

# IntelliShop A Unified AI Recommendation Framework for Adaptive Personalization in Modern Electronic Commerce

R. Swetha<sup>#1</sup>, Nikitha M<sup>\*2</sup>, Agalya S<sup>#3</sup>

# Department Of Artificial Intelligence And Data Science, The Kavery Engineering College, Salem, Tamil Nadu, India  
(Affiliated With Anna University)

[rswetha241201@gmail.com](mailto:rswetha241201@gmail.com), <sup>1</sup>, [nikitha18m@gmail.com](mailto:nikitha18m@gmail.com), <sup>2</sup>, [agalya15405@gmail.com](mailto:agalya15405@gmail.com), <sup>3</sup>

## Abstract

Most online retail platforms still infer customer preference mainly from overt actions such as clicks, cart additions, and purchases. Although these signals are useful, they offer only a partial view of user intent. Repeatedly ignored products, temporary shopping goals, and expired interests are often left unmodeled, which can lead to repetitive suggestions and outdated personalization. This paper introduces IntelliShop, an adaptive recommendation framework for electronic commerce that addresses these gaps through four coordinated components: User Silence Intelligence for learning from non engagement, a Self Learning Taste Reset Mechanism for reducing obsolete preferences, One Tap Mood Selection for session specific control, and Daily Interest Reset for maintaining freshness. A prototype system was implemented and tested in a pilot retail environment. The findings show better recommendation relevance, fewer recurring ignored items, and quicker response to changes in user behavior when compared with conventional recommendation approaches centered primarily on positive interaction data.

*Keywords—Electronic commerce, recommender systems, personalization, implicit feedback, user behavior modeling, adaptive ranking*

## I. INTRODUCTION

The expansion of electronic commerce has fundamentally altered how people discover and purchase products. Retail activity that once relied on physical availability and limited product assortment now takes place in digital marketplaces containing thousands, and sometimes millions, of items. This shift has widened consumer choice, but it has also introduced a new difficulty. When product volume grows beyond what users can reasonably inspect, finding relevant items becomes increasingly dependent on intelligent recommendation.

For this reason, recommendation engines now occupy a central place in online retail systems. They influence what customers notice, how long they browse, and whether they complete a purchase. Yet despite their importance, many current systems are still built on a narrow reading of behavior. They focus strongly on positive signals such as clicks, wishlists, cart additions, and purchases. These actions certainly matter, but they are not enough on their own. A shopper may repeatedly encounter a product and never engage with it. That absence of response is often meaningful, but in many deployed systems it is treated as if nothing happened.

This limitation produces several familiar problems. First, recommendations can become too narrow, reinforcing a

profile that no longer reflects the user's present needs. Second, past interests can remain active for too long, especially when they were tied to a seasonal event, a one time purchase, or a short lived shopping objective. Third, the system may continue surfacing items or categories that users have already shown little interest in, simply because the model lacks a mechanism for learning from silence. These concerns have been noted in earlier recommendation research, especially in work dealing with implicit feedback, temporal variation, and sequential behavior [1]–[4].

The issue becomes even more pronounced in large retail settings where shopping intent changes rapidly. Someone browsing gifts for a holiday, household items after relocation, or discounted apparel during a sale may generate interaction patterns that are highly temporary. If the system treats these signals as long term identity markers, later recommendations become less useful. What the user receives feels personalized on the surface, but in practice it may be tied to an outdated snapshot of behavior.

This paper is motivated by the need for recommendation systems that can adapt more naturally to how people actually shop. A useful recommender should remember past preferences, but it should also recognize fading interest, interpret repeated non response, and respond to immediate session level intent. IntelliShop was

developed with that objective. The architecture combines User Silence Intelligence, which interprets ignored exposures; a Self Learning Taste Reset Mechanism, which weakens stale preference traces; One Tap Mood Selection, which allows lightweight user guidance during a session; and Daily Interest Reset, which introduces controlled novelty into the recommendation stream.

The main contributions of this work are threefold. First, it proposes a unified recommendation architecture that merges long term history, session activity, negative implicit feedback, and periodic refresh behavior in one operational design. Second, it offers a practical strategy for using user silence and preference decay without resorting to hard profile resets. Third, it presents a prototype implementation and pilot evaluation showing that these mechanisms improve responsiveness and reduce repetitive recommendation behavior in a realistic web commerce setting.

## II. LITERATURE REVIEW

Work on recommender systems emerged from the broader challenge of information overload. Resnick and Varian [1] described recommendation as a filtering function that helps users navigate large collections of information. Their contribution established the foundational purpose of the field. However, the commercial environments considered at that time were much less dynamic than modern retail platforms, where repeated exposures, changing intent, and dense catalogs create additional complexity.

Adomavicius and Tuzhilin [2] later broadened the discussion by reviewing the next generation of recommender systems and emphasizing hybrid modeling, richer representations of users, and contextual information. Their survey remains influential because it moved the field beyond simple rating prediction. Even so, it did not directly address how retail systems should deal with repeatedly ignored products or how long dormant interests should continue to affect recommendations.

Collaborative filtering became one of the most widely used techniques in practice, and Su and Khoshgoftaar [3] summarized its memory based and model based forms. These methods infer preference through behavioral similarity and have proved valuable across many domains. Still, collaborative filtering typically performs best when there is enough positive interaction data. It is far less effective at distinguishing items a user has never seen from items the user has seen and consistently chosen not to engage with.

Latent factor methods strengthened collaborative recommendation further. Koren, Bell, and Volinsky [4] demonstrated that matrix factorization can uncover hidden user item relationships and improve prediction quality.

Although this was a major advance, the learned latent structure often reflects past behavior strongly and may preserve preference signals beyond their useful lifetime unless explicit temporal correction is added.

Rendle et al. [5] addressed implicit feedback ranking through Bayesian personalized ranking, a method designed for pairwise learning without explicit ratings. This made recommendation more suitable for click rich environments. However, most unobserved items remain grouped together within such frameworks, which means that lack of interaction is not carefully separated into non exposure and visible but ignored exposure.

Temporal variation was examined more directly by Koren [6], who showed that both users and items change over time. This is particularly relevant in retail, where intent is often seasonal, episodic, or event driven. Even so, temporal modeling on its own cannot tell whether the absence of engagement reflects declining interest, temporary mismatch, or repeated rejection after exposure.

Context aware recommendation introduced another important perspective. Adomavicius et al. [7] argued that recommendation decisions improve when they account for context such as time or situation. Their multidimensional framework remains valuable, especially for domains where context is well defined. Yet in many online stores, context is only partially observable, and users are rarely given easy ways to express their short term shopping purpose directly.

Research has also looked at the structure of recommendation lists rather than prediction alone. Ziegler et al. [8] showed that diversification can reduce redundancy and improve the browsing experience. This helps prevent monotonous recommendation panels, but diversity does not necessarily solve the deeper issue of stale profile influence. A diversified list may still revolve around product families the user has already moved beyond.

Adaptive online learning has also shaped the field. Li et al. [9] applied contextual bandits to personalized news recommendation and demonstrated the value of continuous exploration and feedback driven updating. The idea is attractive for dynamic systems, but retail platforms often face sparse conversions and delayed signals. Purchases are less frequent than page views, which complicates rapid learning from direct reward.

Deep learning later expanded the range of recommendation techniques. Wang et al. [10] integrated representation learning with collaborative modeling, while Hidasi et al. [11] showed that recurrent architectures can capture short session sequences effectively. Covington et al. [12] described an industrial scale deep recommendation pipeline, and He et al. [13] proposed neural collaborative filtering for richer nonlinear interaction modeling. These systems improved ranking capability considerably,

especially at scale. Nevertheless, most of them still rely predominantly on positive interaction traces and do not explicitly incorporate repeated non engagement as a structured negative signal.

Recent work on calibration and sequence awareness offers further insight. Steck [14] argued that recommendation lists should remain aligned with the user's broader preference distribution instead of optimizing only for immediate relevance. Quadrana et al. [15] reviewed sequence aware systems and highlighted the importance of order, recency, and transitions within user behavior. These ideas advance the field, but two practical questions remain open in many electronic commerce settings. The first is how to interpret highly visible but repeatedly ignored items. The second is how to reduce the influence of interests that were once relevant but are no longer active.

The gap in the literature therefore lies not in the absence of strong recommendation methods, but in the lack of integration across several practical signals. Existing studies have improved similarity modeling, contextual reasoning, sequential learning, and deep representation. What is still missing in many retail systems is a unified design that simultaneously handles ignored exposure, adaptive preference decay, direct session guidance, and periodic freshness. IntelliShop addresses this unmet need by bringing these elements together within a single recommendation framework intended for real web commerce environments.

### III. PROPOSED SYSTEM

IntelliShop is designed as a recommendation framework that preserves useful user history without becoming trapped by it. Rather than depending on one static ranking formula, the system maintains a continuously updated user state shaped by long term behavior, current session activity, repeated non engagement, temporary intent, and controlled freshness signals. The model is intended for practical electronic commerce deployments where users browse, scroll, search, compare, and occasionally perform stronger actions such as saving or purchasing products.

In operational terms, IntelliShop receives behavioral data from the storefront, updates user representations in near real time, assembles a candidate set of products, and then applies a staged ranking procedure. The novelty lies in how that ranking is constructed. Positive interactions still contribute strongly, but they no longer dominate without moderation. A product type that was previously important can lose weight if it is no longer reinforced. Repeatedly ignored products are pushed down. Session mood can temporarily redirect the list, and daily freshness prevents the system from becoming too repetitive.

#### A. USER SILENCE INTELLIGENCE

User Silence Intelligence is meant to capture something recommendation systems often miss: non response can carry information. In retail environments, users rarely take the time to explicitly reject a product. More often, they express disinterest by simply not acting. This behavior is easy to overlook if the system tracks only positive engagement.

The module first determines whether an item was meaningfully visible. A product counts as a qualified exposure only if it remained on screen under conditions where the user could reasonably have noticed it. This distinction matters, because a fast scroll past an item should not be interpreted in the same way as a clearly visible product card. Once qualified exposure is established, the absence of any response is recorded as silence. If similar items continue to be shown and continue to receive no engagement, the system increases a negative preference signal not only for the item itself but also for related features such as category, brand, or price band.

Silence is not treated as a permanent judgment. The signal weakens over time, since disinterest can be situational. A user who ignores travel accessories during routine browsing may later return with a specific need. For that reason, the module penalizes persistent repetition rather than banning items outright.

#### B. SELF LEARNING TASTE RESET MECHANISM

The Self Learning Taste Reset Mechanism addresses the problem of inactive preferences remaining too influential. Recommendation profiles should retain meaningful long term structure, but they should also make room for change. This becomes important in retail because user interests often fluctuate with season, occasion, or immediate necessity.

The mechanism monitors how recently a preference has been reinforced and whether related recommendations are still receiving attention. When an interest that was once strong stops attracting engagement, and associated items begin to accumulate silence, its weight is reduced gradually. The process is adaptive rather than fixed. Some users show stable purchase habits over long periods, whereas others move quickly between categories. The system therefore adjusts the rate of decay according to observed behavioral volatility.

A key design choice is that the mechanism does not erase earlier preferences completely. Instead, it lowers their influence until new behavior either confirms or revives them. This allows the system to forget selectively without destroying potentially valuable historical knowledge.

#### C. ONE TAP MOOD SELECTION

Although recommendation engines often infer user intent from behavior alone, many shopping sessions are driven by short lived goals that are difficult to detect immediately. A person might be looking for everyday essentials, promotional bargains, premium goods, or gift ideas. Long term interaction history offers only partial help in such cases.

One Tap Mood Selection gives the user a lightweight method for shaping the active session. A chosen mood is attached to the current browsing context and passed to the ranking layer as a temporary signal. The recommender then increases the priority of items that match that session level objective. For example, a cost conscious mood can favor discounted or lower priced products, while a gifting mood can emphasize items with broad appeal or stronger social validation.

Because the signal is temporary, it does not distort the long term user profile. Once the session ends or the context changes, the effect fades. In this way, the system gains responsiveness without sacrificing profile consistency.

#### D. DAILY INTEREST RESET

Daily Interest Reset was introduced to preserve freshness across repeated visits. The goal is not to rebuild the user profile from scratch every day. Instead, the mechanism ensures that previous browsing patterns do not fully dictate future recommendations when the user returns with a different objective.

At the start of a new activity cycle, the engine allocates a limited portion of recommendation space to products that are new, trending, recently updated, or not yet heavily explored by the user. These items are not inserted randomly. They must still satisfy broad compatibility checks so that personalization is not lost. What the reset provides is breathing room. It allows the system to test adjacent interests, avoid visual repetition, and surface possibilities that would otherwise remain buried under historical preferences.

This feature is especially helpful in reducing recommendation fatigue, a common problem in mature retail platforms where returning users see similar lists again and again.

#### E. UNIFIED RECOMMENDATION LOGIC

The four modules operate within a shared scoring scheme. Each candidate item receives a composite score derived

from several interacting signals. Historical affinity contributes a baseline estimate based on past clicks, saves, cart additions, and purchases. Session relevance captures the user's immediate browsing direction. A mood based adjustment applies when the current session includes explicit mood input. A freshness component introduces selected opportunities for discovery. Against these positive signals, the system applies penalties derived from repeated silence and from preference staleness identified by the reset mechanism.

As a result, the final ranking reflects more than simple interaction accumulation. It combines memory with reinterpretation. The system recognizes what the user has liked, notices what is no longer working, and leaves room for emerging intent. That combination defines IntelliShop's central design philosophy.

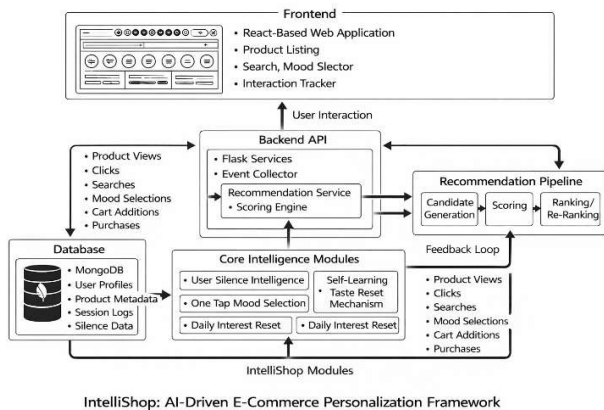
#### IV. SYSTEM ARCHITECTURE

The architecture of IntelliShop can be understood as a sequence of interconnected layers. At the front end sits the interaction layer, where the web application records impressions, searches, clicks, mood choices, cart operations, and purchases. The next layer is responsible for event processing. It validates visibility conditions, records time information, and converts raw behavior into structured events suitable for modeling. The third layer maintains the user profile, including long term preferences, silence statistics, recent activity traces, and active session context. The fourth layer contains the recommendation engine, which handles candidate generation, score computation, and final reordering. The last layer delivers ranked products back to the storefront and records response signals for future updates.

The data flow begins when the user interacts with the store interface. These interactions are sent to the backend through service endpoints. The backend updates the user state, retrieves candidate products from several sources, and executes the ranking logic. Candidate generation can draw on collaborative similarity, product content, popularity trends, and freshness pools. Once assembled, those candidates are processed through the IntelliShop modules, which adjust ranking according to silence patterns, temporal decay, mood signals, and freshness policy.

An advantage of this structure is its modularity. The framework can operate as a standalone recommendation pipeline, but it can also be placed above an existing candidate generator. In other words, retailers do not necessarily need to replace their current recommender entirely. IntelliShop can function as an adaptive ranking layer that strengthens an existing system.

#### Fig. 1. System Architecture Diagram



IntelliShop: AI-Driven E-Commerce Personalization Framework

## IMPLEMENTATION

A prototype version of IntelliShop was developed as a service based web application. The backend used Flask because it offers a compact and flexible environment for API driven development. Separate endpoints were defined for event collection, user state updates, candidate retrieval, and ranking output. This arrangement made it easier to test module behavior independently during development.

The user interface was created with React in order to support dynamic rendering and continuous interaction capture. Product cards, category pages, and session controls were implemented as reusable components. A small mood selection control was incorporated into the browsing interface. React also simplified the integration of visibility tracking, which was necessary for identifying qualified impressions without significantly changing the user experience.

MongoDB served as the primary data store for interaction history, session information, product attributes, and evolving user profiles. Its document based structure proved useful for representing behavioral data that varies across sessions and users. The system also included scheduled background processes for recalculating stale preference weights and refreshing daily novelty pools. Recommendation responses were exposed through REST style endpoints so the prototype could integrate with a broader commerce application.

The implementation was deliberately modest in scale. It was built to validate the design in a realistic environment rather than to simulate a large industrial platform.

## V. ALGORITHMIC DESIGN

The ranking strategy in IntelliShop is based on combining multiple forms of evidence rather than relying on a single predictive signal. Different user actions contribute different levels of support. Purchases receive the highest positive

weight because they reflect clear preference realization. Cart additions and wishlist saves are treated as strong indicators as well, though slightly below purchase intent. Clicks and product views provide weaker but still useful evidence, especially when they recur.

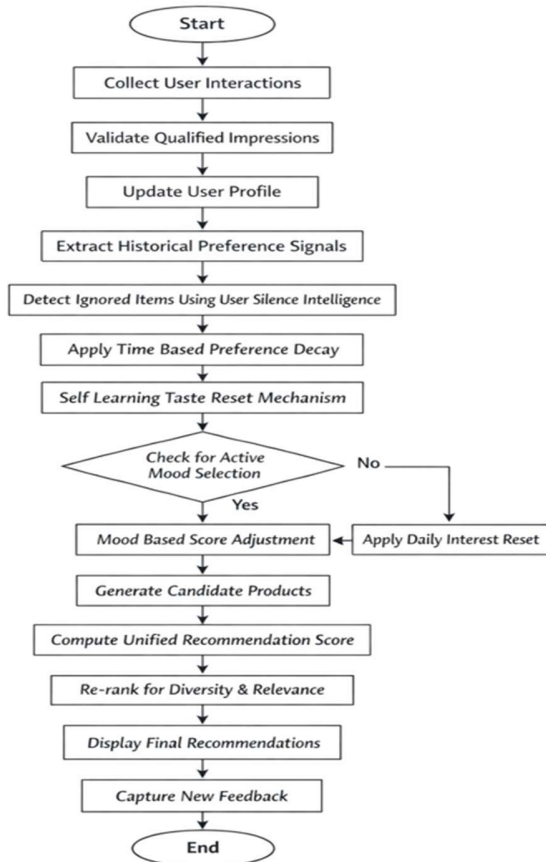
Negative implicit feedback is generated through the silence module. After a product receives qualified exposure, the absence of any action increases its silence count. If the same item, or items sharing similar characteristics, continue to be shown without eliciting engagement, the penalty becomes stronger. The model is intentionally gradual. One ignored display is not enough to suppress an item decisively, since the user may simply have been distracted or browsing casually.

Temporal adjustment is handled through the taste reset mechanism. Interaction signals lose influence when they are not reinforced over time. The decay function is personalized. Users with steady shopping habits experience slower weakening of historical preferences, whereas users with more volatile browsing patterns see faster profile adjustment. This distinction helps the model separate enduring taste from short bursts of activity.

Mood based adjustment acts as a bounded session specific modifier. When a user selects a session mood, the scoring function temporarily boosts items that fit the selected intent. The magnitude of this adjustment is controlled so that recommendations do not become disconnected from the user's broader history. Daily Interest Reset then inserts a limited freshness factor, ensuring that a portion of the ranked list is available for exploration and novelty.

In summary, the algorithm proceeds in stages. It estimates preference from long term and short term positive interactions, subtracts penalties for repeated silence and stale interests, applies session mood and freshness adjustments, and then performs a final reordering step to avoid excessive repetition from the same category or brand. The objective is not merely to maximize immediate click probability, but to sustain usefulness over repeated sessions.

**Fig. 2. Recommendation Decision Flow Diagram**



**WORKFLOW AND PROCESS FLOW**

The system workflow begins when a user arrives at the online storefront and starts browsing. As the session unfolds, the front end captures product visibility, clicks, search activity, and any selected mood input. These records are then transmitted to the backend for processing. At this stage, the system distinguishes between incidental page exposure and meaningful visibility so that later silence based reasoning is grounded in realistic evidence.

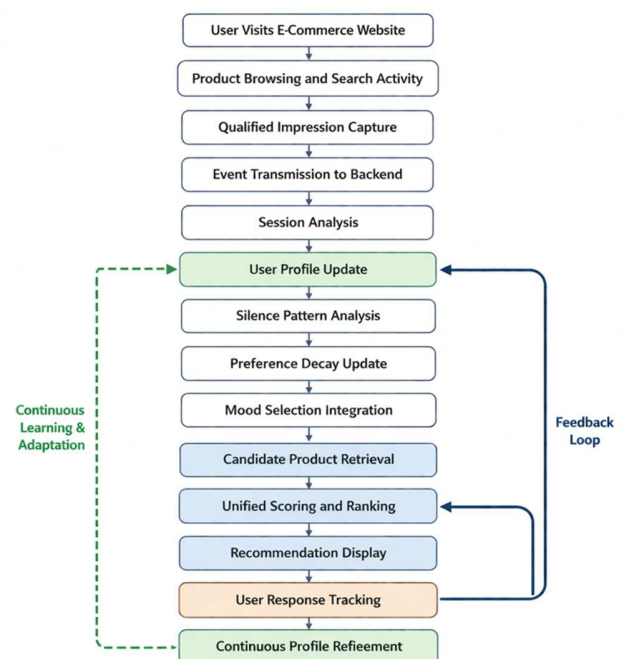
Next, the profile management component updates the user state. It refreshes affinity values, increments silence counters where appropriate, and checks recency patterns to determine whether some preferences should begin to lose weight. If the user has chosen a mood, that information is attached to the live session profile. In parallel, the candidate generation layer retrieves products from several channels, including behavioral similarity, content matching, trend pools, and freshness reserves.

Once the candidates have been assembled, the ranking engine computes a unified score for each item. Historical fit and current session behavior form the base. Mood alignment and freshness may raise selected items, while repeated silence and inactive preferences reduce others.

After scoring, a final list shaping stage prevents the recommendation panel from becoming too concentrated around a single brand or narrow category. The ranked output is then sent back to the user interface, where it becomes part of the next interaction cycle.

The workflow is designed so that most of the complexity remains hidden from the shopper. From the user’s perspective, the system behaves more responsively while requiring little more than ordinary browsing and, when desired, a simple mood selection.

**Fig. 3. System Workflow Diagram**



**RESULTS AND DISCUSSION**

The IntelliShop prototype was assessed in a pilot electronic commerce environment covering fashion, electronics, and home related products. The evaluation dataset contained roughly one and a half thousand users, more than eight thousand items, and six weeks of anonymized interaction records. Initial sessions were used to establish user profiles, while later sessions supported comparative testing. The purpose of the study was not to claim large scale industrial performance, but to determine whether the proposed architecture improves recommendation quality under realistic operating conditions.

Two baseline systems were used for comparison. The first was a popularity driven recommender that ranked items according to overall engagement. The second was a conventional implicit feedback model based on clicks, cart events, and purchases, but without any mechanism for silence interpretation, adaptive preference decay, or

explicit session control. IntelliShop was evaluated against both baselines using user oriented measures including click through behavior, estimated relevance of recommended items, time required to adapt after a shift in interest, and the recurrence of items that had been repeatedly ignored.

The experimental observations were promising. Compared with the popularity baseline, IntelliShop produced clearly stronger first page engagement. Against the click centered baseline, the improvement was smaller but remained consistent. More notably, the proposed framework reduced the repeated appearance of products that had already been shown and ignored in earlier sessions. That reduction matters because recommendation quality is not defined only by higher clicks. It also depends on avoiding visible repetition that weakens user trust.

The prototype adapted more quickly when users moved from one interest area to another. This was especially noticeable in sessions linked to temporary shopping purposes such as holiday gifts, promotional browsing, or short seasonal needs. The mood selection feature was not used in every session, yet the sessions where users did employ it showed better early engagement, suggesting that even minimal explicit control can strengthen fit when intent is situational.

The main evaluation observations are summarized in Table I.

**TABLE I. PERFORMANCE EVALUATION METRICS**

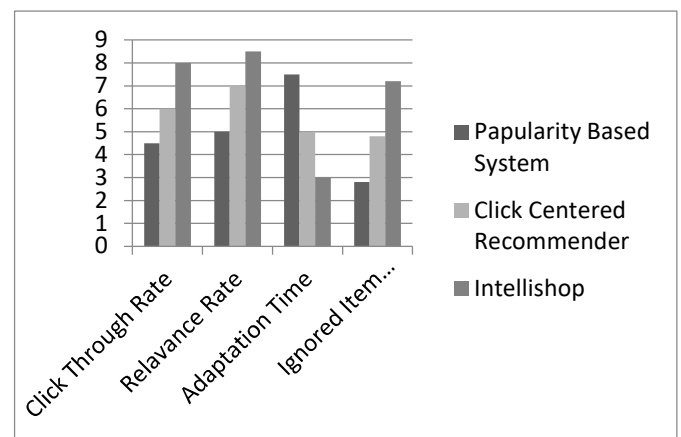
Metric	Description	Result
Click Through Rate	Level of user engagement	Higher
Relevance Rate	Quality of recommendation matching	Strong
Adaptation Time	Speed of profile adjustment	Shorter
Ignored Item Reduction	Decline in repeated irrelevant exposure	Marked

A closer look at module level behavior suggested that User Silence Intelligence contributed most directly to lowering repetition of unwanted products. The Self Learning Taste Reset Mechanism had the strongest influence on adaptation speed, particularly when earlier interests had become inactive. Daily Interest Reset improved list freshness and increased exposure to a wider set of products, although its direct contribution to immediate click rate was less pronounced. The overall improvement came from the

interaction among modules rather than from any single component in isolation.

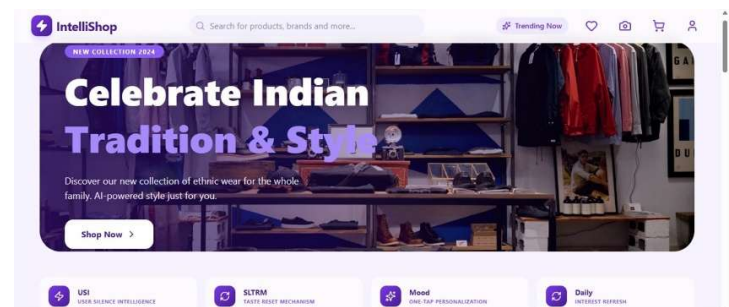
These results should nevertheless be interpreted with care. The study was carried out on a prototype platform, not a full production scale system. In addition, the quality of silence based modeling depends heavily on accurate impression measurement. If visibility tracking is unreliable, non response may be misclassified. Even with these limitations, the findings indicate that the proposed framework offers a credible and practical path toward more adaptive retail recommendation.

**Fig. 4. Performance Comparison Graph**

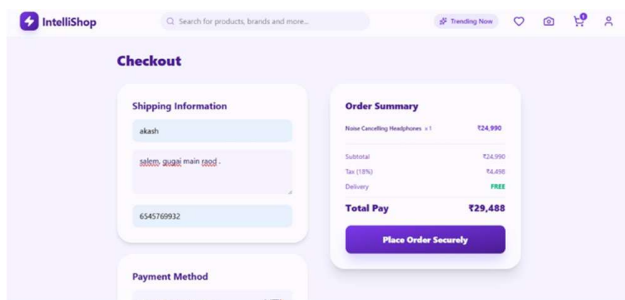


## VI. USER INTERFACE / SCREENSHOTS

**Fig. 5. Main Dashboard / Home Screen**



**Fig. 6. Output / Result Screen**



## ADVANTAGES AND LIMITATIONS

### A. Advantages

A major strength of IntelliShop is that it treats user behavior in a more balanced manner than systems that depend almost entirely on overt positive actions. Because online shoppers often communicate disinterest through inaction rather than explicit rejection, the ability to learn from repeated silence helps reduce unnecessary repetition and improves the perceived relevance of recommendation lists. The taste reset component further enhances adaptability by preventing outdated interests from dominating the profile for too long. This is particularly useful in retail domains where user intent changes with time, season, and occasion.

Another benefit is the inclusion of a lightweight mood based control. Instead of forcing users to edit preferences manually, the system allows quick session level adjustment. The daily reset mechanism contributes a further advantage by preserving freshness and encouraging exploration without abandoning personalization. From a practical standpoint, the architecture is also flexible enough to sit on top of an existing recommendation stack, which lowers the barrier to adoption.

### B. Limitations

The current design also has several constraints. User Silence Intelligence depends on reliable visibility tracking at the interface level, and weak impression logging can reduce the accuracy of negative implicit feedback. New users remain challenging because they offer little behavioral evidence for either positive modeling or silence interpretation. The usefulness of mood selection also depends on whether users choose to engage with it, which may differ across applications and audiences.

In addition, the prototype has been tested at moderate scale only. Its performance under high traffic, low latency production conditions has yet to be studied in depth. Daily

Interest Reset also requires careful calibration. If novelty is emphasized too strongly, recommendation precision may suffer; if it is tuned too lightly, the system may still appear repetitive.

## VII. CONCLUSION AND FUTURE WORK

This study introduced IntelliShop, a unified recommendation framework for adaptive personalization in electronic commerce. The work was motivated by a practical weakness in many existing recommenders: they are effective at learning from clicks and purchases, but far less capable of interpreting ignored items, fading interests, and shifting session intent. To address this limitation, IntelliShop integrates four complementary components: User Silence Intelligence, Self Learning Taste Reset Mechanism, One Tap Mood Selection, and Daily Interest Reset. Rather than presenting a single new prediction model, the paper offers a structured recommendation architecture that combines long term memory, selective forgetting, session awareness, and controlled freshness. The prototype evaluation indicates that this combination can improve recommendation relevance, reduce repeated exposure to ignored products, and shorten the time needed for the system to adjust to changing user needs. The most distinctive aspect of the design is the coupling of silence based learning with adaptive taste decay, since together they allow the system to respond both to present non engagement and to historical obsolescence.

Several directions remain open for future work. More advanced candidate generation and sequence modeling techniques could be incorporated to strengthen recommendation depth. Real time deployment at larger scale should also be explored in order to assess latency, throughput, and online learning behavior. Further extensions may include richer contextual features, stronger privacy aware profile management, and more nuanced user controls beyond a single mood signal. These developments would move recommendation systems closer to the realities of actual shopping behavior, where preferences shift, contexts change, and relevance depends as much on timing as on memory.

## ACKNOWLEDGMENT

The authors would like to thank the development contributors and student participants who supported the pilot environment used for implementation and evaluation.

## REFERENCES

- [1] P. Resnick and H. R. Varian, "Recommender systems," *Communications of the ACM*, vol. 40, no. 3, pp. 56–58, 1997.
- [2] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems A survey of the state

of the art and possible extensions,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 6, pp. 734–749, 2005.

[3] X. Su and T. M. Khoshgoftaar, “A survey of collaborative filtering techniques,” *Advances in Artificial Intelligence*, vol. 2009, pp. 1–19, 2009.

[4] Y. Koren, R. Bell, and C. Volinsky, “Matrix factorization techniques for recommender systems,” *Computer*, vol. 42, no. 8, pp. 30–37, 2009.

[5] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt Thieme, “BPR Bayesian personalized ranking from implicit feedback,” in *Proceedings of the Twenty Fifth Conference on Uncertainty in Artificial Intelligence*, 2009, pp. 452–461.

[6] Y. Koren, “Collaborative filtering with temporal dynamics,” in *Proceedings of the Fifteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2009, pp. 447–456.

[7] G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin, “Incorporating contextual information in recommender systems using a multidimensional approach,” *ACM Transactions on Information Systems*, vol. 23, no. 1, pp. 103–145, 2005.

[8] C. N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen, “Improving recommendation lists through topic diversification,” in *Proceedings of the Fourteenth International Conference on World Wide Web*, 2005, pp. 22–32.

[9] L. Li, W. Chu, J. Langford, and R. E. Schapire, “A contextual bandit approach to personalized news article recommendation,” in *Proceedings of the Nineteenth International Conference on World Wide Web*, 2010, pp. 661–670.

[10] H. Wang, N. Wang, and D. Y. Yeung, “Collaborative deep learning for recommender systems,” in *Proceedings of the Twenty First ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 1235–1244.

[11] B. Hidasi, A. Karatzoglou, L. Baltrunas, and D. Tikk, “Session based recommendations with recurrent neural networks,” in *Proceedings of the International Conference on Learning Representations*, 2016, pp. 1–10.

[12] P. Covington, J. Adams, and E. Sargin, “Deep neural networks for YouTube recommendations,” in *Proceedings of the Tenth ACM Conference on Recommender Systems*, 2016, pp. 191–198.

[13] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T. S. Chua, “Neural collaborative filtering,” in *Proceedings of*

*the Twenty Sixth International Conference on World Wide Web*, 2017, pp. 173–182.

[14] H. Steck, “Calibrated recommendations,” in *Proceedings of the Twelfth ACM Conference on Recommender Systems*, 2018, pp. 154–162.

[15] M. Quadrana, P. Cremonesi, and D. Jannach, “Sequence aware recommender systems,” *ACM Computing Surveys*, vol. 51, no. 4, pp. 1–36, 2018.