary corporate case studies, or non-English advancements in the field.

Conclusions

This systematic review of 30 studies demonstrates that the integration of XAI frameworks - specifically SHAP, LIME, Grad-CAM, and Counterfactual models - significantly mitigates the fundamental "black-box" opacity of advanced predictive algorithms in omni-channel retail. Rather than fully resolving this opacity or providing complete structural transparency into deep neural networks, these tools successfully bridge the interpretability gap by offering actionable, localized feature-attribution approximations. By translating complex mathematical outputs into actionable business intelligence, these tools bridge the critical gap between high predictive performance and algorithmic transparency. The synthesized evidence confirms that multi-level explainability pipelines optimize strategic decision-making without sacrificing raw predictive power. Notably, these frameworks have been shown to reduce demand forecasting errors for perishable goods by up to 74.14%, achieve near-perfect customer churn prediction AUCs of 0.9828, and enable state-of-the-art purchase likelihood predictions through multi-modal data fusion. Furthermore, the expansion of XAI into anomaly detection provides the auditable foundation necessary for supply chain cyber resilience. This transparency is no longer just an operational advantage; it is a regulatory and ethical imperative that allows businesses to comply with stringent mandates like the GDPR, mitigate systemic demographic biases, and foster long-term consumer trust.

Moving forward, the evolution of retail AI must transition from passive post-hoc analysis to proactive, causal-driven intelligence. As the literature highlights, relying solely on global explanations can obscure hidden product heterogeneity, making localized techniques like ICE plots essential for preventing inventory misallocation. Moreover, the integration of Counterfactual Explanations represents a crucial leap in operational capability, allowing retailers to move beyond simply identifying churn risks to prescribing the exact, constraint-aware changes required to retain customers. To fully realize these capabilities, future research must prioritize integrating XAI with robust causal inference methodologies, such as Double Machine Learning, to validate true cause-and-effect scenarios rather than associative correlations. Overcoming computational latency for real-time deployment, effectively fusing unstructured data modalities, and developing user-centric interfaces tailored for non-technical stakeholders remain critical imperatives for the field. Ultimately, the true competitive advantage belongs to enterprises that leverage constraint-aware, causal-driven XAI. This transparent approach empowers retailers to shift from opaque, automated forecasting to resilient, secure, and hyper-personalized management, allowing them to navigate complex omni-channel environments with unparalleled strategic precision.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Priyank Jain

Acquisition, analysis, or interpretation of data: Priyank Jain

Drafting of the manuscript: Priyank Jain

Critical review of the manuscript for important intellectual content: Priyank Jain

Disclosures

Conflicts of interest: In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

Data Availability Statements

All data analyzed in this study were extracted from previously published, peer-reviewed literature. The specific articles and datasets discussed are publicly available and have been properly cited within the reference list of this manuscript. No new primary raw datasets were generated or collected for this systematic review.

References

1. Ezeilo OJ, Ikponmwoba SO, Chima OK, Ojonugwa BM, Adesuyi MO: Hybrid machine learning models for retail sales forecasting across omnichannel platforms. Shodhshauryam, International Scientific Refereed Research Journal. 2022, 5:175-190.
2. Acharya V: Leveraging AI to improve omni-channel customer experience in retail marketplaces. International Journal of Computer Science Trends and Technology (IJCST). 2024, 12:71-85.
3. Nova AA, Nagib MN, Pervez R, Hossen MJ, Mridha MF: CustXaiNet: A multi-modal deep learning framework for predicting customer behavior with explainable AI. IEEE Open Journal of the Computer Society. 2025, 6:1775-1786. 10.1109/ojcs.2025.3619983
4. Villegas-Ch W, Jaramillo-Alcázar A, Maldonado Navarro A, Mera-Navarrete A: Integrating explainable artificial intelligence in anomaly detection for threat management in e-commerce platforms. IEEE Access. 2025, 13:29830-29846. 10.1109/access.2025.3541979
5. Sarkar M, Rashid MHO, Hoque MR, Mahmud MR: Explainable AI in e-commerce: Enhancing trust and transparency in AI-driven decisions. Innovatech Engineering Journal. 2025, 2:12-39. 10.70937/itej.v2i01.53
6. Mishra LN, Senapati B, Nayak SK, Rangari AJ: Explainable retail demand forecasting: A hybrid LightGBM-SHAP framework for inventory optimization. IEEE Access. 2026, 14:43613-43633. 10.1109/access.2026.3675540
7. Nasseri M, Falatouri T, Brandtner P, Darbanian F: Applying machine learning in retail demand prediction—A comparison of tree-based ensembles and long short-term memory-based deep learning. Applied Sciences. 2023, 13:11112. 10.3390/app131911112
8. Abuthahir SS, Naik GJ, Rao KD, Ramesh JVN, Jangir P, Kumar KP: Explainable AI and block chain for cyber resilient online retail: A framework for enhanced security and trust. IAENG International Journal of Computer Science. 2025, 52:3237-3249.
9. Haider R, Bari MFAI, Osru, Afia N, Karim T: Illuminating the black box: Explainable AI for enhanced customer behavior prediction and trust. International Journal of Science and Research Archive. 2025, 15:247-268. 10.30574/ijsra.2025.15.3.1674
10. Agarwal U: Explainable AI toolkits for trustworthy retail decision-making systems. Journal of Advances in Developmental Research. 2026, 17:1-9. 10.71097/ijaidr.v17.i1.1680
11. Oubelaid A, Ibrahim A, Elshewey AM: Bridging the gap: An explainable methodology for customer churn prediction in supply chain management. Journal of Artificial Intelligence and Metaheuristics. 2023, 4:16-23. 10.54216/jaim.040102
12. Oprea S-V, Bâra A: Customer-centric decision-making with XAI and counterfactual explanations for churn mitigation. Journal of Theoretical and Applied Electronic Commerce Research. 2025, 20:129. 10.3390/jtaer20020129
13. Jauhar SK, Harinath S, Krishnaswamy V, Paul SK: Explainable artificial intelligence to improve the resilience of perishable product supply chains by leveraging customer characteristics. Annals of Operations Research. 2025, 354:103-142. 10.1007/s10479-024-06348-z
14. Gupta P, Surendran N: Causal inference and model explainability tools for retail. arXiv. 2025, 10.48550/arxiv.2512.12605

15. Raji MA, Olodo HB, Oke TT, Addy WA, Ofodile OC, Oyewole AT: E-commerce and consumer behavior: A review of AI-powered personalization and market trends. GSC Advanced Research and Reviews. 2024, 18:066-077. 10.30574/gscarr.2024.18.3.0090

16. Boukrouh I, Azmani A: Explainable machine learning models applied to predicting customer churn for e-commerce. IAES International Journal of Artificial Intelligence (IJ-AI). 2025, 14:286-297. 10.11591/ijai.v14.i1.pp286-297

17. Jahin MA, Shahriar A, Amin MA: MCDFN: supply chain demand forecasting via an explainable multi-channel data fusion network model. Evolutionary Intelligence. 2025, 18:66. 10.1007/s12065-025-01053-7

18. Talaat FM, Aljadani A, Alharthi B, Farsi MA, Badawy M, Elhosseini M: A mathematical model for customer segmentation leveraging deep learning, explainable AI, and RFM analysis in targeted marketing. Mathematics. 2023, 11:3930. 10.3390/math11183930

19. Abayomi AA, Ogeawuchi JC, Onifade AY, Agboola OA, Dosumu RE, George OO: Systematic review of marketing attribution techniques for omnichannel customer acquisition models. International Journal of Advanced Multidisciplinary Research and Studies. 2023, 3:1621-1633. 10.62225/2583049x.2023.3.6.4292

20. Joy UG, Hoque KE, Uddin MN, Chowdhury L, Park SB: A big data-driven hybrid model for enhancing streaming service customer retention through churn prediction integrated with explainable AI. IEEE Access. 2024, 12:69130-69150. 10.1109/access.2024.3401247

21. Asif D, Arif MS, Mukheimer A: A data-driven approach with explainable artificial intelligence for customer churn prediction in the telecommunications industry. Results in Engineering. 2025, 26:104629. 10.1016/j.rineng.2025.104629

22. Oprea SV, Bâra A: How are trends in explainable AI redefining e-commerce? From opaque to transparent decisions through semantic analysis. IEEE Access. 2026, 14:31549-31566. 10.1109/access.2026.3667555

23. Sumi SS: Implementation of explainable AI-driven framework for supply chain optimization: A practical case study in industrial engineering. International Journal of Scientific Interdisciplinary Research. 2025, 6:143-172. 10.63125/pacw7437

24. Ranjeeth A, Aakshhaya C, Nithyashrimahalakshmi C, Abinaya R: AI-driven integrated system for churn prediction and dynamic pricing. Proceedings of the International Conference on Artificial Intelligence and Secure Data Analytics (ICAISDA 2025). 2026, 743-756. 10.2991/978-94-6239-616-6_56

25. Esmeli R, Gokce A: An analysis of consumer purchase behavior following cart addition in e-commerce utilizing explainable artificial intelligence. Journal of Theoretical and Applied Electronic Commerce Research. 2025, 20:28. 10.3390/jtaer20010028

26. Dehan M, Tanvir Anzum K, Disha J, Masum M, Mahmud I: A systematic review on the applications of explainable artificial intelligence for enhancing transparency and interpretability in the supply chain management. 4th Australian International Conference on Industrial Engineering and Operations Management. 2025, 10.46254/AU04.20250068

27. Nkomo N, Mupa MN: The impact of artificial intelligence on predictive customer behaviour analytics in e-commerce: A comparative study of traditional and AI-driven models. Iconic Research and Engineering Journal. 2024, 8:432-451.

28. Gupta T, Bansal S: Predictive analytics for customer lifetime value (CLV): Using artificial intelligence to forecast purchasing behavior and churn. European Journal of Advances in Engineering and Technology. 2025, 12:113-120.

29. Chang V, Hall K, Xu QA, Amao FO, Ganatra MA, Benson V: Prediction of customer churn behavior in the telecommunication industry using machine learning models. Algorithms. 2024, 17:231. 10.3390/a17060231

30. Ardhani DA, Tania KD: Knowledge discovery on e-commerce customer churn using interpretable machine learning: A comparative study of SHAP-based classifiers. Journal of Applied Informatics and Computing. 2025, 9:2695-2702. 10.30871/jaic.v9i5.10811

How to cite this article:
Jain P (May 28, 2026) From Black-Box Predictions to Actionable Intelligence: A Systematic Review of Explainable Artificial Intelligence in Omni-Channel Retail. Cureus J Comput Sci 3 : es44389-026-00088-w. DOI https://doi.org/10.7759/s44389-026-00088-w