Explainability in Music Recommender System

Shahrzad Shashaani
shahrzad.shashaani@tuwien.ac.at
TU Wien, Faculty of Informatics
Vienna, Austria

ABSTRACT

Recommendation systems play a crucial role in our daily lives, influencing many of our significant and minor decisions. These systems also have become integral to the music industry, guiding users to discover new content based on their tastes. However, the lack of transparency in these systems often leaves users questioning the rationale behind recommendations. To address this issue, adding transparency and explainability to recommender systems is a promising solution. Enhancing the explainability of these systems can significantly improve user trust and satisfaction. This research focuses on exploring transparency and explainability in the context of recommendation systems, focusing on the music domain. This research can help to understand the gaps in explainability in music recommender systems to create more engaging and trustworthy music recommendations.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; **Music retrieval**.

KEYWORDS

Music Recommender System, Explainability, Explainable AI

ACM Reference Format:
Shahrzad Shashaani. 2024. Explainability in Music Recommender System. In *18th ACM Conference on Recommender Systems (RecSys '24), October 14–18, 2024, Bari, Italy.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3640457.3688028

1 BACKGROUND AND MOTIVATION

The rapid growth of digital music platforms has revolutionized how users discover and consume music. Central to this experience are recommender systems, which capture users' listening habits and preferences and suggest personalized playlists to each user. Despite their effectiveness, a significant challenge remains: the lack of transparency in how recommendations are generated. Users often find it difficult to understand why certain tracks are suggested, leading to mistrust and reduced satisfaction with the system.

Explainability in recommender systems is not just a technical challenge but a critical user experience issue. Providing clear and understandable reasons for recommendations can enhance user trust, engagement, and overall satisfaction. In the context of music, where emotional and personal connections to content are strong, the need for explainable recommendations is even more important. This research seeks to explore how enhancing the transparency and explainability of music recommender systems can address these concerns.

1.1 Problem Statement

Current recommender systems typically function as "black boxes," offering little insight into their decision-making processes. This opacity can lead to user uncertainty, especially when recommendations appear irrelevant or unexpected. The lack of explainability also limits users' ability to refine and personalize their recommendation preferences actively. Users often desire to understand the rationale behind recommendations to make informed decisions and to adjust the system's suggestions to better match their preferences.

Based on [33], music recommender systems face unique challenges compared to other domains due to the complex and subjective nature of music preferences. Unlike typical recommendation tasks, music recommendation must account for a range of contextual and situational factors, such as user mood, activity, and long-term versus short-term listening trends. Additionally, music preferences are highly personal and dynamic, often influenced by social and cultural contexts. This complexity requires sophisticated models that go beyond simple collaborative filtering or content-based approaches commonly used in other domains.

Moreover, music recommender systems must effectively handle sparse and imbalanced data, as users often listen to a small subset of the available catalog, creating a significant popularity bias. This bias can limit the exposure to diverse music, reducing user satisfaction. Addressing this issue necessitates the integration of multi-faceted recommendation strategies, such as hybrid models that combine content-based filtering, collaborative filtering, and contextual information [33].

Rerecommendation, or the process of recommending previously suggested items, adds another layer of complexity. Users often prefer fresh recommendations but may appreciate occasional reminders of past favorites, especially when contextual factors change, such as the user's mood or listening environment. This requires the system to balance between novelty and familiarity, ensuring a satisfying user experience while avoiding redundancy [33].

As a result, explainability in music recommender systems is particularly challenging due to the need for transparent and interpretable explanations that resonate with users on an emotional and psychological level. Furthermore, music is often consumed in varied contexts—work, travel, celebration, relaxation, exercise—which adds another layer of complexity to recommendation systems. This is essential for building trust and enhancing user engagement. Overall, while advancements have been made, significant gaps remain in developing comprehensive and effective explainability frameworks

RecSys '24, October 14–18, 2024, Bari, Italy

ACM ISBN 979-8-4007-0505-2/24/10
https://doi.org/10.1145/3640457.3688028

connected to the music domain, underscoring the need for ongoing research and innovation in this area.

Moreover, explainable music recommender systems benefit multiple stakeholders. For users, enhanced transparency and understanding can lead to higher trust and more tailored music experiences. For artists and music producers, insights into why their music is recommended can inform marketing strategies and audience engagement. For the platforms themselves, better explainability can differentiate their service in a competitive market, improving user retention and satisfaction.

Our research aims to explore frameworks for explainable music recommender systems, focusing on identifying key factors that enhance explainability and examining their impact on user trust and satisfaction. We will review similar approaches in other domains, identify gaps specific to music recommender systems, and propose potential solutions to address these gaps. By doing so, we hope to lay the groundwork for future studies and implementations that can provide personalized and transparent recommendations in the music domain.

1.2 Motivation

The rapid development of digital music platforms has significantly transformed the way users interact with and discover music. Recommender systems play an essential role in this transformation, offering personalized suggestions that provide to individual tastes and preferences. Despite their widespread use and undeniable convenience, these systems often operate as "black boxes," providing little to no insight into how recommendations are generated. This lack of transparency can undermine user trust and satisfaction, as users may find it difficult to understand the rationale behind the suggestions they receive.

Explainability in recommender systems is an emerging field that seeks to address these concerns by making the inner workings of these systems more transparent and understandable. The primary goal is to enhance user trust and engagement by providing clear, comprehensible reasons for recommendations. In the context of music, where personal and emotional connections to content are particularly strong, the need for explainable recommendations is even more critical. Users are more likely to engage with and appreciate a system that not only considers their tastes but also explains the reasoning behind its suggestions. Effective explainability aims to achieve several key objectives [35, 41, 42]:

- **Transparency:** Clarify how the system operates.
- **Trust:** Boost users' confidence in the system.
- **Effectiveness:** Assist users in making informed decisions.
- **Efficiency:** Help users make decisions more quickly.
- **Satisfaction:** Enhance user experience and enjoyment.
- **Fairness:** Be fair and equitable.
- **Responsibility:** Clarify or rationalize decisions or actions.
- **Scrutability:** Enable users to provide feedback and correct the system.
- **Persuasiveness:** Encourage users to try or purchase recommended items.

Previous research has highlighted the importance of explainability in recommender systems across various domains, including e-commerce, movies, and general consumer products [35, 51]. However, there is a notable gap in the literature specifically addressing explainability in music recommender systems. While some studies have explored user perceptions of transparency in music recommendations [14], comprehensive frameworks and methodologies tailored to the unique challenges of music recommendation remain underdeveloped. The motivation for this growing research topic can be summarized as the following:

- Enhancing User Trust and Satisfaction: By providing clear and understandable explanations for recommendations, with the goal of building greater trust between users and the system. This trust is crucial for long-term user engagement and satisfaction.
- Improving User Interaction: Transparent recommendations can empower users to make more informed decisions and refine their preferences actively. This interaction can lead to a more personalized and satisfying user experience.
- Advancing the Field of Music Recommendation: By addressing the specific challenges of explainability in music recommender systems, doing research in this field can contribute to the broader field of recommender systems, offering insights and methodologies that can be applied across various domains.

In conclusion, the lack of explainability in current music recommender systems represents a significant gap in both research and practical application. Our research aim is to bridge this gap by developing a comprehensive understanding of how transparency can be integrated into music recommendation processes, ultimately enhancing user trust, satisfaction, and engagement.

1.3 Research Questions

Here are some questions, that need broader research to answer, mostly in the music recommender system domain:

(1) **What incorporates a good explanation for recommendations, particularly in the music domain?**
Determining what makes an explanation good or sufficient is critical, especially in the highly subjective field of music. This question explores the elements that contribute to effective explanations, such as clarity, relevance, and comprehensibility, and how these elements can be connected to individual user preferences and contexts. By identifying the characteristics of good explanations, we aim to enhance user satisfaction and trust in music recommender systems. Another key consideration is the privacy of the model; thus, we need to find a way to balance transparency and user understanding without revealing the complete recommendation process details.

(2) **How can we measure the effectiveness of explainability in recommender systems within the music domain?**
Measuring the effectiveness of explainability is a complex challenge that requires robust evaluation frameworks. This question aims to establish metrics and methodologies for assessing how well explainable recommendations perform in the music domain. While general evaluation frameworks

exist, different kinds of aims can be considered for explainability, and they can be inconsistent. Therefore, any evaluation needs to use appropriate metrics that match the aim we plan to investigate [36]. By developing and validating new measurement techniques, we seek to provide a comprehensive understanding of explainability's impact in this specific context.

(3) **What are the most effective methods for enhancing the explainability of music recommender systems?** This question is fundamental to our research as it seeks to identify and evaluate the various techniques that can be employed to make music recommender systems more transparent. While there has been significant work on explainability in general recommender systems [35, 51], it can also be domain-specific due to available attributes. So, the unique challenges associated with music recommendations—such as the subjective nature of music preferences—require tailored approaches. Understanding which methods are most effective in this specific context is crucial for developing systems that users can trust and understand.

(4) **How can transparent recommendation processes improve user trust and satisfaction?** User trust and satisfaction are critical components of any successful recommender system. Previous research has shown that explainability can enhance user trust and satisfaction in various domains [12, 41]. However, the specific impact of transparency in music recommender systems remains underexplored. By investigating this relationship, we aim to provide empirical evidence on how explainable recommendations can lead to more positive user experiences, thereby increasing the overall effectiveness of these systems. There is not a direct way to measure the effect of explainability, so in many research, they used the user study [38].

(5) **What is the impact of explainable recommendations on user engagement and interaction with the system?** Engagement and interaction are key metrics for evaluating the success of recommender systems. It can be measured by how the user has control over the recommendations [38]. This question considers how providing explanations for recommendations can influence user behavior, such as the frequency and duration of interaction with the system. While some studies have touched upon this aspect in general settings [49], there is a lack of focused research on the music domain. Our goal is to fill this gap by examining how explainability affects user engagement specifically in music recommender systems.

1.4 Relevance and Importance

The formulation of these research questions is grounded in the recognition that explainability is a multifaceted issue with significant implications for user experience. The questions were derived from a thorough review of the existing literature and identified gaps, particularly the lack of focused research on music recommender systems because of the mentioned particularities of music recommendation. By addressing these questions, our research aims to contribute to the broader understanding of explainability in recommender systems and to provide insights into some of the explainability goals, such as improving user trust, satisfaction, engagement, and overall system transparency and effectiveness.

2 RELATED WORK

In this section, we first review advanced algorithms in recommender systems. We then explore explainability both in the broader context of machine learning and specifically within recommender systems. Finally, we highlight existing approaches and identify areas where further research is needed, particularly in explainability in music recommender systems.

2.1 Approaches in Recommender Systems

Recommender systems typically utilize two main approaches to suggest items: content-based and collaborative filtering methods [9, 26]. Each approach has distinct advantages, leading some researchers to develop hybrid models that combine both methods for enhanced performance [9, 23]. Below, we introduce some advanced approaches for recommendation tasks.

Deep learning techniques have been increasingly applied to recommendation systems. Yang et al. (2019) proposed a personalized recommender system using a time-aware CNN model (TC-PR) for movie recommendations, incorporating user and item features along with temporal context [48]. Saga et al. (2018) integrated CNN with probabilistic matrix factorization, utilizing image shape features for fashion recommendations [30]. Daneshvar et al. (2022) developed RSLC-Net, combining CNN and LSTM architectures to analyze user ratings and visual features, and incorporating social influence data for enhanced movie recommendations [7].

Emotion-based music recommendations have also benefited from deep learning techniques. Joshi et al. (2021) used LSTM, CNN, and CNN-LSTM architectures to detect user emotions from text or facial expressions to suggest songs and playlists [13]. Kim et al. (2021) proposed an advertising video recommendation system that uses CNNs to analyze facial expressions and predict user preferences [15].

Deep learning enhances recommendation algorithms' performance, particularly with large datasets. Wang et al. (2020) developed an LSTM-CNN model for movie recommendations, improving user experience by analyzing behavior data [43]. Van et al. (2013) demonstrated the feasibility of predicting latent factors from music audio for recommendations, showing its effectiveness for new and less popular music [39]. Oramas et al. (2017) addressed the cold-start problem by integrating text, audio, and user feedback data in a deep network architecture [27].

Sachdeva et al. (2018) introduced an attentive neural architecture using item sequences and features to learn user preferences and recommend the next song [29]. Lin et al. (2018) designed a Heterogeneous Knowledge-Based Attentive Neural Network (HK-ANN) for short-term music recommendations using graphical, textual, and visual information [18].

As most of the advanced approaches mentioned above use various deep learning structures, it becomes increasingly difficult to determine the underlying cause of specific recommendations. This makes explainability crucial for stakeholders to maintain confidence and trust in the recommendations they receive. By understanding the reasons behind recommendations, stakeholders can feel more positively about the system's outputs and are more likely to engage with and benefit from the recommendations. Consequently, research has been done on various aspects of explainability to address these challenges.

2.2 Explainability in General

Explainability in Artificial Intelligence (AI) and Machine Learning (ML) has received considerable attention in recent years, focusing on making complex models understandable to users. Various approaches have been proposed to enhance interpretability, transparency, and trust.

Gilpin et al. (2018) provide an overview of the interpretability of machine learning models, discussing different methods to explain AI decisions, which contributes to **Transparency** and **Fairness** [10]. Adadi and Berrada (2018) survey explainable AI (XAI) techniques, categorizing them into post-hoc/intrinsic, global/local, and model-specific/model-agnostic methods, aimed at **Transparency** and **Trust** [2]. Doshi-Velez and Kim (2017) call for a rigorous science of interpretable machine learning, emphasizing the need for standardized evaluation methods, targeting **Transparency** and **Fairness** [8].

Ribeiro et al. (2016) introduce LIME, a technique for explaining the predictions of any classifier, which has become widely adopted in the field. LIME works by approximating the behavior of complex models with interpretable local models, thereby providing insights into individual predictions, addressing **Transparency** and **Trust** [28]. Samek et al. (2017) focus on visualizing deep learning models to make their operations more transparent, introducing methods like heatmaps to highlight influential regions in input data, which aim at **Transparency** and **Trust** [31]. Lipton (2018) discusses the myth of model interpretability, addressing the challenges and misconceptions surrounding this topic and emphasizing the need for clear definitions and expectations, contributing to **Transparency**, **Trust** and **Effectiveness** [19]. Zhang (2020) addresses the challenges of interpretable deep learning, proposing new methods to enhance transparency by combining attention mechanisms with visual explanations, which focus on **Transparency** and **Effectiveness** [50].

Lundberg and Lee (2017) propose SHAP values, a unified approach to interpreting model predictions that have gained significant traction. SHAP values provide a theoretically sound method for attributing the contribution of each feature to a prediction, enhancing the **Transparency** and **Effectiveness** of complex models [21]. Gunning (2019) outlines the DARPA Explainable AI (XAI) program, which aims to produce more explainable models while maintaining high performance, highlighting the need for user-centered explanations, addressing **Transparency**, **Trust**, and **Effectiveness** [11]. Miller (2019) draws insights from social sciences to improve AI explanations, emphasizing the importance of considering user perspectives and cognitive biases in the design of explanations, which contributes to **Transparency**, **Trust**, and **Effectiveness** [25]. Caruana et al. (2015) demonstrate intelligible models for healthcare, showing the practical benefits of interpretable models in critical applications such as predicting patient risk factors, focusing on **Effectiveness** and **Responsibility** [5]. Despite all the research, there has been limited research on what constitutes a good explanation in the context of AI systems. To address this gap, researchers have developed frameworks grounded in philosophy, psychology, and interpretable machine learning to investigate and define the characteristics of effective explanations, which aim at **Transparency** and **Effectiveness** [20].

Most existing approaches focus on implementing explainability in classification tasks, where it is easier to identify common features within a specific class and use them in explanations. However, explainability in recommendation tasks is more complex due to additional factors such as user preferences, item description, context, and short- and long-term trends that must be considered. These factors make the task of providing clear and comprehensive explanations more challenging but also more essential.

2.3 Explainability in Recommender Systems

Explainability in recommender systems has been explored to enhance user trust, transparency, interpretability, and satisfaction [22]. Several studies have proposed different methods to achieve these goals.

Abdollahi and Nasraoui (2017) use explainability for constrained matrix factorization in recommender systems, improving **User Satisfaction** by making recommendations more understandable through the use of explainable factors, thus enhancing **Transparency** [1]. Tintarev and Masthoff (2015) review design and evaluation strategies for explainable recommendations, reviewing how explainability can be added to different recommender approaches, including the importance of **Transparency**, **Scrutability**, **Persuasiveness**, **Effectiveness**, and **Satisfaction** [37]. Zhang and Chen (2020) present a detailed survey on explainable recommendations, identifying key techniques and challenges such as the need for personalized and context-aware explanations, aiming at **Personalization**, **Transparency**, and **Trust** [51]. Zhang et al. (2019) further explore context-aware perspectives that can help to improve explainability [46].

Kouki et al. (2019) introduce personalized explanations for hybrid recommender systems, demonstrating how tailored explanations can improve **User Satisfaction** by aligning recommendations with user preferences and providing clear rationales, enhancing **Transparency** and **Trust** [16]. Schmude et al. (2023) examine the effects of algorithmic decision-making (ADM) systems' transparency on knowledge-sharing decisions, showing how explanations can influence user behavior and **Trust** [34]. Vig et al. (2009) introduced the concept of tagsplanations, which makes recommender systems explainable by predicting tag relevance and preference, providing clear explanations for recommendations, focusing on **Effectiveness** [40].

Wang et al. (2018) propose a multi-task learning approach to generate explainable recommendations using opinionated text data, highlighting the importance of textual explanations in enhancing user understanding, aiming at **Transparency**, **Satisfaction**, and **Effectiveness** [44]. Balog and Radlinski (2019) develop transparent and scrutable user models for personalized recommendation, improving user understanding of the system's logic through clear and interpretable explanations, addressing **Transparency**, **Scrutability**, and **Trust** [4]. Kouki et al. (2020) focus on automatically generating personalized explanations for hybrid recommender systems, enhancing user engagement and trust by providing contextually relevant justifications, focusing on **Personalization** and **Engagement** [17].

Zhang et al. (2018) propose a framework for explainable recommendation systems based on a knowledge graph combined with multi-objective optimization. Their approach leverages the relationships between users and items to construct an explainable candidate list, subsequently optimized for precision, diversity, and explainability. This method is particularly focused on enhancing **Transparency** and **Effectiveness** by clearly showing how recommendations are generated and why certain items are suggested [47].

Abdollahi and Nasraoui (2018) explore the application of knowledge-based explainability in recommender systems, focusing on the integration of explicit knowledge sources to generate understandable explanations for recommendations. Their work aims to enhance **Transparency** and **Trust** by utilizing structured knowledge to rationalize the recommendation process [32].

By understanding the specific aims of these approaches, we can better design and evaluate explainable systems that meet the diverse needs of stakeholders in the music recommendation domain. This understanding highlights the potential gaps in current research and underscores the importance of continued exploration in this field to improve user experience and trust.

2.4 Explainability in Music Recommender Systems

Research specifically focusing on explainability in music recommender systems is relatively limited but growing. Several studies have begun to address the unique challenges of this domain, aiming to enhance transparency, trust, and user satisfaction.

Wang et al. (2019) explore explainable music recommendations via linked data, leveraging semantic relationships between music entities to provide clearer and more meaningful explanations, focusing on **Effectiveness** [45]. Afchar et al. (2021) implement an explainable music recommendation system using matrix factorization and clustering, demonstrating the effectiveness of these techniques in the music domain by clustering similar user preferences and providing explanations based on these clusters, aiming at **User Satisfaction** [3].

Zhao et al. (2018) use user comments to explain music recommendations, showing the potential of natural language explanations for enhancing the transparency of music recommendations in conversational applications, focusing on **User Satisfaction** [53]. Castells et al. (2015) discuss novelty and diversity in recommender systems, which are critical for making music recommendations more engaging and understandable by providing diverse and novel song choices that can be explained in terms of user preferences, addressing **User Satisfaction** [6]. Zhao et al. (2019) propose personalized experiences to explain music recommendations, demonstrating the benefits of tailored explanations that align with individual user tastes and preferences, focusing on **Persuasiveness** and **Effectiveness** [52].

Melchiorre et al. (2022) introduce ProtoMF, a prototype-based matrix factorization approach for explainable recommendations. ProtoMF learns sets of user and item prototypes that capture general consumption characteristics, representing users and items as vectors of similarities to these prototypes. This aims are **Effectiveness**, **Fairness**, and **Transparency** in recommendations by allowing the prediction to be broken down into contributions from distinct prototypes [24].

These studies highlight the unique challenges and opportunities in making music recommender systems more explainable. By addressing these challenges, researchers can enhance user trust and satisfaction, providing a more explainable and engaging experience for users. Most of the approaches are task-specific and not easily adaptable to existing models, indicating a gap that requires further research. Despite the advancements, there remains a significant gap in the literature, indicating a potential area for future research to develop more robust and comprehensive transparent frameworks for the music domain.

3 CONCLUSION

The advantages of recommender systems in improving our daily decisions and optimizing resource use are undeniable. They significantly contribute to saving time and money while enhancing user satisfaction across different domains. However, the ambiguity and uncertainty surrounding their suggestions pose a challenge to user trust and acceptance. By focusing on transparent explainability, we aim to bridge this gap, providing stakeholders with a clearer understanding of how recommendations are generated and fostering greater confidence in these systems. This approach will not only improve user experience but also ensure that the benefits of recommender systems are fully realized and appreciated.

In conclusion, while existing research underscores the significance of explainability in recommender systems, there remains a lack of comprehensive frameworks that effectively operationalize these concepts. Much of the current literature focuses on the theoretical aspects and benefits of explainability, without providing practical solutions that work seamlessly in real-world applications. This paper highlights the necessity of moving beyond conceptual discussions to develop and implement robust, transparent explainability methods. Specifically, our interest lies in enhancing the transparency of recommendations to various stakeholders involved with the system. It is crucial to interpret and communicate the reasoning behind recommendations in a manner that is clear and meaningful to users, developers, and other related parties.

ACKNOWLEDGMENTS

I would like to express my deepest gratitude to my supervisor, Peter Knees, for his invaluable guidance and support throughout this research endeavor. His expertise and guidance have been crucial in shaping this work.

This research was funded in whole or in part by the Austrian Science Fund (FWF) (10.55776/P33526).

REFERENCES

[1] Behnoush Abdollahi and Olfa Nasraoui. 2017. Using explainability for constrained matrix factorization. In *Proceedings of the eleventh ACM conference on recommender systems*. 79–83.

[2] Amina Adadi and Mohammed Berrada. 2018. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE access* 6 (2018), 52138–52160.

[3] Darius Afchar, Alessandro Melchiorre, Markus Schedl, Romain Hennequin, Elena Epure, and Manuel Moussallam. 2022. Explainability in music recommender systems. *AI Magazine* 43, 2 (2022), 190–208.

[4] Krisztian Balog, Filip Radlinski, and Shushan Arakelyan. 2019. Transparent, scrutable and explainable user models for personalized recommendation. In *Proceedings of the 42nd international acm sigir conference on research and development in information retrieval*. 265–274.

[5] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 1721–1730.

[6] Pablo Castells, Neil Hurley, and Saul Vargas. 2021. Novelty and diversity in recommender systems. In *Recommender systems handbook*. Springer, 603–646.

[7] Hirad Daneshvar and Reza Ravanmehr. 2022. A social hybrid recommendation system using LSTM and CNN. *Concurrency and Computation: Practice and Experience* 34, 18 (2022), e7015.

[8] Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* (2017).

[9] G Geetha, M Safa, C Fancy, and D Saranya. 2018. A hybrid approach using collaborative filtering and content based filtering for recommender system. In *Journal of Physics: Conference Series*, Vol. 1000. IOP Publishing, 012101.

[10] Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 80–89.

[11] David Gunning, Mark Stefik, Jaesik Choi, Timothy Miller, Simone Stumpf, and Guang-Zhong Yang. 2019. XAI—Explainable artificial intelligence. *Science robotics* 4, 37 (2019), eaay7120.

[12] Jonathan L Herlocker, Joseph A Konstan, and John Riedl. 2000. Explaining collaborative filtering recommendations. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*. 241–250.

[13] Saurav Joshi, Tanuj Jain, and Nidhi Nair. 2021. Emotion based music recommendation system using LSTM-CNN architecture. In *2021 12th International Conference on Computing Communication and Networking Technologies (ICCCNT)*. IEEE, 01–06.

[14] Iman Kamehkhosh and Dietmar Jannach. 2017. User perception of next-track music recommendations. In *Proceedings of the 25th conference on user modeling, adaptation and personalization*. 113–121.

[15] Gihwi Kim, Ilyoung Choi, Qinglong Li, and Jaekyeong Kim. 2021. A CNN-based advertisement recommendation through real-time user face recognition. *Applied Sciences* 11, 20 (2021), 9705.

[16] Pigi Kouki, James Schaffer, Jay Pujara, John O'Donovan, and Lise Getoor. 2019. Personalized explanations for hybrid recommender systems. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*. 379–390.

[17] Pigi Kouki, James Schaffer, Jay Pujara, John O'Donovan, and Lise Getoor. 2020. Generating and understanding personalized explanations in hybrid recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 10, 4 (2020), 1–40.

[18] Qika Lin, Yaoqiang Niu, Yifan Zhu, Hao Lu, Keith Zvikomborero Mushonga, and Zhendong Niu. 2018. Heterogeneous knowledge-based attentive neural networks for short-term music recommendations. *IEEE Access* 6 (2018), 58990–59000.

[19] Zachary C Lipton. 2018. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue* 16, 3 (2018), 31–57.

[20] Joy Lu, Dokyun Lee, Tae Wan Kim, and David Danks. 2019. Good explanation for algorithmic transparency. *Available at SSRN 3503603* (2019).

[21] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).

[22] Luca Marconi, R Matamoros Aragon, Francesco Epifania, et al. 2023. Discovering the Unknown Suggestion: a Short Review on Explainability for Recommender Systems. In *CEUR WORKSHOP PROCEEDINGS*, Vol. 3463. CEUR-WS, 1–15.

[23] Praveena Mathew, Bincy Kuriakose, and Vinayak Hegde. 2016. Book Recommendation System through content based and collaborative filtering method. In *2016 International conference on data mining and advanced computing (SAPIENCE)*. IEEE, 47–52.

[24] Alessandro B Melchiorre, Navid Rekabsaz, Christian Ganhör, and Markus Schedl. 2022. Protomf: Prototype-based matrix factorization for effective and explainable recommendations. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 246–256.

[25] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence* 267 (2019), 1–38.

[26] Sri Hari Nallamala, Usha Rani Bajjuri, Sarvani Anandarao, D Durga Prasad, and Pragnaban Mishra. 2020. A Brief Analysis of Collaborative and Content Based Filtering Algorithms used in Recommender Systems. In *IOP Conference Series: Materials Science and Engineering*, Vol. 981. IOP Publishing, 022008.

[27] Sergio Oramas, Oriol Nieto, Mohamed Sordo, and Xavier Serra. 2017. A deep multimodal approach for cold-start music recommendation. In *Proceedings of the 2nd workshop on deep learning for recommender systems*. 32–37.

[28] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. " Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.

[29] Noveen Sachdeva, Kartik Gupta, and Vikram Pudi. 2018. Attentive neural architecture incorporating song features for music recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 417–421.

[30] Ryosuke Saga and Yufeng Duan. 2018. Apparel goods recommender system based on image shape features extracted by a CNN. In *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2365–2369.

[31] Wojciech Samek, Thomas Wiegand, and Klaus-Robert Müller. 2017. Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *arXiv preprint arXiv:1708.08296* (2017).

[32] Amina Samih, Amina Adadi, and Mohammed Berrada. 2019. Towards a knowledge based explainable recommender systems. In *Proceedings of the 4th international conference on big data and internet of things*. 1–5.

[33] Markus Schedl, Peter Knees, Brian McFee, and Dmitry Bogdanov. 2022. *Music Recommendation Systems: Techniques, Use Cases, and Challenges*. Springer US, New York, NY, 927–971. https://doi.org/10.1007/978-1-0716-2197-4_24

[34] Timothée Schmude, Laura Koesten, Torsten Möller, and Sebastian Tschiatschek. 2023. On the Impact of Explanations on Understanding of Algorithmic Decision-Making. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 959–970.

[35] Nava Tintarev and Judith Masthoff. 2007. A survey of explanations in recommender systems. In *2007 IEEE 23rd international conference on data engineering workshop*. IEEE, 801–810.

[36] Nava Tintarev and Judith Masthoff. 2012. Evaluating the effectiveness of explanations for recommender systems: Methodological issues and empirical studies on the impact of personalization. *User Modeling and User-Adapted Interaction* 22 (2012), 399–439.

[37] Nava Tintarev and Judith Masthoff. 2015. Explaining recommendations: Design and evaluation. In *Recommender systems handbook*. Springer, 353–382.

[38] Chun-Hua Tsai and Peter Brusilovsky. 2021. The effects of controllability and explainability in a social recommender system. *User Modeling and User-Adapted Interaction* 31 (2021), 591–627.

[39] Aaron Van den Oord, Sander Dieleman, and Benjamin Schrauwen. 2013. Deep content-based music recommendation. *Advances in neural information processing systems* 26 (2013).

[40] Jesse Vig, Shilad Sen, and John Riedl. 2009. Tagsplanations: explaining recommendations using tags. In *Proceedings of the 14th international conference on Intelligent user interfaces*. 47–56.

[41] Alexandra Vultureanu-Albişi and Costin Bădică. 2021. Recommender systems: An explainable AI perspective. In *2021 International conference on innovations in intelligent systems and applications (INISTA)*. IEEE, 1–6.

[42] Alexandra Vultureanu-Albişi and Costin Bădică. 2022. A survey on effects of adding explanations to recommender systems. *Concurrency and Computation: Practice and Experience* 34, 20 (2022), e6834.

[43] Haili Wang, Nana Lou, and Zhenlin Chao. 2020. A personalized movie recommendation system based on LSTM-CNN. In *2020 2nd International Conference on Machine Learning, Big Data and Business Intelligence (MLBDBI)*. IEEE, 485–490.

[44] Nan Wang, Hongning Wang, Yiling Jia, and Yue Yin. 2018. Explainable recommendation via multi-task learning in opinionated text data. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 165–174.

[45] Xiang Wang, Dingxian Wang, Canran Xu, Xiangnan He, Yixin Cao, and Tat-Seng Chua. 2019. Explainable reasoning over knowledge graphs for recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 5329–5336.

[46] Libing Wu, Cong Quan, Chenliang Li, Qian Wang, Bolong Zheng, and Xiangyang Luo. 2019. A context-aware user-item representation learning for item recommendation. *ACM Transactions on Information Systems (TOIS)* 37, 2 (2019), 1–29.

[47] Lijie Xie, Zhaoming Hu, Xingjuan Cai, Wensheng Zhang, and Jinjun Chen. 2021. Explainable recommendation based on knowledge graph and multi-objective optimization. *Complex & Intelligent Systems* 7 (2021), 1241–1252.

[48] Dan Yang, Jing Zhang, Sifeng Wang, and XueDong Zhang. 2019. A time-aware CNN-based personalized recommender system. *Complexity* 2019 (2019), 1–11.

[49] Wenli Yang, Yuchen Wei, Hanyu Wei, Yanyu Chen, Guan Huang, Xiang Li, Renjie Li, Naimeng Yao, Xinyi Wang, Xiaotong Gu, et al. 2023. Survey on explainable AI: From approaches, limitations and Applications aspects. *Human-Centric Intelligent Systems* 3, 3 (2023), 161–188.

[50] Xinyang Zhang, Ningfei Wang, Hua Shen, Shouling Ji, Xiapu Luo, and Ting Wang. 2020. Interpretable deep learning under fire. In *29th {USENIX} security symposium ({USENIX} security 20)*.
[51] Yongfeng Zhang, Xu Chen, et al. 2020. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval* 14, 1 (2020), 1–101.
[52] Guoshuai Zhao, Hao Fu, Ruihua Song, Tetsuya Sakai, Zhongxia Chen, Xing Xie, and Xueming Qian. 2019. Personalized reason generation for explainable song recommendation. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10, 4 (2019), 1–21.
[53] Guoshuai Zhao, Hao Fu, Ruihua Song, Tetsuya Sakai, Xing Xie, and Xueming Qian. 2018. Why you should listen to this song: Reason generation for explainable recommendation. In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*. IEEE, 1316–1322.