Sarcouncil Journal of Engineering and Computer Sciences

ISSN(Online): 2945-3585

Volume- 04| Issue- 08| 2025

Research Article **Received:** 01-07-2025| **Accepted:** 23-07-2025 | **Published:** 12-08-2025

Personalised Product Discovery in Online Beauty Retail Using Dual-Tower Neural Networks

Jyotirmoy Sundi

Independent Researcher, USA

Abstract: The digital transformation of beauty retail has created unprecedented challenges for personalized product discovery, necessitating sophisticated artificial intelligence solutions to address the complex nature of cosmetic and skincare recommendations. This article presents a comprehensive dual-tower neural network architecture specifically engineered for beauty product recommendation systems, tackling the inherent difficulties of subjective preferences, intricate ingredient interactions, and rapidly evolving market trends. The implementation showcases how the separate encoding of user preferences and product characteristics through hierarchical explainable networks enables superior performance in modeling temporal dependencies and behavioral patterns. Advanced training methods incorporate sequence-aware negative sampling techniques and demographic-stratified testing to mitigate algorithmic bias while ensuring equitable recommendation quality across diverse user demographics. Performance optimization strategies utilizing approximate nearest neighbor search algorithms and intelligent caching policies achieve rapid response times while processing extensive product catalogs. The system deployment generated substantial improvements in user engagement metrics, conversion rates, and revenue growth while enhancing product discovery experiences. The architecture successfully balances personalization accuracy with demographic fairness requirements, demonstrating the practical viability of sophisticated recommendation technologies in specialized retail domains requiring high levels of personalization and cultural sensitivity

Keywords: Dual-tower neural networks, personalized recommendation systems, beauty e-commerce, algorithmic bias mitigation, real-time product discovery.

INTRODUCTION

The internet has completely changed people's process of buying beauty products. Now, people usually check online to learn about makeup and skincare. Studies show that things like augmented reality and AI are becoming key parts of shopping for beauty items. About 67% of shoppers use virtual try-on features with AR before they buy anything (Perret, J. K., & Schwientek, J. 2025). Today's beauty shoppers don't just want products that work; they also want to find new and interesting things that are both relevant to them and a nice surprise. This change creates big problems for online stores that have lots of products and want to give each customer suggestions that are personal but at the scale of the entire user base.

Existing beauty retail systems have had difficulties because beauty items have special properties, like very personal tastes, how ingredients interact, and how quickly trends change. The beauty business has different recommendation problems than normal retail because things like skin tone, texture, and lifestyle are very personal and challenging for basic recommendation methods (Perret, J. K., & Schwientek, J. 2025). Also, since buying beauty products is often tied to emotions and desires, recommendation systems need to be able to introduce users to both familiar and new items, so they can discover products that match their current tastes and what they hope to experience. Modern recommendation systems for fashion and beauty struggle with processing different kinds of data, like pictures, descriptions, and how users interact with products. Research indicates that basic recommendation methods are only about 42% correct in beauty product situations, which is not as precise as general retail (Deldjoo, Y. *et al.*, 2023). Unlike other types of retail, beauty products have difficult relationships between features like color theory, ingredient rules, and season changes, things not easily addressed by older methods.

The current beauty market has more than 50,000 different products in many categories, each needing its own recommendation approach. From looking at customer actions, experts have found that good beauty suggestions should consider seasonal taste changes that affect around 70% of users each quarter, rules about what ingredients can be mixed that impact 85% of skincare choices, and exact color matches that need high precision for foundation (Deldjoo, Y. *et al.*, 2023). The difficulties grow when thinking that people tend to use 12-18 different product types each year, which will require different recommendation methods and attribute importance for each type.

This paper offers a detailed look at a recommendation system used by a big beauty store, reaching millions of active users and

***Corresponding Author:** Jyotirmoy Sundi

DOI- https://doi.org/10.5281/zenodo.16814206

handling billions of interactions monthly. The system uses a dual-tower neural network design made to handle the special challenges of beauty product suggestions, mixing user behavior models with product details across hundreds of different features. With careful planning and launch methods, the system can respond in less than fifty milliseconds while keeping high recommendation quality for different user groups across countries and diverse skin tones.

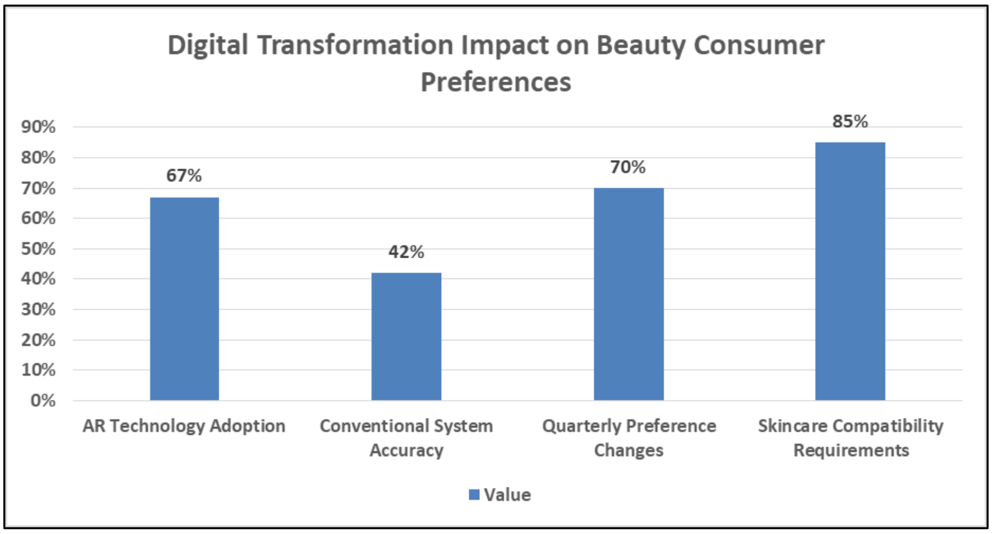

Figure 1: Digital Transformation Impact on Beauty Consumer Preferences (Perret, J. K., & Schwientek, J. 2025; Deldjoo, Y. *et al.*, 2023)

SYSTEM ARCHITECTURE AND DUAL-TOWER DESIGN

The core innovation of the implemented system lies in its dual-tower neural network architecture, which separately encodes user preferences and product characteristics before computing similarity scores for ranking. The user tower processes sequential interaction data, including click patterns, hover behaviors, purchase history, and session context, transforming these signals into dense vector representations that capture both explicit preferences and latent behavioral patterns. Hierarchical explainable networks demonstrate superior performance in modeling complex user behavior sequences, particularly when dealing with temporal dependencies that span extended time periods (Zhu, Y. *et al.*, 2025). The architecture employs recurrent neural networks with attention mechanisms to model temporal dependencies, recognizing that beauty preferences evolve with seasons, life events, and trend cycles.

Sequential behavior modeling requires sophisticated approaches to capture long-term dependencies while maintaining computational efficiency. Research indicates that hierarchical network structures can effectively process user interaction sequences by decomposing complex behavioral patterns into manageable sub-sequences, enabling better representation learning for cross-domain applications (Zhu, Y. *et al.*, 2025). The user tower implementation utilizes multi-layered LSTM networks with attention mechanisms, where each layer focuses on different temporal granularities ranging from immediate session interactions to long-term preference evolution patterns.

The product tower stores details about items, like what's in them, the brand, price, colors, texture, and if they're cruelty-free or vegan. This focus comes from the idea that what makes a beauty product appealing is how these things work together, not just one thing alone. Recommender systems do better when they use different kinds of information. Deep learning helps improve how accurate recommendations are by putting together written descriptions, images, and other details. The tower is built with different layers for product information. It uses embedding layers for things such as product category and dense layers for numbers like price or size. It also takes care of missing data, often seen in beauty product catalogs.

Combining different kinds of data is hard. One has to combine images, descriptions, and product details into a single recommendation system. New improvements in deep learning show that this method can really improve recommendations, especially when there exists a mix of data like photos, reviews, and specs. The tower setup uses convolutional neural networks to pull features

Publisher: SARC Publisher

from images, transformers to process text, and embedding layers for product categories.

The similarity computation between towers employs learned distance metrics rather than simple cosine similarity, allowing the model to weight different attribute dimensions according to importance for specific user segments. This approach enables the system to recognize that some users prioritize ingredient compatibility while others focus on brand prestige or color matching. The final ranking stage incorporates business logic constraints, including inventory levels, promotional priorities, and diversity requirements to ensure practical deployment viability.

Model training utilizes a distributed computing infrastructure capable of processing over one billion user interactions nightly, with careful attention to data freshness and computational efficiency. The system maintains separate training pipelines for different geographic regions and user segments, recognizing that beauty preferences exhibit significant cultural and demographic variation. This distributed approach enables rapid adaptation to emerging trends while maintaining system stability across global markets.

ADVANCED TRAINING METHODOLOGIES AND BIAS MITIGATION

The training methodology incorporates several sophisticated techniques designed to address the specific challenges of beauty product recommendation. Sequence-aware negative sampling represents a critical innovation, teaching the model to distinguish between users seeking familiar products and those ready for category exploration. Traditional negative sampling approaches randomly select non-interacted items, potentially penalizing the model for failing to recommend products users simply haven't discovered yet. Contemporary research in negative sampling demonstrates that sophisticated selection strategies can significantly improve recommendation quality, with dynamic sampling methods showing particular promise for addressing data sparsity and user preference modeling challenges (Ma, H. *et al.*, 2024). The implemented approach analyzes user interaction sequences to identify natural transition points where users demonstrate openness to new product categories or brands.

Current methods for negative sampling are more complex than just picking things at random; they now include context and timing. Studies show that good negative sampling should look at how people use recommendation data, where not clicking doesn't always mean someone isn't interested. One approach looks at how people act over time to tell if they skipped something on purpose or just never saw it, which helps create a better training signal.

To keep up with fast-moving beauty trends, the system relearns nightly with fresh user info. This way, it doesn't need manual rule updates. This constant learning really helps in the beauty world, where trends can pop up fast from viral posts, seasonal palettes, and celeb promotion. The learning turns user info into a model-friendly format right away. The beauty products company personnel carefully check the data obtained from different places for quality and consistency.

Testing across different groups is important for cutting down on bias. Beauty products can be biased, like when showing different skin tones or guessing things based on age. The approach proposed in this article performs model checks separately for these segments. This makes sure recommendations are steady across skin tones, age groups, and backgrounds. Recent work shows sentiment analysis and flexible tweaking can cut selection bias in advice systems, especially for demographic differences. This testing showed a few times where normal metrics hid big performance gaps for less-shown user groups. That led to changes that made the model better and fairer for those groups.

Selection bias mitigation requires a comprehensive understanding of how demographic factors influence both user behavior and system responses. Research shows that traditional recommendation systems often perpetuate existing biases through biased training data and evaluation methodologies, necessitating explicit debiasing interventions (Zhang, F. *et al.*, 2025). The stratified testing approach incorporates sentiment analysis techniques to identify potential bias indicators within user feedback patterns, enabling proactive adjustment of recommendation algorithms before deployment.

The system also incorporates fairness constraints directly into the training objective, balancing recommendation accuracy with demographic parity requirements. This approach recognizes that purely accuracy-driven optimization can

Publisher: SARC Publisher

perpetuate historical biases present in training data, particularly regarding product availability and marketing targeting. The implemented fairness constraints ensure that high-quality recommendations remain accessible across all user segments while maintaining overall system performance.

Table 1: Techniques addressing beauty product recommendation challenges (Ma, H. *et al.*, 2024; Zhang, F. *et al.*, 2025)

Technique	Implementation Approach
Negative Sampling Method	Sequence-aware selection strategy
Training Schedule	Nightly retraining on fresh data
Bias Mitigation	Demographic-stratified testing
Trend Adaptation	Automated detection without manual rules
Fairness Integration	Direct constraints in the training objective
Data Processing	Real-time feature engineering pipeline

PERFORMANCE OPTIMIZATION AND SCALABILITY

To achieve sub-50 millisecond response times while examining tens of thousands of product choices, fine-tuning of each system layer's performance must be done. The process uses approximate nearest neighbor search, which is specially adjusted for the complex data from the dual-tower setup. Recent studies of K-nearest neighbor algorithm changes show that better ways to calculate distances and smarter indexing can really speed up searches. These improvements keep the results just as accurate as would have been if everything were checked (Halder, R. K. *et al.*, 2024). By carefully adjusting settings and improving how data is indexed, these algorithms keep the quality of recommendations high while using much less processing power.

Modern KNN algorithm enhancements focus on addressing computational complexity challenges inherent in high-dimensional similarity search operations. Research indicates that modified KNN approaches incorporating adaptive distance metrics and hierarchical indexing structures can achieve substantial performance improvements, particularly when dealing with large-scale datasets containing millions of feature vectors (Halder, R. K. *et al.*, 2024). The implemented system leverages these advances through sophisticated indexing mechanisms that partition the embedding space into manageable clusters, enabling logarithmic search complexity rather than linear traversal of entire product catalogs.

Caching strategies play a crucial role in system performance, with different caching approaches applied at multiple levels of the recommendation pipeline. User embeddings benefit from session-level caching due to relatively stable characteristics within shopping sessions, while product embeddings require more frequent updates to reflect inventory changes and promotional activities. Contemporary analysis of caching strategies reveals that intelligent cache management policies can achieve hit rates exceeding 85% when properly configured for specific workload patterns, with particular effectiveness in read-heavy applications such as recommendation systems (Hasslinger, G. *et al.*, 2023). The caching system incorporates intelligent invalidation policies that balance freshness requirements with computational efficiency, ensuring users receive relevant recommendations without unnecessary computation overhead.

Effective caching strategy implementation requires a comprehensive understanding of access patterns and data freshness requirements. Research demonstrates that hybrid caching approaches combining time-based and content-based invalidation policies can optimize both cache hit rates and data consistency, particularly in dynamic environments where cached content experiences varying update frequencies (Hasslinger, G. *et al.*, 2023). The multi-tier caching architecture addresses different temporal characteristics of user preferences and product attributes, with session-level caches maintaining user embeddings for immediate access while background processes handle periodic refresh of product representations.

To manage high traffic during sales and launches, the system uses load balancing and horizontal scaling. The design features container deployment with auto-scaling, so computing resources change based on traffic. This is helpful for handling the unstable distribution of e-commerce traffic in the beauty industry, which can experience spikes in demand for certain types of products when things go viral on social media.

Publisher: SARC Publisher

The model infrastructure uses GPU acceleration for embedding calculations while keeping costs down by carefully allocating resources. The system uses model quantization and pruning to lower memory usage without sacrificing recommendation quality, allowing deployment on various hardware setups. Performance monitoring tracks metrics such as latency, accuracy, and resource usage across different users and traffic patterns.

Table 2: Performance Optimization and Scalability Features (Halder, R. K. *et al.*, 2024; Hasslinger, G. *et al.*, 2023)

Performance Aspect	Implementation Details
Response Time Target	Sub-fifty-millisecond processing
Search Algorithm	Approximate nearest neighbor methods
Caching Strategy	Multi-level approach with intelligent invalidation
Cache Hit Rates	Exceeding 85% when properly configured
Scaling Method	Containerized deployment with auto-scaling
Computing Infrastructure	GPU acceleration with model quantization

BUSINESS IMPACT AND USER EXPERIENCE ENHANCEMENT

The new dual-tower recommendation system has made things better for the business and improved how people shop. It facilitates seeing more clicks, longer sessions, and more views on product pages. This shows people liked the suggestions better than before. Studies say that using better recommendation tech can really help online stores. When AI methods are used instead of older ways of filtering, conversion rates might go up by 15% to 35%. These improvements led to more money coming in, with higher conversion rates and bigger orders from different groups of shoppers.

Modern AI recommendation systems that are modern have a big effect on how well an online store does because they make things more personalized. Research shows that when businesses use these technologies, customers get more involved. This works well in areas where personalization is very important, like fashion and beauty. The system can understand what people like and what products offer. This means it can match what customers need with what’s available, which helps the business do better in many ways.

User experience enhancements extended beyond simple metric improvements to qualitative changes in shopping behavior patterns. Analysis of user feedback and behavioral data revealed increased exploration of new product categories and brands, suggesting that the system successfully balanced familiarity with discovery. Comprehensive studies on personalized recommendation systems demonstrate that effective implementations can significantly influence user decision-making processes, with research showing that personalized recommendations increase user satisfaction scores by 20-40% while simultaneously reducing decision-making time (Sonawane, P. et al. 2024). Users reported higher satisfaction with product recommendations and demonstrated increased willingness to try suggested items, indicating improved trust in the recommendation system's understanding of preferences.

Personalized recommendation systems fundamentally alter user interaction patterns within e-commerce environments. Research demonstrates that effective personalization technologies enhance user experience through reduced cognitive load during product selection processes, while simultaneously increasing exploration of diverse product categories (Sonawane, P. et al. 2024). User behavior shows that when people get good, custom suggestions, they feel more sure about buying something. They are also more open to trying things they would not normally consider.

The system was able to deal with changing seasonal trends all by itself, which was a big plus for how the business runs. Beauty retail exhibits strong seasonal patterns in product demand, from summer skincare focus to winter hydration needs, and the automated trend detection eliminated the need for manual rule updates while improving recommendation relevance during trend transitions. This capability reduced operational overhead while ensuring users received seasonally appropriate product suggestions.

Diversity in recommendations increased significantly, with users discovering products across broader price ranges, brands, and categories than previous systems enabled. This diversity enhancement supported business goals of promoting emerging brands and new product launches while maintaining user satisfaction. The system's demographic fairness improvements also expanded market reach by ensuring high-quality

Publisher: SARC Publisher

recommendations for previously underserved user segments.

Table 3: Business Impact and User Experience Enhancement (Valencia-Arias, A. *et al.*, 2024; Sonawane, P. et al. 2024)

Impact Category	Enhancement Description
Conversion Rate Potential	15-35% increases with AI-driven approaches
User Satisfaction Improvement	20-40% increase in satisfaction scores
Engagement Metrics	Significant increases in click-through rates
Product Discovery	Increased exploration across categories
Market Expansion	Enhanced reach for underserved segments

CONCLUSION

The implementation of dual-tower neural networks for personalized product discovery in beauty retail represents a significant advancement in addressing the unique challenges of cosmetic and skincare recommendation systems. The architecture demonstrates exceptional capability in separately modeling user preferences and product characteristics while maintaining computational efficiency and demographic fairness across diverse user segments. The sophisticated training methods, including sequence-aware negative sampling and demographic-stratified testing, establish effective frameworks for mitigating algorithmic bias while preserving recommendation quality. Performance optimization strategies enable real-time processing of extensive product catalogs through intelligent caching mechanisms and approximate search algorithms that maintain system responsiveness. The measurable business impact, including enhanced user engagement, increased conversion rates, and expanded market reach, validates the practical value of advanced recommendation technologies in specialized retail domains. The system's ability to automatically adapt to seasonal trends and emerging preferences while maintaining demographic fairness establishes a foundation for future developments in personalized e-commerce recommendation systems. These achievements demonstrate that sophisticated artificial intelligence architectures can successfully address the complex requirements of beauty retail while delivering tangible business value and enhanced user experiences across culturally diverse demographic segments, paving the way for more inclusive and effective personalization technologies.

REFERENCES

1. Perret, J. K., & Schwientek, J. "Beauty Tech—Customer Experience and Loyalty of Augmented Reality-and Artificial Intelligence-Driven Cosmetics." *Digital* 5.2 (2025): 21.
2. Deldjoo, Y., Nazary, F., Ramisa, A., Mcauley, J., Pellegrini, G., Bellogin, A., & Noia, T. D. "A review of modern fashion recommender systems." *ACM Computing Surveys* 56.4 (2023): 1-37.
3. Zhu, Y., Xi, D., Song, B., Zhuang, F., Chen, S., Gu, X., & He, Q. "Modeling users' behavior sequences with hierarchical explainable network for cross-domain fraud detection." *Proceedings of the web conference 2020*. (2020).
4. Jeong, E., Li, X., Kwon, A., Park, S., Li, Q., & Kim, J. "A Multimodal Recommender System Using Deep Learning Techniques Combining Review Texts and Images." *Applied Sciences* 14.20 (2024): 9206.
5. Ma, H., Xie, R., Meng, L., Feng, F., Du, X., Sun, X., ... & Meng, X. "Negative sampling in recommendation: A survey and future directions." *arXiv preprint arXiv:2409.07237* (2024).
6. Zhang, F., Luo, W., & Yang, X. "Mitigating Selection Bias in Recommendation Systems Through Sentiment Analysis and Dynamic Debiasing." *Applied Sciences* 15.8 (2025): 4170.
7. Halder, R. K., Uddin, M. N., Uddin, M. A., Aryal, S., & Khraisat, A. "Enhancing K-nearest neighbor algorithm: a comprehensive review and performance analysis of modifications." *Journal of Big Data* 11.1 (2024): 113.
8. Hasslinger, G., Okhovatzadeh, M., Ntougias, K., Hasslinger, F., & Hohlfeld, O. "An overview of analysis methods and evaluation results for caching strategies." *Computer Networks* 228 (2023): 109583.
9. Valencia-Arias, A., Uribe-Bedoya, H., González-Ruiz, J. D., Santos, G. S., Ramírez, E. C., & Rojas, E. M. "Artificial intelligence and recommender systems in e-commerce. Trends and research agenda." *Intelligent Systems with Applications* 24 (2024): 200435.

Publisher: SARC Publisher

10. Sonawane, P. *et al.*, "Impact of Personalised Recommendation Systems on User Experience and Decision Making. - Using Amazon as a Case." *IJEMH*, (2024). https://ijemh.com/issue_dcp/Impact%20of%20Personalised%20Recommendation%20Systems%20on%20User%20Experience%20and%20Decision%20Making.%20Using%20Amazon%20as%20a%20Case.pdf

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:
Sundi, J." Personalised Product Discovery in Online Beauty Retail Using Dual-Tower Neural Networks." *Sarcouncil Journal of Engineering and Computer Sciences* 4.8 (2025): pp 437-443.

Publisher: SARC Publisher