# An exploration of recommender systems explanation paradigms: generating and evaluating syntactic, semantic, and generative models with knowledge graphs

## An Extended Abstract

André Levi Zanon
andre.zanon@insight-centre.org
Insight Centre for Data Analytics, School of Computer Science and IT, University College Cork
Cork, Cork, Ireland

Leonardo C. Dutra da Rocha
lcrocha@ufsj.edu.br
Departamento de Ciência da Computação, Universidade Federal de São João del-Rei
São João del Rei, Minas Gerais, Brasil

Marcelo Garcia Manzato
mmanzato@icmc.usp.br
Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo
São Carlos, São Paulo, Brasil

## ABSTRACT

Recommender Systems (RS) generate suggestions for users based on their past interactions and the interactions of other users. However, because RS rely on identifying similarities between users, they may expose users to a limited range of content. In addition, as algorithms have become increasingly complex, users often cannot understand why a recommendation is generated. In this thesis, we explored three approaches—syntactic, semantic, and generative—to generate explanations for recommendations and investigate their potential impact on the diversity of recommendations, with the goal of broadening user interests. Our results show that explanations can play an important role in increasing item diversity, thereby helping mitigate filter bubble effects. We also observed an improvement in the explanation quality across the approaches, from syntactic to semantic and finally to generative models.

## KEYWORDS

Recommender Systems, Explanation, Explanation Evaluation

## 1 INTRODUCTION

Currently, we interact with Recommender Systems (RS) ubiquitously. Entertainment streaming platforms, e-commerce, and social networks use these systems to provide item suggestions for users. To recommend items to a target user, RS find similarities with other users who share similar tastes.

This is why RS create "filter bubbles," where suggested content always aligns with users' tastes, depriving them of diverse debates and opinions. With the increasing complexity of the architecture of these algorithms, RS have also become black boxes for users, meaning that it is not possible to understand why a particular suggestion was generated. Explanations to RS can be created in two ways: intrinsic to the model (alongside) or agnostic (dissociated) from the RS.

One way to generate such explanations is using Knowledge Graphs (KGs). KGs model items and their attributes using a graph where nodes represent users, items, and their attributes, and edges represent relationships between them. For example, the film "Titanic is of the drama genre." This relationship is defined in a KG with two nodes, one representing the film "Titanic" and the other representing the genre "drama," and the edge connecting these nodes represents the relationship "is of the genre."

By expanding this across multiple items and attributes, explanation algorithms for RS extract a set of paths that connect the nodes of items that the user has interacted with to a recommended item node based on shared attributes. Continuing the example, if a user interacted with "Titanic," and the RS recommended "The Godfather", that is also a drama, an explanation path is found between the two items.

In: VII Concurso de Teses e Dissertações (CTD 2025). Anais Estendidos do XXXI Simpósio Brasileiro de Sistemas Multimídia e Web (CTD'2025). Rio de Janeiro/RJ, Brasil. Porto Alegre: Brazilian Computer Society, 2025.

ISSN 2596-1683

## 2 OBJECTIVE AND METHODOLOGY

Based on these paths between interacted and recommended item nodes, the main objective of an agnostic model is to find the most relevant path to show the user as an explanation in order to increase the chance of interaction with the user and increasing aspects such as transparency, persuasion, engagement, efficiency and trust, defined in the literature as explanation goals [2].

In this thesis we explored three main approaches to find the most relevant paths between interacted and item nodes: syntactic, semantic, and generative. Syntactic approaches focus on the number of connections between item nodes; semantic approaches integrate KG structures into a latent space; and generative approaches use language models to compose explanations.

The overall methodology of this study is illustrated in Figure 1. Initially, based on a dataset, we executed a RS to generate recommendations for users. As our process is model-agnostic, we find connections on the KG between nodes of items the user has interacted with and nodes of items recommended by the RS algorithm based on shared attribute nodes using the three paradigms.

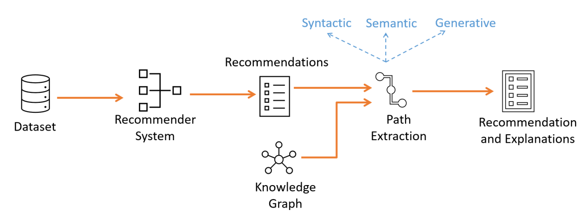


**Figure 1: Project methodology.**

We designed and evaluated explanations generated by each paradigm using online and offline metrics to evaluate them and find correlations between user perception and the selection of explanation paths. Our results indicate that syntactic, semantic, and generative models represent an evolution in offline metrics of explanations and that the diversity and popularity of attributes in the paths affect user perception. Additionally, we found that explanations can be important for increasing item diversity, reducing filter bubbles in RS.

## 3 EXPERIMENTS AND CONTRIBUTIONS

We initially explored syntactic approaches using ontologies and paths between aspects, defined as keywords extracted from users' reviews that characterize an item. Based on these aspects, we built a content-based RS that measured similarity through the ontology. Unlike cosine similarity with simple binary encoding, where each vector position corresponds to the presence of an aspect from the set of all extracted aspects, the ontology-based similarity reduces matrix sparsity and yields good recommendation results while enabling explanation generation.

This work led to two publications: one reporting offline experiments on the algorithm's accuracy [6] and another presenting online experiments with participants to validate its explanations [7]. From the online experiments, we observed that users especially valued explanations when the recommended item was surprising, often due to the similarity between the two aspects. Consequently, providing explanations for surprising recommendations emerges as key to user acceptance.

We then explored KGs to generate explanations based on the ranking of paths between interacted and recommended items using an adaptation of TF-IDF to graphs. TF-IDF is traditionally used to compute the relevance of words in a document; in our scenario, we adapted the formula to graphs, where terms are attributes, and items are documents. The main objective of this research was to analyze the effects of reordering recommendations based on items with more relevant explanations according to the TF-IDF score.

Specifically, we reordered the initial ranking of recommendations generated by the RS according to the scores of all paths between the set of recommended and interacted items. We then compared the accuracy and diversity of the original RS ranking with our reordering method. The results showed that our approach preserved part of the original ranking while increasing the diversity of recommendations. This work also resulted in a publication [4].

We also explored semantic and generative approaches to this problem. In the semantic approach, we designed an algorithm to select an explanation for a recommended item based on the path with the highest similarity to the users. To compute this similarity, we generated semantic representations of both users and items using KG embeddings. After training the KG embedding model, a user embedding was defined as the sum of the embeddings of the items with which the user interacted. In contrast, a path embedding is defined as the sum of the embeddings of all elements (items, attributes, and edges) in the path connecting an interacted item node with the recommended item. Our results showed that, according to three offline explanation metrics defined in the literature, the semantic approach achieved competitive results compared to baseline syntactic approaches.

These offline explanation metrics state that, considering a set of explanations represented as paths in a KG between interacted and recommended items, explanations should connect recently interacted items with recommended items, ensure diversity of attributes across recommendations, and highlight attributes that are popular (i.e., highly connected to many items in the KG) [1]. This line of work resulted in two publications [5] and [3].

For the generative approach, we employed different LLMs to select the best path between interacted and recommended items with respect to the same three offline explanation metrics. Our results showed that LLMs can achieve an optimal balance across these metrics, outperforming both syntactic and semantic approaches. Finally, we conducted a literature review of offline explanation metrics and an online user study to investigate the relationship between the three offline explanation metrics we adopted and the explanation goals. We found that transparency and trust goals are enhanced by showing popular attributes, while engagement and persuasiveness are more strongly linked to attribute diversification in explanations.

## 4 CONCLUSIONS

Our findings suggest a progressive improvement in offline explanation metrics when moving from syntactic to semantic, and finally, to generative models. Moreover, we observed that providing explanations can increase the diversity of recommended items, thereby helping mitigate filter bubbles. For future work, we aim to further investigate the potential of LLMs in generating explanations for RS and to propose new metrics for the offline evaluation of explanations in RS.

## ACKNOWLEDGMENTS

The authors acknowledge CAPES, CNPq, Fapesp (2022/07016-9), AWS and Fapemig for their funding and support of this research.

## REFERENCES


[1] Giacomo Balloccu, Ludovico Boratto, Gianni Fenu, and Mirko Marras. 2022. Post Processing Recommender Systems with Knowledge Graphs for Recency, Popularity, and Diversity of Explanations. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Madrid, Spain) *(SIGIR '22)*. Association for Computing Machinery, New York, NY, USA, 646–656. https://doi.org/10.1145/3477495.3532041

[2] Nava Tintarev and Judith Masthoff. 2007. A Survey of Explanations in Recommender Systems. https://doi.org/10.1109/icdew.2007.4401070

[3] André Zanon, Leonardo Rocha, and Marcelo Manzato. 2024. O Impacto de Estratégias de Embeddings de Grafos na Explicabilidade de Sistemas de Recomendação. In *Proceedings of the 30th Brazilian Symposium on Multimedia and the Web* (Juiz de Fora/MG). SBC, Porto Alegre, RS, Brasil, 231–239. https://doi.org/10.5753/webmedia.2024.241857

[4] André Levi Zanon, Leonardo Chaves Dutra da Rocha, and Marcelo Garcia Manzato. 2022. Balancing the trade-off between accuracy and diversity in recommender systems with personalized explanations based on Linked Open Data. *Knowledge-Based Systems* 252 (2022), 109333. https://doi.org/10.1016/j.knosys.2022.109333

[5] André Levi Zanon, Leonardo Chaves Dutra Da Rocha, and Marcelo Garcia Manzato. 2024. *Model-Agnostic Knowledge Graph Embedding Explanations for Recommender Systems*. Communications in Computer and Information Science, 3–27. https://doi.org/10.1007/978-3-031-63797-1_1

[6] André L. Zanon, Luan Souza, Diany Pressato, and Marcelo G. Manzato. 2020. WordRecommender: An Explainable Content-Based Algorithm based on Sentiment Analysis and Semantic Similarity. In *Proceedings of the Brazilian Symposium on Multimedia and the Web* (São Luís, Brazil) *(WebMedia '20)*. Association for Computing Machinery, New York, NY, USA, 181–184. https://doi.org/10.1145/3428658.3431093

[7] André Levi Zanon, Luan Souza, Diany Pressato, and Marcelo Garcia Manzato. 2022. A user study with aspect-based sentiment analysis for similarity of items in content-based recommendations. *Expert Systems* 39, 8 (2022). https://doi.org/10.1111/exsy.12991