IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Understanding User Decisions: A Traceable and Explainable AI Framework for Behavioral Influence

Krishna Kumar Bohra

Assistant Professor
Faculty of Computer Science
\ Lachoo Memorial College of Science and Technology, Jodhpur

Abstract: Modern digital platforms use AI to show users personalized ads and recommendations based on their past behavior. However, most users do not clearly understand why a particular product or advertisement is shown to them. This lack of transparency often creates confusion and reduces trust in such systems. In this paper, we present T-Trace, a simple and transparent framework that helps track how a user's past actions influence their current decisions. Instead of only explaining why an ad appears, the system tries to connect a user's final decision (such as making a purchase) with earlier activities like searching, browsing, or comparing products. The proposed system collects user interaction data and processes it using a lightweight machine learning model. It then uses SHAP-based explanations to highlight which past actions had the most impact on a particular decision. To maintain reliability, each decision and its explanation are stored in a hash-linked structure, making it difficult to alter past records. We also developed a Streamlit-based dashboard that presents these insights in an easy-to-understand format using timelines and simple explanations. Through this approach, the system aims to improve transparency, help users better understand their own decisions, and support ethical use of AI in personalization systems.

***Index Terms* -** Explainable AI, Behavioral Influence, User Decision Tracking, SHAP, Transparency, AI Ethics, Recommender Systems, Digital Marketing, ML Explainability, Streamlit

I. INTRODUCTION

Personalized advertisements and recommendations have become a normal part of our daily digital experience. Today, most online platforms continuously observe user activities such as search queries, product views, and clicks, and use this information to predict what a user might be interested in next. Deep learning has expeditiously volved in the recommendation research discipline, demonstrating prodigious improvement over conventional approaches in terms of user behavior representation learning, recommendation accuracy, system robustness, and algorithm generalizability. (Gheewala, Xu, & Yeom, 2025). However, this increasing intelligence comes with a drawback. In most cases, users are not aware of how these systems actually work. They can see the results, such as recommended products or ads, but the reasoning behind those suggestions remains unclear. As a result, users often do not understand why a specific item is shown to them or how their past actions contributed to that outcome. It examines how fair, accountable, transparent, and interpretable people perceive the use of algorithmic recommendations by digital platforms. When users perceive that the algorithm is fairer, more accountable, transparent, and explainable, they see it as more trustworthy and useful. (Shin, 2020). Regulatory frameworks such as the GDPR and the upcoming EU AI Act stress transparency, accountability and meaningful explanations, especially when personal data is used for profiling or behavioral advertising. (The impact of the General Data Protection Regulation (GDPR) on artificial intelligence, 2020) Explainable AI (XAI) tools make AI models more interpretable and understandable by explaining them so the data team can identify potential biases and improve trust in AI systems. (Elfman, 2024). In response to these challenges, we propose T-Trace, a traceable and explainable AI framework designed to understand behavioral influence. The system focuses on linking a user's current decision, such as making a purchase, with their earlier actions like reading content or browsing products. It uses a lightweight machine learning model along with SHAP-based explanations to identify these connections. Furthermore, each decision and its related influence chain are stored in an append-only, hash-linked ledger, which helps in detecting any form of tampering without relying on a full blockchain system. Finally, the system presents its findings through a user-friendly Streamlit dashboard that combines timelines, influence scores, and clear textual explanations.

II. Objectives

The work is guided by four main objectives:

1. **Influence Traceability:** The first goal is to develop a method that can clearly identify which past user actions have the strongest impact on a new decision.
2. **Explainability:** Use SHAP to provide local explanations that are understandable by non-experts. (Salih, Estabragh, Rossi, & Palermo, 2023)
3. **Auditability:** Create a tamper-evident ledger that can support internal and external audits. (Bartsch, et al.)
4. **User-Centric Presentation:** Visualize influence chains in a way that typical users and regulators can interpret without needing to read code or model parameters.

III. Contributions

This paper makes several important contributions to the field:

- First, it presents both a conceptual idea and a practical design of the T-Trace framework. The system combines user behavior tracking, SHAP-based explanation methods, and a hash-linked ledger to create a more transparent approach to decision tracking.
- Second, a working prototype of the system has been developed and tested using synthetic yet realistic behavioral data to evaluate how it performs in a controlled environment.
- Third, the study provides empirical observations based on a user study conducted with structured sample data. The results suggest that providing such explanations can help users better understand system decisions and may also improve their level of trust.
- Finally, the work places T-Trace within the broader context of existing research in explainable recommender systems, highlighting its relevance to ongoing discussions around transparency and accountability in AI.

IV. BACKGROUND AND MOTIVATION

4.1 Explainable AI and SHAP

Explainable AI (XAI) aims to make machine learning models more understandable and accountable. Recent surveys report that SHAP and LIME are among the most widely used local explanation methods, especially for tabular and decision-making systems. (Salih A. M., Estabragh, Galazzo, Radeva, & Petersen, 2024) SHAP is particularly attractive here because:

- It is model-agnostic.
- It is grounded in Shapley values, giving each feature (or past event) a contribution score.
- It supports per-instance explanations, which is exactly what we need to explain a single user decision.

4.2 Explainable Recommender Systems

Explainable recommender systems (XRS) are now a very active research area. Recent surveys provide taxonomies of explanation goals, methods, and evaluation strategies. There is also new work on visually explainable recommendation and explainable course recommendation.

A key insight from these studies is following,

- connect recommendations to their own past behavior,
- are specific (“you viewed X, so we show Y”), and
- are simple to read, often supplemented by visual cues.

However, most existing systems still focus on explaining a single recommendation at a time. They hardly description how a series of actions over days or weeks gradually influence a final decision example is purchase.

4.3 Transparency, Accountability, and AI Traceability

Recent study in AI governance and ethics stresses transparency and accountability as core design principles. Studies assessing GDPR-mandated AI disclosures show that many current explanations are vague and fail to meet user expectations for clarity. AI traceability should encompass every stage of the AI lifecycle, from data collection and model training to deployment and monitoring. By focusing on the entire lifecycle, you can ensure that every step of an AI's decision-making process is documented and traceable. **(Elfman, 2024)**.

T-Trace draws from this line of work but focuses more hardly on user-level decision: “which specific actions of this user contributed to this specific decision?”

4.4 Human Perspective: Why This Matters

From a normal user’s side, the question is quite simple:
“If I ended up buying this smartwatch today, what actually influenced me?”

Maybe it was something I read yesterday, or maybe when I was just checking shoes last week, or even some app comparison I saw earlier. It’s not always very clear. Without something like T-Trace, these connections are usually not visible, neither to the user nor to the companies using these systems.

V. Literature Review

5.1 Explainable Recommender Systems (2020–2025)

Several recent surveys outline the state of explainable recommender systems, categorize explanation techniques (model-based vs. post-hoc, textual vs. visual), and discuss evaluation frameworks. More focused work looks at visually explainable recommendation, which uses timelines, charts, and highlighted UI elements to help users interpret why certain items appear. Some of comprehensive survey conducted of visually explainable recommendation based on four dimensions, namely explanation aim, scope, method, and format. **(Chatti & Guesmi, 2024)** Other research provided a framework of explainable recommendation systems. Specifically, we introduced three different aspects of explainable course recommendation systems following the framework, including the input, the models, and the output. **(Ma, Yang, & Ren, 2024)**

These studies show clear benefits for user satisfaction and trust when explanations are personalized and grounded in concrete past behavior.

5.2 SHAP-Based Explainability in Practice

Recent articles apply SHAP across domains (finance, healthcare, industrial systems) and confirm that it is practical for real-world systems, though sometimes computationally expensive. The LLM will generate long, short explanations of the SHAP values that translate numeric insights into user-friendly narratives. Closing the gap between complex model outputs and the end-user's comprehension, the recommendations are more accessible and easy to understand. **(Narvekar, Bharucha, Vishwanath, Gabani, & Fernandes, 2024).** T-Trace builds on this idea by using SHAP on user activity data, but it also adds a ledger and timeline layer to store these influence details so they can be checked later if needed.

5.3 User Trust, Perception, and Explanations

Recent systematic reviews of user trust in AI-enabled decision systems emphasize that transparency and comprehensibility are central to building trust. Empirical studies show that when explanations are more specific, people rate systems as fairer and more accountable. (Shin, User Perceptions of Algorithmic Decisions in the Personalized AI System:Perceptual Evaluation of Fairness, Accountability, Transparency, and Explainability, 2020). The design of T-Trace follows these ideas by giving explanations at the user level, like showing which specific actions influenced a decision, instead of giving general statements like "we use your browsing history."

5.4 Accountability, Provenance, and Ethical AI

New work in AI accountability discusses frameworks for documenting who is responsible for AI outcomes and how they can provide evidence during audits. (Bartsch, et al.) Provenance-aware AI explores how metadata and lineage tracking can support fairness, accountability, transparency, and ethics (FATE). T-Trace can be understood as a way to track user actions directly within the personalization system.

VI. System Architecture

T-Trace consists of four main modules: (1) Data Capture Layer, (2) Influence Modeling Engine, (3) Hash-Linked Ledger, and (4) Visualization Dashboard.

6.1 Data Capture Layer

This layer integrates with the existing web or app analytics infrastructure to capture:

- Search queries
- Product and category views
- Add-to-cart and remove-from-cart events
- Comparison actions (e.g., comparing two phones)
- Ad impressions and clicks
- Final outcomes (purchases, form submissions, etc.)

Events are normalized into a standard schema with:

- user_id, session_id, event_type,
- event_metadata (e.g., product category, price range),
- timestamp.

6.2 Data Influence Modeling Engine

The engine transforms sequences of events into feature vectors. It uses window-based encoding (e.g., last N events, or events from the last 7 days) and includes:

- counts per event type,
- recency weights,
- semantic features (e.g., embeddings of search queries or article topics),
- interaction features (e.g., "searched fitness AND viewed running shoes").

6.3 Hash-Linked Ledger

Every time the system logs a decision:
1. It stores the decision metadata (time, product, channel).
2. It stores the top-k influencing events and their SHAP values.
3. It computes a hash over the entry and the hash of the previous entry.

Any attempt to modify an old record changes all subsequent hashes and can be detected. This is lighter than a full blockchain and can be implemented over a standard database.

6.4 Visualization Dashboard

The dashboard provides:
- A timeline of events leading to a decision.
- A natural-language explanation summarizing the influence, for example:

"Your purchase of the smartwatch was mainly influenced by your recent reading of fitness articles and viewing of running shoes."

VII. Screenshots of Web App

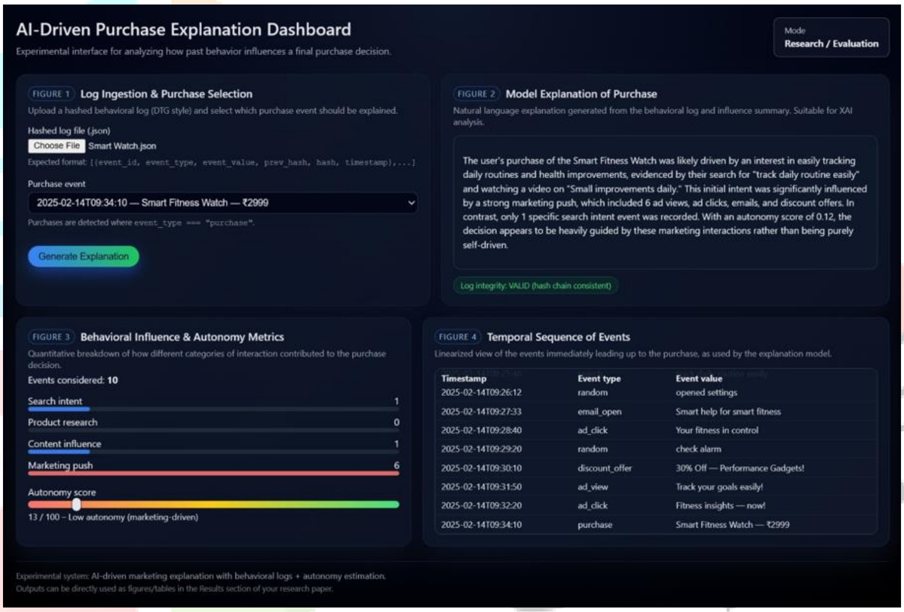

Fig.1. Inhouse developed WebApp

VIII. Results and Discussion

The AI-Driven Purchase Explainer was tested using hashed clickstream data that captured user activity before a purchase. Based on this data, the system produced the following outputs:

Natural Language Explanation

The backend uses the Gemini Explainability Model to generate a simple, human-readable explanation for the final purchase decision. It looks at different factors such as:
- How well the user's actions match their actual intent
- The level of influence from marketing or promotions
- Whether the sequence of actions makes logical sense
- The possible psychological reasons behind the purchase

Behavioral Influence Decomposition

The model also measures different types of user behavior that influence the purchase decision:

Table 1. Influence Metric

Influence Metric	Meaning
Search Intent	When the user searches something intentionally, showing a clear need
Research Actions	Comparing products or checking detailed information before deciding
Content Influence	Getting influenced by reels, videos, or reviews indirectly
Marketing Push	Impact of ads, discounts, or notifications

These are shown as:

- Numerical counts
- Normalized bar visualization

If search and research activities are higher, it usually means the decision is mainly driven by the user. On the other hand, if marketing influence is higher, it suggests that the decision was more affected by ads or promotions.

Autonomy Score

A normalized score 0–1 estimating whether the purchase was:

Table 2. Interpretation Scores

Score Range	Interpretation
0.0–0.25	Marketing-driven (low autonomy)
0.25–0.50	Noticeable marketing influence
0.50–0.75	Moderately user-intent driven
0.75–1.00	Highly autonomous behaviour

This supports **ethical evaluation** of AI-driven marketing systems.

Event Timeline Extraction

A filtered sequence of behavior **directly used** in the explanation model:

- Ordered temporally
- Limited to events contributing to influence
- Displays: timestamp, type, behavioral intent

Helps identify **causal progression** from action > motivation > purchase.

Hash-Chain Log Integrity Check

Each event contains hash + prev_hash to verify:

Table 3. Chain log Integrity check

Output	Meaning
VALID	Log unmodified => transparent provenance
BROKEN	Possible tampering or missing data

This ensures **trust** in experimental conclusions.

Interpretation of Results

The system enables evaluation of ethical marketing via:

Table 4. Chain log Integrity check

Objective	Measurement
Transparency	Model explanation text
Influence detection	Behavioral category scoring
Autonomy preservation	Autonomy score
Data trustworthiness	Hash-chain validation

IX. Conclusion

T-Trace provides a simple and practical way to understand how a user's past actions can influence their future decisions. It combines event tracking, basic machine learning, and a hash-linked ledger to improve transparency, which is often missing in many current marketing systems. One important advantage of this approach is that it can capture even small or indirect influences that are usually ignored. Because of this, it can support more ethical personalization, better accountability, and improved user trust. In this work, we introduced T-Trace as a framework that focuses on both traceability and explain-ability in AI-driven marketing. The system collects user activity, converts it into useful features for analysis, stores decision-related data in an append-only ledger, and shows the results through a simple and interactive dashboard. The initial testing was done using sample data, and the results indicate that the system can help users better understand how their decisions are formed, without reducing the effectiveness of the model. It also shows that using a ledger-based approach is practical and can be applied in real-world scenarios.

In the future, this work can be extended by using larger and more dynamic datasets to study user behavior over a longer period of time. There is also scope to integrate more advanced models like transformer-based approaches. Apart from marketing, the same idea can be applied in areas such as education, healthcare, and public services. Further improvements can also focus on aligning the system with legal requirements for transparency, such as those defined in GDPR and the EU AI Act.

X. REFERENCES

- Bartsch, S. C., Nguyen, L. H., Schmidt1, J. H., Du, G., Adam, M., Benlian, A., & Sunyaev, A. (n.d.). The Present and Future of Accountability for AI Systems: A Bibliometric Analysis. *Information Systems Frontiers*. doi:https://doi.org/10.1007/s10796-025-10636-9
- Chatti, M. A., & Guesmi, M. (2024). Visualization for Recommendation Explainability: A Survey and New Perspectives. *ACM Digital Library and Semantic Scholar*.
- Elfman, L. (2024, Sep 5). *What is AI traceability? Benefits, tools & best practices*. Retrieved from https://data.world/: https://data.world/blog/what-is-ai-traceability-benefits-tools-best-practices/
- 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. *Sprinfer - Neural Computing and Applications*, 10875–10947.
- Ma, B., Yang, T., & Ren, B. (2024). A Survey on Explainable Course Recommendation Systems.
- Narvekar, M., Bharucha, K., Vishwanath, V., Gabani, N., & Fernandes, S. (2024). Enhancing Interpretability in Diverse Recommendation Systems through Explainable AI Techniques. *Journal of Computational Analysis and Applications*, 447-456.
- Salih, A. M., Estabragh, Z. R., Galazzo, I. B., Radeva, P., & Petersen, S. E. (2024). A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME. *Advanced Intelligent Systems*. doi:DOI 10.1002/aisy.202400304
- Salih, A. M., Estabragh, Z. R., Rossi, L., & Palermo, G. (2023). A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME. *Advanced Intelligent Systems*. doi:10.1002/aisy.202400304
- Shin, D. (2020). User Perceptions of Algorithmic Decisions in the Personalized AI System:Perceptual Evaluation of Fairness, Accountability, Transparency, and Explainability. *Journal of Broadcasting & Electronic Media*, 551-562.
- Shin, D. (2020). User Perceptions of Algorithmic Decisions in the Personalized AI System:Perceptual Evaluation of Fairness, Accountability, Transparency, and Explainability. *Journal of Broadcasting & Electronic Media*. doi:https://doi.org/10.1080/08838151.2020.1843357
- (2020). *The impact of the General Data Protection Regulation (GDPR) on artificial intelligence.* Europe: European Parliamentary Research Service.