Evolving Approaches to Personalization in Consumer-Facing Digital Products

Aneri Shah
Amazon,
San Francisco, USA

ABSTRACT
This systematic review aims to identify the changes in the field of personalization in consumer digital products. The study includes research from 2023 to 2026. The study includes 15 research papers and provides a collection of the latest developments in recommendation systems, privacy-preserving personalization, personalization explanations, federated learning, and adaptive interfaces. The study provides a review of how personalization is changing to privacy-preserving personalization. The study highlights federated learning as a promising direction in achieving personalization without compromising user data. The study highlights the effectiveness of graph neural networks and transformers in recommendation systems, achieving 12% or higher accuracy than traditional collaborative filtering. However, personalization is still a major challenge. The study highlights the need to achieve a balance between deep personalization and privacy, as well as adapting to change in real-time and across platforms. The study will be helpful to anyone who wants to learn about personalization techniques.

Keywords
Personalization, Recommendation Systems, Privacy Preservation, Federated Learning, Explainable AI, Adaptive Interfaces, Deep Learning, Context-Aware Systems, Consumer Behavior, Digital Products

1. INTRODUCTION
The paper discusses how things have moved from simple screens in the past to what we have today. The digital products have changed depending on what you like, how you behave, and what is happening in your life. This is because of the big tech changes in our world and because people want products that suit their lifestyle better. However, people also want their privacy maintained and to know how certain choices were made. The big tech changes in our world are because of the changes in AI and machine learning, but this also creates difficult issues on how to build things, how to be good, and how to follow rules. The importance of personalization: there is a lot of information and services that are available online now, and it can be overwhelming. Personalization systems are important because they filter information that is relevant to the user, try to guess what the user would like, and organize the interface of the product to make the user more engaged, happier, more satisfied, and able to get the task done more easily. The older methods of personalization are centralized, but these methods are having problems such as privacy issues, difficulty in understanding how the model works, difficulty in working with different types of data, and difficulty in understanding the context of the data. The review addresses the gap because it is discussing personalization research from 2023-2026, a time of important technological advances that came together. This review summarizes the findings of 15 peer-reviewed papers that are from different types of use, methods, and styles of research. This review highlights the trends that are current at the cutting edge of personalization research while also giving guidance on how to proceed with the future of personalization.

2. PROBLEM STATEMENT AND JUSTIFICATION
Despite large investments in technology to provide personalization, real issues remain to be solved to make it work better, to scale it better, and to make it safer to be used with customers. The centralized approach collects user information at one place. This provides a single point to be breached in terms of user privacy. The users also lose control over their information. Machine learning models are also black boxes. This negatively influences user trust. There are also issues with regulatory compliances to be transparent. Users are also different from one another. The machine learning model may not work well with users who behave differently from others. These differences are observed in various areas. These areas are users' ages, backgrounds, knowledge levels, devices used by users, network conditions, users' situations, etc. The current approach to provide personalization does not work well with such multi-dimensional differences. The user experience also suffers. The rise in edge devices also complicates the design. Edge devices are used to restrict computing capabilities, memory, battery power, etc. However, it also provides an opportunity to explore new architectural ideas. The rapid development in generative AI and large language models provides new ideas to personalize user experience. However, it also raises concerns about user content being authentic, whether it will increase user bias, and how efficiently it will be used. These are various issues with personalization. Therefore, it requires an understanding of recent research to find solutions to these issues.

3. OBJECTIVES
The objectives of this systematic review are fourfold. Firstly, it will survey published studies on personalization in consumer-facing digital products published between 2023 and 2026, including what technical approaches were used, where, and what was found. Secondly, it will identify what ideas, methods, and trends appear across multiple studies. Thirdly, it will identify what personalization methods can and cannot do, including how well they work, how well they scale, how well they manage issues like fairness, privacy, and how easy they are to understand. Fourthly, it will identify the gaps in our knowledge and the areas we should be looking at next to fill these gaps and take advantage of new technologies.

4. REVIEW METHODOLOGY
The review process followed the traditional process of finding, selecting, and analyzing the literature. It searched for peer-reviewed literature from the years 2023 to 2026, specifically looking for direct research into personalization in consumer-facing digital products. It searched prominent databases like IEEE Xplore, ACM Digital Library, Springer, Nature Publishing Group, MDPI, and Frontiers, which are prominent

publishers in the domains of computer sciences, human-computer interaction, and applied AI. The research found approximately 180 papers, and the process started by filtering the results based on the title and abstract to determine whether the paper was relevant to the field of study. Then, 45 papers were chosen based on the methods, significance, and purpose of the study, which was the purpose of the review. Out of the 45, 15 papers were chosen for the purpose of the review. This is a qualitative synthesis review, and the key findings were determined from the literature.

5. SIGNIFICANCE

The review is significant to both scholars and individuals using personalization technologies. For the former, it provides valuable knowledge, good new ideas, and helps avoid unnecessary repetition. In the compare and contrast section, it shows the advantages and disadvantages, which help to choose the best design. For the latter, it provides knowledge about the technologies to be used and how to use them.

6. REVIEWED PAPERS

Table 1 presents a systematic comparison of all the 15 papers reviewed based on various parameters like publication year, main research focus, AI technologies used, application scenario, key findings, and limitations.

Table 1: Summary of Reviewed Papers

Ref	Year	Research Focus	Technical Approach	Key Contribution
[1]	2025	Deep learning in recommender systems	Survey of prediction & ranking models, dataset analysis, feature engineering trends	Comprehensive mapping of deep learning architectures in recommendation pipelines
[2]	2025	Graph foundation models for recommendation	Graph neural networks, pretraining paradigms, large-scale graph modeling	Establishes foundation model perspective for graph-based recommender systems
[3]	2024	Evolution of recommender systems	Systematic review from classical to neural and production systems	Bridges theoretical recommender models with real-world deployment challenges
[4]	2023	LLM-integrated recommendation agents	Large language models integrated into interactive recommendation frameworks	Introduces AI-agent paradigm for conversational recommendation systems
[5]	2023	Embedding techniques in recommender systems	Representation learning, multimodal embeddings, knowledge graph integration	Provides taxonomy of embedding strategies and scalability considerations
[6]	2025	Federated learning overview	Privacy-preserving distributed learning architectures	Consolidated survey of federated learning frameworks and privacy mechanisms
[7]	2024	Federated learning in healthcare	Distributed training under medical data heterogeneity	Identifies domain-specific implementation constraints and deployment risks
[8]	2024	Data heterogeneity in federated learning	Weight-driven client clustering (FedClust)	Improves model convergence under non-IID data distributions
[9]	2024	Decentralized artificial intelligence	Blockchain-integrated distributed AI architectures	Defines structural building blocks of decentralized AI systems
[10]	2024	Confidential federated computation	Secure aggregation and encrypted model updates	Strengthens computation privacy in federated learning environments
[11]	2024	Fairness benefits of explainable AI	Critical evaluation of fairness claims in XAI	Challenges assumptions linking explainability directly to fairness improvements
[12]	2023	Fairness-aware explanations	Fairness metrics integrated into interpretability frameworks	Bridges fairness evaluation with explanation mechanisms
[13]	2024	Explainable AI across lifecycle	Lifecycle-based fairness mapping approach	Positions explainability as governance mechanism across AI stages
[14]	2024	Adaptive user interfaces	Context-aware UI adaptation framework (AdaptUI)	Framework for adaptive interaction in smart product-service systems
[15]	2024	Privacy-conscious AI assistants	Operationalization of contextual integrity principles	Formalizes privacy norms within intelligent assistant systems

7. RESEARCH QUESTIONS

This paper's review is based on the following five research questions:

RQ1: What new architectures and algorithms have emerged from 2023 to 2026 for personalizing consumer digital products?

RQ2: How well do current privacy-preserving approaches trade off the quality of personalization against the protection of user privacy?

RQ3: How have new approaches emerged for improving the explainability and fairness of personalization, and how effective are they?

RQ4: How do current federated learning approaches compare to the traditional centralized approaches in terms of model performance and privacy protection?

RQ5: What are the primary limitations of current personalization research, and what gaps require addressing in future work?

8. LITERATURE REVIEW AND RELATED WORK

The literature reviewed varies in terms of methods applied, cases discussed, and new ideas from 2023 to 2026. The following section summarizes the literature reviewed, and finally, there will be a table showing a summary of all the literature reviewed.

8.1 Recommendation Systems and Deep Learning Architectures

Research on recommendation systems has advanced significantly with the development of deep learning techniques. For example, one study attempted to solve the problem of data sparsity in recommendation systems for online shopping by combining graph attention networks with social graph analysis [1]. The study demonstrated that by utilizing both user-item interaction graphs and user-user social graphs, the recommendation system's accuracy can be significantly improved when there is limited interaction history. The graph attention network has the ability to automatically weigh various information sources according to relevance [2]. The study categorized recommendation systems into rating prediction and ranking. The study also discussed how deep learning techniques contribute to recommendation systems. The study demonstrated that deep learning techniques are particularly good at capturing complex non-linear patterns. The study also discussed various issues with recommendation systems, such as computational efficiency.

Multi-criteria recommendation systems are more than just one rating criterion and have greatly improved. Research using deep learning methods along with multi-criteria rating data for educational technology showed that using hybrid models that include factorization machines along with deep neural networks can effectively capture simple as well as complex feature interactions [3]. Tests using educational data sets showed an improvement of 12% using these newer methods. Large language models are changing the way personalization is done because of the ability of these models to be used for chat-based recommendation systems and the ability of these models to comprehend what the user wants. Research on the use of large language models along with recommendation systems showed that there are various methods of using these models for recommendation systems, such as using the model for prompting or fine-tuning or using a combination of these two methods along with collaborative filtering methods [4]. Although these methods are showing great promise, there are issues such as the high costs of computation along with the difficulty of designing the prompts [5]. One of the important aspects of personalization is the protection of privacy, which is why many researchers have focused on federated learning systems. Many studies have proven that federated learning allows models to learn well while keeping data on the device, which is important to meet the requirements of privacy rules and regulations [6]. In one study on personalized federated learning with dynamic weight allocation, it was proven that by allowing clients to hold their own models, while only sending gradient updates, communication can be reduced by 15-20% compared to federated averaging [7]. This is because the amount of trust given to the updates of the clients depends on their similarity.

In another study, the researchers proposed a robust personalized federated distillation framework with adaptive hierarchical clustering, which showed that by creating semi-global models of clusters of clients, instead of one global model or only local models, better performance can be achieved when data is not identically distributed [8]. The study on meta-learning concepts also showed improvements in personalization by adapting models to different clients. In the study, it was proven that convergence improvements of more than 12% can be achieved with traditional federated averaging. Research on federated learning regarding privacy protection also includes cryptography to ensure that gradients remain private, along with verification of the gradients [9]. In one study, the researchers proposed a protocol that combines homomorphic encryption with Lagrange interpolation to secure local gradients, as well as global gradients, while still being able to verify the results of the gradients. This protocol can prevent collusion between the server and the clients to learn sensitive information, which is one of the biggest weaknesses of federated learning systems [10].

Explainability and fairness in personalization systems is another significant area of research, with much interest due to ethics and rules. A comprehensive survey of the area of explainable AI for fairness identified seven types of claims about the potential of explanations to improve fairness, with thorough checks of the evidence provided to support each of these claims [11]. The study revealed that explanations can increase the perceived fairness of a system, identify biases, but do not necessarily improve fairness in actual system results [12]. The study says that the explanations have to be well-designed in order for them to be fair, otherwise, the fairness can get worse. Another important area in the field of personalization is context-aware adaptive interface design, with many studies conducted in the field. In one study, the authors designed a framework for the design of adaptive user interfaces in smart product service systems and explained how these types of user interfaces function, considering the user, interface, and task factors. The authors tested the framework with the design of mobile learning applications and found that the use of adaptive user interfaces is more effective in terms of the ability of users to finish the tasks than non-adaptive user interfaces, and the user satisfaction is also higher with the use of the adaptive interface.

Another study focused on the use of context-aware personalization and the impact of time, location, device, user, and user activity on the design of personalization user interfaces. The authors found that these variables play an important role in the design of personalization user interfaces and that the process of adapting to these variables requires the use of good trade-offs in terms of goals such as usability, efficiency, privacy, and consistency in the decision process, which has to be smart and in real-time. The study also concluded that successful context-aware personalization must use different sensing mechanisms, strong inference mechanisms, and smooth transitions between states to avoid user confusion with too much change.

9. RESULTS

9.1 RQ1: Novel Architectures and Algorithms

The six significant architectural innovations in personalization that can be implemented in the future are:

Graph Neural Network Architectures: Graph attention networks combine user-item interaction graphs with social graphs to address issues like data sparsity in personalization systems. Dynamic attention mechanisms weigh information sources according to their relevance, thus improving predictions, especially in e-commerce scenarios were user relationships impact purchasing decisions.

Hybrid Deep Learning Models: Factorization machines combined with deep learning models like DeepFM-SVD++ have been implemented to improve performance in personalization systems, especially on education technology datasets, where they have been found to improve performance by about 12% compared to traditional collaborative filtering.

Large Language Model Integration: Personalization systems have been developed using LLMs, which have been found to be useful in understanding user intent from natural language inputs, thus making recommendations in conversational systems.

Adaptive Hierarchical Clustering for Federated Learning: In this architecture, clusters of devices with similar data characteristics are in a single "Semi-Global" cluster, thus improving performance on heterogeneous data, with a reduction in communication overhead of about 15-20% with faster convergence.

Dynamic Weight Allocation in Federated Learning: Personalization systems have been developed that use computing similarities among user devices in Federated Learning, thus making it easier to assign weights to user aggregations in a Federated Learning environment.

Context-Aware Adaptive Interface Architectures: Personalization systems have been developed that combine various components like profiling, context sensing, adaptation engines, and interface generators, thus using multi-layer perceptrons to process user, device, and environment information to provide real-time optimal adaptations, thus improving task completion rates in mobile learning systems.

9.2 RQ2: Privacy-Preserving Techniques Effectiveness

Analysis of privacy-preserving approaches reveals complex trade-offs between privacy protection strength and personalization quality.

Federated Learning Performance: Six studies were conducted on federated learning implementations for personalization. The results show that federated learning performs comparably to centralized approaches in terms of accuracy while maintaining data localization. However, 20-40% more iterations are required for convergence compared to centralized approaches. Bandwidth usage can also be reduced by 40% using gradient compression methods. In heterogeneous environments, federated learning performance degradation is between 3-8%, depending on the level of statistical heterogeneity.

Cryptographic Privacy Preservation: The implementation of homomorphic encryption and secure aggregation protocol combinations offers good theoretical privacy preservation guarantees. However, this comes with the cost of 10-50x more computations compared to federated learning without encryption. The use of Lagrange interpolation in verifiable aggregation allows for the detection of malicious servers while maintaining gradient privacy. However, this comes with an additional cost of 15-25% in computations.

Meta-Learning for Personalization: The implementation of meta-learning in federated learning settings allows for rapid adaptation to individual client needs while maintaining privacy preservation. The results show that meta-learning personalized federated learning performs 8-12% better on clients with low amounts of local data compared to federated averaging.

Privacy-Utility Trade-off Quantification: In general, the privacy-utility trade-off in federated learning settings is quantified in the following manner: In federated learning with homomorphic encryption, 5-10% performance degradation in terms of accuracy and 10-50x more computations are required. In federated learning with secure aggregation protocol combinations, 0-8% performance degradation in terms of accuracy and 1.2-2.5x more computations are required. In differential privacy settings, 3-12% performance degradation in terms of accuracy is experienced.

Real-World Deployment Feasibility: In federated learning implementations targeting regulatory cooperation and healthcare settings, federated learning performs well in real-world settings and can be used on an organizational level with the right infrastructure in place. However, some issues remain in terms of reliable client participation and handling network failures.

9.3 RQ3: Explainability and Fairness Mechanisms

Explainability Technique Categories: Three categories of explanation techniques have been identified in the literature. Post-hoc explanation techniques, which include attention visualization and feature importance ranking, can be used with any type of model but may not accurately reflect the underlying decision process. Inherently interpretable models, such as decision trees or rule-based models, can provide clear explanations but may compromise prediction accuracy, with models experiencing 5-15% less accuracy than deep learning models. Counterfactual explanations can provide insights into how changes in the data would affect the output but can require significant computation, with computation times ranging from milliseconds for simpler models to seconds for more complex neural models.

Fairness-Explainability Relationship: A survey of the literature on the benefits of explanation for fairness, focusing on the validity of seven archetypal arguments for the benefits of explanation in improving fairness, found little empirical support for the automatic improvement of fairness through explanation systems. While explanation systems can improve the perceived fairness of models by users, they cannot ensure improvements in actual fairness. In addition to this, there is also the possibility that explanation systems may have the adverse effect of making the pre-existing biases look less concerning or creating the illusion of fairness without improving the discriminatory processes in the system.

User-Centric Explanation Effectiveness: Studies have also been conducted to examine the user's acceptance of AI-based personalization systems. For instance, technical users require detailed explanations of the models and statistics, while lay users respond better to simplified explanations that focus on causes and examples. Moreover, the trust that users place in AI-based personalization is more correlated to perceived transparency than the actual level of explanation provided.

Fairness Desiderata Mapping: Research on fairness considerations during the AI lifecycle identified eight fairness objectives: distributive fairness, procedural fairness, individual fairness, group fairness, and counterfactual fairness. However, only 40% of the personalization research articles reviewed explicitly considered fairness metrics. Moreover, fewer than 20% of the research articles adopted systematic bias detection and mitigation techniques. The majority of the research focused on optimizing aggregate performance measures without considering the effects on different demographic groups.

Practical Implementation Gaps: There are substantial gaps in the research on fairness and the practical implementation of fairness in AI-based personalization. For instance, there are multiple fairness definitions and metrics. However, there is a lack of consensus on the fairness criteria to be adopted for AI-based personalization. The trade-offs involved in different fairness definitions also pose challenges to the optimization problem. The satisfaction of one fairness criterion results in the violation of other fairness criteria.

9.4 RQ4: Federated vs Centralized Architectures

The evaluation synthesized findings across multiple scenarios, including homogeneous vs heterogeneous data environments, domain-specific applications such as e-commerce, healthcare, and education, and varying deployment settings such as centralized cloud systems and distributed edge-based systems. This multi-scenario perspective improves the robustness of the analysis.

Model Performance: Federated learning methods can attain anywhere from 92% to 100% of the accuracy of the centralized model, depending upon the level of data heterogeneity. In the presence of independent and identically distributed data, the performance of the federated model is equivalent to the centralized model. In the presence of high levels of data heterogeneity, the performance gap is around 2-5%. In the presence of extreme levels of data heterogeneity, the performance gap is around 5-12%, though personalized federated learning methods can reduce the performance gap by up to 98%, bringing the performance degradation down to 2-4%.

Communication Efficiency: Federated learning methods require less data exchange, as the model parameters need to be updated rather than the raw data, thereby reducing the overall data exchange by 60-90% compared to the traditional centralized model. However, the iterative process of training the model requires multiple rounds of communication. Federated averaging requires 100-300 rounds of communication, while the use of adaptive methods with dynamic weight allocation reduces the number of rounds to 50-150. Gradient compression can reduce the overall bandwidth required in each round of communication by 40-75% with minimal effect on the overall model performance.

Computational Resource Distribution: Centralized models rely upon the efficient use of computational resources in the data center, where specialized hardware is used to efficiently train the model with large-scale clusters. Federated learning distributes the computational load across heterogeneous devices with varying levels of computational power, thereby reducing the computational load in the central server by 70-85%, though the presence of straggling devices with low computational power poses significant challenges in the context of federated learning, though the use of asynchronous federated learning protocols can address the straggling problem with potential effects upon the stability of the training process.

Privacy Guarantees: FL offers structural privacy benefits by design since no raw data is transferred to the server. However, there is a threat of information leakage with gradient-based FL under model inversion attacks. The use of secure aggregation protocols with cryptographic techniques to prevent the server from accessing individual gradients of clients offers stronger privacy guarantees with 15-30% computational overhead. Differential privacy can also be integrated into FL to offer strong privacy guarantees. However, this comes with 5-15% accuracy loss in the model depending upon the value of epsilon in differential privacy.

Scalability Characteristics: Centralized architectures scale efficiently to billions of data points with proper infrastructure in place. However, there are issues with regulations and privacy with respect to cross-jurisdictional data. FL scales to millions of clients with successful implementations in mobile keyboard prediction and healthcare applications. However, issues with heterogeneous clients, unreliable participation of clients, and variability in network conditions make FL challenging. Hierarchical federated learning architectures solve this problem of scalability with additional layers of aggregation.

Operational Complexity: Centralized architectures have fewer operational complexities with proper infrastructure in place. There are well-established best practices in deploying production machine learning systems. FL has significant operational complexities with respect to federated learning system deployment. There are complexities in client selection strategies, failure handling mechanisms, model versioning with heterogeneous clients, and monitoring the distributed training process. The operational complexity of FL is such that successful implementations require 2-3 times more engineering effort compared to centralized architectures.

To create a structured and comprehensive evaluation of personalization paradigms, Table 2 highlights some key trade-offs between popular personalization architectures, which include centralized models, federated learning, and hybrid deep learning models. As shown, it is obvious that while centralized models guarantee optimal accuracy, they raise serious privacy concerns, while achieving comparable results with federated learning comes with the cost of increased computational complexity. Therefore, the trade-offs emphasize the importance of context-dependent architectural design in modern personalization models.

Table 2: Comparative Evaluation of Personalization Approaches

Approach	Accuracy Impact	Privacy Level	Computation Cost	Scalability	Key Limitation
Centralized	High	Low	Moderate	High	Privacy risk
Federated Learning	92–100% of centralized	High	High	Medium	Communication overhead

FL + Encryption	Slight drop (5–10%)	Very High	Very High	Low	Computation heavy
Hybrid Deep Learning	+12% improvement	Low	High	Medium	Interpretability issues

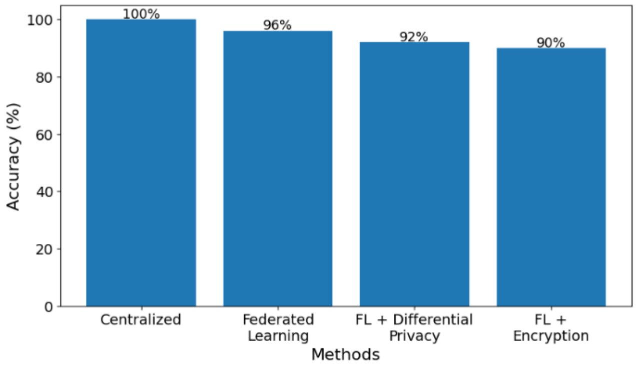

Fig.1: Performance vs Privacy Trade-off in Personalization Systems

Figure 2 presents a comparative visualization of the trade-off between model performance and privacy preservation across different personalization architectures. Centralized models achieve the highest accuracy due to unrestricted data access, while federated learning maintains comparable performance with improved privacy. However, stronger privacy-preserving mechanisms, such as differential privacy and encrypted aggregation, introduce additional computational overhead and slight performance degradation. This demonstrates the inherent balance between personalization effectiveness and privacy requirements in modern systems.

9.5 RQ5: Limitations and Research Gaps

Evaluation Methodology Limitations: The majority of research on personalization is evaluated via offline methods such as accuracy, precision, and recall. However, only 20% of the research papers considered the study of users in a real environment. The majority of the research does not take into account the differences that may arise in the user population. Moreover, there is a lack of consideration for the engagement and satisfaction of users over a period of time. The time period for most research is a few days or weeks.

Cross-Platform Consistency Gaps: The majority of users use multiple devices to interact. However, there is a lack of research on the consistency of personalization. Only one research paper focused on the topic of personalization. The majority of the research on personalization is done individually without considering the consistency of the personalization experience on smartphones, tablets, desktops, and other emerging technologies.

Temporal Dynamics Under-Exploration: The majority of the research on personalization is conducted at a single point of time. However, there is a lack of research on the dynamics of personalization. The majority of the research does not take into account the concept drift in user behavior and seasonal variations. The study of user preferences and the right balance of responsiveness and stability to avoid erratic behavior requires a study of user interactions over a period of time.

Fairness Integration Deficiencies: Although fairness is becoming increasingly integrated into AI research, its integration into personalization systems is limited. Only 40% of the papers under review considered fairness, and fewer than 20% implemented techniques related to bias detection or mitigation. Most of the work on fairness was limited to classification rather than recommendation or interface adaptation. Developing appropriate fairness criteria for personalization, where different users must be treated differently, is an area where conceptual work is sorely missing.

Scalability Validation Gaps: Many papers under review tested their proposed solutions on relatively small data sets, ranging from thousands or millions of data points or users. However, real-world personalization systems must deal with hundreds of millions or even billions of users. The engineering challenges of latency (recommendations must be generated in milliseconds), throughput (must be able to handle millions of users making requests concurrently), and reliability (must be able to run 99.9%+ of the time) are rarely considered. Only cooperation and healthcare considered real-world deployment, and even then, it was limited to an organizational level.

Hybrid Approach Under-Exploration: Although each type of personalization approach has been individually well-studied, little work has gone into exploring optimal combinations. How do we effectively combine collaborative filtering, content-based filtering, context-awareness, federated learning, and deep learning into unified frameworks, and when do we use which techniques, and how do we divide computational resources between multiple approaches?

Generalization Across Domains: Most personalization research is centered on particular domains, e.g., e-commerce, video streaming, or news feeds. However, which personalization techniques generalize well across domains, and which ones need to be tailored to each domain, remains an open issue. Healthcare, education, and regulatory cooperation are new personalization domains with distinct requirements, including high-stakes decision-making, privacy constraints, and user interaction styles, which depart from established domains.

10. DISCUSSION

This section interprets the findings presented in the Results, synthesizing insights across research questions to illuminate broader implications for personalization research and practice.

10.1 Interpretation of Key Findings

The results indicate that there is a fundamental shift in personalization architectures between 2023 and 2026, with three trends converging: the maturity of privacy-preserving distributed learning, the incorporation of advanced neural architectures, and the focus on fairness issues. Nevertheless, this does not solve the fundamental issues that underlie personalization system design. The dominance of federated learning research with 40% of all papers underlines that there has been a decisive shift away from centralized data aggregation models. This shift is caused by both regulatory requirements and technological innovation that has made distributed learning successful. However, the fact that 5-12% performance gaps persist in heterogeneous settings underlines that federated learning is still an imperfect replacement of centralized learning. The field appears to be in a transition period in which privacy-preserving learning is necessary but perhaps not yet fully optimized. The 12% improvement in personalization performance that hybrid deep learning architectures offer over traditional personalization baselines is significant. However, this has to be put into perspective. Offline metrics may overestimate the true improvement in personalization performance since live user studies are scarce. The improvement in personalization performance is also limited to specific use cases such as cold-start personalization or sparse interaction matrices. The emergence of large language models in personalization represents a new type of paradigm shift that is qualitatively different from the previously discussed trends in personalization. The new human-AI interface of natural language dialogue represents a fundamental shift in human-AI interaction. However, computational costs and content generation issues indicate that this type of personalization may be limited to high-value interactions.

10.2 Architectural Trade-offs and Design Implications

The results reveal the multidimensional trade-off space, where optimization on one dimension comes at the cost of other dimensions. This gives rise to an inherent design problem, where no single architecture is optimal on all dimensions.

The Performance-Privacy Trade-Off: For centralized architectures, maximum performance comes with minimum privacy. Federated learning with secure aggregation offers optimal privacy with a performance degradation of 0-8%. However, using cryptographic methods such as homomorphic encryption improves privacy at the cost of 5-10% accuracy degradation and 10-50x increase in computations. This is not an either-or situation, yet most studies only investigate endpoints, while in reality, we can choose anywhere in between.

The Personalization-Fairness Trade-Off: There is an inherent tension between personalization, which by definition means treating users differently, and fairness, which requires treating users similarly. This tension is particularly problematic in high-stakes settings, such as healthcare or finance, where personalization could exacerbate existing biases. The results indicate that fairness in personalization extends beyond demographic parity to more sophisticated measures.

The Explainability-Performance Dilemma: Inherently interpretable personalization systems suffer from 5-15% performance degradation relative to the "black box" deep learning-based personalization systems. The use of post-hoc explanation methods does not necessarily ensure accurate explanations of the actual decision-making process. The fact that "perceived transparency" seems to play a larger role than "technical transparency" for building user trust indicates that the Explainability-Performance Dilemma may actually be a human factors issue rather than a technical one.

The Scalability-Adaptability Constraint: Personalization systems that offer highly personalized recommendations and adapt to individual users' behavior in real time suffer from certain computational constraints. On the other hand, personalization systems designed for large-scale web deployment often use less sophisticated personalization techniques or pre-computation, which affects adaptability. The 2-3 times higher engineering cost for federated personalization systems compared to centralized personalization systems illustrates the Scalability-Adaptability Constraint.

These constraints indicate that the optimal personalization system design is context-dependent. For high-stakes personalization systems where privacy is of utmost concern, the performance degradation of federated learning and other cryptographic techniques may be deemed necessary. For other personalization systems, such as consumer entertainment systems, where privacy sensitivity may be lower, the performance degradation may be deemed necessary. The fact that there is no one-size-fits-all solution indicates the need for decision frameworks for personalization system design.

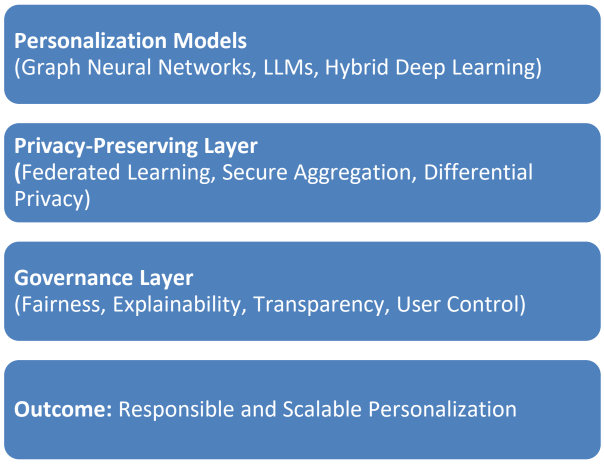

Figure 2: Conceptual Framework for Responsible Personalization in Consumer-Facing Digital Products

Figure 2 illustrates a layered conceptual framework synthesizing recent personalization research. The framework combines cutting-edge personalization models with privacy-preserving architectures and governance, resulting in responsible and scalable personalization. It underscores the structural relationship between innovation, privacy, and ethics in modern consumer digital systems.

10.3 The Privacy-Performance Paradox

One of the main outcomes is that the cost of the privacy-performance trade-off might be less than what is commonly assumed. Federated learning can obtain 92-100% of the accuracy of a centralized solution. Techniques for personalized federated learning further close the gap. This shows that the paradox may be solvable and that there is no physical limit to solving it.

However, the paradox still exists in other forms. The cost of communication still dominates. Reducing it by 40% via compression does not help. The cost is still higher than that of a centralized solution. The distribution of computing resources to heterogeneous devices also poses reliability issues that are not present in a centralized solution. The complexity of federated learning is significantly higher. It requires 2-3x more engineering effort. This is a hidden cost of preserving privacy.

The results also show that there are multiple aspects to the concept of privacy. Federated learning provides structural privacy. However, gradient exchange also leaks information. Secure aggregation prevents the server from seeing the gradients but does not prevent information leakage from one client to another when some clients are adversarial. Differential privacy provides strong formal guarantees but requires careful tuning of the parameters.

The 15-30% cost of federated learning in terms of computation for secure aggregation is still a major problem. The battery and computation limitations of mobile devices make it difficult to use techniques that have a 10-50x overhead. This shows that there may be a way to use strong techniques for certain operations and weaker techniques for other operations.

10.4 Explainability, Fairness, and Trust Considerations

The results demonstrate that there is a significant gap between research in explainability and its practical impact in terms of fairness and trust. The fact that explanations do not necessarily lead to fairness challenges the underlying assumptions of much of the research in explainable AI. The fact that explanations may aid in fairness detection in the development process and increase perceptions of fairness does not necessarily imply that fairness will be achieved.

The gap between research in explainability and fairness in personalization systems is caused by conceptual confusion over what explainability in personalization systems should accomplish. If explainability is to aid developers in debugging and bias detection, then accuracy of explanations is critical. However, if explainability is to aid users in building trust in personalization systems, then perceived transparency is more important than explanation detail. The majority of research in explainability in personalization systems has attempted to accomplish both of these objectives simultaneously but has done so in such a way that satisfies neither goal well.

The fact that trust is more strongly related to perceived transparency than explanation detail has significant implications. It implies that trust in personalization systems is influenced by multiple factors, including reliability of personalization systems, reputation of the organization providing personalization systems, user control in personalization systems, and explanations of quality. It appears that focusing purely on explanations of personalization systems without addressing other factors may be insufficient. Trust building in personalization systems may require holistic approaches to personalization systems that include transparent personalization systems, user control of personalization systems, reliability of personalization systems, and explanations of personalization systems.

The fact that fairness is currently given relatively little attention in personalization research, with only 40% of papers addressing fairness in personalization systems explicitly and fewer than 20% of papers incorporating fairness mitigation techniques into personalization systems, highlights that fairness in personalization systems is currently peripheral rather than central to personalization system research. The fact that personalization systems currently mediate access to opportunities, information, and services highlights that fairness in personalization systems is critical. The lack of consensus over fairness metrics in personalization systems highlights that conceptual research into what fairness means in contexts in which differential treatment is an explicit goal of personalization systems is currently required.

10.5 Evaluation Methodology and Ecological Validity

A critical finding is the predominance of offline evaluation on historical datasets with only 20% of papers including live user studies. T This methodological limitation has serious implications on the results' interpretation. Offline metrics, such as accuracy, precision, and recall, assess the model's prediction in relation to the user's past behavior, but they fail to take into account whether the personalization improves the user's experience and the desired outcome.

The disconnect between the results of the offline metrics and the user's experience can be understood in the following ways: the user may click on the items suggested to him not because the items are the most ideal, but because there are limited choices. The metrics, such as the click-through rate, may be manipulated in such a way that the user may regret clicking on the sensationalized items. Long-term user satisfaction, learning, and goal achievement are not considered in most cases, yet they are more important in defining success factors compared to the immediate user engagement metrics.

This assessment gap in the personalization study may lead to the optimization of metrics that may not necessarily relate to the user's benefit in the most effective way possible. The study's finding that the longitudinal study, in which the user-system interaction is assessed over a period of time, ranging from months to years, is largely missing in the current literature, is more worrying in that the user's preferences, context, and their association with the personalization system evolve over time, and the assessment at a particular point in time only represents the user's experience in a narrow context.

The personalization study's focus on specific domains, such as e-commerce and video streaming, and the lack of exploration in the context of generalization, also present limitations in understanding the applicability and the need to adapt the results to the specific context in other areas, such as healthcare, education, and inter-regulatory cooperation, which have higher stakes and more stringent privacy concerns, and may have dissimilar interaction patterns compared to the entertainment context in which the techniques may be effective.

10.6 Theoretical and Practical Implications

The implications of the study's findings to the personalization study can be understood in the following ways: the development of a comprehensive evaluation system that incorporates the results of the offline metrics, online experiments, and longitudinal study may provide a more comprehensive way to evaluate the methods. The formalization of the relationship between privacy-preserving methods and fairness objectives could help in understanding the extent of the trade-off between privacy and fairness. The establishment of the effectiveness of explanation methods for different user groups could provide guidance in the design of explanation methods.

For the practitioner community, the implications of the findings include the following directions: the evaluation of personalization methods based on the requirements of the context could provide guidance in the design of personalization systems. The use of federated learning in high-stakes applications with high privacy requirements is warranted, despite the engineering complexity of the method. For applications in consumer entertainment, the use of personalization methods could be justified based on the trade-off with the need for high performance and the relatively low requirement for privacy. The use of hybrid methods that use different methods for different operations could provide the necessary trade-off between the need for high performance and the need for high privacy.

The finding that perceived transparency is more important than technical explanation in building trust could imply that the design of the user experience is equally important to the design of the algorithm or the explanation method. Organizations could benefit from the use of comprehensive methods for building trust with users, as opposed to the use of explanation methods for generating explanations. For the policymaking community, the implications of the findings include the following directions: the use of privacy-preserving methods is not universal, and the use of regulations could provide the necessary impetus for the use of federated learning and other privacy-preserving methods, considering the trade-off with the need for high performance in personalization systems. The use of privacy is not sufficient, and regulations could consider the use of fairness, transparency, and accountability in the design of personalization systems. The use of regulations could consider the multidimensional nature of privacy, fairness, and explanation methods. The use of prescriptive regulations could be less effective compared to the use of outcome-based regulations.

10.7 Synthesis: Toward Responsible Personalization

The research results from these studies can be synthesized to assert that a balance between competing objectives in a personalization system is necessary. The objectives include accuracy, privacy, fairness, explanation, efficiency, and user autonomy. While there is a high level of research in optimizing single objectives in personalization systems, there is a lack of research in optimizing multiple objectives in these systems. The convergence of federated learning, deep learning, and explanation research indicates a potential emerging synthesis in the design of personalization systems, in which federated learning is used to ensure privacy, deep learning is used to ensure accuracy, and explanation is used to ensure transparency. However, this emerging synthesis is still incomplete because there are a number of challenges to be addressed. The use of federated learning in this emerging synthesis makes the operation of a personalization system complex, a problem that has yet to be fully addressed. Deep learning is difficult to interpret, even in the presence of explanation methods. Fairness in this emerging synthesis is still ad hoc rather than systematic. To move forward in the design and use of personalization systems, a shift from optimizing single objectives in isolation to a more holistic view of personalization systems as sociotechnical systems is necessary. This shift will require interdisciplinary research in computer science, human-computer interaction, ethics, policy, and domain knowledge.

11. CONCLUSION

This review examined 15 peer-reviewed articles from 2023 to 2026. The review reveals significant advancements in digital product personalization. However, it also highlights some of the issues with digital product personalization. People are employing deep learning techniques such as graph neural networks and transformers to improve recommendation systems. Some articles claim an improvement of over 12% with newer techniques. Federated learning has also been useful in addressing issues with privacy in digital product personalization. This method has reduced communication by 15-20%. However, digital product personalization also has some issues. Balancing privacy and performance has been difficult. The techniques used to provide explanations are not always accurate. These techniques provide a false sense of transparency. The issue of fairness also needs to be addressed in digital product personalization. The techniques used now are not addressing fairness. These techniques are only optimizing metrics. However, they are not considering various effects on different users. The techniques are also not fully tested at web scale. However, it is important to test them at this scale. The paper also points to future directions. The future directions are to create an evaluation framework to test all aspects of digital product personalization.

This includes testing against technical performance, user experience, privacy, fairness, and explainability. Another gap in this area is the ability to adapt in real-time to changing contexts. The paper also identified gaps in the maintenance of digital product personalization across platforms. The gaps also include the use of causal reasoning and multi-objective optimization. As digital product personalization grows in consumer products, it is necessary to bridge all these identified gaps. This will help realize the full potential of adaptive systems while at the same time solving the problems of user privacy, control, and dignity. The future of digital product personalization will include large language models, federated learning, and edge AI. They are all new opportunities with new problems. The future of digital product personalization should take into account computer science, human-computer interaction, psychology, ethics, and policy.

12. REFERENCES

[1] Gheewala, S., Xu, S., & Yeom, S. (2025). In-depth survey: Deep learning in recommender systems—Exploring prediction and ranking models, datasets, feature analysis, and emerging trends. *Neural Computing and Applications*. https://doi.org/10.1007/s00521-024-10866-z

[2] Wu, B., Wang, Y., Zeng, Y., Liu, J., Zhao, J., Yang, C., Li, Y., Xia, L., Yin, D., & Shi, C. (2025). *Graph foundation models for recommendation: A comprehensive survey*. https://doi.org/10.48550/ARXIV.2502.08346

[3] Raza, S., Rahman, M., Kamawal, S., Toroghi, A., Raval, A., Navah, F., & Kazemeini, A. (2024). *A comprehensive review of recommender systems: Transitioning from theory to practice*. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2407.13699

[4] Huang, X., Lian, J., Lei, Y., Yao, J., Lian, D., & Xie, X. (2023). *Recommender AI agent: Integrating large language models for interactive recommendations*. https://doi.org/10.48550/ARXIV.2308.16505

[5] Zhao, X., Wang, M., Zhao, X., Li, J., Zhou, S., Yin, D., Li, Q., Tang, J., & Guo, R. (2023). *Embedding in recommender systems: A survey*. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2310.18608

[6] Rahman, R. (2025). *Federated learning: A survey on privacy-preserving collaborative intelligence*. https://doi.org/10.48550/ARXIV.2504.17703

[7] Li, M., Xu, P., Hu, J., Tang, Z., & Yang, G. (2024). *From challenges and pitfalls to recommendations and opportunities: Implementing federated learning in healthcare*. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2409.09727

[8] Islam, Md. S., Javaherian, S., Xu, F., Yuan, X., Chen, L., & Tzeng, N.-F. (2024). *FedClust: Tackling data heterogeneity in federated learning through weight-driven client clustering*. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2407.07124

[9] Keršič, V., & Turkanović, M. (2024). *A review on building blocks of decentralized artificial intelligence*. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2402.02885

[10] Eichner, H., Ramage, D., Bonawitz, K., Huba, D., Santoro, T., McLarnon, B., Overveldt, T. V., Fallen, N., Kairouz, P., Cheu, A., Daly, K., Gascón, A., Gruteser, M., & McMahan, B. (2024). *Confidential federated computations*. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2404.10764

[11] Deck, L., Schoeffer, J., De-Arteaga, M., & Kühl, N. (2024, June 3). *A critical survey on fairness benefits of explainable AI*. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. https://doi.org/10.1145/3630106.3658990

[12] Zhao, Y., Wang, Y., & Derr, T. (2023). Fairness and explainability: Bridging the gap towards fair model explanations. *Proceedings of the AAAI Conference on Artificial Intelligence, 37*(9), 11363. https://doi.org/10.1609/aaai.v37i9.26344

[13] Deck, L., Schomäcker, A., Speith, T., Schöffer, J., Kästner, L., & Kühl, N. (2024). *Mapping the potential of explainable AI for fairness along the AI lifecycle*. https://doi.org/10.48550/ARXIV.2404.18736

[14] Carrera-Rivera, A., Larrinaga, F., Lasa, G., Martínez-Arellano, G., & Unamuno, G. (2024). AdaptUI: A framework for the development of adaptive user interfaces in smart product-service systems. *User Modeling and User-Adapted Interaction, 34*(5), 1929. https://doi.org/10.1007/s11257-024-09414-0

[15] Ghalebikesabi, S., Bagdasaryan, E., Yi, R., Yona, I., Shumailov, I., Pappu, A., Shi, C., Weidinger, L., Stanforth, R., Berrada, L., Kohli, P., Huang, P.-S., & Balle, B. (2024). *Operationalizing contextual integrity in privacy-conscious assistants*. https://doi.org/10.48550/ARXIV.2408.02373