Journal of Science
Basic and Applied Number. 21 June 2026

DOI: https://doi.org/10.36602/jsba.2026.21.49

AI-Augmented Human–Computer Interaction: Advances, Challenges, and Future Directions

Anwar Alhenshiri[1] and Wesam Elturjman[1]
[1]Department of Computer Science, Faculty of Information Technology, Misurata University
*Corresponding author: Anwar Alhenshiri, alhenshiri@it.misuratau.edu.ly

Submission date: 30. 10.2025 Acceptance date: 26 .3 .2026 Electronic publishing date: 2.4.2026

Abstract: This paper surveys recent advances in applying Artificial Intelligence (AI) to Human–Computer Interaction (HCI), with an emphasis on adaptive systems, personalization, and intelligent interfaces. AI techniques, particularly natural language processing and reinforcement learning, enable systems to understand user intent, adapt to context, and improve usability across domains such as healthcare, education, and recommender systems. The survey highlights core approaches, their strengths and limitations, and outlines open challenges, offering a roadmap for integrating AI into HCI more effectively.

Keywords: HCI, AI, Reinforcement Learning, Adaptability, Context, Integration.

Introduction:

The rapid growth of Artificial Intelligence (AI) has fundamentally reshaped Human–Computer Interaction (HCI), moving beyond static interfaces toward systems that are adaptive, personalized, and capable of understanding complex user behaviors. Modern users expect digital systems not only to respond efficiently but also to anticipate their needs, interpret intent, and deliver tailored experiences. These expectations have been fueled by advances in machine learning and natural language processing (NLP), which enable intelligent interfaces across domains such as healthcare, education, and personalized recommendation (Norman, 2013; Shneiderman et al., 2016; Devlin et al., 2019). Figure 1 shows the evolution stages of HCI.

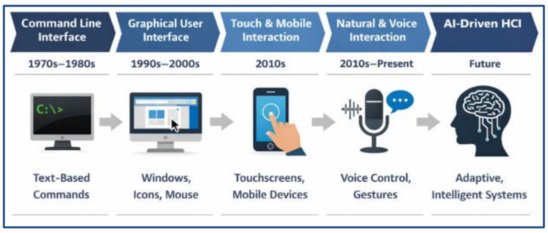

Figure 1. *Evolution of Human–Computer Interaction*

Recent breakthroughs have highlighted the potential of AI to enhance personalization. Transformer-based models such as BERT and GPT provide rich contextual embeddings that allow for accurate query disambiguation and semantic understanding (Lee and Chen, 2024). Reinforcement learning (RL) techniques extend this capability by enabling continuous adaptation, using real-time signals such as clicks, dwell time, and query reformulations to optimize ranking strategies (Miller and White, 2023; Yin et al., 2021). Complementary approaches such as graph-based user modeling capture long-term preferences and contextual dependencies, ensuring that systems remain responsive not only within a session but also across extended periods (Taylor and King, 2022; Pu et al., 2025).

Together, these advances suggest the possibility of a new generation of AI-augmented HCI systems that can dynamically balance personalization, adaptability, and usability.

Despite this progress, important limitations persist. Current systems often struggle with ambiguous or multimodal queries, limiting their ability to capture user intent accurately (Wang and Liu, 2023). Overreliance on implicit feedback such as clicks may introduce bias and fail to reflect genuine satisfaction (Williams and Davis, 2022). Moreover, the integration of adaptive learning raises questions of fairness, transparency, and trust, particularly in high-stakes settings such as clinical decision support or digital learning platforms (Hall and Young, 2023). These challenges underscore a critical research problem: How can AI techniques be systematically integrated into HCI to achieve robust personalization while ensuring transparency, fairness, and long-term user satisfaction?

This paper addresses this gap by proposing a survey of methods that can be integrated into a unified framework for intelligent HCI.

Specifically, we examine three pillars: (i) NLP-

Journal of Science Basic and Applied **Number. 21 June 2026**

driven semantic understanding for intent detection, (ii) reinforcement learning-based adaptive ranking for real-time personalization, and (iii) user modeling through graph-based representations for long-term preference tracking. By synthesizing progress across these domains, this survey highlights how their integration can overcome existing limitations and establish more effective, user-centered systems (Green and Black, 2024; Agarwal et al., 2022).

The remainder of the paper is structured as follows. Section 2 explains the literature review methodology. Section 3 shows the study selection process. Section 4 provides theoretical and technological foundations of AI in HCI. Section 5 reviews methods for intent detection and semantic modeling. Section 6 explores reinforcement learning and counterfactual approaches for adaptive personalization. Section 7 discusses user modeling techniques for long-term personalization. Section 8 surveys application domains where these methods are applied. Section 9 outlines evaluation strategies and benchmarks. Section 10 presents open research challenges, and Section 11 concludes with directions for future work.

Literature Review Methodology

This survey follows a structured literature review approach to ensure comprehensive and unbiased coverage of research on AI-enhanced Human–Computer Interaction (HCI). Relevant studies were collected from major academic databases, including IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, and Google Scholar, to ensure high-quality peer-reviewed sources.

The search process employed combinations of keywords such as “AI in HCI,” “intelligent user interfaces,” “personalized systems,” “intent detection,” “reinforcement learning,” “user modeling,” and “multimodal interaction”. Boolean operators (AND, OR) were used to refine and expand the search.

To maintain relevance and quality, studies were included if they were published between 2015 and 2025, focused on AI techniques in HCI, and provided clear methodological or experimental contributions. Studies were excluded if they were not directly related to HCI or personalization, lacked sufficient technical detail, were duplicates, or were not peer-reviewed.

Study Selection Process

The study selection process followed a systematic review protocol inspired by PRISMA guidelines to ensure transparency and reproducibility. Initially, relevant studies were identified through comprehensive searches across multiple academic databases using predefined keywords. Duplicate records were then removed, and titles and abstracts were screened to exclude studies not directly related to AI in HCI or personalization. The remaining articles were assessed in full text against the inclusion and exclusion criteria to ensure their relevance, methodological rigor, and quality. Finally, the selected studies were included for detailed analysis and comparative discussion in this survey. The overall selection process is summarized in Figure 2.

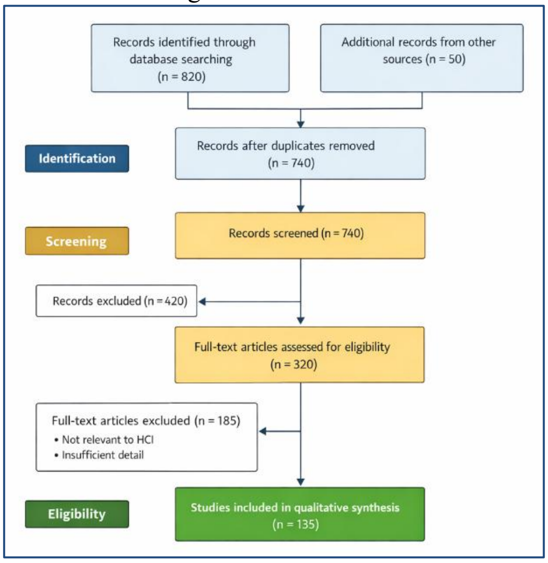

Figure 2. *Summary of Key Studies Reviewed*

Background and Foundations

The integration of Artificial Intelligence (AI) into Human–Computer Interaction (HCI) is built upon well-established theoretical and technological foundations. Classical HCI emphasized usability and cognitive principles, while modern approaches increasingly leverage machine learning to deliver adaptive and personalized experiences. This section reviews the key foundations that underpin AI-enhanced HCI, including advances in deep learning and natural language processing (NLP) for semantic understanding, reinforcement learning (RL) for adaptive personalization, and user modeling techniques for capturing long-term preferences. Together, these approaches form the technical basis for intelligent, context-aware systems. The evolution of AI-HCI integration is depicted in Figure 3.

Journal of Science Basic and Applied **Number. 21 June 2026**

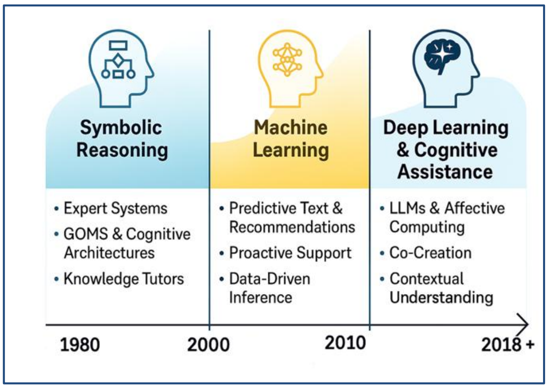

Figure 3. *Evolution of AI-HCI Integration*

Human–Computer Interaction and the Role of AI

Human–Computer Interaction (HCI) traditionally emphasized usability, cognitive ergonomics, and task efficiency (Norman, 2013; Shneiderman et al., 2016). While these foundations remain relevant, the field has shifted toward intelligent interfaces capable of adapting to user behavior and context. This transformation has been accelerated by advances in Artificial Intelligence (AI), particularly in machine learning and data-driven personalization. AI now enables systems not only to interpret user actions but also to predict preferences, anticipate intent, and respond adaptively across multiple application domains (Smith and Brown, 2023).

Deep Learning and Natural Language Processing

Deep learning has provided powerful tools for semantic modeling. Neural networks, and more recently transformer-based architectures such as BERT (Devlin et al., 2019) and GPT variants, allow systems to capture rich contextual representations of language, making it possible to disambiguate queries and understand intent with higher accuracy (Lee and Chen, 2024). These models significantly outperform traditional methods such as TF–IDF and BM25, enabling more natural and human-like interaction. However, they also pose challenges in terms of computational cost, interpretability, and generalization to unseen tasks (Thakur et al., 2021).

Reinforcement Learning for Adaptive Personalization

Reinforcement Learning (RL) has emerged as a complementary technique for real-time adaptation. Unlike supervised models trained offline, RL agents learn dynamically from feedback, using implicit signals such as clicks, dwell time, or query reformulations to refine ranking strategies (Miller and White, 2023; Yin et al., 2021). Counterfactual learning-to-rank methods further address the issue of biased interaction logs, ensuring more robust and fair personalization (Agarwal et al., 2018; Agarwal et al., 2022). While promising, RL introduces new challenges including reward design, data sparsity, and balancing short-term versus long-term user satisfaction.

User Modeling and Personalization

Personalization in HCI depends on effective user modeling. Early approaches relied on explicit profiles or demographic data, but these were limited in capturing evolving preferences. Recent advances in Graph Neural Networks (GNNs) and embedding-based models have enabled richer representations of user interests, linking behaviors, content, and context into unified profiles (Taylor and King, 2022; Pu et al., 2025). Such models allow for personalization that extends beyond immediate sessions, though they raise concerns regarding privacy, fairness, and scalability.

Beyond structural modeling, researchers are increasingly focusing on dynamic and context-aware personalization. Hybrid systems now combine short-term session signals—such as clicks, dwell time, or recent queries—with long-term historical data to balance recency and stability in recommendations (Wu et al., 2021). This dual-layered modeling approach is particularly effective in domains like news and music recommendation, where user interests change rapidly yet remain anchored in broader preferences. At the same time, new methods in privacy-preserving personalization, such as federated learning and differential privacy, are emerging to mitigate risks of sensitive data exposure. These developments signal a move toward personalization strategies that are not only more accurate and adaptive but also ethically responsible and socially aware.

To summarize, deep learning, NLP, reinforcement learning, and advanced user modeling provide the technical foundation for intelligent HCI. Yet, despite progress, gaps remain in seamlessly integrating these components into unified systems that are accurate, adaptive, and trustworthy. This survey builds on these foundations to review recent work and propose directions for bridging these gaps.

Intent Detection and Semantic Modeling

Understanding user intent lies at the heart of intelligent HCI. Because queries and interactions are often short, ambiguous, or context-dependent, systems must go beyond surface-level keyword matching to capture deeper semantic meaning. Advances in natural language processing (NLP)

Journal of Science Basic and Applied **Number. 21 June 2026**

and deep learning have made it possible to model intent more accurately, enabling systems to disambiguate queries, infer goals, and deliver contextually relevant results. This section reviews the evolution of intent detection methods, from traditional keyword-based models to modern transformer-based and multimodal approaches, highlighting their impact on adaptive interaction.

Importance of Intent Understanding

A central challenge in Human–Computer Interaction is the ability to accurately interpret user intent. Queries and commands are often short, ambiguous, or context-dependent, making intent detection essential for delivering relevant and personalized results. For example, the query “jaguar” may refer to the animal, the car brand, or a sports team, depending on the user’s context. Accurate intent modeling is therefore critical for minimizing ambiguity and improving user satisfaction (Williams and Davis, 2022; Wang and Liu, 2023).

Early Approaches

Traditional approaches relied on rule-based systems and keyword-matching models such as TF–IDF and BM25. While computationally efficient, these methods ignored semantic relationships between terms, often failing in cases of synonymy, polysemy, or nuanced user goals (Smith and Brown, 2023). Early statistical learning models attempted to capture patterns in query logs, but their effectiveness was limited by feature engineering constraints.

Deep Learning for Intent Detection

The advent of deep learning introduced more powerful tools for capturing semantic meaning. Recurrent neural networks (RNNs) and convolutional neural networks (CNNs) were applied to query classification, intent recognition, and dialogue systems (Clark and Gardner, 2017). These models improved performance over shallow methods but struggled with long-term dependencies and scalability.

Transformer-based Models

Transformer architectures have revolutionized semantic modeling by introducing attention mechanisms that capture contextual relationships across entire sequences. BERT (Devlin et al., 2019) and subsequent transformer-based models such as GPT variants provide rich embeddings that significantly enhance intent detection and disambiguation (Lee and Chen, 2024). These models have been successfully applied in search engines, chatbots, and recommender systems. However, they present challenges in terms of computational cost, interpretability, and domain generalization (Thakur et al., 2021).

Multimodal Intent Understanding

Recent research has expanded intent detection beyond text to include voice, image, and behavioral signals. For example, integrating audio and visual cues improves disambiguation in multimodal search scenarios, where user intent is expressed through a combination of modalities (Wang and Liu, 2023). Such approaches align closely with modern HCI trends that emphasize natural, immersive, and multimodal interactions.

Open Challenges

Despite these advances, several gaps remain. Current models still face difficulties in handling context shifts across sessions, explaining predictions in a user-friendly way, and reducing reliance on large-scale labeled datasets. Moreover, the integration of semantic models into real-time adaptive systems requires balancing accuracy with computational efficiency. These challenges highlight the need for hybrid approaches that combine deep semantic modeling with user-specific contextualization, an idea further developed in later sections of this survey.

Reinforcement Learning and Adaptive Ranking

Personalized interaction in HCI requires systems that can adapt dynamically to user feedback rather than relying solely on static, pre-trained models. Reinforcement Learning (RL) provides a natural framework for this adaptation, treating ranking and personalization as sequential decision-making problems where user interactions serve as feedback signals. By continuously updating policies based on clicks, dwell time, or query reformulations, RL enables systems to refine recommendations in real time. This section reviews RL-based approaches to adaptive ranking, including recent advances in counterfactual learning-to-rank, and examines their potential and limitations for intelligent HCI.

The Need for Adaptive Ranking

In information-rich environments, static ranking strategies often fail to reflect the dynamic nature of user preferences. While traditional supervised learning methods can optimize relevance offline, they are limited in adapting to real-time feedback. Users’ interactions—clicks, dwell time, scrolling, or query reformulations—represent valuable implicit signals of satisfaction. Leveraging these signals requires adaptive ranking mechanisms that can learn continuously during the interaction process (Radlinski and Joachims, 2008; Turner and Scott, 2022).

Journal of Science Basic and Applied Number. 21 June 2026

Reinforcement Learning in Personalization

Reinforcement learning (RL) provides a natural framework for adaptive personalization by modeling ranking as a sequential decision-making problem. The system observes a state (e.g., user query and context), selects an action (e.g., ranking order), and receives rewards based on user feedback (e.g., clicks or dwell time). Over time, the agent learns policies that maximize cumulative reward, aligning system outputs with user satisfaction (Miller and White, 2023).
Several RL algorithms have been applied in this context. Deep Q-Networks (DQN) and Policy Gradient methods enable optimization of ranking functions in high-dimensional spaces, while contextual bandit models provide lightweight solutions for balancing exploration (trying new results) and exploitation (reinforcing successful results). Research demonstrates that RL-based ranking can outperform static methods, particularly in session-based recommendation and personalized search (Yin et al., 2021; Hou et al., 2023).

Counterfactual Learning-to-Rank

A key limitation of RL approaches is their dependence on biased feedback logs. For example, position bias (users clicking higher-ranked items regardless of relevance) can skew the learning process. To address this, counterfactual learning-to-rank (LTR) techniques introduce methods to estimate and correct for bias. By applying inverse propensity weighting and related techniques, counterfactual LTR enables more reliable learning from historical data (Agarwal et al., 2018; Oosterhuis and de Rijke, 2021; Agarwal et al., 2022). This makes it possible to combine the benefits of offline training with unbiased evaluation of ranking functions.

Challenges and Research Opportunities

While RL and counterfactual methods hold promise, several challenges remain. Reward design is non-trivial, as implicit signals like clicks may not always reflect genuine relevance. Data sparsity limits the effectiveness of RL in domains with fewer interactions. Additionally, RL agents may optimize for short-term engagement at the expense of long-term satisfaction, raising ethical and practical concerns. Addressing these challenges requires hybrid frameworks that combine RL with robust user modeling, fairness-aware objectives, and interpretable decision processes.

User Modeling and Long-Term Personalization

Personalization in HCI depends on accurately modeling users over time. Beyond short-term intent, effective systems must capture evolving interests and contextual patterns across sessions. Recent advances in embeddings and graph-based methods provide richer user representations, enabling adaptive and sustainable personalization. This section reviews the progression from early profile-based models to modern deep and graph-based approaches for long-term personalization.

Importance of User Modeling

Personalization in Human–Computer Interaction (HCI) depends fundamentally on the system’s ability to model users accurately. Effective user modeling allows systems to go beyond one-time interactions and deliver consistent, adaptive experiences across sessions. This is particularly critical in domains such as healthcare, e-learning, and recommender systems, where preferences and needs evolve over time. By maintaining and updating models of user interests, systems can provide continuity, relevance, and engagement that static approaches fail to achieve (Ricci et al., 2015; Aggarwal, 2016).

Early Approaches

Traditional approaches to user modeling relied on explicit profiles—such as user-provided ratings, demographic attributes, or preference lists—or on collaborative filtering methods that inferred preferences from user–item interactions. While effective in domains with dense interaction data, these methods were limited by their reliance on either user willingness to provide input or on large-scale historical logs. They often failed to capture dynamic changes in preference and struggled with the cold-start problem, where little prior information was available (Williams and Davis, 2022).

Embedding-Based User Representations

The rise of deep learning enabled user modeling through latent embeddings, where users and items are projected into a shared vector space. Matrix factorization techniques, followed by neural embeddings, made it possible to capture subtle patterns of similarity between users and content (Koren et al., 2009; Wu et al., 2020). Such models form the basis of many modern recommendation systems, providing compact yet powerful representations of user behavior. However, embedding-based approaches typically emphasize short-term interactions, often neglecting long-term preference evolution.

Graph-Based and Hybrid Approaches

Journal of Science Basic and Applied Number. 21 June 2026

Recent work has explored the use of graph neural networks (GNNs) to represent users, items, and contextual relationships within a unified graph structure. By leveraging heterogeneous interactions (e.g., user–user, user–item, item–attribute), GNN-based methods can capture richer contextual dependencies and improvepersonalization in sparse data environments (Taylor and King, 2022; Pu et al., 2025). Hybrid systems further combine short-term session features with long-term graph-based embeddings, allowing for personalization that balances recency and stability. These methods have shown strong results in domains such as news recommendation, music streaming, and personalized search (Wu et al., 2021).

Challenges in User Modeling

Despite their promise, modern user modeling techniques face persistent challenges. Privacy concerns arise from storing and processing detailed user histories, necessitating approaches such as differential privacy or federated learning. Fairness and bias must be addressed to avoid reinforcing harmful stereotypes or unequal treatment of users. Finally, scalability remains a pressing issue, as models must handle millions of users and items in real-time environments. Addressing these challenges requires developing user models that are not only accurate but also transparent, ethical, and efficient.

Table 1. summarizes the main AI techniques used in HCI, highlighting their advantages, limitations, and suitable application domains.

Table 1. A Comparison of AI Techniques in HCI

Technique	Advantages	Disadvantages	Suitable Applications
NLP (Transformers)	Strong semantic understanding; context-aware interaction	High computational cost; limited interpretability	Search engines, chatbots, virtual assistants
Reinforcement Learning (RL)	Real-time adaptation; learns from user feedback	Complex reward design; requires large interaction data	Personalized ranking, adaptive recommendation
Graph-Based Models (GNNs)	Captures long-term preferences; models relationships effectively	Scalability challenges; high complexity	Social networks, long-term recommendation systems
Embedding-Based Models	Efficient and s		

Application Domains

Healthcare

Healthcare is one of the most prominent domains for AI-enhanced HCI, where intelligent systems assist clinicians in diagnosis, treatment planning, and patient engagement. NLP models enable the interpretation of clinical notes and patient queries, improving decision support tools and conversational health assistants (Esteva et al., 2019). Reinforcement learning has been explored in treatment recommendation and adaptive patient monitoring, where real-time feedback can guide interventions (Gottesman et al., 2019). User modeling further supports personalization by tailoring health information and recommendations to individual patient histories. Challenges in this domain include ensuring interpretability, trust, and data privacy, as errors or biases may have life-critical consequences.

Education

In digital learning environments, personalization is key to enhancing student engagement and success. AI-driven tutoring systems use intent detection to interpret learner queries, while reinforcement learning adapts instructional content in response to student performance (Woolf et al., 2013). User modeling, including graph-based approaches, allows for the tracking of learner progress over time, supporting adaptive pathways through course material. These systems improve outcomes by aligning content delivery with the learner's evolving knowledge state, though concerns remain about fairness, explainability, and potential over-reliance on automation.

Recommender Systems and Entertainment

Recommender systems in music, video, and e-commerce are perhaps the most widely recognized application of AI in HCI. Platforms like Spotify and Netflix use embeddings and graph-based modeling to capture user preferences, while reinforcement learning fine-tunes recommendations through real-time feedback signals such as skips or completions (Zhou et al., 2018). NLP techniques enhance these systems by enabling natural language queries and content tagging. While such systems are highly effective in improving engagement and discovery, they raise concerns about filter bubbles, fairness, and transparency in the recommendation process.

Broader Interactive Systems

Beyond these sectors, AI-enhanced HCI is increasingly applied in domains such as customer service, where chatbots use NLP to handle support requests, and in smart environments, where adaptive interfaces personalize interactions with IoT devices. In each case, the integration of intent understanding, adaptive feedback, and user modeling enables richer, more human-like interactions.

Summary

Across healthcare, education, recommender systems, and broader interactive contexts, a consistent pattern emerges: effective AI-enhanced HCI relies on the interplay of semantic understanding, adaptive feedback mechanisms, and long-term user modeling. However, all domains face common challenges related to transparency, ethical deployment, and balancing personalization with broader social objectives such as fairness and inclusivity.

Evaluation Methodologies and Benchmarks

Evaluating AI-enhanced HCI systems is challenging, as success depends not only on accuracy but also on adaptability, transparency, and user satisfaction. Traditional metrics such as precision and recall provide useful baselines, but they must be complemented by user-centric measures and standardized benchmarks to capture interactive performance. This section reviews system-level metrics, user-focused evaluations, and recent benchmark datasets that support the study of adaptive and personalized interaction.

Importance of Evaluation

Evaluating AI-enhanced HCI systems presents unique challenges. Unlike traditional information retrieval, where success can be measured against fixed relevance judgments, HCI involves dynamic, context-dependent, and user-centered interactions. Systems must not only retrieve relevant information but also adapt over time, build trust, and sustain engagement. This requires a combination of quantitative metrics and qualitative assessments to capture both system performance and user experience (Turner and Scott, 2022).

System-Centric Metrics

Classical evaluation relies on precision and recall, which measure retrieval accuracy and coverage. More advanced metrics such as Mean Average Precision (MAP) and Normalized Discounted Cumulative Gain (NDCG) account for ranking quality by rewarding systems that place highly relevant items near the top of results (Järvelin and Kekäläinen, 2002). In adaptive systems, these metrics are often calculated at cut-off points (e.g., Precision@K, Recall@K) to reflect the limited depth of user attention (Turner and Scott, 2022). While valuable, these metrics cannot fully capture user satisfaction in interactive, evolving scenarios.

User-Centric Metrics

To complement system-centric measures, user-centric evaluation focuses on outcomes such as engagement, trust, transparency, and satisfaction. Implicit feedback signals—including click-through rates, dwell time, scrolling depth, and query reformulations—serve as proxies for user satisfaction (Radlinski and Joachims, 2008). Explicit feedback, gathered through surveys or controlled studies, provides additional insight into perceived usefulness and fairness. Evaluating trust and transparency is particularly important in domains such as healthcare and education, where users must rely on system decisions with confidence (Hall and Young, 2023).

Benchmarks and Datasets

Recent efforts have led to the creation of large-scale benchmarks that facilitate systematic evaluation across domains. The BEIR benchmark (Thakur et al., 2021) provides a heterogeneous suite of retrieval datasets for testing generalization in information retrieval. In personalization, the MIND dataset (Wu et al., 2020) and PENS dataset (Ao et al., 2021) offer large-scale user–item interaction logs for evaluating recommender systems. These benchmarks enable reproducibility and cross-system comparison, but they remain limited in capturing real-time adaptation and longitudinal personalization. Table 2 summarizes the key studies reviewed in this survey, highlighting their techniques, application domains, main contributions, and limitations.

Table 2. A Comparison of Studies Reviewed

Study	Technique	Application Domain	Key Contribution	Limitation
Devlin et al. (2019)	BERT (NLP)	Search, NLP	Contextual embeddings	High computational cost
Agarwal et al. (2022)	Counterfactual LTR	Ranking systems	Bias correction	Complex implementation
Yin et al. (2021)	RL-based search	Personalized search	Adaptive ranking	Data sparsity
Wu et al. (2020)	Embeddings	News recommendation	Large-scale dataset (MIND)	Short-term focus
Pu et al. (2025)	GNN	Recommendation	Multi-scale preference modeling	Scalability

Gaps and Future Directions

While progress in evaluation frameworks has been substantial, important gaps remain. Current metrics often overemphasize short-term interactions, neglecting long-term satisfaction and fairness. Many benchmarks fail to incorporate multimodal signals or cross-domain transfer, both of which are central to next-generation HCI. Future evaluation should focus on dynamic, longitudinal, and user-centric benchmarks, integrating both offline metrics and online user studies to provide a more holistic understanding of system performance.

Open Research Challenges

Journal of Science Basic and Applied Number. 21 June 2026

Despite notable progress in applying AI to HCI, current systems face limitations that hinder their ability to deliver truly adaptive, transparent, and user-centered experiences. While deep learning, reinforcement learning, and user modeling have each advanced personalization, their integration into unified, scalable frameworks remains incomplete. Moreover, as these systems move into sensitive domains such as healthcare, education, and digital platforms, the demands on trust, fairness, and accountability become even more pressing.

These realities point to several unresolved research challenges. Systems must become more interpretable and transparent so that users understand and trust their outputs. They must address bias and fairness issues arising from interaction data. The tension between real-time adaptation and long-term user satisfaction requires new strategies in reinforcement learning. The growth of multimodal interaction raises challenges in aligning heterogeneous signals, while privacy and ethical concerns intensify as more personal data is used for personalization. Finally, scalability and efficiency remain barriers to deploying advanced models at real-world scale.

The following subsections discuss these challenges in detail, outlining why they matter and how they shape the future of AI-enhanced HCI.

Interpretability and Transparency

As AI-driven HCI systems become more complex, their decision-making processes often operate as black boxes. Users and stakeholders require explanations for why certain results are ranked or why specific recommendations are made (Doshi-Velez and Kim, 2017). Without interpretability, trust in adaptive systems is undermined, particularly in high-stakes domains such as healthcare or education. Developing explainable AI (XAI) approaches that balance transparency with usability remains a critical challenge.

Fairness and Bias Mitigation

Personalized systems frequently learn from biased user interaction data, leading to position bias, popularity bias, or demographic bias (Agarwal et al., 2018; Oosterhuis and de Rijke, 2021). Left unaddressed, these biases risk reinforcing stereotypes and creating unequal access to information. Fairness-aware ranking algorithms and debiasing techniques must be further developed and standardized to ensure equitable outcomes across diverse user populations.

Real-Time Adaptation and Long-Term Balance

Reinforcement learning has shown promise in adapting rankings based on real-time feedback. However, systems often over-optimize for short-term engagement (e.g., clicks, dwell time) while neglecting long-term satisfaction and well-being (Miller and White, 2023). Balancing immediate feedback with persistent preference modeling is an unresolved research problem. Future work must design reward functions and evaluation metrics that incorporate short- and long-term user outcomes simultaneously.

Multimodal and Context-Aware Interaction

Modern HCI increasingly involves multimodal inputs—voice, gesture, image, and text combined in a single interaction (Wang and Liu, 2023). While multimodal AI models are emerging, integrating them into real-time adaptive systems remains complex. Challenges include aligning heterogeneous data streams, handling missing modalities, and ensuring consistent personalization across contexts. Research in multimodal fusion and context-aware modeling is needed to bridge this gap.

Privacy, Security, and Ethical Concerns

Personalization depends on detailed user modeling, often requiring large-scale behavioral and contextual data. This raises concerns regarding data privacy, consent, and security (Hall and Young, 2023). Emerging solutions such as federated learning and differential privacy offer promising directions, but they must be adapted to the unique requirements of interactive, adaptive systems. Ethical frameworks that guide data collection, storage, and usage are urgently needed.

Scalability and Deployment

While many AI techniques perform well in controlled experiments, deploying them at scale—where millions of users interact in real time—remains a major challenge. Issues of computational cost, latency, and robustness limit the applicability of advanced models such as transformers in production systems (Thakur et al., 2021). Developing lightweight, efficient architectures that retain accuracy without sacrificing speed is an open area of investigation.

2026-2030 Research Roadmap

Between 2026 and 2030, human–AI interaction is expected to move from simple tool use toward genuine co-evolution, where systems and users learn and adapt together. To achieve this, the

Journal of Science Basic and Applied Number. 21 June 2026

research community has outlined a roadmap built around five interconnected priorities.

The first priority is transparency. In the past, explain-ability was treated as an add-on, but in the coming years it must become a central design principle. Systems should be able to explain not only individual decisions but also the logic behind their adaptive behaviour over time. Explanations must be tailored to user expertise and context, adjusting in depth and form to ensure that people can truly understand how and why the system is acting. This shift requires new evaluation methods that assess comprehension, trust, and the alignment of user and system mental models.

Alongside transparency comes the need for co-adaptive intelligence. Instead of simply learning about users, systems must learn with them, evolving through long-term interaction. Such systems will detect shifts in user goals, preferences, and knowledge, while making their own learning process visible so that users can guide it. This transforms adaptation into a collaborative process, where safeguards are necessary to prevent misalignment or over-personalisation.

Equally vital is the integration of ethical intelligence into the very fabric of system design. Rather than relying on external oversight, ethics must be embedded in adaptive algorithms themselves. Systems should be built to respect autonomy, privacy, fairness, and accountability, operating with explicit consent and providing audit trails of decision-making. Importantly, ethical design should be participatory, ensuring that diverse user groups have a voice in shaping how systems adapt and what counts as responsible behaviour.

The roadmap also emphasises multimodal and context-aware interaction. Future systems must be able to interpret complex signals—such as speech, gesture, gaze, emotion, and environmental cues—so they can respond appropriately to users' states and situations. These capabilities will allow AI to act not only efficiently but also empathetically, recognising when users are tired, frustrated, or distracted. Seamless continuity across devices and contexts will further support fluid interaction, creating a sense of partnership rather than fragmented exchanges.

Finally, the roadmap highlights the need for scalable evaluation infrastructures. Traditional measures like task completion time are insufficient for adaptive systems that change over long periods. Instead, researchers must develop benchmarks, datasets, and simulation environments that capture the dynamics of adaptation, trust, and ethical alignment. This will make it possible to compare systems, reproduce findings, and ensure accountability across the field.

Taken together, these five directions outline a vision of human–AI systems that are transparent, collaborative, ethically aware, sensitive to context, and rigorously evaluated. By embedding these principles, the next generation of AI will not simply perform tasks more efficiently—it will grow alongside its users, creating interactions that are truly human-centred.

Conclusion

This survey has examined how Artificial Intelligence (AI) is reshaping Human–Computer Interaction (HCI) through advances in semantic modeling, reinforcement learning, and user personalization. We highlighted how NLP-driven intent detection, reinforcement learning for adaptive ranking, and graph-based user modeling together enable systems that anticipate needs, adapt continuously, and evolve with user behavior. At the same time, challenges remain around transparency, fairness, real-time adaptation, multimodality, privacy, and scalability. Addressing these issues is essential for building intelligent systems that are accurate, trustworthy, and ethically aligned.

References:

[1] Aggarwal, C. C. (2016). *Recommender Systems: The Textbook*. Springer.

[2] Agarwal, A., Takatsu, K., Zaitsev, I., & Joachims, T. (2018). A general framework for counterfactual learning-to-rank. *arXiv preprint arXiv:1805.00065*.

[3] Agarwal, A., Zaitsev, I., & Joachims, T. (2022). Counterfactual learning to rank for utility-maximizing metrics. In *Proceedings of SIGIR 2022*. ACM.

[4] Ao, X., Zong, Y., et al. (2021). PENS: A dataset and framework for personalized news recommendation. In *ACL 2021*.

[5] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. NAACL-HLT* (pp. 4171–4186).

[6] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.

[7] Esteva, A., et al. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24–29.

[8] Gottesman, O., et al. (2019). Guidelines for reinforcement learning in healthcare. *Nature Medicine*, 25(1), 16–18.

Journal of Science
Basic and Applied **Number. 21 June 2026**

[9] Green, S., & Black, R. (2024). An integrated AI framework for personalized information retrieval. *Journal of Artificial Intelligence Research*, 70, 500–520.

[10] Hall, G., & Young, N. (2023). Applying design science research in information systems engineering. *Communications of the ACM*, 65(1), 90–9