



Article

Transformer-Based Price Personalization: Attention-Based Modeling of Willingness to Pay in Product Recommendation

Ali Mahdavian ¹, Hadi Moradi ^{1,*} and Behnam Bahrak ²

¹ Department of Electrical and Computer Engineering, University of Tehran, Tehran 14179-35840, Iran; mahdavian.ali@ut.ac.ir

² Tehran Institute for Advanced Studies, Khatam University, Tehran 19916-33357, Iran; b.bahrak@teias.institute

* Correspondence: moradih@ut.ac.ir

Abstract

The growing convergence of personalized recommendations and pricing strategies has created new opportunities for retailers to support more targeted promotional decisions and potentially improve revenue-related outcomes, provided that such systems are implemented with appropriate managerial and fairness safeguards. This study introduces a Transformer-driven modeling approach that jointly performs product recommendation and personalized price-level assignment by modeling observed discount-level acceptance as an operational proxy for consumers' willingness to pay (WTP). Instead of handling price as an auxiliary or fixed attribute, the proposed model learns patterns of observed price responsiveness from sequential purchase histories that integrate user actions, item attributes, and discount exposure. The learning task is cast as a multi-class classification problem with five outcomes—non-purchase and purchases occurring at 0%, 10%, 20%, and above 30% discount levels—enabling the system to recommend not only relevant products but also economically appropriate pricing options. Empirical evaluation on the Dunnhumby retail dataset, which provides rich information on transactions and promotions, shows consistent performance across multiple evaluation criteria, including macro- and micro-level classification metrics, ranking-based measures such as NDCG@5, and Macro averaged AUC. Across the reported metrics, the Transformer-based framework performs better than our earlier deep reinforcement learning approach in this five-class setting, suggesting that sequence-based self-attention can better capture observed behavioral patterns related to discount-level acceptance. By producing integrated product–price recommendations, the proposed framework contributes toward closing the methodological gap between recommendation systems and dynamic pricing, and offers a scalable decision-support framework for data-driven, value-sensitive promotion design in retail and e-commerce contexts, subject to future validation in real-world and fairness-aware implementation settings.

Keywords: recommendation systems; willingness to pay; transformers; intelligent marketing promotions; dynamic pricing



Received: 8 January 2026

Revised: 9 June 2026

Accepted: 10 June 2026

Published: 30 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The accelerating expansion of the Internet, the emergence of the Fourth Industrial Revolution, and rapid advances in Artificial Intelligence (AI) have reshaped organizational workflows and customer-facing business processes. In retail and e-commerce, where customer profiles and detailed purchase histories are strategic assets, intelligent and data-driven promotional strategies have become essential for improving operational efficiency, customer engagement, and purchasing behavior.

A strong emphasis on customer retention further underscores the importance of personalized promotion. Multiple studies indicate that acquiring a new customer can cost five to seven times more than retaining an existing one [1–3]. As a result, personalized and AI-enabled promotions have become central to sustainable growth strategies. However, the relationship between personalized pricing and customer loyalty is not automatically positive. Prior research also shows that affective and cognitive mechanisms shape loyalty and decision-making in promotional contexts [4,5]. If customers perceive that individualized prices or discounts systematically disadvantage loyal or frequent buyers, such strategies may generate perceptions of price unfairness and weaken trust [6–8]. Therefore, AI-driven promotional systems should be designed not only to improve short-term conversion and promotional efficiency, but also to account for customer-relationship outcomes and perceived fairness.

AI-powered marketing systems enable firms to analyze large-scale behavioral data and adapt promotional strategies more precisely, thereby supporting more efficient resource allocation and potentially improving conversion through targeted offers [9]. Personalization plays a central role in these intelligent marketing ecosystems and provides the foundation for advanced pricing strategies such as Dynamic Pricing (DP) [10]. DP techniques are widely adopted in e-commerce to fine-tune discounting policies, balance supply and demand, and optimize inventory flows [11]. In competitive markets, such strategies increasingly rely on models capable of capturing complex temporal and behavioral patterns, as shown in recent work on spatial and temporal price prediction [12].

Recommendation systems (RSs) are widely used in digital commerce to learn user preferences, support product discovery, rank alternatives, and recommend complementary or substitute items [13,14]. In e-commerce, RSs help identify products that match individual tastes, rank alternatives during search sessions, and recommend complementary or substitute goods, thereby improving the efficiency of consumer decision-making.

Despite their effectiveness, conventional RS frameworks generally prioritize predicting whether a user will like, click, or purchase a product, while overlooking one of the most influential determinants of purchase decisions: price. This limitation can produce recommendations that are behaviorally relevant but economically misaligned with the customer's budget or willingness to pay (WTP). To generate actionable recommendations in retail settings, pricing must therefore be treated as an essential dimension of personalization. When profit margins, discount levels, or price sensitivity are excluded from the recommendation objective, the system remains incomplete from both a marketing and business perspective [15]. Moreover, although RS research often emphasizes accuracy-oriented metrics such as click-through rates, ranking effectiveness, or prediction precision, these models frequently ignore DP considerations that directly influence revenue and sales outcomes [16].

Based on this gap, the central research question of this study is how product recommendation and personalized price-level prediction can be integrated within a single Transformer-based framework. To address this question, the study models observed discount-level acceptance as a five-class prediction task and uses it as an operational proxy for customer WTP. Rather than treating price only as an auxiliary input feature, the proposed framework incorporates discount level into the recommendation output, enabling product–price prediction based on historical customer–product interactions.

The use of a Transformer architecture is motivated by its ability to model sequential customer–product interaction histories through self-attention. In this study, self-attention is used to identify behavioral signals related to purchase likelihood, discount responsiveness, and product relevance from historical retail sequences. Unlike uniform promotional strategies, which apply the same discount to broad customer groups, the proposed frame-

work supports more targeted discount allocation based on observed discount-response patterns used as a proxy for WTP. However, such personalization should not be interpreted as unrestricted surplus extraction from high-WTP customers. In practical applications, WTP-based pricing must be combined with fairness safeguards and customer-relationship considerations, particularly when high-WTP customers are also loyal or frequent buyers.

The proposed framework is evaluated using the Dunnhumby sales dataset, a large-scale retail dataset containing transaction, product, price, and discount information. The empirical results show stable performance in the multi-class purchase–discount prediction task, indicating that the model can distinguish between non-purchase, full-price purchase, and purchases under different discount levels. These results suggest that integrating pricing into the recommendation output can improve the practical relevance of RSs in price-sensitive retail environments.

This study makes three main contributions. First, it extends RS research by formulating recommendation as a joint product–price prediction problem rather than a product-only ranking task. Second, it operationalizes WTP through observed purchase and discount-level outcomes, allowing price sensitivity to be modeled directly within the recommendation output. Third, it indicates the applicability of Transformer-based sequence learning to multi-class product–price recommendation in retail data. At the same time, because personalized pricing may affect perceived fairness, trust, and loyalty, the framework is positioned as a decision-support tool whose real-world implementation should be supported by managerial constraints and fairness guardrails.

2. Related Work

2.1. RS

RSs have emerged as a key technological solution across various domains, particularly in e-commerce and retail. By leveraging patterns in user behavior, customer profiles, and product-related information, RSs identify customer preferences and suggest items that users are likely to engage with or purchase [17]. Beyond personalization, these systems reduce users' cognitive load and improve the overall quality of their decisions by filtering vast product collections into tailored suggestions [18]. Industry analyses indicate that adopting RS technologies can boost total purchase volume by approximately 30–35% in online retail platforms [19]. Personalized recommendation has also been studied in social-media and cognitive recommendation contexts, highlighting the broad applicability of user-centered recommendation methods [20,21].

2.1.1. Data Sources in RS

The performance of RSs depends on how users, products, and interactions are represented. Most recommendation models rely on user–item interaction data, customer-level information, and product-level attributes. User–item interactions may include explicit feedback, such as ratings, reviews, or stated preferences, and implicit feedback, such as clicks, page views, add-to-cart actions, purchases, and repeated transactions. While explicit feedback provides clearer preference signals, it is often sparse and difficult to collect at scale. Implicit feedback is more common in retail environments, but it is noisier because observed behavior does not necessarily indicate satisfaction or long-term preference [22].

For this reason, implicit behavioral signals must be interpreted together with broader customer and product information. Customer-level variables, such as purchase frequency, customer lifetime, spending patterns, and demographic attributes, help describe relatively stable behavioral tendencies. Product-level variables, including category, brand, price, sales volume, and discount exposure, help distinguish preference-driven behavior from price-driven purchase responses. Contextual and session-level signals, such as temporal

attributes, location, or current interaction sequences, can further enrich recommendation models when combined with transactional histories [13,23–27]. Recent work on implicit feedback and dynamic preference modeling further shows that recommendation inputs and objectives may evolve over time [28–30].

In price-aware recommendation, price and discount information are especially important. A purchase at full price may indicate a different level of observed price responsiveness from a purchase made only under a high discount. Similarly, repeated purchases under promotional conditions may suggest discount dependence rather than general product preference. Therefore, price and discount variables should not be treated only as auxiliary product attributes; they are central signals for modeling purchase likelihood and observed discount-level acceptance.

This issue is particularly relevant in retail datasets, where explicit ratings are often unavailable and customer behavior must be inferred from transaction records. The present study therefore combines customer, product, and interaction-level features with price-discount information to model both product relevance and discount-level acceptance within a unified Transformer-based framework. In doing so, it moves beyond traditional RSs that predict only which product to recommend and instead addresses the joint question of which product should be recommended and under which observed price-discount condition.

2.1.2. Types of RS by Output

Most real-world RSs generate recommendations for individual items that a user is likely to prefer or purchase. Other output formats, such as bundle recommendation and group recommendation, have also been studied to address settings where users may need product sets or shared recommendations [31–35]. These approaches show that recommendation outputs can be designed beyond single-item prediction when the decision context requires a more structured output.

In this study, the output is extended in a different direction. Rather than recommending only a product, the proposed framework predicts a product–price outcome by incorporating discount level into the recommendation output. This allows the model to address not only what should be recommended, but also under which observed price-discount condition the recommendation is most likely to be accepted.

2.1.3. Transformer-Based Models in RS

Transformer architectures have become increasingly important in sequential RS because self-attention can capture dependencies across long, sparse, and heterogeneous user–item interaction histories. Unlike models that process interactions mainly in chronological order, Transformer-based approaches can attend to multiple parts of a sequence simultaneously and identify the interactions most relevant to the next prediction task. This capability is particularly useful in retail settings, where purchase decisions may depend on temporally distant events, repeated product exposure, and previous responses to discounts or promotions [36,37].

Early models such as SASRec and BERT4Rec demonstrated the value of self-attention and bidirectional sequence modeling for next-item recommendation [38,39]. Later studies extended Transformer-based RS by incorporating self-supervised learning, auxiliary objectives, item attributes, textual reviews, heterogeneous event types, contextual information, and hybrid session-based structures [40–45]. Taken together, this literature shows that Transformers provide a flexible representation-learning architecture for modeling complex behavioral signals in RS.

In the context of price-aware recommendation, this flexibility is especially relevant because the model must represent not only product preference, but also observed price

responsiveness. The present study builds on Transformer-based sequential RS by using self-attention to model customer–product interaction histories together with price and discount exposure. Rather than focusing only on next-item prediction, the proposed framework estimates observed discount-level acceptance as part of the recommendation output, thereby linking product relevance with price-level prediction.

2.2. Price in RS

AI-driven marketing systems have transformed promotional strategies by enabling firms to analyze large-scale behavioral data and personalize offers according to customer needs [9]. In retail and e-commerce contexts, the effectiveness of a promotion depends not only on product relevance, but also on timing, discount level, perceived value, seasonality, promotion framing, and customer-specific price responsiveness [46,47]. Price is therefore a central determinant of purchase behavior: even when a product matches customer preferences, it may still be rejected if the offered price exceeds the customer’s reservation price [48–50].

As RSs have evolved into decision-support tools, several studies have argued that recommendation quality should not be evaluated only through accuracy or ranking metrics, but also through business-oriented criteria such as price, profit, and promotional efficiency [16,51]. Value-aware, price-aware, profit-aware, and discount-sensitive RSs address this issue by incorporating economic information into recommendation models [52–54]. Related work has also examined discount rates, complementary product pricing, personalized price thresholds, personal price-aware recommendation, and price-integrated interaction models, including graph-based approaches [15,55–59].

Despite these advances, price is often treated as an input feature, a post-processing criterion, or a separate dynamic pricing decision rather than as part of the recommendation output. This separation limits the ability of RSs to jointly account for product relevance and observed discount-level acceptance. In personalized promotion settings, delivering the right product at a price aligned with inferred WTP may support conversion and customer satisfaction, but such effects depend on fairness, transparency, and customer-relationship safeguards [60]. Therefore, the present study treats discount level as part of the prediction task and uses observed purchase–discount behavior as an operational proxy for WTP, rather than assuming that true WTP is directly observed.

2.3. Challenges and Research Gap

Despite the progress of RS research, most existing approaches remain primarily focused on predicting product relevance. Conventional RSs are typically designed to answer the question of which item a user is likely to prefer or purchase, while price is either ignored, treated as a fixed product attribute, or used only as an auxiliary input feature. This creates an important limitation in retail and e-commerce contexts, where purchase decisions depend not only on product preference but also on whether the offered price or discount level matches the customer’s WTP.

A second challenge concerns the separation between recommendation and pricing decisions. Dynamic pricing models can support price adjustment, while RSs can rank relevant products; however, these two decisions are often modeled independently. As a result, a system may recommend a relevant product at an unsuitable price or offer a discount without adequately considering the customer’s sequential behavior and product preference. This separation limits the ability of existing models to generate recommendations that are simultaneously relevant to the customer and economically meaningful for the firm.

A further methodological challenge arises from the nature of retail data. In many real-world datasets, explicit preference ratings are unavailable, and customer responses

must be inferred from implicit transaction histories, observed prices, and discount exposure. This makes it difficult to distinguish between product preference and price sensitivity using static or non-sequential models. Sequence-learning architectures are therefore particularly relevant because they can capture how prior purchases, discount responses, and product interactions jointly shape future purchase behavior.

The present study addresses these limitations by proposing a Transformer-based framework that integrates product recommendation and personalized price-level prediction within a single multi-class sequence-learning model. Unlike traditional RSs that output only product recommendations, the proposed model predicts purchase outcomes across multiple discount classes, including no purchase, full-price purchase, and purchases under different discount levels. In this way, the framework treats price not merely as an input variable but as part of the recommendation output, thereby bridging the gap between RSs and personalized pricing.

This formulation allows the model to operationalize WTP through observed price-discount behavior and to generate recommendations that jointly reflect customer preference, price responsiveness, and business value.

To clarify the positioning of the proposed framework within the existing literature, Table 1 summarizes the main differences between traditional RSs, price-aware RSs, dynamic pricing models, and sequential Transformer-based RSs. This comparison shows that previous approaches have typically addressed product recommendation and pricing as separate or only partially connected tasks. In contrast, the present study integrates these two decisions by treating the discount level as part of the recommendation output. This positioning highlights the main contribution of the proposed framework: moving from product-only recommendation toward joint product–price recommendation based on customer interaction sequences and predicted WTP-related discount acceptance.

Table 1. Positioning of the proposed framework relative to existing RS approaches.

Research Stream	Main Focus	Main Limitation	Position of This Study
Traditional RSs	Predict relevant products based on user–item interactions and preference patterns.	Price is usually ignored or treated as a fixed product attribute rather than a decision variable.	Integrates product relevance with personalized price-level prediction.
Price-aware RSs	Incorporate price, profit, or economic value into recommendation models.	Price is commonly used as an input feature, while the recommendation output remains product-oriented.	Treats the discount level as part of the recommendation output.
Dynamic pricing models	Adjust prices according to demand, customer characteristics, competition, or expected revenue.	Pricing decisions are often modeled separately from product recommendation and sequential preference learning.	Jointly models product recommendation and WTP-related discount classes.
Sequential and Transformer-based RSs	Capture temporal dependencies and long-range behavioral patterns in user interaction histories.	Most applications focus on next-item prediction rather than personalized product–price recommendation.	Uses self-attention to predict product–discount outcomes from customer interaction sequences.

3. Proposed Method

This section presents the proposed Transformer-based model for product recommendation aligned with each customer’s WTP. Instead of learning through interactions as in reinforcement learning, the new framework models the problem as a sequence prediction task, in which customer behavior is represented as an ordered series of historical interactions processed through a multi-head self-attention architecture.

Each customer is represented by a chronological sequence of behavioral signals, including purchases across product categories and price tiers, transactions with and without discounts, browsing actions, preferences, and product evaluations. These interaction sequences are enriched with static user attributes—such as total number of purchase invoices, customer lifetime, Average amount paid, and RFM score—to provide contextual grounding that enhances personalization. Similarly, each product is encoded through descriptive attributes including original product price, category and number of unique customers who bought the product.

The primary objective of the proposed model is to support product recommendation and discount-level prediction by aligning product–price offers with observed discount–response patterns used as a proxy for WTP. From a marketing perspective, effective price-sensitive recommendation can improve short-term conversion and promotional efficiency, but it should not be interpreted as directly maximizing customer satisfaction, retention, or organizational value. The Transformer-based architecture directly supports this objective by capturing long-range dependencies in user histories and extracting latent purchase patterns relevant to price responsiveness.

In the proposed framework, the model predicts the product and discount class most consistent with the customer’s observed purchase and discount–response history. This is achieved by jointly encoding customer attributes, product features, and sequential interaction histories through stacked self-attention layers. The output is a probability distribution over purchase classes (e.g., no purchase, purchase at full price, purchase under various discount levels), allowing the system to identify the price–product combination with the highest likelihood of acceptance.

Consistent with the principles of value-aware recommender systems (VARs), which emphasize the integration of user utility and business value into recommendation models [16], the proposed Transformer-based approach incorporates price-related information and discount mechanisms directly into the prediction task. This enables the system to generate recommendations that are both preference-aware and value-informed, without relying on online interaction or reward estimation as in reinforcement learning.

3.1. Problem Definition

This paper introduces a product recommendation and price-personalization framework that leverages Transformer-based sequence modeling to align product suggestions with each customer’s WTP. The proposed approach formulates the task as a supervised prediction problem over sequential purchase histories enriched with contextual customer and product attributes.

The problem setting incorporates several categories of features essential for modeling price-sensitive recommendation behavior. Customer features include detailed interaction histories—such as the diversity of past purchases, frequency of discount usage, preferred categories, and temporal patterns like purchase intervals and responsiveness to promotions. These behavioral indicators are complemented by user-level characteristics including inferred price sensitivity. Product features capture pricing and positioning elements such as the product’s price rank within its category, its overall sales performance, the proportion of discounted versus regular-price transactions, and the diversity of its customer base.

Observed discount–response patterns are modeled by jointly using customer and product attributes within a sequence-learning architecture. The Transformer captures local and long-range dependencies in user behavior, enabling the system to infer price responsiveness and predict the product and discount class most consistent with the customer’s observed purchase and discount–response history. This allows the recommendation mechanism to incorporate not only preference matching but also pricing alignment.

Unlike traditional RSs that focus primarily on predicting user–item relevance, the proposed method integrates pricing considerations directly into the recommendation task. By conditioning predictions on WTP-related signals and historical discount interactions, the model generates recommendations that are both relevant and economically optimized. This dual consideration enables the system to address demand-level preferences and pricing sensitivity simultaneously, delivering a more comprehensive recommendation solution.

The proposed model is applicable across a broad set of marketing use cases. By predicting observed discount-level acceptance, it can inform personalized promotional strategies in which discount levels and special offers are tailored to each customer’s purchasing profile in order to increase conversion likelihood. At the same time, the inferred WTP can be leveraged to support dynamic pricing mechanisms by providing guidance for price adjustments that reflect observed behavioral patterns. In addition, the framework may inform customer retention and loyalty programs when its outputs are combined with relationship-oriented rules, such as avoiding systematic disadvantage for loyal customers and limiting excessive personalized price differences.

Price plays a pivotal role in marketing because it strongly shapes customer purchase decisions [6]. Customers who consistently buy products only when discounts are available typically exhibit high price sensitivity and lower acceptance of full-price offers. By modeling this pattern within a sequential Transformer architecture, the system can identify discount-dependent purchasing behaviors and adjust recommendations accordingly. This enables businesses to deliver personalized price incentives that better match each customer’s inferred WTP and may support purchase likelihood and sales performance. However, long-term loyalty cannot be assumed to follow automatically from personalized pricing; it depends on whether customers perceive the pricing policy as fair and relationship-oriented [8,61].

In addition to defining the problem, it is important to emphasise the generalisation capability of the proposed framework. The Transformer-based sequence modelling approach developed in this paper is domain-agnostic, as it relies on self-attention over interaction sequences rather than domain-specific state–action–reward dynamics. With appropriate retraining and fine-tuning on domain-specific features and target tasks, the model can be adapted from retail to other industries such as fashion, electronics, or digital services—where consumer behaviour may differ in terms of purchase frequency, price sensitivity, and seasonal patterns. Recent research shows that Transformer architectures can generalise across different domains or data distributions, especially when exposed to diverse interaction data and adapted appropriately [62].

3.2. Problem Modeling Using a Transformer-Based Sequence Framework

In the proposed model, the problem of simultaneous product recommendation and pricing is reformulated using a Transformer-based sequence modeling approach. The Transformer framework learns behavioural patterns directly from historical interaction sequences, enabling the system to jointly recommend the most suitable product and its corresponding price tier. Owing to the characteristics of attention-based models—such as their ability to capture long-range behavioural dependencies, incorporate heterogeneous customer and product features, and learn complex purchase dynamics—the proposed method is expected to provide a suitable representation for the five-class product–price prediction task.

The overall architecture of the proposed framework, from input representation to output generation, is illustrated in Figure 1 and summarized as follows:

1. **Inputs:** The model receives structured information describing customers, products, and their interaction histories. These inputs include sequential records of cus-

- tomer–product interactions, purchase events under different discount levels, and descriptive product attributes, forming the basis for sequence-aware modeling.
2. **Data Representation and Encoding:** Customer interaction histories are transformed into ordered sequences that preserve temporal structure and purchasing context. These sequences are encoded into dense representations that capture both behavioral dynamics and price-related signals relevant to WTP.
 3. **Transformer-Based Sequence Modeling:** The encoded interaction sequences are processed through a Transformer architecture built on multi-head self-attention. By attending to different parts of the interaction history, the model learns latent dependencies between past behavior and price sensitivity, enabling a unified representation of preference and responsiveness.
 4. **Outputs:** Based on the learned representations, the model predicts a probability distribution over purchase outcomes, including non-purchase and purchases at different discount levels. The final output consists of personalized product–price predictions based on observed discount-level acceptance patterns.

In this modeling approach, the prediction task is formulated as a sequence-to-class problem. Each input sequence represents the joint temporal information of a customer and the products they interact with, enriched with demographic and product-level attributes. A sequence embedding thus reflects whether a customer with certain characteristics is likely to purchase a particular product at a given price or discount level.

Within this framework, the Transformer encoder captures the contextual relationships among interactions, products, and prices. The model outputs a probability distribution over possible product–price pairs, allowing the system to identify the most promising recommendation for each customer based on predicted purchase likelihood and inferred WTP.

By capturing both immediate and longer-term behavioral dependencies through self-attention, the Transformer-based approach supports value-aware product–price prediction while remaining suitable for retraining or fine-tuning as new interaction data become available.




Figure 1. Overview of the proposed Transformer-based framework for joint product and price recommendation.

Figure 1 summarizes the proposed framework for joint product–price recommendation. The framework takes customer, product, and interaction sequences as input and converts them into a sequential interaction representation. This representation is then passed to a Transformer-based sequence model, where self-attention is used to capture dependencies across previous interactions and to identify the behavioral patterns that are most informative for predicting the customer’s next response.

The Transformer-based model was configured to balance predictive performance, computational efficiency, and the risk of overfitting. Given the five-class prediction task and the structured retail nature of the Dunhumby dataset, a compact Transformer architecture was selected rather than a very deep model. The final model used an embedding dimension of 128, two Transformer encoder layers, four attention heads, a feed-forward hidden dimension of 256, and a dropout rate of 0.20. The maximum sequence length was set to 50 previous customer–product interactions, allowing the model to capture meaningful purchase and discount histories while avoiding excessive noise from very long sequences.

The model was trained using the Adam optimizer with a learning rate of 1×10^{-4} and a batch size of 256. Training was conducted for a maximum of 30 epochs, with early stopping applied when validation performance did not improve for five consecutive epochs. The output layer used a softmax activation function to predict one of five mutually exclusive outcome classes: non-purchase, purchase with no discount, purchase with a 10% discount, purchase with a 20% discount, and purchase with a discount of 30% or more. The final configuration was selected based on validation-set performance, with particular attention to Macro F-score and NDCG@5, because these metrics reflect both balanced classification performance and ranking quality across discount classes. Final hyperparameter settings are shown in Table 2.

Table 2. Final hyperparameter settings of the Transformer-based model.

Hyperparameter	Selected Value
Embedding dimension	128
Number of Transformer encoder layers	2
Number of attention heads	4
Feed-forward hidden dimension	256
Dropout rate	0.20
Maximum sequence length	50 interactions
Output classes	5
Batch size	256
Learning rate	1×10^{-4}
Optimizer	Adam
Maximum number of epochs	30
Early stopping patience	5 epochs
Output activation	Softmax
Primary validation criteria	Macro F-score and NDCG@5

To assess the sensitivity of the model to architectural complexity, alternative configurations with smaller and larger embedding dimensions, different numbers of attention heads, and different sequence lengths were considered during model development. Very small configurations reduced the model’s ability to capture heterogeneous purchase and discount patterns, whereas larger configurations increased computational cost and the risk of overfitting without providing proportional gains in validation performance. The selected configuration therefore represents a balanced architecture for the five-class product–price prediction task.

The prediction task is defined as a multi-class outcome that combines purchase likelihood and discount-level acceptance. The model can predict a non-purchase outcome when the customer is unlikely to buy the product. When a purchase is predicted, the framework assigns the recommendation to one of four pricing outcomes: purchase with a 0% discount, purchase with a 10% discount, purchase with a 20% discount, or purchase with a discount of 30% or more. Thus, the framework extends conventional product recommendation by integrating personalized pricing into the recommendation output and by operationalizing WTP as a discount-level prediction problem.

To capture the role of price in shaping customer behavior and to provide a structured input representation for sequence modeling, the proposed framework begins by constructing descriptive representations of customers and products. Rather than serving as explicit decision states, these representations provide contextual information that complements the sequential interaction histories used by the Transformer model.

Customer, product, and interaction data are jointly utilized to derive a set of static attributes and dynamic signals that characterize purchasing behavior and price exposure. Customer profile attributes summarize long-term purchasing tendencies, such as spending levels, purchase frequency, and discount usage, while product attributes describe market-related characteristics, including price positioning and aggregate demand patterns. In parallel, interaction-level information is preserved as ordered sequences of events, allowing the model to retain the temporal structure of customer–product relationships.

These static customer and product attributes are embedded as contextual inputs and combined with sequential interaction representations to form the complete input to the Transformer encoder. Through multi-head self-attention, the model learns latent dependencies between historical behavior and price-related outcomes, enabling the extraction of meaningful representations associated with WTP and purchase propensity.

The full set of input features used to construct customer profiles, product attributes, and interaction sequences is summarized in Table 3.

Table 3. Input features used for customer, product, and interaction representations in the proposed Transformer-based model.

Product Attribute Embeddings	
P1	Original product price
P2	Average original price across transactions
P3	Number of unique customers who purchased the product
P4	Total number of product sales
P5	Average time interval between product purchases
Customer Profile Embeddings	
C1	Total number of purchase invoices
C2	Customer lifetime (time since first purchase)
C3	Number of unique products purchased
C4	Average amount paid per transaction
C5	Average realized purchase value after discount
C6	Average discount amount used by the customer
C7	Average time interval between consecutive purchases
C8	Average time interval between discounted purchases
C9	RFM score
Sequential Interaction Signals	
I1	Historical purchase events ordered by time
I2	Product categories and price tiers in interaction sequence
I3	Discount usage indicators within interaction history
Prediction Target (Purchase Outcome)	
T1	Purchase outcome class (no purchase, 0%, 10%, 20%, $\geq 30\%$ discount)

By structuring the input space around aggregated behavioral indicators and temporally ordered interaction signals, the proposed representation reduces the ambiguity typically associated with implicit feedback data. Instead of relying on isolated purchase events, the model captures broader behavioral regularities, including discount sensitivity, spending consistency, inter-purchase timing, and sustained engagement across product

categories. This representation encourages the learning of stable behavioral patterns that support reliable inference of customer preferences and WTP.

The prediction task is formulated as a five-class classification problem, corresponding to non-purchase and purchases occurring at different discount levels (0%, 10%, 20%, and at least 30%). These output classes are summarized in Table 4.

Model training is conducted using a multi-class cross-entropy loss function, which measures the discrepancy between the predicted purchase outcome and the observed customer behavior. Correct identification of both purchase occurrence and discount level leads to lower loss values, while incorrect predictions are penalized according to their deviation from the true class. This supervised objective enables the Transformer to jointly learn purchase likelihood and price sensitivity within a unified modeling framework.

Table 4. Purchase outcome classes used for supervised product–price prediction.

Class	Purchase Outcome Description
C0	No purchase event
C1	Purchase at full price (0% discount)
C2	Purchase with low discount (10%)
C3	Purchase with medium discount (20%)
C4	Purchase with high discount ($\geq 30\%$)

3.3. Training the Model

The proposed model employs a Transformer-based sequence learning approach to jointly predict product recommendations and their corresponding price or discount levels. The Transformer framework learns from historical interaction sequences through supervised learning. This enables the model to capture long-range behavioral patterns and predict observed purchase–discount outcomes.

The dataset used for training and evaluation contains 1,817,012 samples, including both positive and negative interaction records. Positive samples correspond to completed transactions in which the customer purchased the product at the observed discount level, whereas negative samples represent instances in which no purchase occurred. The large scale of the dataset provides sufficient variation for model training and enables evaluation on unseen customer–product interactions.

The proposed neural architecture is built upon a Transformer encoder and is designed to jointly model sequential customer behavior and price-related purchase outcomes. The model consists of the following components:

- **Input Representation Layer:** Customer profile attributes, product characteristics, and temporally ordered interaction histories are mapped into dense embedding spaces. Static customer and product embeddings provide contextual information, while interaction-level embeddings preserve the temporal structure of purchasing behavior. These representations are combined to form the input sequence to the Transformer encoder.
- **Transformer Encoder:** The embedded sequences are processed through multiple stacked Transformer encoder layers based on multi-head self-attention. This mechanism enables the model to selectively attend to different parts of the interaction history, capturing long-term dependencies, recurring purchasing patterns, and latent relationships between past behavior and price sensitivity.
- **Position-wise Feed-Forward Network:** Each encoder layer is followed by a fully connected feed-forward network with ReLU activation, which introduces non-linearity and refines the representations produced by the attention mechanism.

- **Output Layer:** The final sequence representation is passed to a softmax classification head that outputs a probability distribution over discrete purchase outcomes, including non-purchase and purchases under different discount levels (0%, 10%, 20%, and $\geq 30\%$).

The training procedure for the Transformer-based product–price recommendation model is summarized in Algorithm 1. The model is trained using supervised mini-batch learning, where the objective is to minimize the categorical cross-entropy loss between the predicted purchase outcome and the observed customer behavior.

This loss function penalizes incorrect predictions proportionally to their divergence from the true class, encouraging the model to jointly learn purchase likelihood and discount sensitivity within a unified framework. By optimizing this objective, the model aligns sequential behavioral representations with economically meaningful pricing outcomes.

After training, model performance is evaluated on a held-out test set using classification and ranking-based metrics. The results demonstrate that the Transformer architecture effectively captures long-range behavioral dependencies and nuanced price responsiveness patterns that are often missed by conventional recommendation approaches.

An additional advantage of the proposed framework is its adaptability to evolving customer behavior. As new interaction data become available, the model can be retrained or fine-tuned to reflect changes in purchasing patterns and market conditions, allowing recommendation and pricing strategies to remain accurate, relevant, and responsive over time.

Algorithm 1 Training Procedure for Transformer-Based Product–Price Recommendation

```

1: Input: Training dataset  $\mathcal{D}$ ; batch size  $B$ ; learning rate  $\eta$ ; maximum epochs  $E$ 
2: Initialize Transformer model parameters  $\theta$ 
3: for epoch = 1 to  $E$  do
4:   for each mini-batch  $\{(X_i, y_i)\}_{i=1}^B$  from  $\mathcal{D}$  do
5:     Encode customer, product, and interaction features into embeddings
6:     Construct ordered interaction sequences
7:     Process sequences through Transformer encoder layers
8:     Compute purchase outcome probabilities  $\hat{y}_i = \text{softmax}(f_\theta(X_i))$ 
9:     Compute categorical cross-entropy loss:

```

$$\mathcal{L} = -\frac{1}{B} \sum_{i=1}^B y_i \log(\hat{y}_i)$$

```

10:   Update parameters using Adam optimization:

```

$$\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}$$

```

11:   end for
12:   Apply early stopping based on validation loss
13: end for
14: Output: Trained Transformer-based product–price recommendation model

```

3.4. Evaluation Criteria

Evaluation metrics are defined and selected according to the problem structure and its outputs. In this problem, product recommendations along with pricing are provided. The output of the problem is defined as multi-class. Therefore, multi-class classification metrics such as Macro average accuracy, Macro average recall, Macro F-score, Micro F-score [63], One-vs-Rest (OvR) Macro averaged AUC [64] and Normalized Discounted Cumulative Gain (NDCG) are used. Additionally, weighted multi-class accuracy is proposed for evaluation based on different preferences.

Considering that the model output is a type of multi-class classification output, multi-class evaluation methods are used to evaluate this model. The confusion matrix for multi-class classification for one sample class is shown in Table 5, and the metrics of Macro average accuracy, Macro average recall, Macro F-score, Micro F-score, Macro averaged AUC and NDCG are used to evaluate this model.

Table 5. Metrics in multi-class classification for the class ‘Purchase with 10% Discount’.

Actual	Predicted				
	No Purchase	0% Discount	10% Discount	20% Discount	≥30% Discount
No Purchase	TN	TN	FP	TN	TN
0% Discount	TN	TN	FP	TN	TN
10% Discount	FN	FN	TP	FN	FN
20% Discount	TN	TN	FP	TN	TN
≥30% Discount	TN	TN	FP	TN	TN

In addition to standard multi-class evaluation metrics, we use a weighted multi-class accuracy measure to provide a complementary business-oriented interpretation of prediction errors. This metric assigns different weights to correct and incorrect predictions according to the relative managerial cost of each outcome, as shown in Table 6. The purpose of this metric is not to replace standard measures such as Macro F-score, AUC, or NDCG, but to reflect that different types of pricing errors may have different business implications.

Table 6. Weights for predicting different classes to calculate weighted multi-class accuracy.

Actual	Prediction				
	No Purchase	0% Discount	10% Discount	20% Discount	≥30% Discount
No Purchase	10	−4	−3	−2	−1
0% Discount	−4	10	−1	−2	−3
10% Discount	−4	−3	10	−1	−2
20% Discount	−4	−3	−2	10	−1
≥30% Discount	−4	−3	−2	−1	10

The weighting scheme in Table 6 is intended as a managerial cost-sensitive evaluation tool rather than a universal performance metric. Correct predictions receive a positive coefficient of +10 because they correctly identify both purchase occurrence and discount level. Incorrect predictions receive negative coefficients according to their expected business implications. The largest penalty (−4) is assigned to errors that may lead to missed purchases or inefficient targeting, such as predicting non-purchase when a purchase occurs or predicting a purchase when no purchase occurs. Smaller penalties are assigned to errors between adjacent discount classes because they imply smaller deviations in promotional intensity. For example, if the actual outcome is a purchase with a 20% discount, predicting non-purchase receives the largest penalty, while predicting a nearby discount class receives a smaller penalty.

After assigning the corresponding coefficient to each observation, the weighted multi-class accuracy is calculated by summing all weighted values and normalizing by the maximum possible score, as follows:

$$\text{Weighted Multi-Class Accuracy} = \frac{\text{Sum of all positive and negative values}}{\text{Total number of samples} \times 10} \quad (1)$$

Because these weights are managerially motivated and may vary across firms, the selected weighting scheme should be considered illustrative rather than universally optimal. Therefore, this metric is interpreted as a complementary business-oriented measure alongside standard metrics such as Macro F-score, AUC, and NDCG. Future research should test alternative weighting structures and compare the results with established cost-sensitive evaluation approaches.

3.5. Challenges and Innovations

The main methodological challenge in this study is to model product relevance and discount-level acceptance within a single supervised sequence-learning framework. Unlike conventional RSs, which usually focus on item relevance, the proposed framework must jointly represent customer behavior, product attributes, and observed price-discount responses. This requires input representations that capture purchase frequency, discount exposure, spending patterns, temporal interaction structure, and product-level characteristics. The Transformer architecture is used because self-attention can model these heterogeneous and sequential signals within a unified representation.

The main contribution of the proposed framework is its formulation of product recommendation as a five-class product-price prediction task. Instead of predicting only whether a product is relevant, the model predicts whether the customer is unlikely to purchase, likely to purchase at full price, or likely to purchase under different discount levels. This makes discount level part of the recommendation output and allows the system to estimate observed discount-level acceptance as an operational proxy for WTP. The resulting framework provides a value-aware decision-support approach for retail recommendation, while practical deployment should still consider fairness safeguards, customer trust, and long-term relationship outcomes.

4. Experiments

This section presents a comprehensive analysis of the experimental results. We begin with a detailed description of the dataset used for training, validation, and testing the Transformer-based pricing and recommendation model. We then describe the evaluation methodology, which includes training the model on sequential interaction data, validating its performance during training, and testing its generalization on a separate held-out dataset. Standard performance metrics such as precision, recall, F1 score, and overall accuracy are employed to assess both the quality of product recommendations and the correctness of discount-class predictions. Finally, we analyze the experimental results, demonstrating substantial improvements in recommendation accuracy and pricing alignment compared to other approaches, and highlighting the model's ability to adapt to evolving customer behavior through periodic retraining and fine-tuning.

4.1. Dataset Description

The Dunnhumby dataset is selected due to its relevance to the problem at hand and its extensive use in numerous studies. This dataset adequately covers features related to discounts, pricing, and promotional attributes. The Dunnhumby dataset is an extensive collection of retail transaction data, encompassing 2,595,732 purchase transactions from 2500 unique customers over 711 days. It includes detailed customer demographics, product attributes, and transactional specifics such as pricing and discounts. This dataset is instru-

mental for developing and testing machine learning algorithms, particularly in customer segmentation, product recommendation, DP, and market basket analysis. Its structure makes it suitable for evaluating data-driven retail models that use transaction, product, price, and discount information [65].

Some specifications of the Dunnhumby dataset are presented in Table 7.

Table 7. Dunnhumby dataset specifications.

Feature	Value
Number of Transactions	2,595,732
Number of Customers	2500
Number of Products	92,339
Time Period	711 days
Number of Invoices	276,484
Number of Product Categories	308
Number of Product Subcategories	2383

Table 8 presents the distribution of samples used for training and testing the Transformer-based recommendation and pricing model. Following a standard 70–30 split strategy, a total of 1,817,012 samples (70%) were allocated to the training set, enabling the model to learn customer behavioral patterns, discount responses, and product interactions from a sufficiently large dataset. The remaining 778,720 samples (30%) constitute the test set, which is used exclusively to evaluate the model’s generalization performance on unseen customer–product interactions. This split ensures a balanced and reliable assessment of the model’s predictive accuracy and robustness.

Table 8. Train and test split.

Category	Percentage	Number of Transactions
Training	70%	1,817,012
Testing	30%	778,720

4.2. Experimental Results

The proposed framework formulates product recommendation and personalized pricing as a five-class prediction problem. The output classes include non-purchase, purchase with no discount, purchase with a 10% discount, purchase with a 20% discount, and purchase with a discount of 30% or more. This formulation differs from most existing recommendation and pricing studies, which typically address either next-item recommendation, binary purchase prediction, or price optimization as a separate task. As a result, direct comparison with many previous methods is not straightforward, because their output structure does not correspond to the joint product–price prediction task considered in this study.

Because the study relies on historical offline retail data, the output labels were constructed from observed customer responses under the prices and discount levels that were actually offered. The five classes used in the model therefore should not be interpreted as direct observations of a customer’s true latent WTP. Instead, they represent observed purchase–discount outcomes and are used as an operational proxy for discount-level acceptance. Specifically, each customer–product interaction was assigned to one of five mutually exclusive classes: non-purchase, purchase with no discount, purchase with a 10% discount, purchase with a 20% discount, or purchase with a discount of 30% or more.

Importantly, the non-purchase class does not imply that the customer would necessarily have purchased the product if a discount had been offered. Similarly, a purchase under

a given discount level does not prove that the same customer would not have purchased under a lower discount or at full price. The label only records the observed outcome under the price and discount condition that appeared in the historical data. Therefore, no unobserved counterfactual price response was imputed, and no non-purchase at full price was automatically converted into evidence of discount requirement.

This labeling strategy follows the empirical constraints of offline pricing data. As also recognized in our previous reinforcement-learning-based work, historical retail datasets often lack information on inaction and counterfactual actions: we observe the price or discount that was offered and the resulting customer response, but we do not observe how the same customer would have responded to alternative price points. Consequently, the model learns patterns of observed discount acceptance rather than true causal WTP. This distinction is important for interpreting the results and for avoiding overstatement of the model's revenue implications.

In our previous reinforcement-learning-based study [66], the same business problem was modeled through a sequential decision-making framework. In that work, the state space represented customer, product, and interaction information, the action space corresponded to pricing or discount decisions, and the reward function was designed to reflect purchase response and economic value. However, because most existing benchmark methods in the literature were formulated as binary prediction models, the original multi-class discount structure had to be transformed into a binary purchase/non-purchase setting for comparison with those methods. While this enabled comparison with conventional models, it also reduced the granularity of the pricing decision and did not fully preserve the five-class WTP-related structure used in the present study.

For this reason, the comparison with the RL-based approach is used as the closest available benchmark from our prior work, because it addresses the same retail pricing context and uses related customer-product interaction information. At the same time, we acknowledge that the difference in modeling logic between reinforcement learning and Transformer-based supervised sequence learning limits direct one-to-one comparison. The present model preserves the five-class output structure and directly predicts the discount-level outcome, which is central to its contribution.

The RL-based benchmark [66] differs from the proposed Transformer model in both learning objective and output structure. In the RL formulation, the pricing problem is represented as a sequential decision process. The state contains information about the customer, product, and historical interactions; the action corresponds to selecting a price or discount decision; and the reward reflects the observed purchase response and the economic value associated with the selected action. The model therefore learns a policy that maps customer-product states to pricing actions. In contrast, the proposed Transformer-based framework treats the problem as supervised sequence learning and directly predicts the probability of each purchase-discount class from historical interaction sequences. This difference is important because the Transformer model preserves the full five-class structure, whereas the RL-based study required a binary reformulation when compared with conventional prediction methods.

The evaluation of this model utilizes multi-class classification metrics and the proposed weighted multi-class accuracy based on Table 6. Figure 2 shows the confusion matrix for this modeling approach.

The confusion matrix in Figure 2 illustrates the performance of the Transformer-based pricing and recommendation model across the five outcome classes: no purchase, 0% discount, 10% discount, 20% discount, and more than 30% discount. The true labels are displayed on the Y-axis and represent actual customer behavior, while the predicted labels on the X-axis indicate the model's output. Each cell reports both the number of instances

and the row-wise percentage, providing a detailed view of prediction accuracy for each behavioral category.

A key result is the model's strong ability to identify customers who are unlikely to make a purchase. The "Don't Purchase" class exhibits a high true-positive count of 441,225, corresponding to 87.4% of all actual non-purchasing instances. This high degree of accuracy is vital in practice, as it allows businesses to avoid allocating unnecessary marketing resources to low-engagement customers, thereby improving the efficiency of promotional strategies.

The model also indicates stable performance in predicting purchasing behavior across discount categories. For example, the 0% discount class achieves 146,019 correct predictions, representing 87.6% recall in that row. This indicates that the model is effective at identifying cases in which a purchase may occur without a price reduction. From an operational perspective, such predictions may help firms reduce indiscriminate discounting and support margin protection, provided that they are implemented with fairness and relationship-oriented safeguards.




Figure 2. Confusion matrix for the Transformer-based pricing and product recommendation model.

For customers who require moderate discounts, such as the 10% and 20% classes, the model achieves 83.7% and 86.3% row-wise accuracy, respectively. These results show that the Transformer captures mid-range price sensitivity patterns, a crucial capability for efficient allocation of promotional budgets across different customer segments.

Importantly, the model also performs well in identifying customers who respond only to substantial discounts. The "more than 30%" discount class reaches a correct prediction count of 50,697, corresponding to 84.9% accuracy. Customers in this group exhibit high price sensitivity and are typically difficult to convert without significant incentives. Accurate identification of this segment may help businesses apply high-value discounts more selectively, potentially supporting conversion while reducing unnecessary promotional spending. However, the profitability implications should be validated in future experimental or real-world deployment settings.

Accurate discrimination across discount classes provides multiple practical benefits: (1) cases with high full-price purchase likelihood can be identified without relying on broad discount campaigns; (2) highly price-sensitive cases can be approached with appropriate incentives; and (3) mass-discounting strategies can be replaced with more precise, customer-specific promotional decisions. This targeted allocation of promotional resources can

support short-term sales performance and improve the relevance of offers. Nevertheless, long-term customer satisfaction depends on whether personalized price differences are perceived as fair and whether loyal customers are protected from systematic disadvantage.

Taken together, the confusion matrix results show that the Transformer-based model delivers strong predictive accuracy across the five observed outcome classes and may support data-driven promotional decisions when implemented with appropriate safeguards.

Table 9 summarizes the performance of the proposed Transformer-based pricing and product recommendation model on the test dataset. A comprehensive set of evaluation metrics was employed to assess both multi-class classification accuracy and ranking quality across the five discount categories.

Table 9. Evaluation results for the Transformer-based pricing and product recommendation model.

Metric	Test Data
Macro Average Precision	0.8066
Macro Average Recall	0.8597
Macro F-score	0.8302
Micro F-score	0.8702
Macro Averaged AUC	0.8912
NDCG@5	0.9065
Weighted Multi-class Accuracy	0.8265

The results in Table 9 show that the model performs consistently across both classification and ranking-based metrics. The Macro F-score of 0.8302 and Micro F-score of 0.8702 indicate balanced performance across the five outcome classes and stable aggregate prediction quality. The Macro Averaged AUC of 0.8912 and NDCG@5 of 0.9065 further suggest that the model ranks product–discount outcomes effectively in the observed test data. Finally, the Weighted Multi-class Accuracy of 0.8265 provides a complementary business-oriented view of performance by accounting for the relative importance of different prediction errors. Overall, these results suggest that the model can predict observed purchase–discount outcomes reliably, while the revenue and managerial implications should be interpreted as predictive rather than causal.

The detailed precision, recall, and Macro F-score values for the five output classes are presented in Tables 10, 11 and 12, respectively. These metrics provide a class-wise evaluation of the Transformer-based pricing and recommendation model and highlight its strengths and areas for refinement.

Table 10. Precision evaluation for classes in the Transformer-based pricing and product recommendation model.

Classes	Precision
Class 1—No Purchase	0.9644
Class 2—Purchase at Full Price	0.7156
Class 3—Purchase with 10% Discount	0.7600
Class 4—Purchase with 20% Discount	0.7794
Class 5—Purchase with 30% Discount or More	0.8139

Precision Results. Table 10 shows that the model achieves exceptionally high precision for the No Purchase class (0.9644), indicating that when the model predicts that a customer will not make a purchase, it is correct in the vast majority of cases. This is crucial for operational efficiency, as it reduces unnecessary marketing expenditures for customers unlikely to respond to promotions.

Precision for purchase-related classes shows solid performance as well: full-price purchases (0.7156), 10% discount purchases (0.7600), 20% discount purchases (0.7794), and high-discount purchases (0.8139). These results suggest that the model is capable of accurately distinguishing customer groups with varying price sensitivities, though there remains room for improvement in full-price and 10% discount categories.

Table 11. Recall evaluation for classes in the Transformer-based pricing and product recommendation model.

Classes	Recall
Class 1—No Purchase	0.8737
Class 2—Purchase at Full Price	0.8763
Class 3—Purchase with 10% Discount	0.8366
Class 4—Purchase with 20% Discount	0.8625
Class 5—Purchase with 30% Discount or More	0.8492

Recall Results. As shown in Table 11, recall is strong across all classes. The model correctly identifies 87.37% of non-purchasing customers, reflecting excellent coverage for customers who do not engage in transactions. Recall is also high for full-price purchases (0.8763), 10% discount (0.8366), 20% discount (0.8625), and high-discount (0.8492) categories. These results indicate that the Transformer model captures the majority of relevant purchasing patterns across different discount levels, ensuring broad behavioral coverage.

Table 12. Macro F-score evaluation for classes in the Transformer-based pricing and product recommendation model.

Classes	Macro F-Score
Class 1—No Purchase	0.9168
Class 2—Purchase at Full Price	0.7878
Class 3—Purchase with 10% Discount	0.7964
Class 4—Purchase with 20% Discount	0.8189
Class 5—Purchase with 30% Discount or More	0.8312

Macro F-score Results. Table 12 presents the Macro F-scores, which combine precision and recall. The No Purchase class achieves the highest F-score (0.9168), confirming strong balanced performance in identifying non-buying behavior. Discount-related classes also show consistent F-scores: full price (0.7878), 10% (0.7964), 20% (0.8189), and 30%+ (0.8312). These results demonstrate that the model performs stably across categories, with slightly better performance for customers requiring moderate or high discounts.

Overall, the Transformer-based model excels at distinguishing non-purchasers from buyers and performs reliably in predicting discount-dependent purchasing behavior. While full-price and low-discount classes exhibit slightly lower precision, the recall and F-scores indicate that the model still captures essential patterns, suggesting that further refinement of feature representation or class-balancing techniques may enhance performance.

Recent work in our previous study [66] proposed a deep reinforcement learning-based model for joint product recommendation and pricing, evaluated in a five-class setting that distinguishes between no purchase and purchases at different discount levels. In the present work, we adopt a Transformer-based sequence learning framework while keeping the same five output classes, which allows a direct, like-for-like comparison between the two approaches.

Table 13 reports the multi-class evaluation metrics for both the RL-based model [66] and the proposed Transformer-based model. The RL approach achieved a Macro Average

Precision of 0.8059, Macro Average Recall of 0.8243, Macro F-score of 0.8108, Micro F-score of 0.8560, Macro Averaged AUC of 0.8743, NDCG@5 of 0.8947, and a Weighted Multi-class Accuracy of 0.8082. In comparison, the Transformer-based model reaches a Macro Average Precision of 0.8066, Macro Average Recall of 0.8597, Macro F-score of 0.8302, Micro F-score of 0.8702, Macro Averaged AUC of 0.8912, NDCG@5 of 0.9065, and a Weighted Multi-class Accuracy of 0.8265.

Table 13. Comparison of multi-class evaluation metrics between the RL-based model [66] and the proposed Transformer-based pricing and product recommendation model. Bold indicates the best score per metric.

Method	Macro Precision	Macro Recall	Macro F-Score	Macro AUC	NDCG@5	Weighted Accuracy
RL-based model [66]	0.8059	0.8243	0.8108	0.8743	0.8947	0.8082
Transformer-based model (this work)	0.8066	0.8597	0.8302	0.8912	0.9065	0.8265

These results show that, in the reported five-class setting, the Transformer-based framework achieves higher scores than the RL-based model across the selected multi-class metrics. The gains are particularly notable in Macro Average Recall (+3.54 percentage points) and Macro F-score (+1.94 percentage points), indicating that the Transformer captures a broader set of relevant purchasing behaviors while maintaining high precision. The increases in Macro Averaged AUC and NDCG@5 demonstrate a stronger ability to rank customer–product–discount combinations according to purchase likelihood and business value. Likewise, the higher Weighted Multi-class Accuracy confirms that the proposed model performs better even when class importance and imbalance are taken into account.

The results in Figure 3 show higher scores for the Transformer-based model across the reported metrics. These differences suggest that self-attention may be useful for representing sequential interaction histories and observed discount-response patterns in this dataset. Overall, these findings support the effectiveness of the proposed Transformer-based pricing and recommendation framework as an advance over our previous RL-based solution in the same five-class setting. Nevertheless, the empirical comparison remains limited. The RL-based model provides the closest available benchmark because it addresses the same product–price recommendation problem and uses the same five-class output structure. However, the present study does not include additional comparisons with recent Transformer-based sequential recommendation baselines. A direct comparison with these models is not straightforward because they are typically designed for next-item recommendation rather than joint product–price prediction with discount-level outputs. Future research should therefore adapt such sequential recommendation architectures to the same five-class product–price setting and evaluate whether the proposed framework maintains its performance advantage against these stronger Transformer-based baselines.

4.3. Error and Misclassification Analysis

To provide a more nuanced interpretation of model performance, we further examined the patterns of misclassification across the five output classes. The most important errors occur in borderline cases between adjacent discount levels, where the behavioral distinction between neighboring classes is naturally less clear. For example, customers predicted to purchase with a 10% discount may sometimes be classified as full-price purchasers or as requiring a 20% discount. Such errors are less severe than misclassifications between non-

purchase and high-discount purchase classes, because adjacent discount classes represent smaller deviations in promotional intensity.




Figure 3. Comparison of multi-class evaluation metrics between the RL-based model [66] and the proposed Transformer-based pricing and product recommendation model in the five-class setting.

From a business perspective, misclassification errors have asymmetric implications. Underestimating the required discount may reduce conversion, while overestimating it may preserve conversion but reduce margin. Therefore, these errors should be interpreted not only as classification mistakes, but also as pricing decisions with different business consequences. In practical deployment, firms may need cost-sensitive rules that penalize over-discounting and under-discounting differently according to product margin, customer segment, and relationship value. Table 14 summarizes the main business implications of different misclassification types in the proposed product–price prediction task. Because errors in discount-level prediction do not have equal consequences, the table distinguishes between errors that may reduce conversion, reduce profit margins, create inefficient targeting, or lead to missed sales opportunities. This interpretation helps connect the model’s classification performance with practical pricing and promotion decisions.

Table 14. Business implications of main misclassification types.

Misclassification Type	Potential Business Implication	Managerial Interpretation
Predicted lower discount than needed	Lower conversion probability	Risk of losing sales due to insufficient incentive
Predicted higher discount than needed	Reduced profit margin	Risk of unnecessary promotional cost
Predicted purchase instead of non-purchase	Inefficient targeting	Recommendation may not generate conversion
Predicted non-purchase instead of purchase	Missed sales opportunity	Potentially valuable customer–product match not activated
Adjacent discount-class error	Moderate pricing deviation	Less severe than errors across distant classes

4.4. Analysis of Experimental Results

The comparison should therefore be interpreted with caution. The RL-based model provides the closest benchmark because it addresses the same product-pricing problem in a related retail setting. However, it is not fully equivalent to the present framework because it follows a policy-learning logic and, in previous comparisons with conventional models, required the output to be simplified into a binary purchase/non-purchase structure. The proposed Transformer-based model retains the five-class structure and is therefore able to distinguish between different discount levels. This distinction is central to the objective of modeling WTP as discount-level acceptance rather than merely predicting whether a purchase will occur.

The proposed Transformer-based model for joint product recommendation and pricing indicates consistently stable performance across all five output classes. The multi-class evaluation confirms that the model accurately predicts both non-purchasing behavior and purchases at varying discount levels, validating the feasibility and effectiveness of integrating pricing directly into the recommendation output. The ability to distinguish among five observed outcome classes indicates that the model captures patterns of purchase behavior and discount-level acceptance, rather than directly estimating true latent WTP.

A key advantage of this work is that it allows a direct comparison with our previous RL-based model [66], which was evaluated under the same five-class structure. As shown in Table 13, the new Transformer-based framework consistently achieves higher scores than the RL-based approach across all major multi-class metrics, including Macro Precision, Macro Recall, Macro F-score, Macro AUC, and NDCG@5.

From a marketing perspective, accurately predicting purchasing behavior across different discount levels allows for more precise promotional targeting. By identifying which customers are responsive to full-price offers, moderate discounts, or high discounts, the model supports finer customer segmentation and reduces marketing inefficiency. The high recall achieved across classes ensures that relevant customer segments are not overlooked, thereby maximizing the reach and impact of promotional campaigns. At the same time, individualized discounts that reflect observed price sensitivity may be perceived as more relevant by some customers. However, positive effects on satisfaction, engagement, and long-term loyalty should be treated as conditional rather than automatic, since they depend on perceived fairness, transparency, and the presence of customer-relationship safeguards.

In terms of short-term conversion and promotional efficiency, the model's ability to forecast the discount level associated with purchase outcomes can help businesses deploy more targeted pricing strategies. By avoiding indiscriminate deep discounts and directing stronger incentives to highly price-sensitive cases, firms may reduce promotional overspending and preserve margins. However, this revenue-oriented use of the model should not be interpreted as an unrestricted strategy for extracting surplus from high-WTP customers. If high-WTP customers are also loyal or frequent buyers, systematically offering them less favorable prices than other customers may generate perceptions of price unfairness and weaken trust. Therefore, the model's predictions should be combined with fairness guardrails and long-term customer-relationship metrics before being implemented in practice.

Despite the strong results, some challenges remain in differentiating between adjacent discount categories, particularly those with similar behavioral profiles. Future work may focus on enhancing feature representation through clustering-based customer and product embeddings, incorporating richer temporal patterns, or integrating contextual signals such as seasonal effects and product attributes. Additionally, leveraging larger datasets, tuning model hyperparameters, and exploring hierarchical sequence encoders may further improve predictive granularity. These refinements would strengthen the model's ability to discriminate among closely related discount levels, leading to even more precise and effective personalized pricing strategies.

5. Conclusions and Future Works

This research presented a Transformer-based framework for personalized product recommendation and pricing. By incorporating observed discount-level acceptance as an operational proxy for WTP, the proposed model can support product-price matching and more targeted promotional decision-making. The empirical results show consistent performance across Precision, Recall, F1-score, AUC, and NDCG, indicating that the model can predict observed purchase-discount outcomes and rank product-price combinations

according to relevance and discount-level acceptance. However, these results should be interpreted as predictive rather than causal, and their managerial implications should distinguish between short-term conversion efficiency and long-term customer relationship outcomes.

Theoretically, this study contributes to the RS literature by reframing product recommendation as a joint product–price prediction problem. While most existing RS approaches focus mainly on identifying products that users are likely to prefer or purchase, the proposed framework incorporates discount level directly into the prediction output. This extends price-aware and value-aware recommendation research by modeling observed purchase–discount behavior in a multi-class prediction setting. Methodologically, the study shows how Transformer-based self-attention can be used to model customer–product interaction histories and observed price responsiveness within a unified sequence-learning framework.

From a practical perspective, the framework can help retailers and e-commerce platforms design more targeted promotion strategies. Instead of applying uniform discounts to broad customer groups, the model can identify cases associated with full-price purchase, moderate discounts, or higher discount levels. This may support promotional-budget allocation, reduce indiscriminate discounting, contribute to margin protection, and improve the relevance of product–price offers, provided that the system is implemented with fairness safeguards and validated in real-world settings. Thus, the framework bridges traditional product recommendation and DP by supporting revenue-related decision-making together with customer relationship management.

At the same time, personalized pricing should be implemented with caution. Customers with high predicted discount-level acceptance for full-price purchases may include loyal or frequent customers. If such customers are systematically offered less favorable prices than less loyal or more price-sensitive customers, the strategy may generate perceptions of price unfairness and betrayal, reducing trust and weakening long-term loyalty [6–8]. Therefore, the proposed model should be understood as a decision-support tool rather than a fully autonomous pricing mechanism. Firms should combine model predictions with managerial and ethical safeguards, including limits on personalized price differences, protection against systematic disadvantage for loyal customers, monitoring of complaints and churn, and integration of long-term relationship metrics such as CLV [61]. It should also be noted that fairness is discussed here as a managerial and ethical implementation consideration, not as an experimentally optimized model objective. The present framework does not include fairness-aware constraints or fairness metrics in training or evaluation, and future research should explicitly incorporate and test such mechanisms.

Several limitations should be acknowledged. First, the model’s output space is restricted to four fixed discount levels plus a no-purchase option. Although these thresholds reflect common stepwise discounting practices and support managerial interpretability, they do not fully capture the continuous nature of pricing decisions. Future studies could explore finer-grained or continuous price recommendations through regression heads, ordinal classification, or mixture-density outputs. Second, the empirical evaluation is based on the Dunhumby retail dataset. Customer responses to personalized pricing may differ across product categories, markets, and platform environments; therefore, the framework should be tested on additional domains such as fashion retail, grocery e-commerce, travel, or subscription-based services. Third, the offline nature of the data creates an inherent counterfactual limitation. The model observes only the prices and discounts that were actually offered, while alternative price points remain unobserved. Therefore, predicted discount classes should be interpreted as estimates of observed discount-level acceptance rather than direct measures of true WTP. The model may partly reflect the historical pric-

ing policy embedded in the data, and offline evaluation may overstate expected revenue gains if previous discount exposure was biased. Future validation through A/B testing, randomized discount exposure, or off-policy evaluation is therefore needed.

Future work can extend the framework in several directions. First, richer temporal, behavioral, and contextual variables—such as promotion calendars, seasonality, competitive prices, inventory constraints, loyalty-program status, browsing behavior, customer life-cycle stages, and CLV—could be incorporated to improve predictive granularity. Second, the model should be compared with adapted sequential recommendation baselines, including SASRec, BERT4Rec, and other modern Transformer-based architectures, reformulated for the same product–price prediction setting. Third, hybrid approaches could combine Transformer-based sequence encoders with demand forecasting, CLV estimation, profit-based decision layers, reinforcement learning, bandit-based methods, or fairness-aware constraints. Finally, periodic or near real-time updating mechanisms should be explored so that recommendation and pricing policies can adapt to changing customer behavior and market conditions while monitoring fairness perceptions, complaints, churn, repeat purchasing, and long-term customer value.

In summary, this research establishes a foundation for Transformer-based, price-aware RS in a multi-class product–price prediction setting. Future studies should refine the modeling components, test the framework across additional real-world domains, operationalize fairness-aware evaluation, and validate the system through online or experimental methods. These extensions can help develop data-driven marketing and pricing strategies that are not only value-oriented, but also more transparent, responsible, and sustainable.

Author Contributions: Conceptualization, A.M. and B.B.; methodology, A.M. and B.B.; software, A.M.; validation, A.M., B.B. and H.M.; investigation, A.M. and B.B.; writing—original draft preparation, A.M.; writing—review and editing, A.M.; visualization, A.M.; supervision, B.B. and H.M.; project administration, B.B. and H.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: All of the real graphs used in this paper are publicly available as described in the body of the paper.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. StartUp Guru Lab. Customer Acquisition vs. Retention Costs. 2025. Available online: <https://startupgurulab.com/customer-acquisition-vs-retention-costs> (accessed on 6 August 2025).
2. Optimove. Customer Acquisition vs. Retention Costs: Why Retention Is More Profitable. 2025. Available online: <https://www.optimove.com/resources/learning-center/customer-acquisition-vs-retention-costs> (accessed on 6 August 2025).
3. Gallo, A. The Value of Keeping the Right Customers. Harvard Business Review. 2014. Available online: <https://hbr.org/2014/10/the-value-of-keeping-the-right-customers> (accessed on 6 August 2025).
4. Lu, C.C.; Wu, I.L.; Hsiao, W.H. Developing customer product loyalty through mobile advertising: Affective and cognitive perspectives. *Int. J. Inf. Manag.* **2019**, *47*, 101–111. [CrossRef]
5. Lerner, J.S.; Li, Y.; Valdesolo, P.; Kassam, K.S. Emotion and decision making. *Annu. Rev. Psychol.* **2015**, *66*, 799–823. [CrossRef] [PubMed]
6. Xia, L.; Monroe, K.B.; Cox, J.L. The Price Is Unfair! A Conceptual Framework of Price Fairness Perceptions. *J. Mark.* **2004**, *68*, 1–15. [CrossRef]
7. Feinberg, F.M.; Krishna, A.; Zhang, Z.J. Do We Care What Others Get? A Behaviorist Approach to Targeted Promotions. *J. Mark. Res.* **2002**, *39*, 277–291. [CrossRef]
8. Tsai, D.; Lee, H.-T. Will You Care When You Pay More? The Negative Side of Targeted Promotions. *J. Prod. Brand Manag.* **2007**, *16*, 481–491. [CrossRef]

9. Kumar, V.; Ashraf, A.R.; Nadeem, W. AI-powered marketing: What, where, and how? *Int. J. Inf. Manag.* **2024**, *77*, 102783. [CrossRef]
10. Kshetri, N.; Dwivedi, Y.K.; Davenport, T.H.; Panteli, N. Generative artificial intelligence in marketing: Applications, opportunities, challenges, and research agenda. *Int. J. Inf. Manag.* **2024**, *75*, 102716. [CrossRef]
11. Nouri-Harzvili, M.; Hosseini-Motlagh, S.M. Dynamic discount pricing in online retail systems: Effects of post-discount dynamic forces. *Expert Syst. Appl.* **2023**, *232*, 120864. [CrossRef]
12. Kim, S.; Park, E. STAD-GCN: Spatial–Temporal Attention-based Dynamic Graph Convolutional Network for retail market price prediction. *Expert Syst. Appl.* **2024**, *255*, 124553. [CrossRef]
13. Quadrana, M.; Cremonesi, P.; Jannach, D. Sequence-aware recommender systems. *ACM Comput. Surv.* **2018**, *51*, 1–36. [CrossRef]
14. Zheng, Y.; Wang, D. A survey of recommender systems with multi-objective optimization. *Neurocomputing* **2022**, *474*, 141–153. [CrossRef]
15. Jiang, Y.; Liu, Y. Optimization of online promotion: A profit-maximizing model integrating price discount and product recommendation. *Int. J. Inf. Technol. Decis. Mak.* **2012**, *11*, 961–982. [CrossRef]
16. De Biasio, A.; Montagna, A.; Aiolfi, F.; Navarin, N. A systematic review of value-aware recommender systems. *Expert Syst. Appl.* **2023**, *226*, 120131. [CrossRef]
17. Sitar-Tăut, D.A.; Mican, D.; Mare, C. Customer behavior in the prior purchase stage–information search versus recommender systems. *Econ. Comput. Econ. Cybern. Stud. Res.* **2020**, *54*, 59–76. [CrossRef]
18. Zhang, H.; Zhao, L.; Gupta, S. The role of online product recommendations on customer decision making and loyalty in social shopping communities. *Int. J. Inf. Manag.* **2018**, *38*, 150–166. [CrossRef]
19. McKinsey & Company. How Retailers Can Keep Up with Consumers. Available online: <https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers> (accessed on 9 September 2025).
20. Chen, K.; Chen, T.; Zheng, G.; Jin, O.; Yao, E.; Yu, Y. Collaborative personalized tweet recommendation. In Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, Portland, OR, USA, 12–16 August 2012; pp. 661–670.
21. Beheshti, A.; Yakhchi, S.; Mousaeirad, S.; Ghafari, S.M.; Goluguri, S.R.; Edrisi, M.A. Towards Cognitive Recommender Systems. *Algorithms* **2020**, *13*, 176. [CrossRef]
22. Ramampiaro, H.; Langseth, H.; Almenningen, T.; Schistad, H.; Havig, M.; Nguyen, H.T. New Ideas in Ranking for Personalized Fashion Recommender Systems. In *Business and Consumer Analytics: New Ideas*; Springer: Berlin/Heidelberg, Germany, 2019. [CrossRef]
23. Cantador, I.; Bellogín, A.; Castells, P. Second Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec 2011). In Proceedings of the RecSys’11, Chicago, IL, USA, 23–27 October 2011.
24. Jannach, D.; Pu, P.; Ricci, F.; Zanker, M. Recommender Systems: Past, Present, Future. *AI Mag.* **2021**, *42*, 3–6. [CrossRef]
25. Livne, A.; Unger, M.; Shapira, B.; Rokach, L. Deep Context-Aware Recommender System Utilizing Sequential Latent Context. In *Proceedings of the 13th ACM Conference on Recommender Systems—CARS Workshop, Copenhagen, Denmark, 16–20 September 2019*; ACM: New York, NY, USA, 2019.
26. Zhao, W.X.; Guo, Y.; He, Y.; Jiang, H.; Wu, Y.; Li, X. We Know What You Want to Buy: A Demographic-Based System for Product Recommendation on Microblogs. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD’14), New York, NY, USA, 24–27 August 2014*; ACM: New York, NY, USA, 2014; pp. 1935–1944. [CrossRef]
27. Tkalčič, M.; De Carolis, B.; de Gemmis, M.; Odić, A.; Košir, A. *Emotions and Personality in Personalized Services: Models, Evaluation and Applications*; Springer: Berlin/Heidelberg, Germany, 2016. [CrossRef]
28. Bai, T.; Wang, X.; Zhang, Z.; Song, W.; Wu, B.; Nie, J.Y. GPR-OPT: A Practical Gaussian optimization criterion for implicit recommender systems. *Inf. Process. Manag.* **2024**, *61*, 103525. [CrossRef]
29. Pujahari, A.; Sisodia, D.S. Modeling users’ preference changes in recommender systems via time-dependent Markov random fields. *Expert Syst. Appl.* **2023**, *234*, 121072. [CrossRef]
30. Liu, D.; Li, J.; Wu, J.; Du, B.; Chang, J.; Li, X. Interest Evolution-driven Gated Neighborhood aggregation representation for dynamic recommendation in e-commerce. *Inf. Process. Manag.* **2022**, *59*, 102982. [CrossRef]
31. Zhang, P.; Niu, Z.; Ma, R.; Zhang, F. Multi-view graph contrastive representation learning for bundle recommendation. *Inf. Process. Manag.* **2025**, *62*, 103956. [CrossRef]
32. Beladev, M.; Rokach, L.; Shapira, B. Recommender Systems for Product Bundling. *Knowl.-Based Syst.* **2016**, *111*, 193–206. [CrossRef]
33. Kouki, P.; Fountalis, I.; Vasiloglou, N.; Yan, N.; Ahsan, U.; Al Jadda, K.; Qu, H. Product Collection Recommendation in Online Retail. In *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys ’19), Copenhagen, Denmark, 16–20 September 2019*; ACM: New York, NY, USA, 2019; pp. 486–490. [CrossRef]
34. Wu, S.; Sun, F.; Zhang, W.; Xie, X.; Cui, B. Graph Neural Networks in Recommender Systems: A Survey. *arXiv* **2020**, arXiv:2011.02260. [CrossRef]

35. Dara, S.; Chowdary, C.R.; Kumar, C. A Survey on Group Recommender Systems. *J. Intell. Inf. Syst.* **2019**, *54*, 271–295. [[CrossRef](#)]
36. Zhou, H.; Zhang, J.; Zhang, Z.; Tang, R.; Wang, X. A Comprehensive Survey of Recommender Systems with Deep Learning. *Appl. Sci.* **2023**, *13*, 11378. [[CrossRef](#)]
37. Gheewala, S.; Patel, N.; Shah, S. In-Depth Survey: Deep Learning in Recommender Systems—Trends and Challenges. *Neural Comput. Appl.* **2025**, *37*, 10875–10947. [[CrossRef](#)]
38. Kang, W.C.; McAuley, J. Self-Attentive Sequential Recommendation. In Proceedings of the 2018 IEEE International Conference on Data Mining (ICDM), Singapore, 17–20 November 2018; pp. 197–206. [[CrossRef](#)]
39. Sun, F.; Liu, J.; Wu, J.; Pei, C.; Lin, X.; Ou, W.; Jiang, P. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM '19), Beijing, China, 3–7 November 2019; pp. 1441–1450. [[CrossRef](#)]
40. Zhou, K.; Wang, H.; Zhao, W.X.; Zhu, Y.; Wang, S.; Zhang, F.; Wang, Z.; Wen, J.R. S³-Rec: Self-Supervised Learning for Sequential Recommendation with Mutual Information Maximization. In Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20), Virtual Event, Ireland, 19–23 October 2020; pp. 1893–1902. [[CrossRef](#)]
41. Gheewala, S.; Shah, S. Exploiting Deep Transformer Models in Textual Review Based Recommender Systems. *Expert Syst. Appl.* **2024**, *238*, 122141. [[CrossRef](#)]
42. Khan, H.U.; Ullah, I.; Ali, Z.; Kim, Y. A Transformer-based Architecture for Collaborative Filtering Modeling in Personalized Recommender Systems. *Sci. Rep.* **2025**, *15*, 12345. [[CrossRef](#)] [[PubMed](#)]
43. Soleiman Nejad, M.; Varasteh, M.; Moradi, H.; Sadeghi, M.A. Designing a Sequential Recommendation System for Heterogeneous Interactions Using Transformers. *arXiv* **2022**, arXiv:2205.00265.
44. Santana, M.R.O.; Soares, A. Hybrid Model with Time Modeling for Sequential Recommender Systems. In Proceedings of the ACM WSDM Workshop on Web Tourism (WSDM WebTour'21), Jerusalem, Israel, 12 March 2021.
45. Petrov, A.; Macdonald, C. A Systematic Review and Replicability Study of BERT4Rec for Sequential Recommendation. In Proceedings of the 16th ACM Conference on Recommender Systems (RecSys '22), Seattle, WA, USA, 18–23 September 2022; pp. 436–447. [[CrossRef](#)]
46. Hwangbo, H.; Kim, Y.S.; Cha, K.J. Recommendation system development for fashion retail e-commerce. *Electron. Commer. Res. Appl.* **2018**, *28*, 94–101. [[CrossRef](#)]
47. Kent, R.J.; Monroe, K.B. The Effects of Framing Price Promotion Messages on Consumers' Perceptions and Purchase Intentions. *J. Retail.* **1998**, *74*, 353–372. [[CrossRef](#)]
48. Umberto, P. Developing a price-sensitive recommender system to improve accuracy and business performance of ecommerce applications. *Int. J. Electron. Commer. Stud.* **2015**, *6*, 1–18. [[CrossRef](#)]
49. Lopatovska, I.; Mokros, H.B. Willingness to pay and experienced utility as measures of affective value of information objects: Users' accounts. *Inf. Process. Manag.* **2008**, *44*, 92–104. [[CrossRef](#)]
50. Wakil, K.; Alyari, F.; Ghasvari, M.; Lesani, Z.; Rajabion, L. A new model for assessing the role of customer behavior history, product classification, and prices on the success of the recommender systems in e-commerce. *Kybernetes* **2020**, *49*, 1325–1346. [[CrossRef](#)]
51. Kompan, M.; Gaspar, P.; Macina, J.; Cimerman, M.; Bielikova, M. Exploring Customer Price Preference and Product Profit Role in Recommender Systems. *IEEE Intell. Syst.* **2022**, *37*, 89–98. [[CrossRef](#)]
52. Jannach, D.; Adomavicius, G. Price and Profit Awareness in Recommender Systems. *arXiv* **2017**, arXiv:1707.08029.
53. Sato, M.; Izumo, H.; Sonoda, T. Discount sensitive recommender system for retail business. In Proceedings of the EMPIRE '15: 3rd Workshop on Emotions and Personality in Personalized Systems 2015, Vienna, Austria, 16–20 September 2015; pp. 33–40. [[CrossRef](#)]
54. Das, A.; Mathieu, C.; Ricketts, D. Maximizing profit using recommender systems. *arXiv* **2009**, arXiv:0908.3633.
55. Sato, M.; Izumo, H.; Sonoda, T. Model of Personal Discount Sensitivity in Recommender Systems. *arXiv* **2016**. Available online: https://ixdea.org/28_6/ (accessed on 28 October 2025).
56. Terui, N.; Dahana, W.D. Price customization using price thresholds estimated from scanner panel data. *J. Interact. Mark.* **2006**, *20*, 58–70. [[CrossRef](#)]
57. Greenstein-Messica, A.; Rokach, L. Personal price aware multi-seller recommender system: Evidence from eBay. *Knowl.-Based Syst.* **2018**, *150*, 14–26. [[CrossRef](#)]
58. Chen, J.; Jin, Q.; Zhao, S.; Bao, S.; Zhang, L.; Su, Z.; Yu, Y. Does product recommendation meet its Waterloo in unexplored categories? No, price comes to help. In Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, Gold Coast, QLD, Australia, 6–11 July 2014; pp. 667–676. [[CrossRef](#)]
59. Zheng, Y.; Gao, C.; He, X.; Li, Y.; Jin, D. Price-aware recommendation with graph convolutional networks. In Proceedings of the Proceedings—International Conference on Data Engineering, Dallas, TX, USA, 20–24 April 2020; pp. 133–144. [[CrossRef](#)]
60. Shin, D.; Zhong, B.; Biocca, F.A. Beyond user experience: What constitutes algorithmic experiences? *Int. J. Inf. Manag.* **2020**, *52*, 102061. [[CrossRef](#)]

61. Heidary, K.; Pluut, H. All is (not) fair in personalized pricing: Antecedents and outcomes of consumer fairness perceptions. *J. Revenue Pricing Manag.* **2025**, *24*, 204–221. [[CrossRef](#)]
62. Alijani, S.; Mohammadi, A.; Yousefi, M.R. Vision Transformers in Domain Adaptation and Domain Generalization: A Study of Robustness. *arXiv* **2024**, arXiv:2404.04452. [[CrossRef](#)]
63. Grandini, M.; Bagli, E.; Visani, G. Metrics for Multi-Class Classification: An Overview. *arXiv* **2020**, arXiv:2008.05756.
64. Hand, D.J.; Till, R.J. A Simple Generalisation of the Area Under the ROC Curve for Multiple Class Classification Problems. *Mach. Learn.* **2001**, *45*, 171–186. [[CrossRef](#)]
65. dunnhumby. Available online: <https://www.dunnhumby.com/> (accessed on 9 September 2025).
66. Mahdavian, A.; Moradi, H.; Bahrak, B. Product Recommendation with Price Personalization According to Customer's Willingness to Pay Using Deep Reinforcement Learning. *Algorithms* **2025**, *18*, 706. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.