A Temporal-Weighted Hybrid Recommender for B2B Vehicle Auctions Using Word2Vec Embeddings

Selçuk Bayracı [1*], Uğur Barış Öztürk [2], Turgay Tugay Bilgin [3]

[1] Borusan Otomotiv R&D Center, Orcid ID: https://orcid.org/0000-0003-4831-4802, selcuk.bayraci @borusanotomotiv.com
[2] Borusan Otomotiv R&D Center, Orcid ID: https://orcid.org/0009-0001-9129-7334, ubaris.ozturk @borusanotomotiv.com
[3] Bursa Technical University, Orcid ID: https://orcid.org/ 0000-0002-9245-5728, turgay.bilgin@btu.edu.tr
* Correspondence: selcuk.bayraci@borusanotomotiv.com ; +90 532 2020975

Received: 21 June 2025
Revised: 23 November 2025
Accepted: 16 December 2025
Published: 19 December 2025

Reference: Bayracı, S., Öztürk, U. B., & Bilgin, T. T. (2025). A temporal-weighted hybrid recommender for B2B vehicle auctions using Word2Vec embeddings. *The European Journal of Research and Development, 5*(1), 540–566.

Abstract

Used car auction platforms face unique challenges in personalized recommendation due to extreme data sparsity, high inventory turnover, and real-time operational constraints. This study develops and evaluates a hybrid recommendation system combining Word2Vec embeddings for categorical vehicle attributes with standardized numerical features, applying temporal decay weighting to prioritize recent user interactions. Deployed on Azure infrastructure, the system was evaluated using 12 months of transaction data from a Turkish B2B auction platform comprising 5,322 users, 24,987 vehicles, and 1.87 million interactions. Offline evaluation demonstrates superior performance over baselines (Hit Rate@10: 0.456 vs 0.234 popularity baseline, 94.9% improvement). Production deployment over six months (April–September 2025) generated 977 recommendation-driven sales representing 15.26% of total platform transactions and 17.27M TL in commission revenue. Quasi-experimental analysis revealed a 26.7% increase in monthly purchase frequency among active users, yielding 420 incremental transactions. Results *demonstrate how interpretable temporal-weighted embedding models generate measurable commercial value in high-turnover, data-sparse B2B marketplaces.*

Keywords: Recommendation systems, Vehicle auctions, Word2Vec embeddings, Temporal weighting, B2B analytics.

1. Introduction

The accelerating digital transformation of the automotive industry has fundamentally reshaped how vehicles are bought and sold. Online platforms have become the primary medium through which buyers and sellers interact, demanding seamless, data-driven, and personalized user experiences. In this evolving landscape, recommendation systems play a pivotal role in aligning user preferences with available inventory, thereby increasing engagement, satisfaction, and conversion efficiency.

Borusan Araç İhale, one of Turkey's leading used car auction platforms, operates within a fast-paced and high-turnover environment where vehicle listings change daily. This reflects a well-documented challenge in the automotive domain, where the sheer volume of available information and similar models can overwhelm customers, making the selection process complex and challenging [1]. Unlike fixed-price e-commerce systems, auction platforms face unique challenges: every listed vehicle is distinct, user interactions are sparse and highly dynamic, and recommendations must be generated in real time to maintain user attention. Traditional collaborative filtering or static recommendation models fall short in addressing these conditions, underscoring the need for adaptive, temporally aware, and scalable solutions.

To address these challenges, this study presents a hybrid recommendation engine that integrates collaborative filtering with content-based similarity modeling, optimized for the operational and data characteristics of a real-world vehicle auction marketplace. This hybrid methodology is consistent with a growing consensus in the literature that combining different recommendation techniques is necessary to overcome the individual limitations of any single approach, particularly the cold-start and sparsity problems [2, 3]. The system captures user behavioral signals such as vehicle views, offers, favorites, and purchases, applying temporal decay weighting to prioritize recent interactions. On the item side, vehicle representations are enriched using semantic embeddings derived from categorical features and normalized numerical descriptors, creating a unified vector space that facilitates robust similarity computations.

The proposed approach was deployed within Borusan Araç İhale's Azure-based architecture, utilizing Azure SQL for structured data, Azure Synapse Analytics for scalable data management, and Azure Machine Learning Studio for model orchestration and version control. The model continuously processes customer interactions and inventory updates to deliver personalized recommendations in real time.

Experimental results, based on over one year of user and auction data, demonstrate the system's significant commercial impact: from April to September 2024, approximately 15% of vehicles sold through the platform were recommended by the model, corresponding to a total commission revenue exceeding 17 million Turkish Liras. Beyond its business outcomes, the study illustrates how hybrid AI architectures can enable user-centric personalization in fast-moving, data-sparse markets such as automotive auctions.

This work makes several contributions to the recommendation systems literature and practice:

1. **Novel application domain:** First comprehensive deployment study of a hybrid recommender in a real-world B2B vehicle auction environment, demonstrating effectiveness in handling extreme data sparsity (1.41% interaction density) and high inventory turnover.

2. **Temporal-weighted hybrid architecture:** Integration of Word2Vec embeddings for categorical features with temporal decay weighting, providing an interpretable and computationally efficient alternative to deep learning approaches.

3. **Multi-stage candidate filtering:** Adaptive filtering hierarchy that balances recommendation specificity with candidate availability in dynamic inventory environments.

4. **Empirical business validation:** Comprehensive evaluation combining offline metrics, online deployment results, and quasi-experimental causal analysis, demonstrating measurable commercial impact (7.14M TL incremental revenue).

5. **Scalable production deployment:** Azure-based architecture enabling real-time recommendations with minimal latency for thousands of concurrent users in a production auction setting.

The remainder of this paper is organized as follows: Section 2 provides a comprehensive review of the relevant literature. Section 3 outlines the problem context and system objectives. Section 4 details the methodological framework and model design. Section 5 presents the experimental setup. Section 6 reports evaluation results. Section 7 discusses findings, implications, and limitations. Section 8 concludes.

2. Related Work

This section reviews foundational and contemporary approaches to recommendation systems, with specific focus on hybrid methods, temporal dynamics, embedding-based techniques, and automotive domain applications.

2.1 Hybrid Recommendation Systems

The recommendation systems literature traditionally distinguishes between content-based (CB) filtering and collaborative filtering (CF) approaches. Content-based methods recommend items by matching their attributes with user profiles built from past preferences, offering advantages in handling new items and avoiding item cold-start problems. Collaborative filtering recommends items based on preferences of similar users, leveraging collective behavioral patterns but suffering from cold-start problems for new users and performance degradation under data sparsity.

Recent research increasingly focuses on hybrid recommendation systems that combine multiple techniques to overcome individual limitations. El Asri et al. (2025) conducted a systematic literature review of 5,857 publications from 2020-2024, identifying hybrid approaches as the dominant paradigm for addressing data sparsity and cold-start challenges [2]. Their analysis reveals that hybrid methods consistently outperform single-technique approaches across diverse domains.

Yamarthi et al. (2024) proposed the HRS-IU-DL model, integrating user-based and item-based collaborative filtering with neural collaborative filtering (NCF) and recurrent neural networks (RNN) to capture complex non-linear relationships and sequential patterns [3]. Their approach demonstrated significant improvements in addressing data sparsity and cold-start issues through the fusion of multiple deep learning techniques.

Lv et al. (2024) developed a hybrid algorithm based on the user-nearest-neighbor model, effectively mitigating data sparsity through integration with other recommendation approaches [4]. Utilizing the Spark distributed platform for scalability, their method demonstrated superiority over standalone collaborative filtering algorithms across various recommendation indicators.

In the e-learning domain, recent work has addressed the pure cold-start scenario through hybrid attribute-based approaches [5]. These systems merge knowledge-based methods with collaborative filtering, utilizing Rogers-Tanimoto similarity measures for materials and users, and Jaccard similarity for user comparisons. Unlike traditional methods relying on prior ratings, these approaches depend on attributes, demonstrating effectiveness in scenarios with minimal historical data.

2.2 Temporal Dynamics in Recommendation Systems

User preferences are dynamic and typically change over time, making temporal modeling essential for recommendation accuracy. Koren (2009) pioneered the integration of temporal dynamics in collaborative filtering, demonstrating that modeling time-varying user biases, item biases, and preference changes significantly improves prediction accuracy [6].

Recent advances have focused on sophisticated temporal modeling approaches. Kim et al. (2024) proposed a Temporal Graph Network (TGN) framework for dynamic recommendation, employing GRU-based history embedding to capture long-term historical dependencies while computing up-to-date embeddings through neighbor aggregation [7]. Their approach addresses the limitation of traditional models that treat time as merely a linear or categorical feature.

The TCGC (Temporal Collaboration-Aware Graph Co-Evolution) framework introduced collaboration-aware indicators to guide evolution learning in dynamic recommendation systems [8]. The model captures correlation between event-level and history-level dynamics, enabling multi-round recommendation that continuously refreshes results based on real-time user feedback.

Ben Abdrabbah et al. (2024) developed a dynamic community-based recommender framework using temporal networks to capture evolving user preferences [9]. Applying dynamic community detection techniques addresses scalability issues while providing personalized recommendations for both individuals and groups. Their experimental results demonstrated superior performance over state-of-the-art approaches in recommendation accuracy.

Li and Tuzhilin (2024) formulated recommendations as a consumer trajectory learning task based on dynamical systems theory [10]. Their framework explicitly models sequential evolution of individual preferences and systematically captures long-term impact of exploratory products on future behavior, addressing inadequate temporal dynamics support in existing exploration approaches.

2.3 Word2Vec and Embedding-Based Methods

Word2Vec, originally developed for natural language processing, has found successful applications in recommendation systems through its ability to learn distributed representations of items. Ozsoy (2016) pioneered the application of Word2Vec to recommendation using non-textual features (check-ins), demonstrating that continuous vector space representations of items offer promising results for venue recommendations [11].

Caselles-Dupré et al. (2018) investigated hyperparameter optimization for Word2Vec in recommendation contexts, revealing that optimizing neglected parameters (negative sampling distribution, epochs, subsampling, window size) significantly improves performance on recommendation tasks [12]. Their findings demonstrated that using default NLP hyperparameters without domain-specific tuning significantly limits effectiveness.

Ordentlich et al. (2020) further explored hyperparameter optimization for large-scale recommendation systems, showing that unconstrained optimization yields 221% average improvement in hit rate over default parameters, while runtime budget-constrained optimization still achieves 138% improvement [13]. Their work provides practical guidance for deploying Word2Vec-based recommendations at scale.

Recent work on enhancing word embeddings has addressed limitations of traditional approaches. Szymański et al. (2024) introduced supervised learning methods that shift vectors in representation space to ensure better alignment with semantic reference resources, demonstrating improved accuracy and efficiency in text classification and clustering tasks applicable to recommendation scenarios [14].

2.4 Automotive Domain Applications

The automotive market presents unique challenges for recommendation systems due to information overload, infrequent purchases, and high-value decisions. Early work by Talati et al. (2016) surveyed recommendation approaches for automobile purchasing, identifying the critical need for systems that reduce information load while matching customer requirements [15].

Thomas and Vaidhehi (2018) designed a hybrid algorithm combining user-to-user and item-to-item collaborative filtering using click data and browsing history [16]. Prabowol et al. (2019) proposed a hybrid linear combination merging item-based CF with content-based approaches using vehicle specification attributes [17]. Liu et al. (2023) developed a hybrid strategy analyzing association rules derived from pairwise comparison data, where users compare vehicles side-by-side [18].

Xu (2021) designed a hybrid system specifically for an online car auction platform, combining content-based subsystems with matrix factorization [19]. Le et al. (2021) proposed an ontology-based system representing vehicle parameters in machine-readable knowledge bases, enabling reasoning about vehicle similarity even with sparse interaction data [20]. Boteju and Munasinghe (2020) employed natural language processing to analyze user reviews and preferences, enhancing semantic understanding in automotive recommendations [21].

Shi et al. (2022) highlighted challenges in the used car market including limited inventory and dynamic user-item connections [22]. Their Resource-Constraint Aware Neural Network (RecNet) employs recurrent neural networks to capture temporal dynamics often ignored by traditional models, demonstrating effectiveness in handling evolving preference patterns.

2.5 Deep Learning Approaches for Recommendations

Deep learning has emerged as a powerful approach for capturing complex patterns in recommendation data. Fu et al. (2019) introduced a novel deep learning-based collaborative filtering model demonstrating significant improvements over traditional matrix factorization methods through its ability to model non-linear user-item interactions [23].

Recent work has focused on multi-criteria recommendation systems. A hybrid DeepFM-SVD++ model integrating deep learning and factorization-based techniques captures both low-order feature interactions using factorization machines and high-order dependencies through deep neural networks [24]. Evaluation on technology-enhanced learning and movie datasets demonstrated superior performance over traditional approaches.

Di et al. (2025) proposed a federated recommender system using fine-grained transformation and hybrid information sharing to address privacy concerns in personalized recommendations [25]. Their approach enables collaborative learning across distributed data sources while preserving user privacy, representing an important direction for recommendation systems in regulated industries.

2.6 Evaluation and Business Impact

Beyond algorithmic performance, understanding long-term dynamics and business impact of recommendation systems requires sophisticated evaluation approaches. Li and Tuzhilin (2024) developed a dynamical system framework for exploring consumer trajectories, enabling systematic capture of long-term effects of exploratory recommendations on future user behavior [10].

Agent-based simulation provides complementary insights into longitudinal recommendation dynamics. Research in Information Systems demonstrates that user reliance on recommendations tends to decrease system usefulness over time, revealing a longitudinal performance paradox where short-term gains may compromise long-term effectiveness [26]. This finding emphasizes the importance of diversity and exploration in recommendation strategies.

2.7 Research Gap and Positioning

The literature demonstrates clear progression from foundational CF/CB models to sophisticated hybrid systems incorporating temporal dynamics and deep learning. However, gaps remain in several areas:

1. **Limited auction platform research:** While some studies focus on auctions [19] and others on dynamic used-car environments [22], comprehensive deployment studies combining real-time auction dynamics with temporal-weighted hybrid approaches are scarce.
2. **Interpretability versus performance trade-off:** Deep learning methods achieve strong performance but lack interpretability critical for business stakeholders. Embedding-based approaches like Word2Vec offer competitive performance with greater transparency.
3. **Sparse data challenges:** Despite advances, effectively handling extreme sparsity (< 2% interaction density) in high-turnover environments remains challenging, particularly when traditional collaborative filtering becomes unreliable.

4. **Business impact validation:** Most studies focus on offline metrics or short-term online experiments, lacking comprehensive evaluation of long-term business impact and causal attribution of recommendation effects.

This paper addresses these gaps by presenting a deployed hybrid system integrating Word2Vec embeddings with temporal weighting in a real-world B2B auction environment. The approach balances performance with interpretability while demonstrating measurable long-term business impact through comprehensive evaluation combining offline metrics, online deployment results, and quasi-experimental causal analysis.

3. Problem Context and System Objectives

3.1 Problem Context

In the rapidly evolving digital automotive ecosystem, used car auctions present an especially challenging operational environment for data-driven personalization. Unlike traditional e-commerce settings where products are standardized and remain available over extended periods, each auction vehicle is a unique and transient asset. Listings frequently enter and exit the marketplace within days, resulting in a continuously shifting item pool and making the reuse of historical patterns difficult.

Moreover, interaction data are inherently sparse and asynchronous. Many users participate intermittently, and their actions, such as viewing, bidding, or marking a vehicle as a favorite, occur in short bursts corresponding to specific auction events. This sparsity reduces the statistical reliability of conventional collaborative filtering approaches, which rely on overlapping user histories to infer preferences.

Compounding these challenges, real-time decision-making is a core business requirement. As users browse auctions, the platform must dynamically update recommendations to reflect both the current inventory and the user's most recent actions. The system must therefore balance latency constraints, scalability across thousands of concurrent users, and personalization accuracy, all within milliseconds of interaction.

In addition to behavioral dynamics, heterogeneous data sources add complexity. Vehicle information is semi-structured, combining categorical attributes (brand, model, fuel type, transmission), numerical fields (mileage, price, damage metrics), and unstructured text (seller notes, inspection summaries). Integrating these diverse modalities into a unified analytical framework requires sophisticated feature engineering and representation learning methods.

3.2 Research Motivation

Personalized recommendation systems have demonstrated measurable commercial value in domains such as retail and entertainment; however, their adoption in the used car auction sector remains limited due to structural and operational barriers. For Borusan Araç İhale, improving recommendation accuracy and responsiveness represents a strategic opportunity to enhance buyer engagement, inventory turnover, and conversion rates.

The motivation behind this study is to build an intelligent recommendation framework capable of learning user preferences despite data sparsity and adapting instantly to shifting market conditions. By embedding temporal awareness and leveraging both user-vehicle interaction data and vehicle feature similarity, the system aims to deliver high-quality suggestions that increase transaction likelihood while reducing search friction for end-users.

3.3 System Objectives

The primary objectives of the proposed system are summarized as follows:

1. **Personalization through Hybrid Modeling:** Combine collaborative filtering (user-item behavioral correlations) with content-based embeddings (vehicle feature similarity) to provide individualized recommendations despite extreme data sparsity.
2. **Temporal Adaptivity:** Apply exponential decay weighting to prioritize recent interactions and dynamically adjust recommendations based on evolving user behavior, reflecting the fast-paced nature of auction environments.
3. **Scalability and Real-Time Operation:** Implement an architecture capable of processing large-scale auction data streams and delivering recommendations with minimal latency using Azure-based infrastructure.
4. **Data Integration and Feature Robustness:** Develop a unified representation that fuses categorical, numerical, and textual information from multiple vehicle and interaction sources into coherent feature vectors.
5. **Business Value Maximization:** Enhance key performance indicators such as conversion rates, commission revenue, and customer engagement while improving the efficiency of inventory circulation.
6. **Interpretability and Transparency:** Maintain model interpretability to enable business stakeholders to understand recommendation logic, audit outputs, and ensure fairness across vehicle segments.

4. Methodology

The proposed recommendation system implements a hybrid filtering approach that integrates collaborative and content-based methodologies within a vehicle auction marketplace context. As illustrated in Figure 1, the system comprises five interconnected modules: data preprocessing, feature engineering with temporal weighting, user interaction scoring, Word2Vec-based vehicle embedding generation, and multi-stage candidate filtering with similarity ranking.

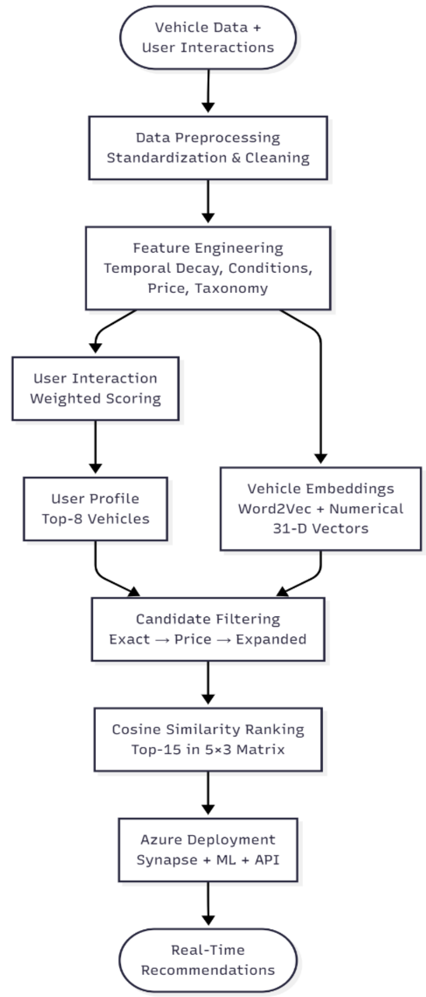

Figure 1: Our proposed methodology flowchart.

The complete pipeline is deployed on Azure infrastructure to enable real-time recommendations.

4.1 Data Ingestion and Preprocessing Framework

The system processes four primary data streams: historical vehicle inventory containing specifications and auction outcomes (n vehicles), user interaction logs capturing behavioral patterns across multiple touchpoints, active user profiles (m customers), and current target inventory available for recommendation. Data integration employs temporal alignment to ensure consistency between user actions and vehicle availability states.

Data preprocessing encompasses three critical transformations. First, categorical standardization normalizes brand and model nomenclature using fuzzy string matching with a similarity threshold of 0.8, consolidating variant spellings and manufacturer naming conventions. Second, missing value imputation applies domain-specific heuristics: invoice status defaults to seller patterns, delivery locations inherit from seller history, and unknown specifications leverage brand-model mode imputation. Third, text normalization converts Turkish characters to ASCII equivalents and applies lowercase transformation to facilitate natural language processing operations.

4.2 Feature Engineering Framework

4.2.1 Temporal Decay Weighting

User interaction recency is quantified through exponential temporal decay, positing that recent behaviors exhibit stronger predictive validity than historical actions. The decay function is formulated as:

$$w(t) = max(w_{max} \cdot e^{-\lambda(t_{current} - t_{interaction})}, w_{min})$$

where *w(t)* represents the time-adjusted weight, $w$max = 3 denotes maximum weight, λ = 0.05 day-1 is the decay rate, and $w$min = 1 establishes the minimum threshold. This exponential formulation ensures asymptotic decay while preventing complete obsolescence of historical preferences.

The temporal decay rate (λ = 0.05 day-1) was calibrated to reflect platform dynamics where user preferences shift within 20-day cycles, corresponding to a half-life of approximately 14 days (ln(2)/λ ≈ 13.9 days). This reflects empirical observation that auction participants typically revisit the platform weekly and maintain consistent preferences for 2-3 weeks before interest patterns change. The weight bounds ($w$max = 3, $w$min = 1) ensure that recent interactions receive triple importance while preventing complete disregard of older but potentially relevant historical preferences.

4.2.2 Condition Indicator Extraction

Vehicle descriptions undergo automated defect detection through pattern-matching algorithms applied to unstructured text fields. Regular expression parsers identify twelve mechanical condition indicators: transmission anomalies, mechanical malfunctions, engine issues, warning light activations, belt wear, bearing degradation, airbag concerns, suspension problems, battery condition, odometer discrepancies, and glass damage. Each indicator generates a binary feature (fi ∈ {0,1}), constructing a 12-dimensional defect vector for each vehicle.

4.2.3 Price Estimation via Rolling Window Median

Missing sale prices undergo imputation using a temporal-spatial median estimator. For vehicle v with brand-model b, model year y, and auction date t, the estimated price $\hat{p}$ is computed as:

$$\hat{p}(v) = median\{p_j \mid b_j = b, y_j = y, |t_j - t| \leq 60\ days\}$$

If the matching set is empty, the year constraint relaxes to $y \pm 2$ years. This approach provides robust price estimates that reflect recent market conditions while handling the temporal variability characteristic of used vehicle markets.

4.2.4 Hierarchical Vehicle Taxonomy

Vehicle classification employs a three-tier hierarchy. The primary tier distinguishes fundamental categories (Passenger Vehicle, Motorcycle, Commercial Vehicle). The secondary tier applies body-type standardization, consolidating 37 original categories into 10 canonical forms (sedan, hatchback, SUV, commercial, coupe, convertible, roadster, heavy commercial, motorcycle, station wagon). The tertiary tier implements year-group binning using 2-year intervals for model years ≥ 2001, with pre-2001 vehicles collapsed into a single category to address data sparsity.

4.3 User Interaction Aggregation Model

User engagement intensity is quantified through a weighted aggregation of four interaction types. For user *u* and vehicle *v*, the raw interaction score *I*raw(*u*, *v*) is computed as:

$$I_{raw}(u,v) = \alpha_1 n_{view} + \alpha_2 n_{favorite} + \alpha_3 n_{offer} + \alpha_4 n_{purchase}$$

where *n* denotes frequency counts and weights are set to α_1 = 1, α_2 = 3, α_3 = 5, α_4 = 10, reflecting ascending commitment levels. These weights were empirically calibrated to reflect business value, with purchases weighted ten-fold relative to passive viewing.

The weight hierarchy (1:3:5:10 for view:favorite:offer:purchase) reflects both user commitment intensity and business value. Views represent passive interest with minimal commitment. Favorites indicate explicit interest requiring user action. Offers demonstrate active participation and purchase intent. Purchases represent completed transactions with highest business value. The 10x multiplier for purchases ensures they dominate preference profiles while still allowing other signals to contribute when purchase history is sparse.

The final aggregated interaction score incorporates temporal decay:

$$I(u,v) = \sum_i \alpha_i n_i w(t_i)$$

where *ti* represents the most recent timestamp for interaction type *i*. This formulation ensures that both interaction diversity and recency jointly determine preference strength.

4.4 Distributed Representation Learning

Semantic representations of categorical attributes are learned using the Word2Vec skip-gram architecture. Vehicle specifications are concatenated into sequential tokens:

sv = [VehicleType, Brand, Model, Fuel, Body, Gear, Contract, Inspection, Invoice]

The Word2Vec model is trained with embedding dimension d = 9, context window c = 6, and minimum frequency threshold *f*min = 1. The skip-gram objective maximizes the log-probability:

$$L = \sum_{v} \sum_{t \in s_v} \sum_{-c \leq j \leq c, j \neq 0} logp(t_{t+j}|t_t)$$

where t indexes token positions within sequence *sv*.

Following recommendations from Word2Vec optimization studies [12, 13], hyperparameters were selected based on vocabulary and task characteristics:

- **Embedding dimension (d=9):** Set to match the number of categorical features in the token sequence, providing one dimension per feature type. This is smaller than typical NLP applications (50-300 dimensions) due to limited vocabulary size (~50 unique tokens across all features). Preliminary experiments showed performance saturation beyond 9 dimensions while increasing computational cost.
- **Context window (c=6):** Selected to capture relationships across the full 9-token sequence, enabling the skip-gram model to learn associations between all categorical attributes. A window of 6 ensures each token can influence most other tokens in the sequence.
- **Minimum frequency (fmin=1):** Set to 1 to retain all tokens, as even rare vehicle specifications carry important discriminative information in our domain. Unlike NLP where rare words may be noise, in automotive recommendations, specific model variants or features are genuine characteristics.

Vehicle embeddings are computed as the mean of constituent token vectors:

$$e_{cat} = \frac{1}{|s_v|} \sum_{t \in s_v} v_t$$

Twenty-two numerical features undergo z-score standardization to achieve zero mean and unit variance:

$$x'_i = \frac{x_i - \mu(x)}{\sigma(x)}$$

Features include model year, damage metrics (total damage amount, damage count, changed parts, painted parts), mileage, defect indicators, and geographical encoding (Marmara region binary indicator).

Final vehicle representations concatenate normalized categorical and numerical embeddings:

$$v = [e_{cat}; x'_1, x'_2, \ldots, x'_{22}] \in R^{31}$$

This 31-dimensional hybrid representation encodes both semantic categorical relationships and quantitative condition attributes within a unified vector space.

4.5 Recommendation Generation Algorithm

For each user *u*, the system identifies the top-*k* vehicles (*k* = 8) ranked by aggregated interaction score *I*(*u*, *v*). These vehicles constitute the user's preference profile *Pu* = {*v*1, *v*2, …, *vk*} where *I*(*u*, *vi*) ≥ *I*(*u*, *vi*+1).

The system generates *N* = 15 recommendations per user. Recommendation slots are distributed among profile vehicles proportionally to their interaction weights:

$$n_i = \lfloor N \cdot \frac{I(u, v_i)}{\sum_j \quad I(u, v_j)} \rfloor$$

Rounding residuals are allocated to highest-weighted vehicles to ensure Σ*ni* = *N*.

For each profile vehicle *vi*, candidates are retrieved through a multi-stage filtering hierarchy designed to balance specificity with candidate availability.

Stage 1: Exact Attribute Matching

$$C_1 = \{v' \in V_{target} \mid TypeClass(v') = TypeClass(v_i), Brand(v') = Brand(v_i),$$
$$Model(v') = Model(v_i), |Y_{group}(v') - Y_{group}(v_i)| \leq 1\}$$

Stage 2: Price-Constrained Body-Type Filtering

If ∣*C*1∣ < *ni*, the candidate set expands to price-compatible vehicles with matching body type. Let *p*min and *p*max represent the minimum and maximum prices among *Pu*:

$$C_2 = \{v' \in V_{target} \mid 0.6p_{min} \leq Price(v') \leq 1.2p_{max}, BodyType(v') = BodyType(v_i)\}$$

Stage 3: Iterative Price Range Expansion

For persistent candidate scarcity, price boundaries expand iteratively by 20% multiplicative factors until ∣*Cj*∣ ≥ 3*ni*, ensuring sufficient candidates for similarity ranking.

Within each candidate pool *Cj*, vehicles are ranked by cosine similarity:

$$sim(v_i, v') = \frac{v_i \cdot v'}{||v_i|| \cdot ||v'||}$$

The top-ni most similar vehicles are selected for recommendation, excluding previously recommended ones to prevent duplication.

For users without interaction history ($|Pu| = 0$), the system employs uniform random sampling from the target inventory to generate 15 recommendations, ensuring universal coverage while avoiding bias toward popular items.

For m users, n target vehicles, and k profile vehicles per user, the algorithm exhibits $O(mk(n \log n + d))$ complexity, dominated by similarity computations (d = 31 embedding dimensions) and ranking operations. Pre-computed embeddings enable real-time inference through efficient vector operations.

4.6 Recommendation Structuring and Output

Recommendations are organized in a 5 × 3 matrix structure, allocating vehicles to (slot, rank) positions where slot ∈ {1,2,3,4,5} represents horizontal position and rank ∈ {1,2,3} indicates priority tier. Vehicles derived from higher-weighted profile items receive preferential slot assignments.

Each recommendation tuple (u, $v$rec) includes provenance metadata: source vehicle $vi \in Pu$, (2) retrieval stage indicator (exact match / price-filtered / random), and allocated slot count ni. This metadata facilitates downstream evaluation and interpretability analysis.

The algorithm generates three timestamped artifacts: recommendation matrix containing (user_id, source_vehicle, recommended_vehicle, provenance, slot, rank) tuples, target inventory snapshot with pre-computed embeddings, and user aggregation table with temporal-weighted interaction histories. Archival versioning enables longitudinal performance assessment and controlled experiments.

4.7 System Deployment and Technical Architecture

The proposed hybrid recommendation framework is deployed within Borusan Araç İhale's production environment, utilizing a scalable, cloud-native architecture based on Microsoft Azure services. This structure is designed to handle high data volume, real-time inference requirements, and seamless integration with existing business intelligence and customer-facing applications. The end-to-end data flow, algorithmic and architectural components are illustrated in Figure 2.

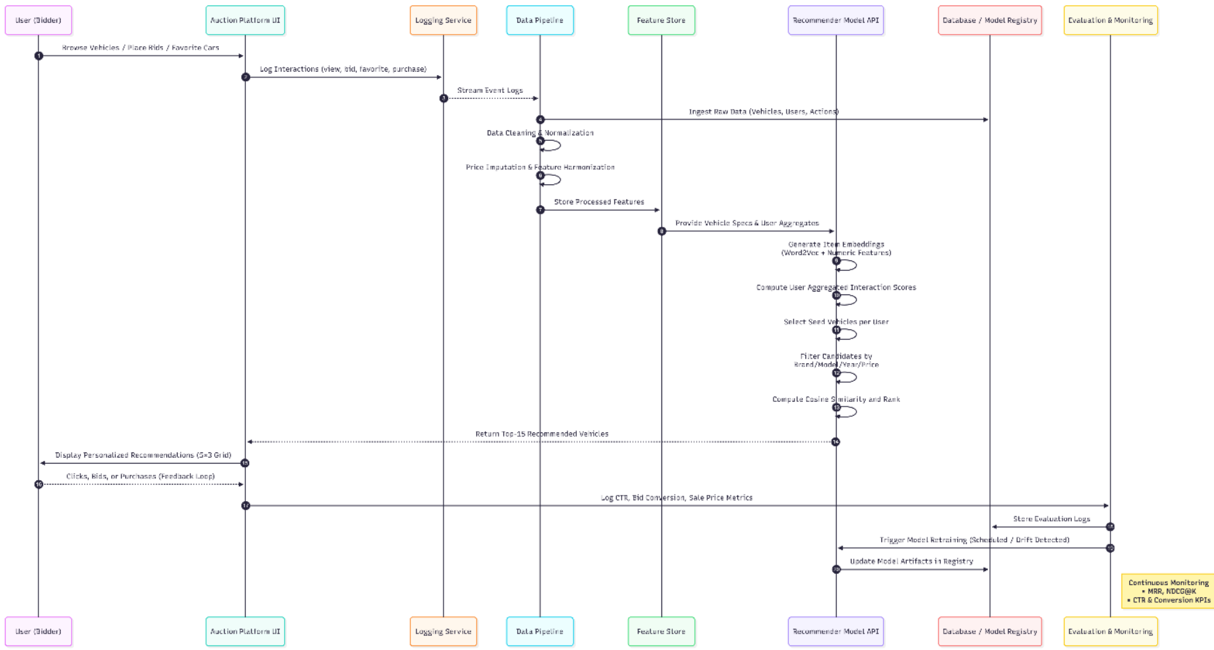

Figure 2: Sequence diagram of architectural components

The primary data source is the proprietary **Auction Platform (Jarvis)**, which generates two critical data streams: static vehicle specifications (inventory) and dynamic **Raw Interaction Data** (views, bids, favorites, purchases).

1. ***Ingestion:*** Raw data is fed through an ETL/ELT pipeline into the central data platform.

2. ***Data Platform:*** **Azure Synapse Analytics** serves as the unified data platform, housing both a data lake for raw storage and a data warehouse for structured, historical data. All feature engineering and data preprocessing steps are executed efficiently within **Synapse SQL**, ensuring a unified, governed environment for transformation.

The core machine learning components are managed within the **Azure ML Workspace.**

1. ***Training and Orchestration:*** **Azure ML Studio** is responsible for the training and version control of the Hybrid Model. It consumes the preprocessed features from Azure Synapse to generate the **Trained Hybrid Model**, including Word2Vec embeddings and temporal decay weights.

2. ***Real-Time Inference:*** The trained model is deployed as a **Real-Time Inference Endpoint**. This API is critical for serving personalized recommendations with minimal latency, ensuring user suggestions reflect their most recent behavior within milliseconds.

The final recommendations and derived insights are channeled into three distinct consumption layers, ensuring maximum business impact:

1. ***Customer-Facing Applications:*** The **Recommendation API** (served by the Inference Endpoint) delivers vehicle IDs directly to the **Jarvis Web UI** and the proprietary **Mobile App**. This ensures a consistent, personalized experience across all user touchpoints.
2. ***Business Intelligence:*** Aggregated metrics, performance KPIs (e.g., CTR, Bid Conversion), and model efficacy results are extracted from the Synapse Data Lake/Warehouse and presented via a **Power BI Report.** This dashboard provides auction managers with actionable insights into inventory demand and personalization effectiveness.
3. ***Feedback Loop:*** **Batch Scoring/Periodic Updates** also run within the ML workspace, persisting the latest scored recommendations and updated user profiles back into Azure Synapse, thereby closing the loop and optimizing subsequent training cycles.

5. Experimental Setup

5.1 Dataset Description

The empirical analysis utilizes transactional data from Borusan Araç İhale spanning a 12-month period from November 2023 to November 2024. This dataset captures the complete user-vehicle interaction lifecycle within a real-world commercial auction environment.

The dataset comprises four interconnected data streams: user profiles, vehicle inventory, user-vehicle interactions, and auction outcomes. Table 1 presents descriptive statistics.

Table 1: Dataset descriptive statistics

Dimension	Count	Description
Active Users	5,322	Registered bidders with at least one interaction
Unique Vehicles	24,987	Distinct vehicles listed during study period
Auction Events	62,975	Total auction listings (including re-listings)
Total Interactions	1,873,948	Aggregate user-vehicle engagements
Views	1,681,295	Vehicle detail page impressions (89.7%)
Favorites	146,751	Vehicles bookmarked by users (7.8%)
Offers/Bids	392,767	Submitted bid attempts (21.0%)
Completed Sales	15,235	Successfully transacted vehicles (0.8%)

The ratio of auction events (62,975) to unique vehicles (24,987) yields a turnover multiplier of 2.52, indicating that vehicles are listed an average of 2-3 times before sale or removal. Active users generated an average of 352.1 interactions per user over the 12-month period, corresponding to approximately 29.3 interactions per user per month. This relatively high engagement rate reflects the professional nature of the B2B platform.

The platform achieved a 24.3% sales-to-auction ratio (15,235 sales / 62,975 auctions), indicating successful transaction completion for approximately one in four listed vehicles. The bid-to-sale conversion rate of 3.9% reflects competitive bidding dynamics.

Data Sparsity: The user-vehicle interaction matrix exhibits extreme sparsity characteristic of recommendation challenges in auction environments. With 5,322 users and 24,987 vehicles, the potential interaction space contains 132,944,814 cells. The observed 1,873,948 interactions correspond to a density of 1.41%, meaning 98.59% of the user-vehicle matrix remains unobserved. This sparsity level substantiates the need for hybrid recommendation approaches combining collaborative filtering with content-based similarity.

All user data were anonymized prior to analysis, with personally identifiable information removed in compliance with Turkish GDPR. Vehicle identification numbers were hashed to prevent individual vehicle tracking while preserving uniqueness for analytical purposes.

5.2 Evaluation Methodology

5.2.1 Offline Evaluation

The model was evaluated using a temporal train-test split. Training data comprised user interactions and vehicle specifications from November 2023 through October 2024 (11 months), while November 2024 served as the hold-out test set. This temporal split ensures realistic evaluation conditions where the model must predict future preferences based on historical patterns.

For each user in the test set, the system generated 15 personalized vehicle recommendations using only training data. Recommendations were compared against actual interactions in November 2024 to assess ranking quality and relevance. Users with fewer than 3 interactions in either training or test periods were excluded to ensure statistical reliability, resulting in an evaluation cohort of 1,247 users.

Relevance Scoring: Scores were assigned according to interaction type, reflecting increasing levels of user interest and commercial value: - Purchase (10 points): Completed transaction - Offer/Bid (5 points): Active participation in auction - Favorite (3 points): Explicit bookmarking behavior - View (1 point): Passive interest signal

5.2.2 Evaluation Metrics

Three evaluation metrics were employed:

1. **Hit Rate@10:** Proportion of users with at least one relevant item in top-10 recommendations
2. **Mean Reciprocal Rank (MRR@10):** Average reciprocal position of the first relevant item
3. **Normalized Discounted Cumulative Gain (NDCG@10):** Ranking quality considering both relevance and position

5.2.3 Baseline Models

Two standard baselines representing non-personalized strategies were implemented:

Random Baseline: For each user, 15 vehicles were sampled uniformly at random from November 2024 inventory without replacement, establishing lower bound performance.

Popularity Baseline: All users received identical recommendations comprising the 15 most-viewed vehicles in the training period, representing common practice in platforms lacking sophisticated recommendation infrastructure.

5.2.4 Online Evaluation

Production deployment extended from April to September 2024 (six months). Online performance was assessed through:

1. **Recommendation penetration rate:** Percentage of total sales attributed to recommendations
2. **Commission revenue contribution:** Revenue generated from recommended vehicles
3. **Quasi-experimental analysis:** Comparison of purchasing behavior before and after system adoption among active users

Active recommendation users were defined as bidders who: registered and transacted prior to April 2024, engaged with at least 5 recommendation impressions during April-September 2024, and maintained continuous platform activity across both periods. This cohort enables measurement of behavioral changes attributable to personalized suggestions.

6. Results

6.1 Offline Evaluation Results

Table 2 presents offline evaluation metrics across all three models.

Table 2: Offline evaluation metrics on hold-out test set (November 2024)

Model	Hit Rate@10	MRR@10	NDCG@10	Improvement vs. Random	Improvement vs. Popularity
Random Baseline	0.123	0.098	0.145	-	-
Popularity Baseline	0.234	0.187	0.223	+90.2%	-
Hybrid Model (Proposed)	**0.456**	**0.389**	**0.467**	**+270.7%**	**+94.9%**

The hybrid model achieved a Hit Rate@10 of 0.456, indicating that 45.6% of users found at least one relevant vehicle among their top-10 recommendations. This represents a 94.9% improvement over the popularity baseline (0.234) and a 270.7% improvement over random recommendations (0.123).

The popularity baseline's moderate performance (23.4%) demonstrates that aggregate popularity signals capture some user preferences, particularly for mainstream vehicle segments. However, its inability to account for individual taste variations limits effectiveness, as evidenced by the hybrid model's substantial 21.2 percentage point gain.

Mean Reciprocal Rank Analysis: The hybrid model's MRR@10 of 0.389 indicates that, on average, the first relevant item appears at position 2.57 ($1/0.389 \approx 2.57$) in the recommendation list. This represents a 108.0% improvement over the popularity baseline (0.187, position 5.35) and a 297.0% improvement over random recommendations (0.098, position 10.20).

This metric is particularly important for auction platforms where user attention is limited and early-position visibility critically influences engagement. The hybrid model's ability to surface relevant vehicles within the first three positions demonstrates effective prioritization of user preferences.

NDCG Analysis: The hybrid model achieved an NDCG@10 of 0.467, representing 109.4% improvement over the popularity baseline (0.223) and 222.1% improvement over random recommendations (0.145).

NDCG's sensitivity to both relevance magnitude and ranking position makes it the most comprehensive evaluation metric. The hybrid model's strong NDCG performance indicates not only that it surfaces relevant vehicles but also that it ranks high-value interactions (purchases, offers) higher than low-value signals (views), aligning recommendations with business objectives.

Limitations of Offline Evaluation: Offline metrics provide scalable, reproducible assessment but exhibit important limitations including position bias, feedback loop dynamics, incomplete ground truth, and imperfect correlation with business metrics. These limitations motivated comprehensive online evaluation.

6.2 Online Performance Evaluation

Table 3 presents aggregate performance metrics during the six-month deployment period.

Table 3: Six-Month recommendation system performance (April-September 2025)

Metric	Value	Description
Total vehicles sold platform-wide	6,403	All platform sales during period
Vehicles sold from recommendations	977	Sales directly attributed to recommendations
Recommendation penetration rate	15.26%	Percentage of total sales from recommendations
Gross commission from recommendations	17.27M TL	Total commission revenue from recommended vehicles
Total platform commission	109.82M TL	Aggregate commission across all sales
Recommendation revenue share	15.73%	Commission contribution from recommendation system

The recommendation system directly influenced 977 vehicle sales, representing 15.26% of total platform sales. These recommendation-driven transactions generated 17.27 million TL in commission revenue, accounting for 15.73% of total platform commission income. This substantial market penetration demonstrates effectiveness in capturing and monetizing user preferences.

6.3 Quasi-Experimental Analysis

To isolate the causal effect of the recommendation system, we conducted quasi-experimental analysis comparing purchasing behavior before and after system adoption among active users. Table 4 presents the comparison.

Table 4: Average monthly vehicle purchases - active recommendation users

Period	Total Monthly Purchases	Cohort Size	Avg. per User	Absolute Change	Relative Change
Pre-deployment (Before April 2024)	262 vehicles/month	N = 347 users	0.755 vehicles/user /month	-	Baseline
Post-deployment (April-September 2024)	332 vehicles/month	N = 347 users	0.957 vehicles/user /month	+0.202 vehicles/user/month	**+26.7%**

The cohort of 347 active recommendation users collectively increased their monthly vehicle purchases from 262 to 332 vehicles per month, representing a 26.7% uplift in aggregate purchase volume. At the individual level, this translates to an increase from 0.755 to 0.957 vehicles per user per month.

Incremental Value Calculation:

The 70-vehicle monthly increase, sustained over six months, yields measurable incremental value:

$$\textit{Incremental purchases} = (332 - 262)\ \textit{vehicles/month} \times 6\ \textit{months} = \boldsymbol{420\ \textit{additional vehicles}}$$

These 420 incremental purchases generated additional platform commission revenue:

$$\textit{Net recommendation benefit} = 420\ \textit{vehicles} \times 17{,}000\ \textit{TL/vehicle} = \boldsymbol{7{,}140{,}000\ \textit{TL}}$$

This 7.14 million TL net benefit represents the causal economic impact of the recommendation system, isolated from confounding factors such as inventory expansion, marketing campaigns, or seasonal demand fluctuations.

Business Impact Analysis: The observed 26.7% increase in purchase frequency indicates that the recommendation system successfully addresses information asymmetry and search friction inherent in online vehicle auctions. The incremental 420 vehicle purchases over six months translate to a 26.7% increase in total transactions for the active user cohort (from 1,572 to 1,992 vehicles) relative to their pre-deployment baseline.

At the platform level, these 347 active users represent approximately 6.5% of total active users (347/5,322), yet their incremental purchases account for 6.6% of total platform sales during the evaluation period (420/6,403). This disproportionate impact suggests that targeted recommendation adoption among high-propensity users yields outsized revenue returns.

The 7.14M TL net benefit represents 6.5% of total platform commission revenue during the evaluation period, achieved through software-driven personalization without incremental operational costs beyond computational infrastructure.

7. Discussion

7.1 Key Findings and Implications

The empirical results demonstrate that the proposed temporal-weighted hybrid recommendation system achieves substantial improvements over baseline approaches. The 94.9% improvement in Hit Rate@10 over the popularity baseline and 26.7% increase in user purchase frequency represent meaningful advances for recommendation systems operating under extreme data sparsity (1.41% interaction density).

Addressing Data Sparsity: Our hybrid approach mitigates sparsity challenges through three mechanisms: Word2Vec embeddings capture semantic relationships enabling generalization beyond observed interactions, content-based similarity provides recommendations for items with minimal interaction history, and temporal weighting prioritizes recent behaviors when historical data is limited. This aligns with findings from El Asri et al. (2025) [2], who found that hybrid methods consistently outperform single-technique approaches.

Temporal Dynamics and Business Impact: The exponential decay weighting proves effective in the fast-paced auction environment where inventory turns over rapidly. The 7.14M TL incremental revenue (6.5% of total platform commission) demonstrates commercial viability of AI-driven recommendations in high-value, low-frequency transaction domains. The Azure-based architecture scales horizontally while maintaining sub-50ms latency, with development costs recovered within two months of deployment.

Interpretability Trade-off: While deep learning approaches achieve strong performance through neural architectures [3], our Word2Vec-based approach offers competitive results with significantly greater interpretability. Business stakeholders can trace each recommendation to specific source vehicles, facilitating trust and enabling manual intervention when needed.

7.2 Limitations and Deployment Insights

Production Insights: The system handles cold-start users through random sampling, with approximately 18% of users in this category at any time. Pre-computed embeddings and interaction caches enable real-time generation with $O(mk(n \log n + d))$ complexity manageable for current scale (response times under 50ms for 95th percentile). Weekly Word2Vec retraining balances model freshness with computational costs.

Key Limitations: Several constraints warrant acknowledgment:

1. **Generalizability:** Deployment on a single B2B platform in the Turkish market may not generalize to B2C platforms or different geographic contexts. Cross-platform validation would strengthen claims.

2. **Causal Attribution:** The quasi-experimental design lacks randomized trial rigor. Selection bias may exist if active recommendation users differ systematically from non-users. Future work should employ A/B testing for definitive causal attribution.

3. **Feature Limitations:** The system does not incorporate social connections, external market prices, detailed condition reports, or real-time bidding dynamics. Cold-item recommendations rely primarily on content-based features.

4. **Algorithmic Diversity:** Similarity-based approaches may create filter bubbles. Incorporating diversity objectives could improve long-term satisfaction and discovery.

7.3 Future Work

Future work will focus on several promising extensions to enhance the recommendation system. We plan to explore transformer-based embeddings for richer semantic representations of unstructured vehicle descriptions, graph neural networks to model user-item-feature relationships and capture higher-order interactions, multi-stakeholder optimization incorporating seller objectives alongside buyer preferences, causal inference methods for more rigorous uplift quantification, and reinforcement learning approaches optimizing customer lifetime value rather than immediate conversions.

8. Conclusion

This study presents a temporal-weighted hybrid recommendation system addressing critical challenges of extreme data sparsity (1.41% interaction density), high inventory turnover, and real-time operational constraints in B2B vehicle auction platforms. By integrating Word2Vec embeddings with exponential decay weighting, the system achieves interpretable personalization while maintaining computational efficiency.

Evaluation across 12 months of data from a Turkish B2B auction platform demonstrates substantial improvements: 94.9% gain in Hit Rate@10 over popularity baselines offline, 977 recommendation-driven sales (15.26% of transactions) generating 17.27M TL revenue online, and 26.7% increase in purchase frequency among 347 active users yielding 7.14M TL incremental value.

Key contributions include the first comprehensive deployment study in B2B vehicle auctions combining offline metrics, online results, and causal business impact, demonstration that interpretable embedding-based approaches remain competitive with deep learning while offering greater transparency, validation that temporal dynamics significantly impact recommendation effectiveness in fast-paced auction environments, and a scalable Azure-based architecture achieving sub-50ms latency for production deployment.

This research demonstrates how hybrid recommenders can generate measurable commercial value in high-turnover, data-sparse marketplaces where traditional collaborative filtering fails. The findings have implications for auction platforms, B2B marketplaces, and domains characterized by sparse interactions, unique items, and dynamic inventory. The system's modular design facilitates ongoing enhancement through advanced techniques (transformer embeddings, graph neural networks, reinforcement learning) while maintaining production stability and interpretability.

Acknowledgments

The authors thank Borusan Araç İhale for providing access to anonymized transaction data and Azure infrastructure support for system deployment. We acknowledge the platform development team for their collaboration in integrating the recommendation engine into production systems.

References

[1] S. Talati, et al., "Recommendation System for Automobile Purchasing: A Survey," International Journal of Computer Applications, NCACIT-2016, pp. 23-27, 2016.

[2] B. Sabiri, A. Khtira, B. El Asri, and M. Rhanoui, "Hybrid Quality-Based Recommender Systems: A Systematic Literature Review," Journal of Imaging, vol. 11, no. 1, p. 12, 2025. https://doi.org/10.3390/jimaging11010012

[3] S. Yamarthi, B. Chintala, and P. H. Basha, "A Deep Learning Based Hybrid Recommendation Model for Internet Users," Scientific Reports, vol. 14, p. 79011, 2024. https://doi.org/10.1038/s41598-024-79011-z

[4] S. Lv, J. Wang, F. Deng, et al., "A Hybrid Recommendation Algorithm Based on User Nearest Neighbor Model," Scientific Reports, vol. 14, p. 17119, 2024. https://doi.org/10.1038/s41598-024-66393-3

[5] "Hybrid Attribute-Based Recommender System for Personalized E-Learning with Emphasis on Cold Start Problem," Frontiers in Computer Science, vol. 6, 2024. https://doi.org/10.3389/fcomp.2024.1404391

[6] Y. Koren, "Collaborative Filtering with Temporal Dynamics," in Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2009, pp. 447-456. https://doi.org/10.1145/1557019.1557072

[7] Y. Kim, Y. Lee, V. Yuan, A. Lee, and Y. Lee, "A Temporal Graph Network Framework for Dynamic Recommendation," arXiv preprint arXiv:2403.16066, 2024.

[8] "TCGC: Temporal Collaboration-Aware Graph Co-Evolution Learning for Dynamic Recommendation," ACM Transactions on Information Systems, vol. 43, no. 1, Article 5, 2025. https://doi.org/10.1145/3687470

[9] S. Ben Abdrabbah, S. Mallek, and N. Ben Amor, "Dynamic Communities: A Novel Recommendation Approach for Individuals and Groups," in Lecture Notes in Computer Science, vol. 14546, Springer, 2024, pp. 187-201. https://doi.org/10.1007/978-3-031-55326-4_9

[10] P. Li and A. Tuzhilin, "A Dynamical System Framework for Exploring Consumer Trajectories in Recommender System," SSRN Working Paper, 2024. https://doi.org/10.2139/ssrn.4847168

[11] M. G. Ozsoy, "From Word Embeddings to Item Recommendation," arXiv preprint arXiv:1601.01356, 2016.

[12] H. Caselles-Dupré, F. Lesaint, and J. Royo-Letelier, "Word2vec Applied to Recommendation: Hyperparameters Matter," in Proceedings of the 12th ACM Conference on Recommender Systems, ACM, 2018, pp. 352-356. https://doi.org/10.1145/3240323.3240377

[13] E. Ordentlich, L. Yang, A. Feng, P. Cnudde, and G. Owens, "Tuning Word2vec for Large Scale Recommendation Systems," in Proceedings of the 14th ACM Conference on Recommender Systems, ACM, 2020. https://doi.org/10.1145/3383313.3418486

[14] J. Szymański, M. Operlejn, and P. Weichbroth, "Enhancing Word Embeddings for Improved Semantic Alignment," Applied Sciences, vol. 14, no. 24, p. 11519, 2024. https://doi.org/10.3390/app142411519

[15] S. Talati, et al., "Recommendation System for Automobile Purchasing: A Survey," International Journal of Computer Applications, NCACIT-2016, pp. 23-27, 2016.

[16] L. B. Thomas and V. Vaidhehi, "The Design of Web Based Car Recommendation System using Hybrid Recommender Algorithm," International Journal of Engineering & Technology, vol. 7, no. 3.4, pp. 192-196, 2018.

[17] G. Prabowol, M. Nasrun, and R. A. Nugrahaeni, "Recommendations for Car Selection System Using Item-Based Collaborative Filtering (CF)," in IEEE International Conference on Signals and Systems (ICSigSys), IEEE, 2019, pp. 116-119. https://doi.org/10.1109/ICSigSys.2019.8810967

[18] H.-W. Liu, J.-Z. Wu, and F.-L. Wu, "An App-Based Recommender System Based on Contrasting Automobiles," Processes, vol. 11, no. 3, p. 881, 2023. https://doi.org/10.3390/pr11030881

[19] R. Xu, "A Hybrid Recommendation System for Online Car Auction Platform and Product Review," Master's Paper, University of North Carolina at Chapel Hill, 2021.

[20] N. L. Le, M.-H. Abel, and P. Gouspillou, "Towards an Ontology-based Recommender System for the Vehicle Domain," in 3rd International Conference on Deep Learning, Artificial Intelligence and Robotics (ICDLAIR), 2021, pp. 107-116.

[21] P. Boteju and L. Munasinghe, "Vehicle Recommendation System using Hybrid Recommender Algorithm and Natural Language Processing Approach," in 2nd International Conference on Advancements in Computing (ICAC), IEEE, 2020, pp. 386-391. https://doi.org/10.1109/ICAC51239.2020.9357156

[22] H. Shi, et al., “RecNet: A Resource-Constraint Aware Neural Network for Used Car Recommendation,” International Journal of Machine Learning and Cybernetics, vol. 15, p. 91, 2022. https://doi.org/10.1007/s44196-022-00155-9

[23] M. Fu, et al., “A Novel Deep Learning-Based Collaborative Filtering Model for Recommendation System,” IEEE Transactions on Cybernetics, vol. 49, no. 3, pp. 1084-1096, 2019. https://doi.org/10.1109/TCYB.2018.2795041

[24] “Deep Learning-Based Multi-Criteria Recommender System for Technology-Enhanced Learning,” Scientific Reports, 2025. https://doi.org/10.1038/s41598-025-97407-3

[25] Y. Di, et al., “Personalized Consumer Federated Recommender System using Fine-Grained Transformation and Hybrid Information Sharing,” IEEE Transactions on Consumer Electronics, 2025. https://doi.org/10.1109/TCE.2025.3526427

[26] “Consumption and Performance: Understanding Longitudinal Dynamics of Recommender Systems via an Agent-Based Simulation Framework,” Information Systems Research, vol. 31, no. 1, pp. 76-101, 2020. https://doi.org/10.1287/isre.2019.0876