

Why-Not Explainable Graph Recommender

Hervé-Madelein Attolou¹, Katerina Tzompanaki¹, Kostas Stefanidis², and Dimitris Kotzinos¹

¹ETIS, CY Cergy Paris University, ENSEA, CNRS, UMR8051

²Tampere University



Acknowledgments

This work is partially supported by the National French Agency, under the EXPIDA - ANR-22-CE23-0017 project.



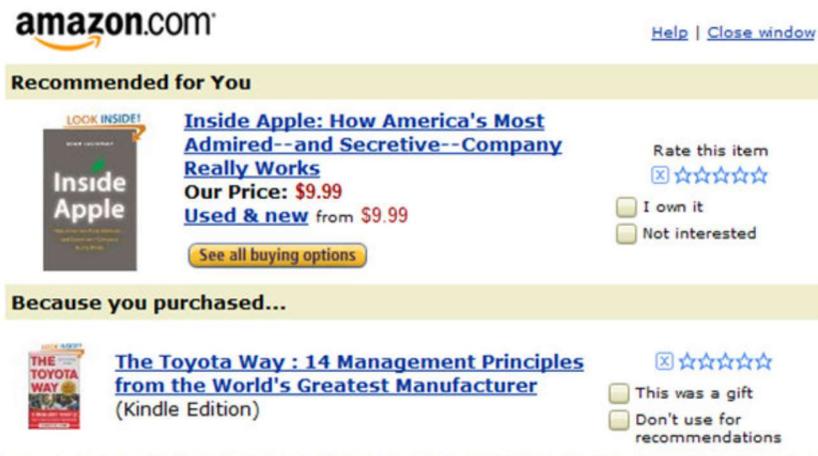
40th IEEE International Conference on
Data Engineering
Utrecht, Netherlands | 13th - 17th May

2024

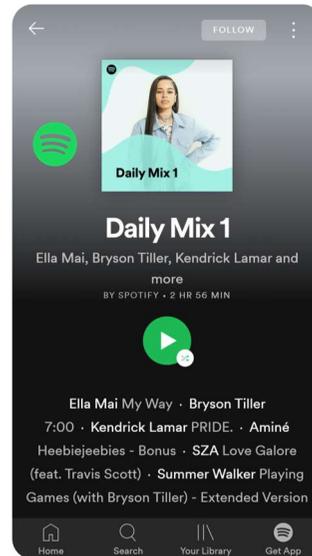


Recommendation Systems

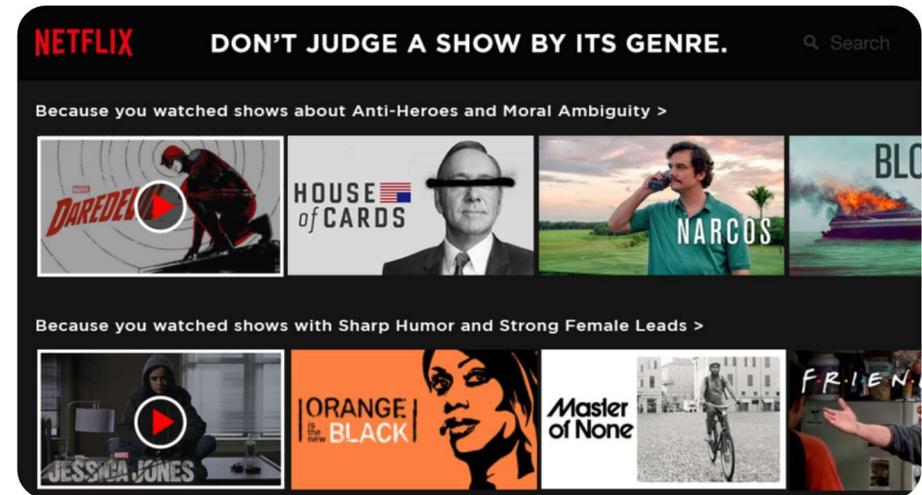
Recommendation systems are personalizing our web experience.



What to buy



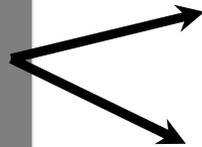
What to listen



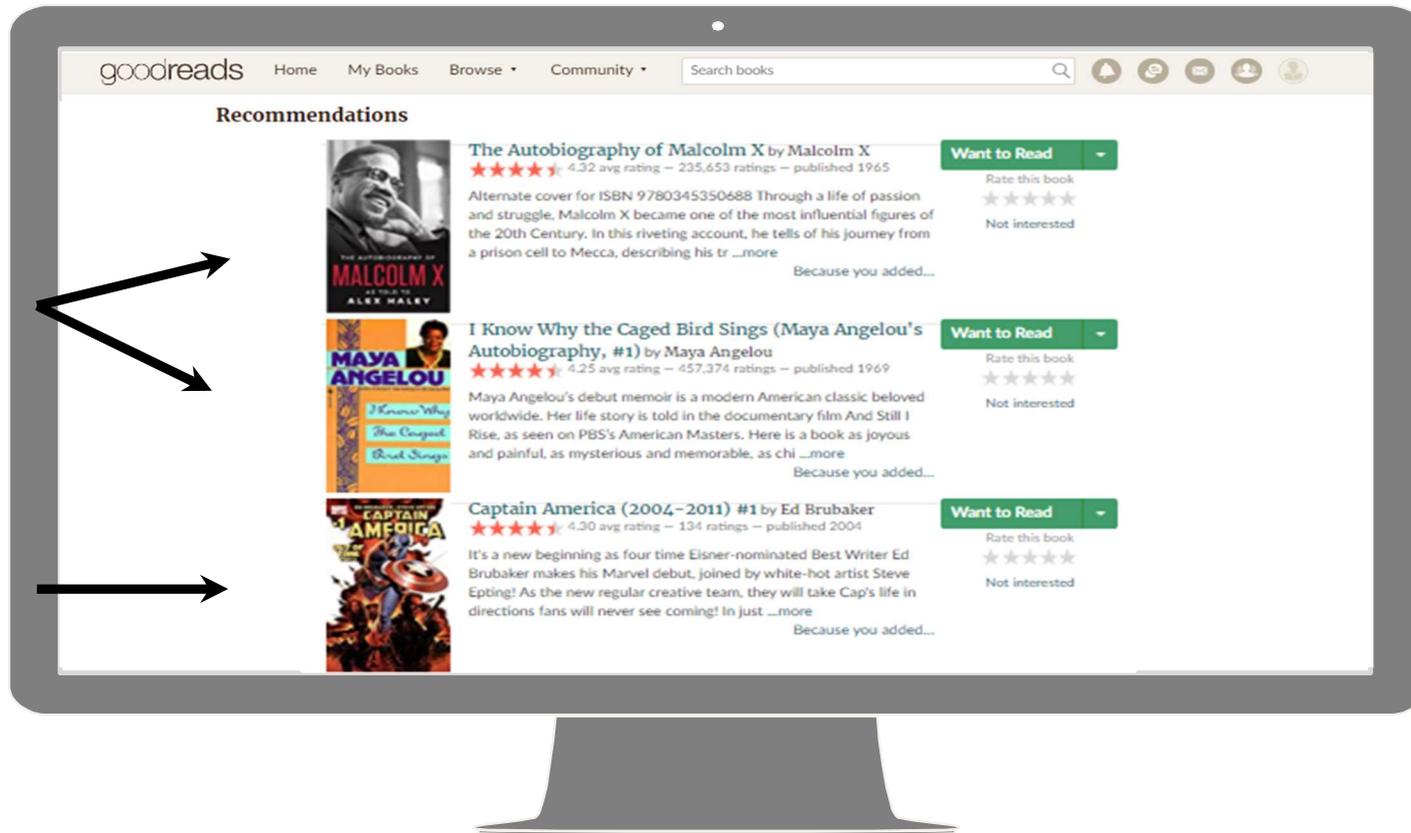
What to watch

Recommendation Systems

Biography

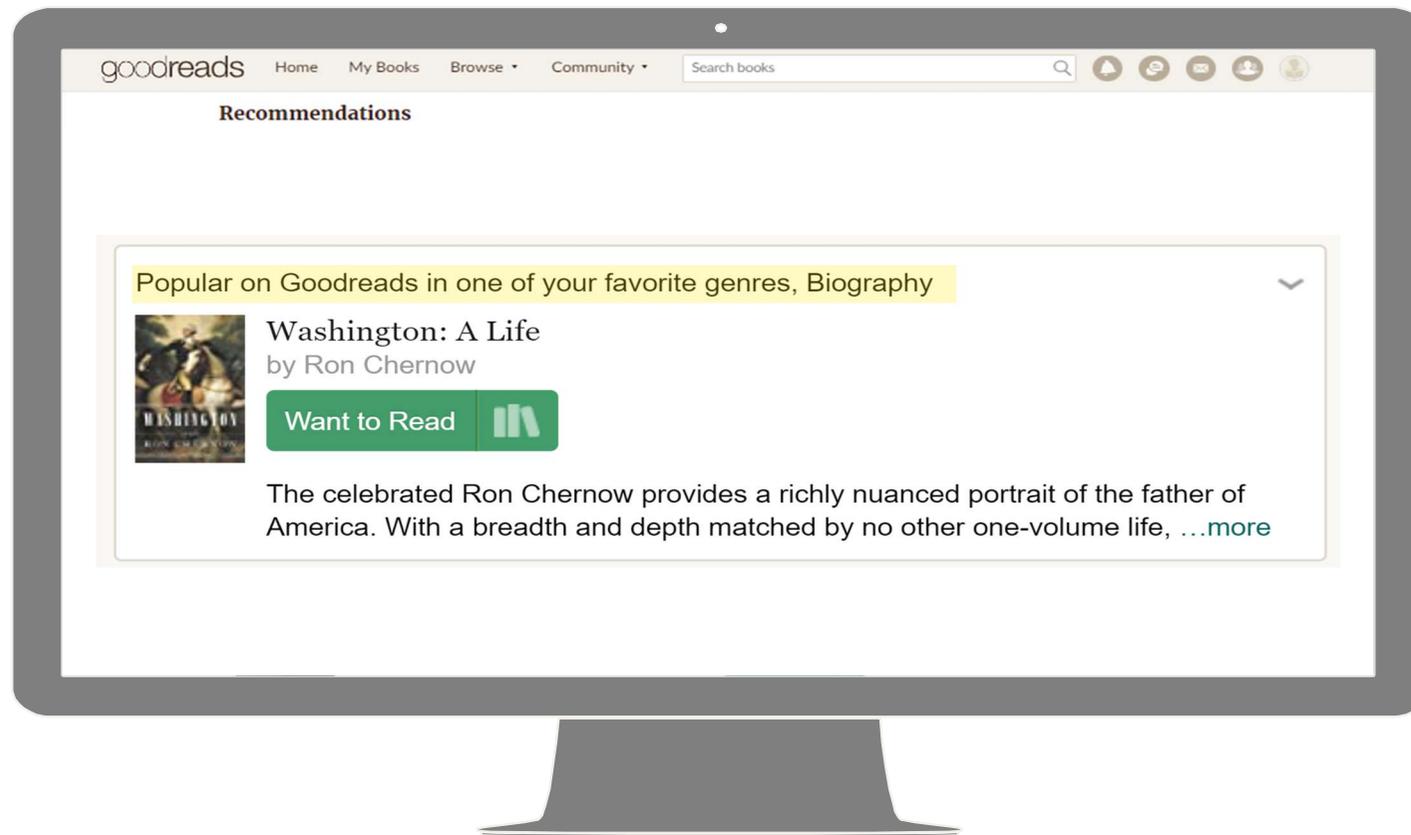


Comics



Book recommendations from Goodreads when logging in with an Amazon account.

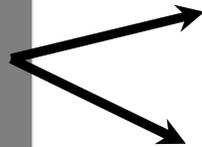
Explainable Recommendation Systems



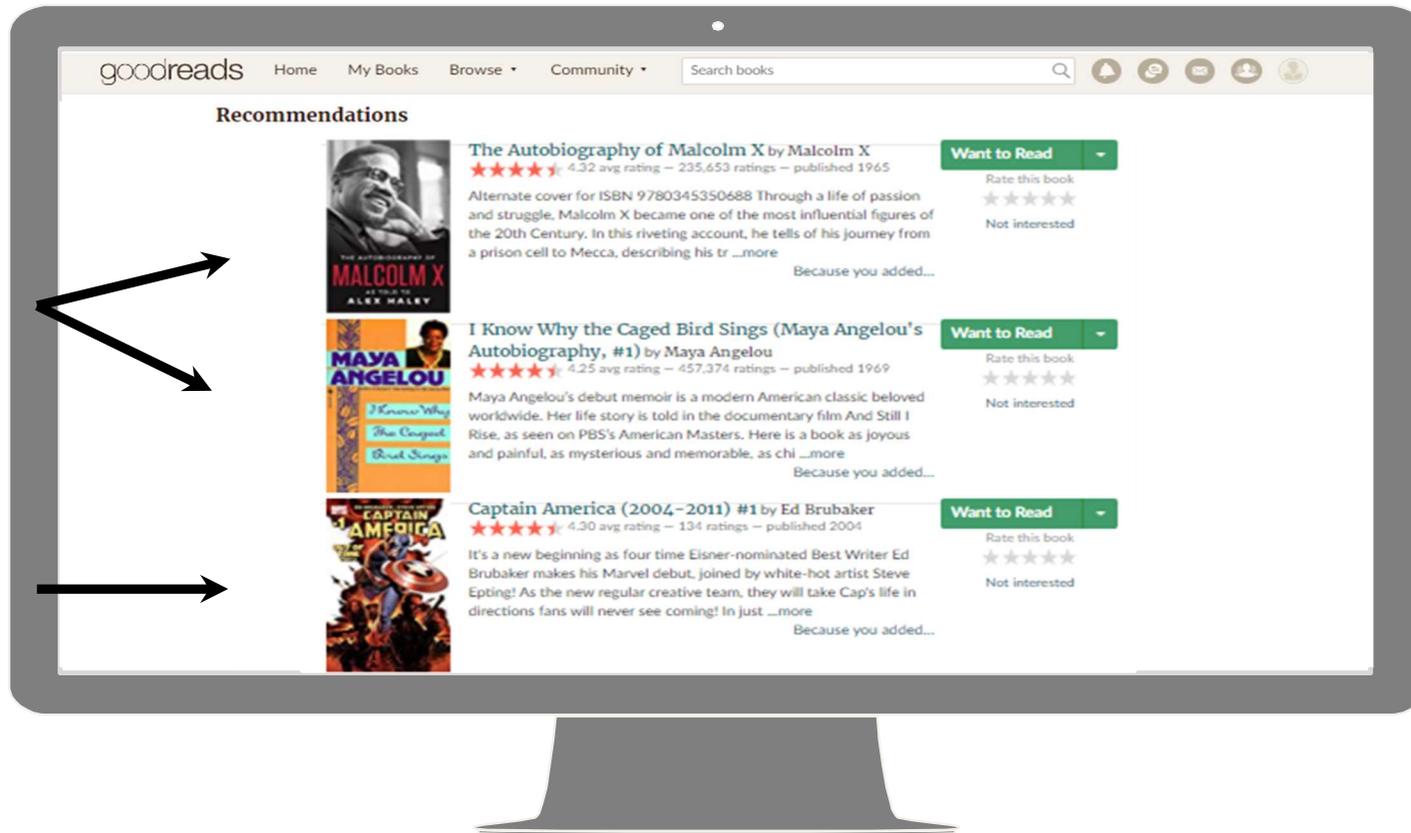
Example **explanation** for a recommendation.

Explainable Recommendation Systems

Biography



Comics



Why-Not Question:

Why don't I get recommended any Computer Science books **?⁵**

Why-Not Explanations

Why-Not Question:

- **Why** an item or category of items is **not** in the recommendation list ?
- **Why** an item or category of items is **not** ranked in a higher position ?

Why-Not Explanations:

- **Explain the absence** of certain options from a user's recommendation list.
- **Provide more information** to the system designer.
- Promote decision **transparency** and strengthen fairness.
- **Detect discriminative** systems and biases in the original data.

Related & Background Work



Related work

Recommenders

- Collaborative Filtering [12]
- Content-Based Filtering [8]
- Context-Based Recommenders [17]
- ...

Explainable Recommenders

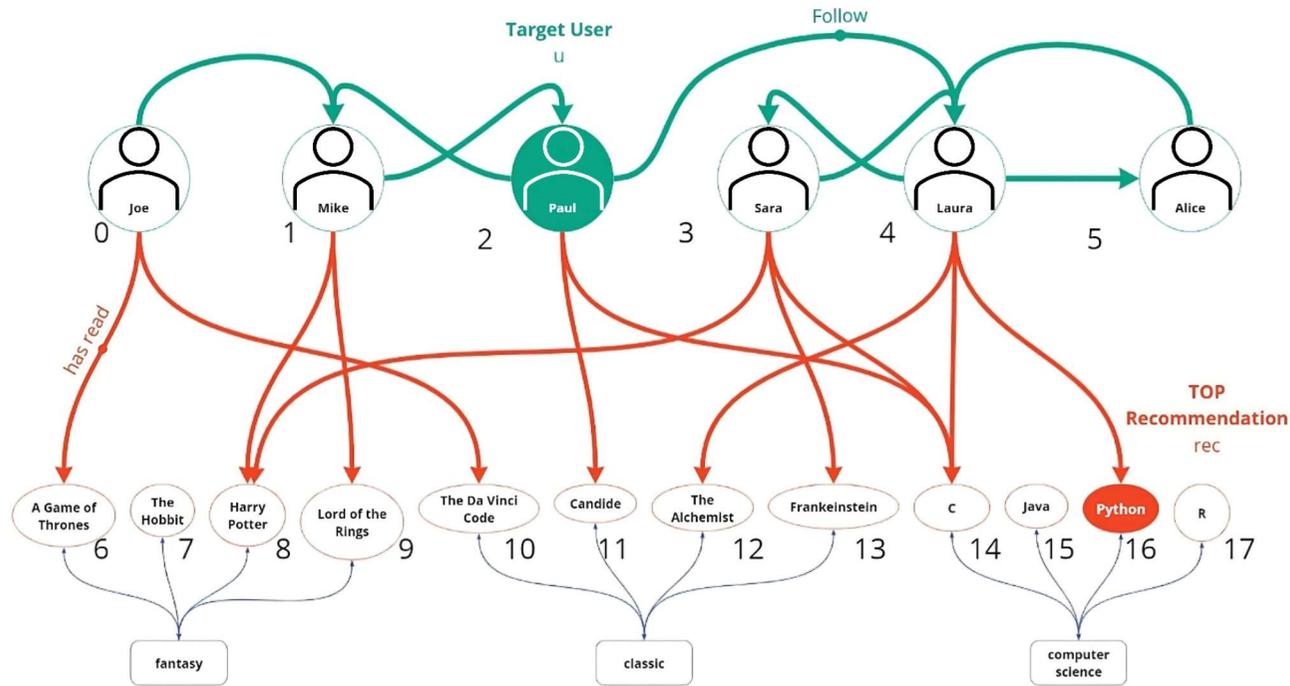
- Explainable Matrix Factorization [1]
- Rule Mining [21]
- Graph Based [5]
- Black-Box Explanations [20]
- Counterfactual Explanation for ML
- ...

Why-Not Questions

- Why-Not Questions In Databases [6, 14],
- Top-K Spatial Queries [7]
- Workflow Analysis [4]
- ...

- Why-Not Questions In Collaborative Filtering [13]
- **Why-Not Explainable Recommenders**

Background: Heterogeneous graph



An **heterogeneous graph** $G = (V, E, \theta)$ consists of a set of nodes V , a set of edges $E \subseteq V \times V$, and a mapping θ from each node and each edge to their types, such that $\theta_V : V \rightarrow T_V$ and $\theta_E : E \rightarrow T_E$ with $|T_V| + |T_E| > 2$.

In this work a heterogeneous graph generally contains at least two node types:

- **users** $U \in T_V$
- **items** $I \in T_V$

Background: Recommendation Algorithm: Personalized PageRank (PPR)

$$PPR(u, \cdot) = \alpha \sum_{l=0}^{\infty} (1 - \alpha)^l e_u W^l$$

The distribution of a random walk in G , starting at u :

- with probability α : teleports to a set of seed nodes $\{s\}$
- with probability $1 - \alpha$: continues the walk to a randomly chosen outgoing edge.

$$\text{rec} = \arg \max_{i \in I \setminus N_{\text{out}}(u)} PPR(u, i)$$

The item i with the highest $PPR(u, i)$ will be recommended.

α → teleportation probability

s → single seed

e_s → the one-hot vector

W → the transition matrix

Problem
Definition
&
Solution

Why-Not Explanations for Graph Recommender Systems

Problem Definition

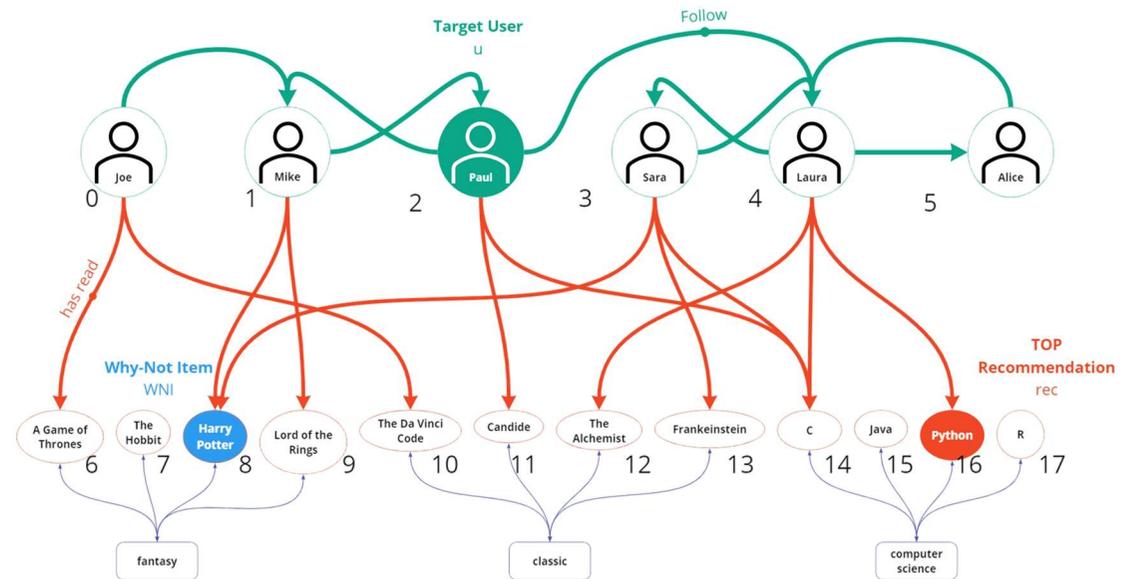
Given:

- a graph $G = (V, E, \theta)$,
- a user u ,
- a Why-Not Item **WNI**,
- and an initial recommendation *rec*

a **Why-Not explanation** is the set of edges $A^* \subseteq A, A \in \{A^+, A^-\}$, with:

- $A^+ = \{a^+ \mid a^+ = (u, i) \notin E, i \in I\}$, and
- $A^- = \{a^- \mid a^- = (u, i) \in E, i \in I\}$

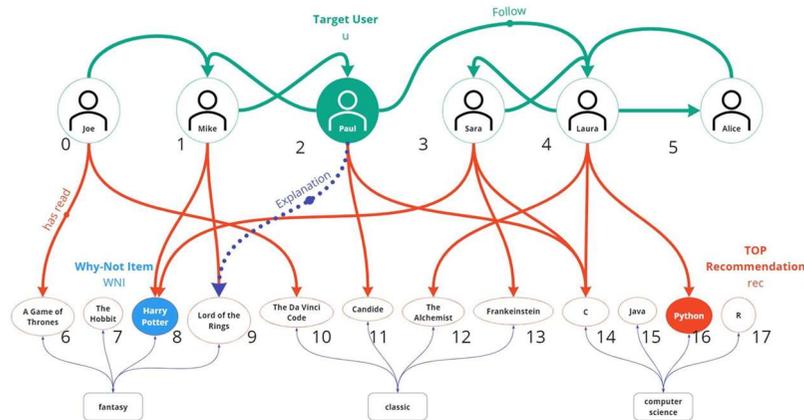
such that $G' = (V, E', \theta)$, with $E' = E \cup A^+ \setminus A^-$ generates **WNI** as the top-1 recommendation.



Example of Why-Not explanations

Why Harry Potter is *not* my recommendation?

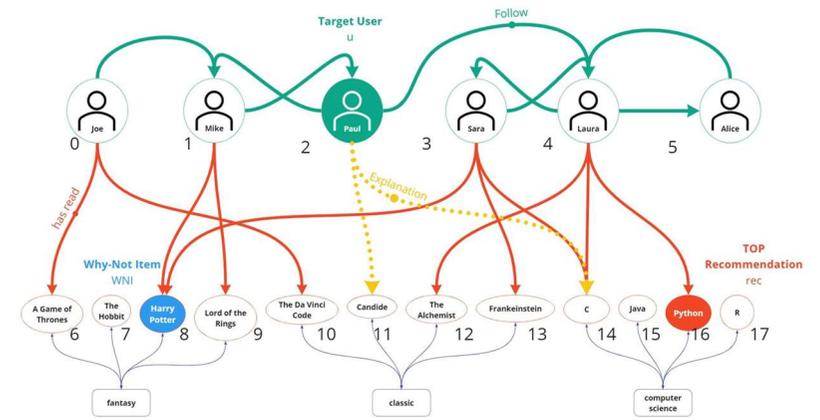
Add Mode
Suggests **NEW ACTIONS**



$$A^+ = \{(2,9)\}$$

“If Paul reads *Lord of the Rings*, the recommendation will be *Harry Potter*”

Remove Mode
Pinpoints **PAST ACTIVITY**



$$A^- = \{(2,11), (2,14)\}$$

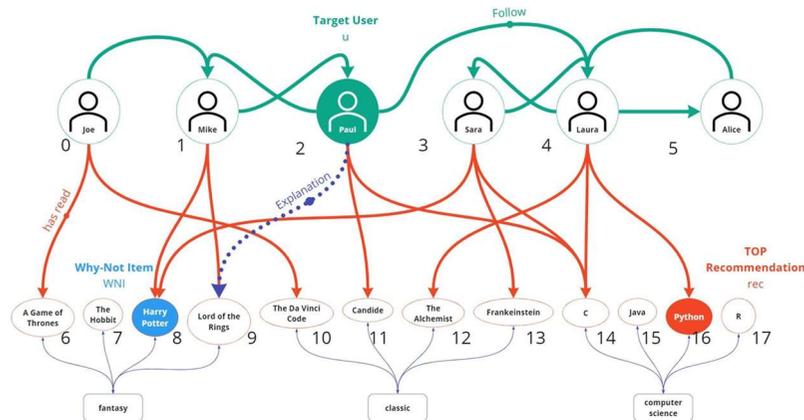
“If Paul had not read *Candide* and *C* the recommendation, would have been *Harry Potter*”

Example of Why-Not explanations

Why Harry Potter is not my recommendation?

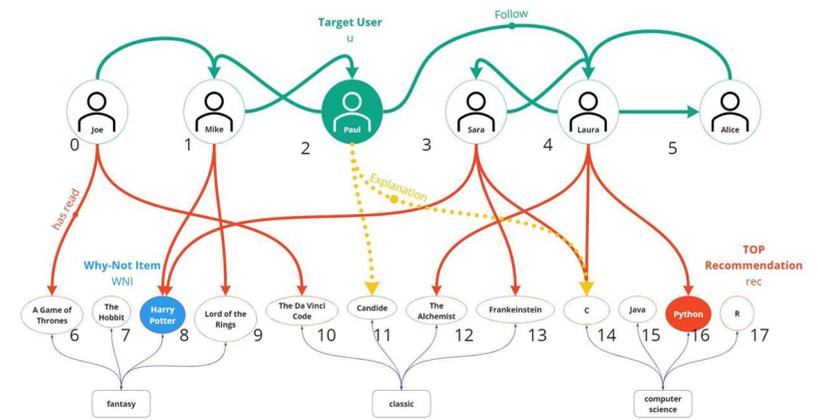
Add Mode

Suggests **NEW ACTIONS**



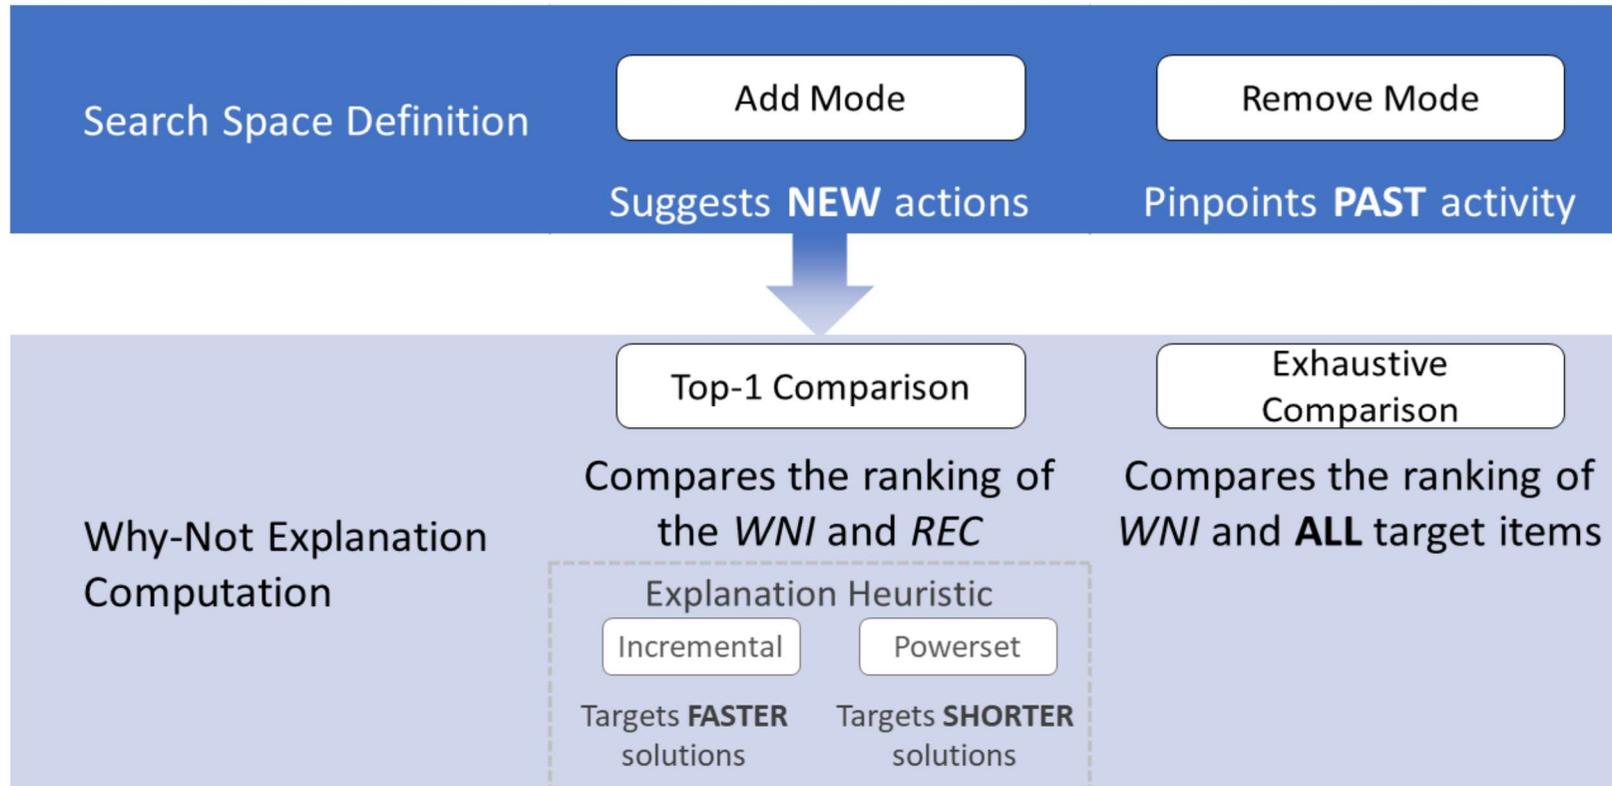
Remove Mode

Pinpoints **PAST ACTIVITY**



1. List the **existing/potential** neighbors by their contribution to the current recommendation compared to Why-Not Item.
2. **Disconnect/Connect** a set of neighbors so that the recommendation changes to the Why-Not Item.

EMiGRe: a Framework for Explaining Missing Graph Recommendations



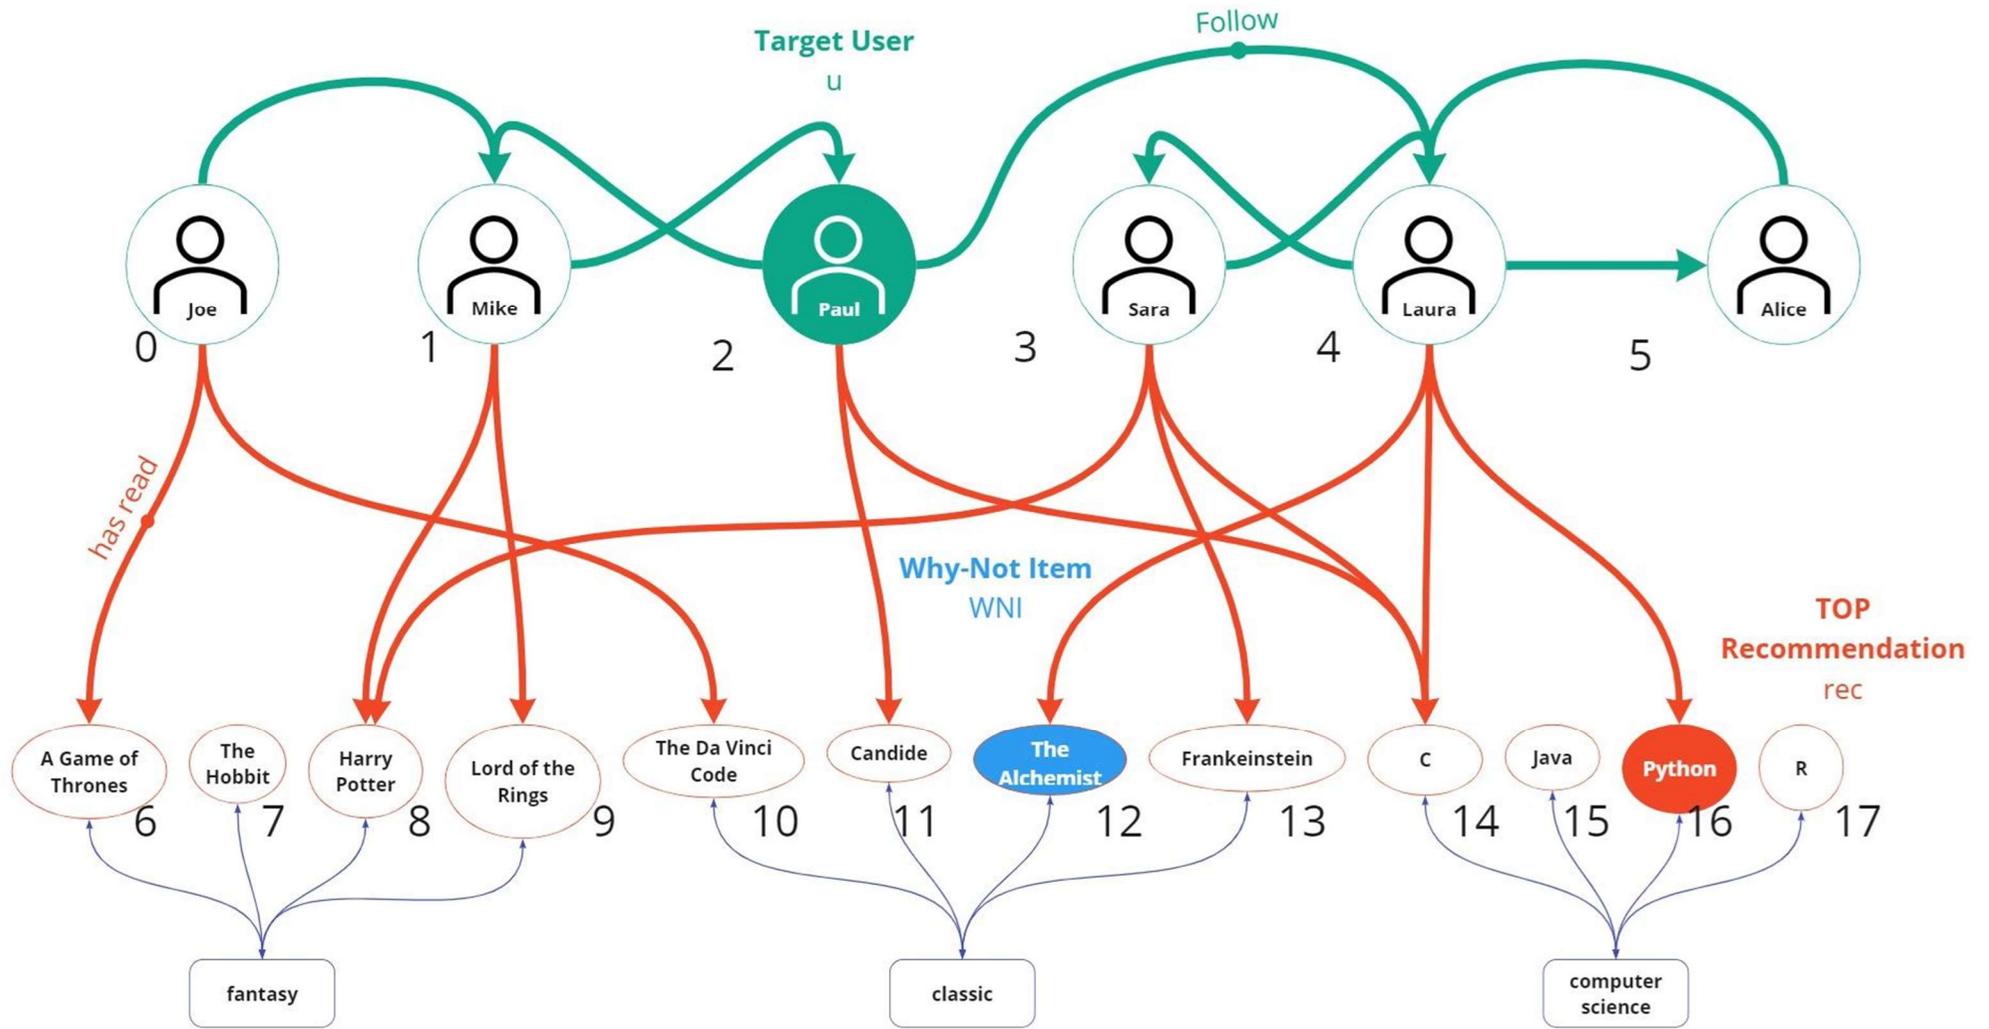
$$contribution_{add}(n_i) = PPR(n_i, WNI | A) - PPR(n_i, rec | A)$$

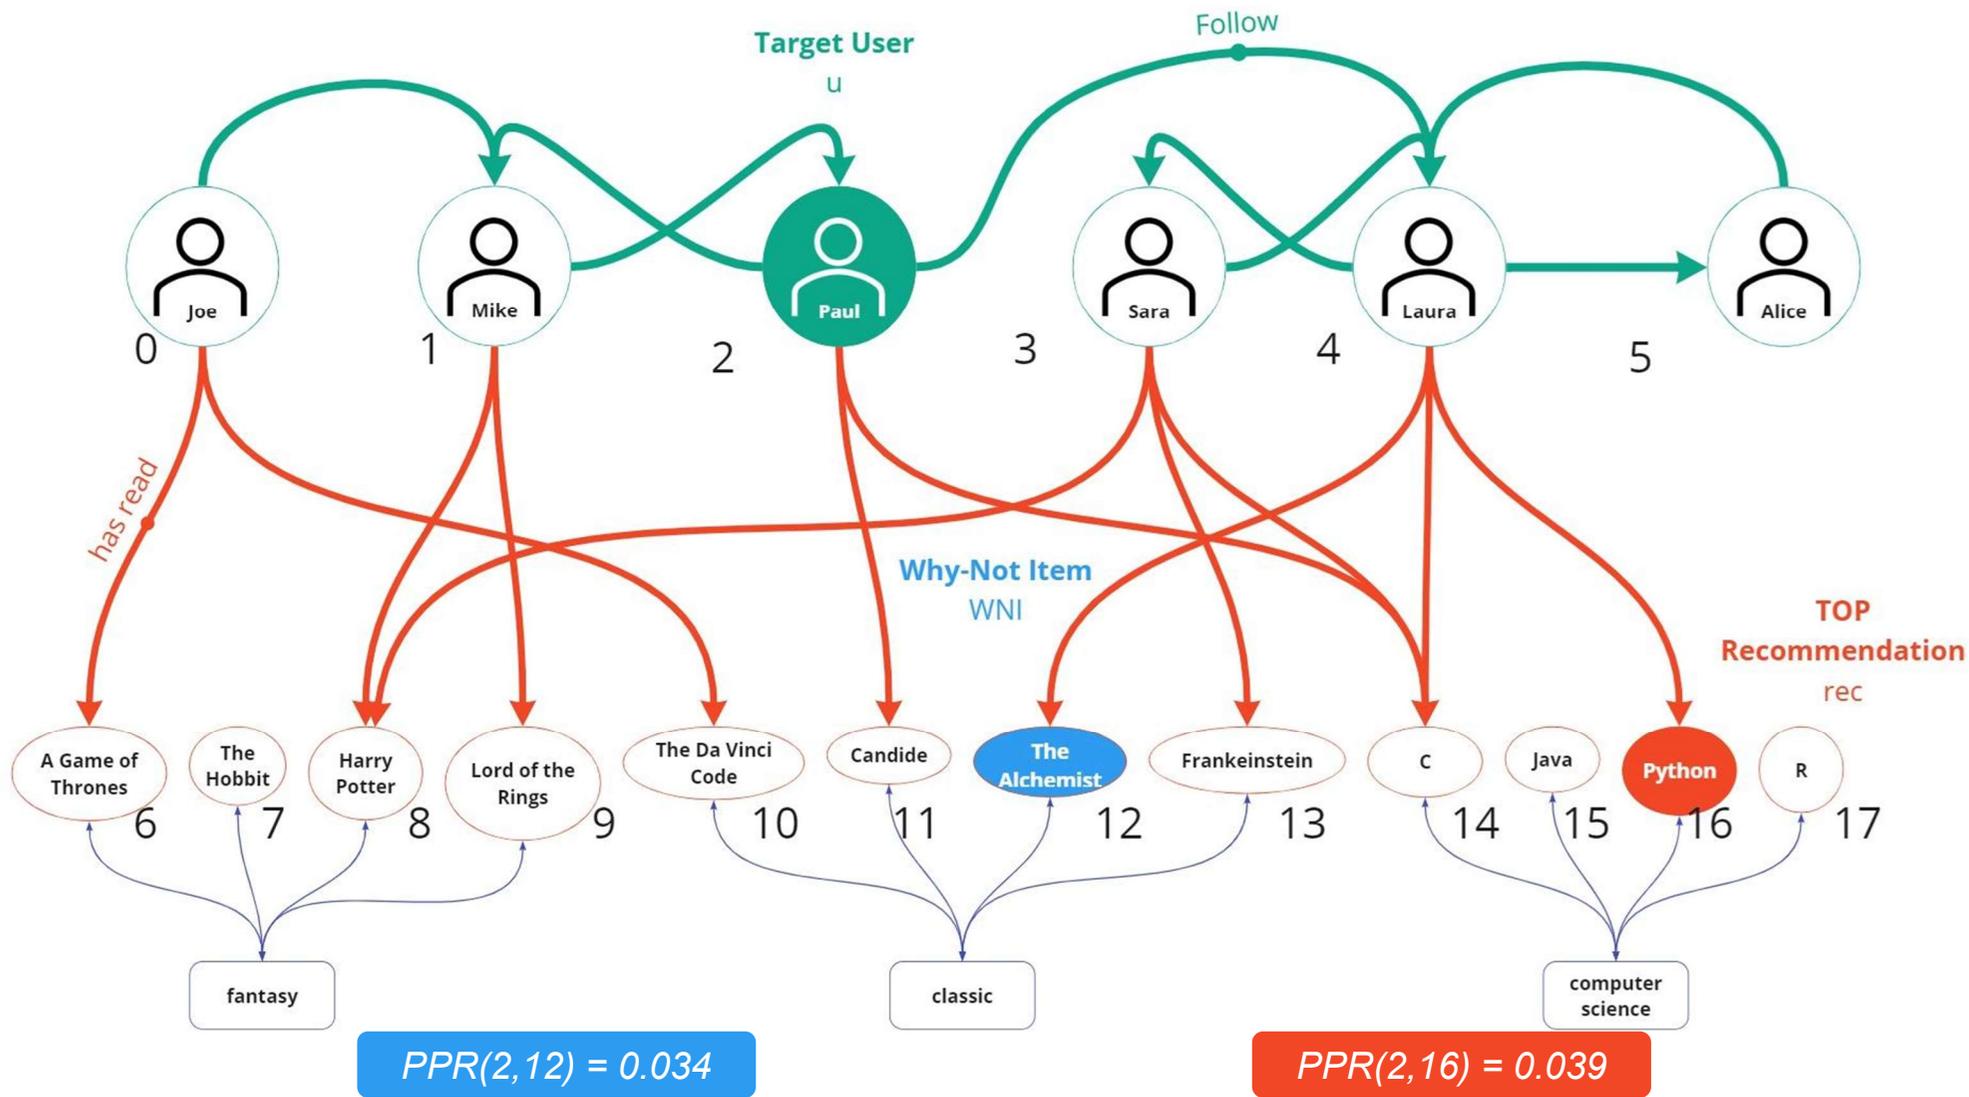
$$contribution_{rmv}(n_i) = W(u, n_i) \cdot (PPR(n_i, rec | A) - PPR(n_i, WNI | A))$$

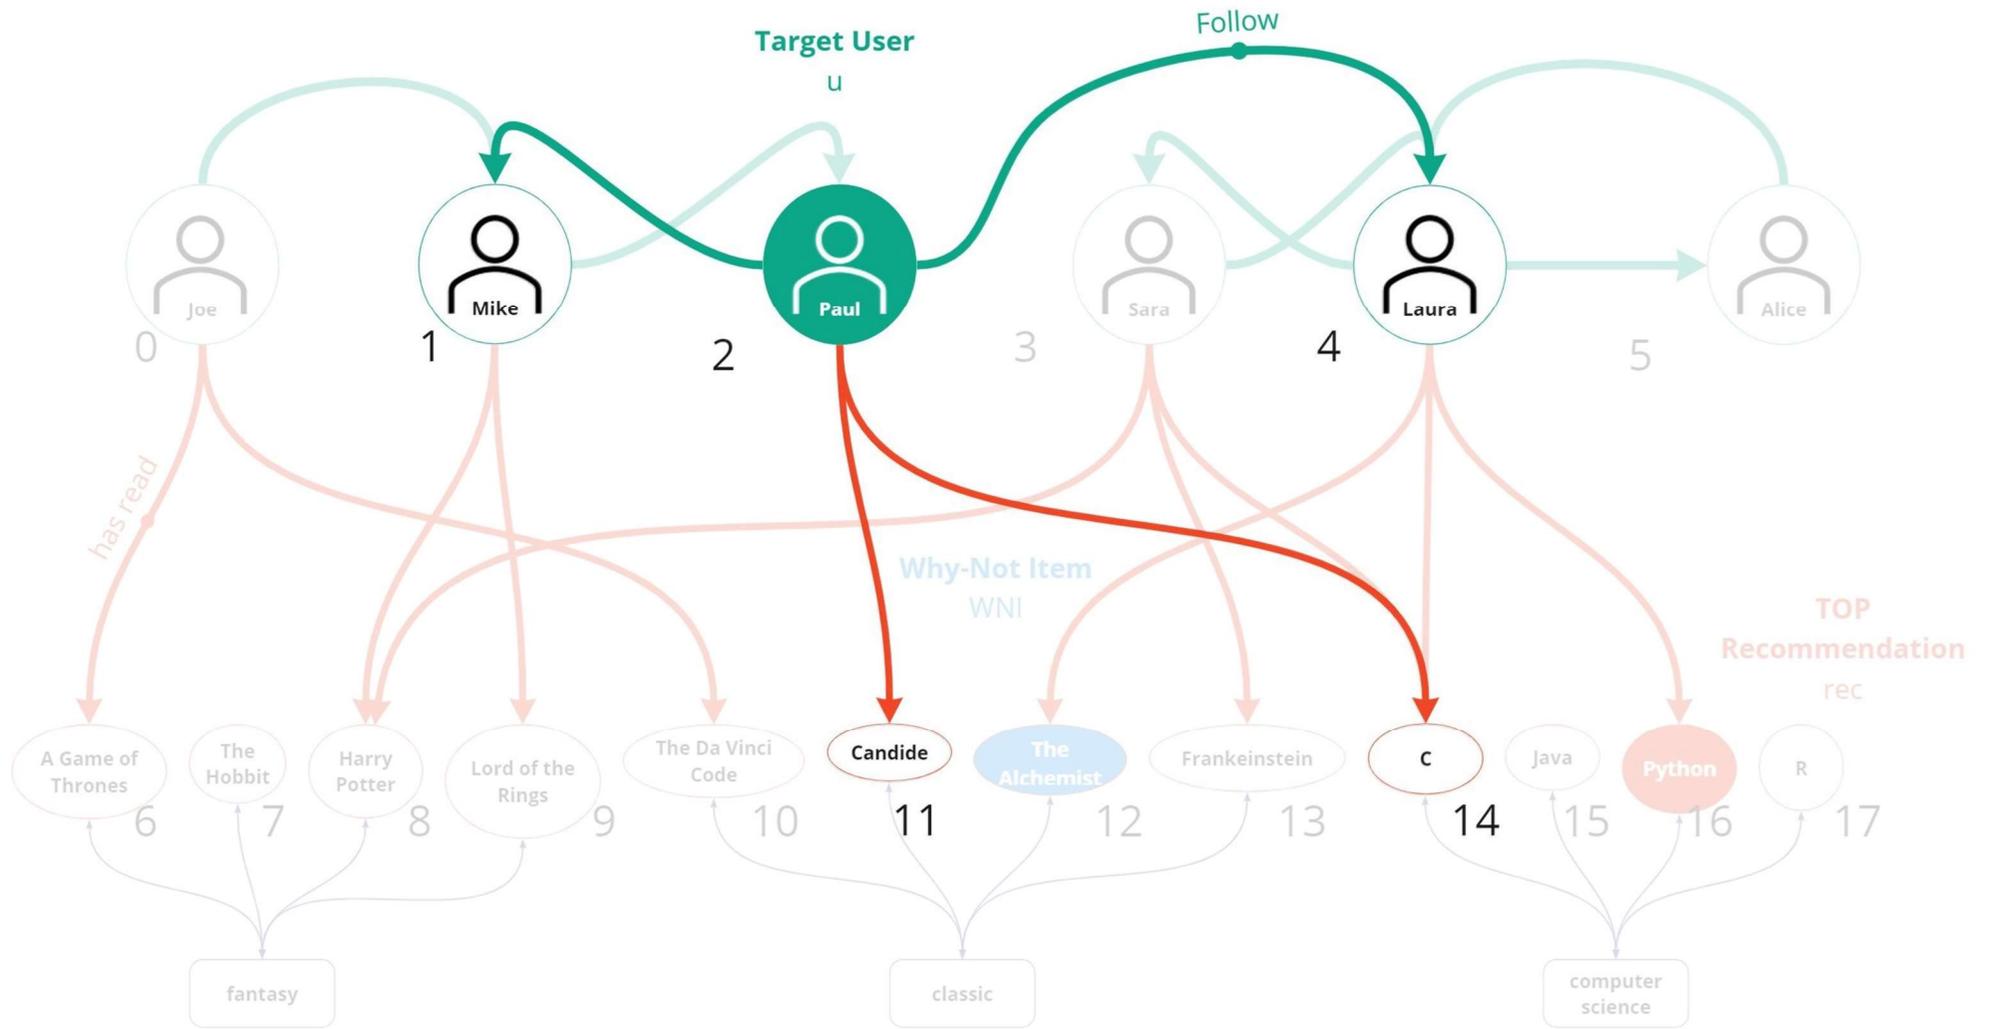
Example

— — —

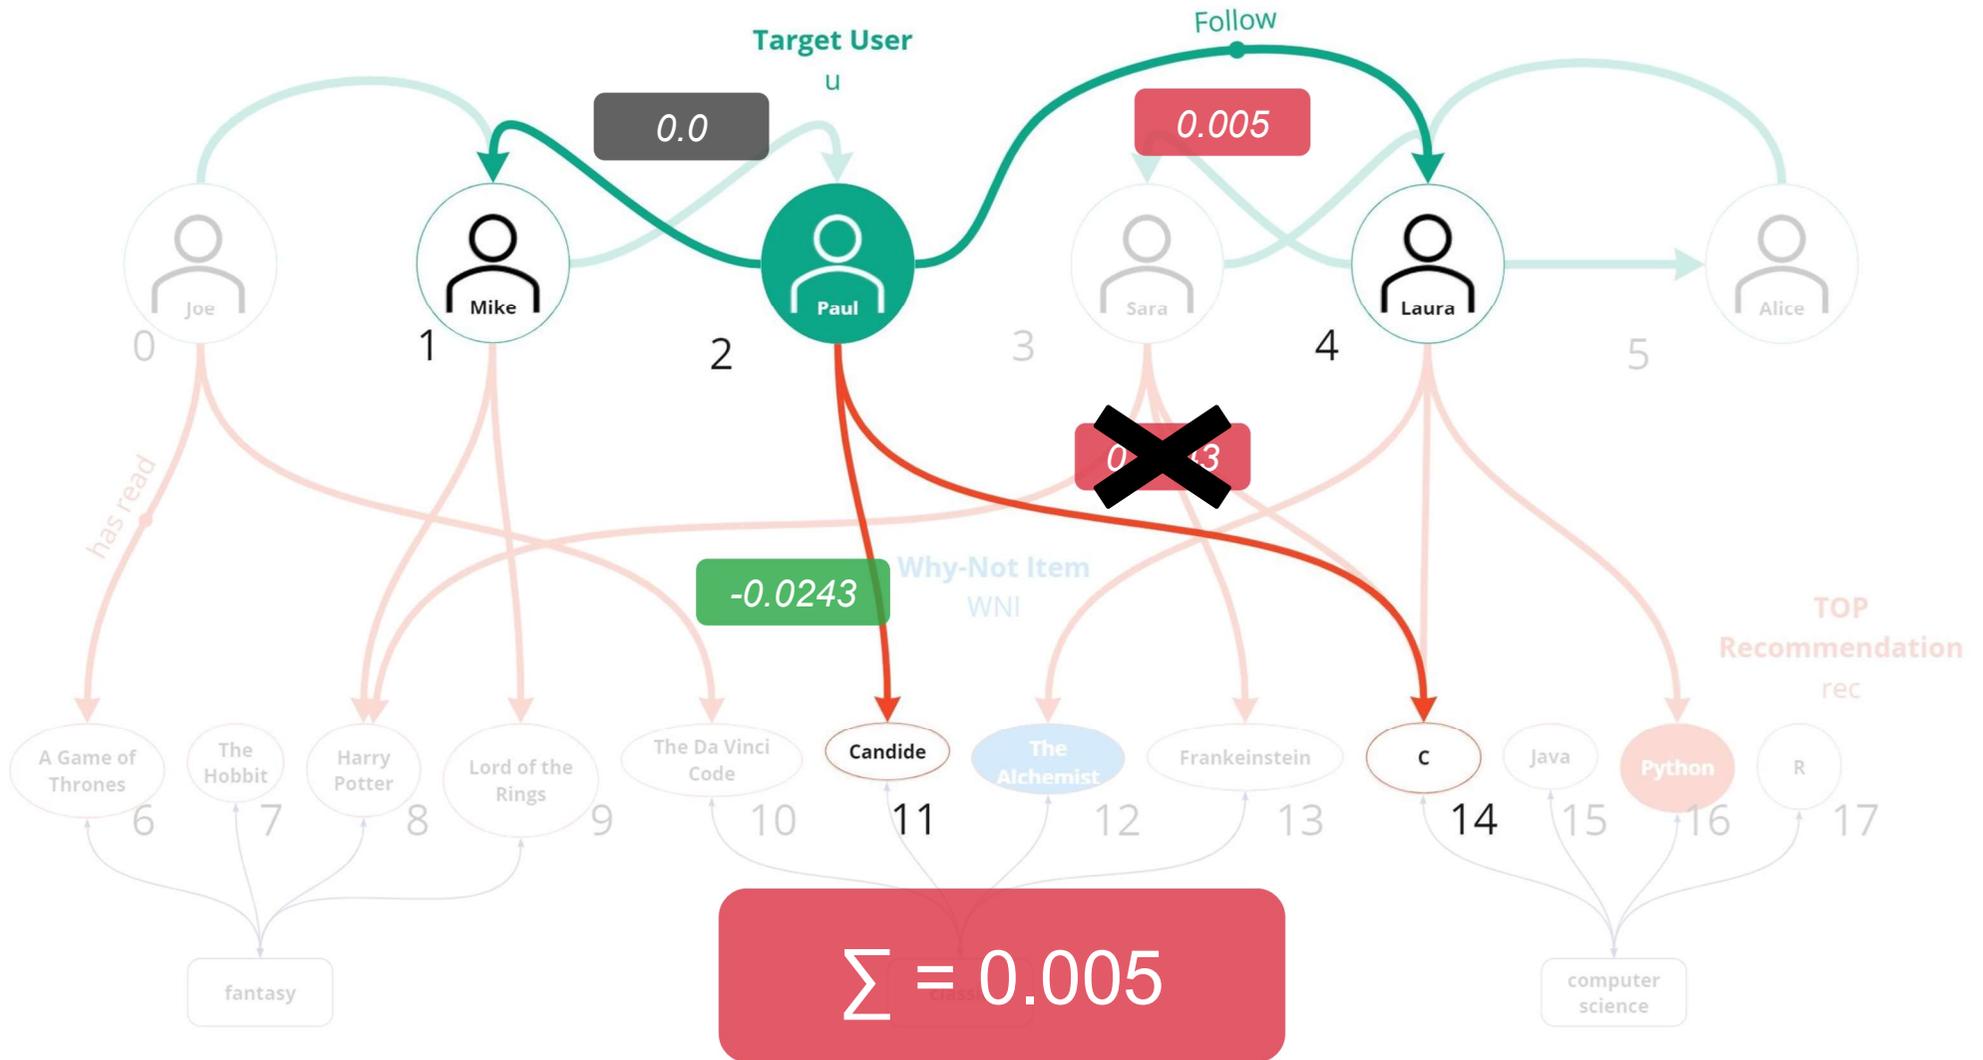
Incremental heuristic in Remove Mode

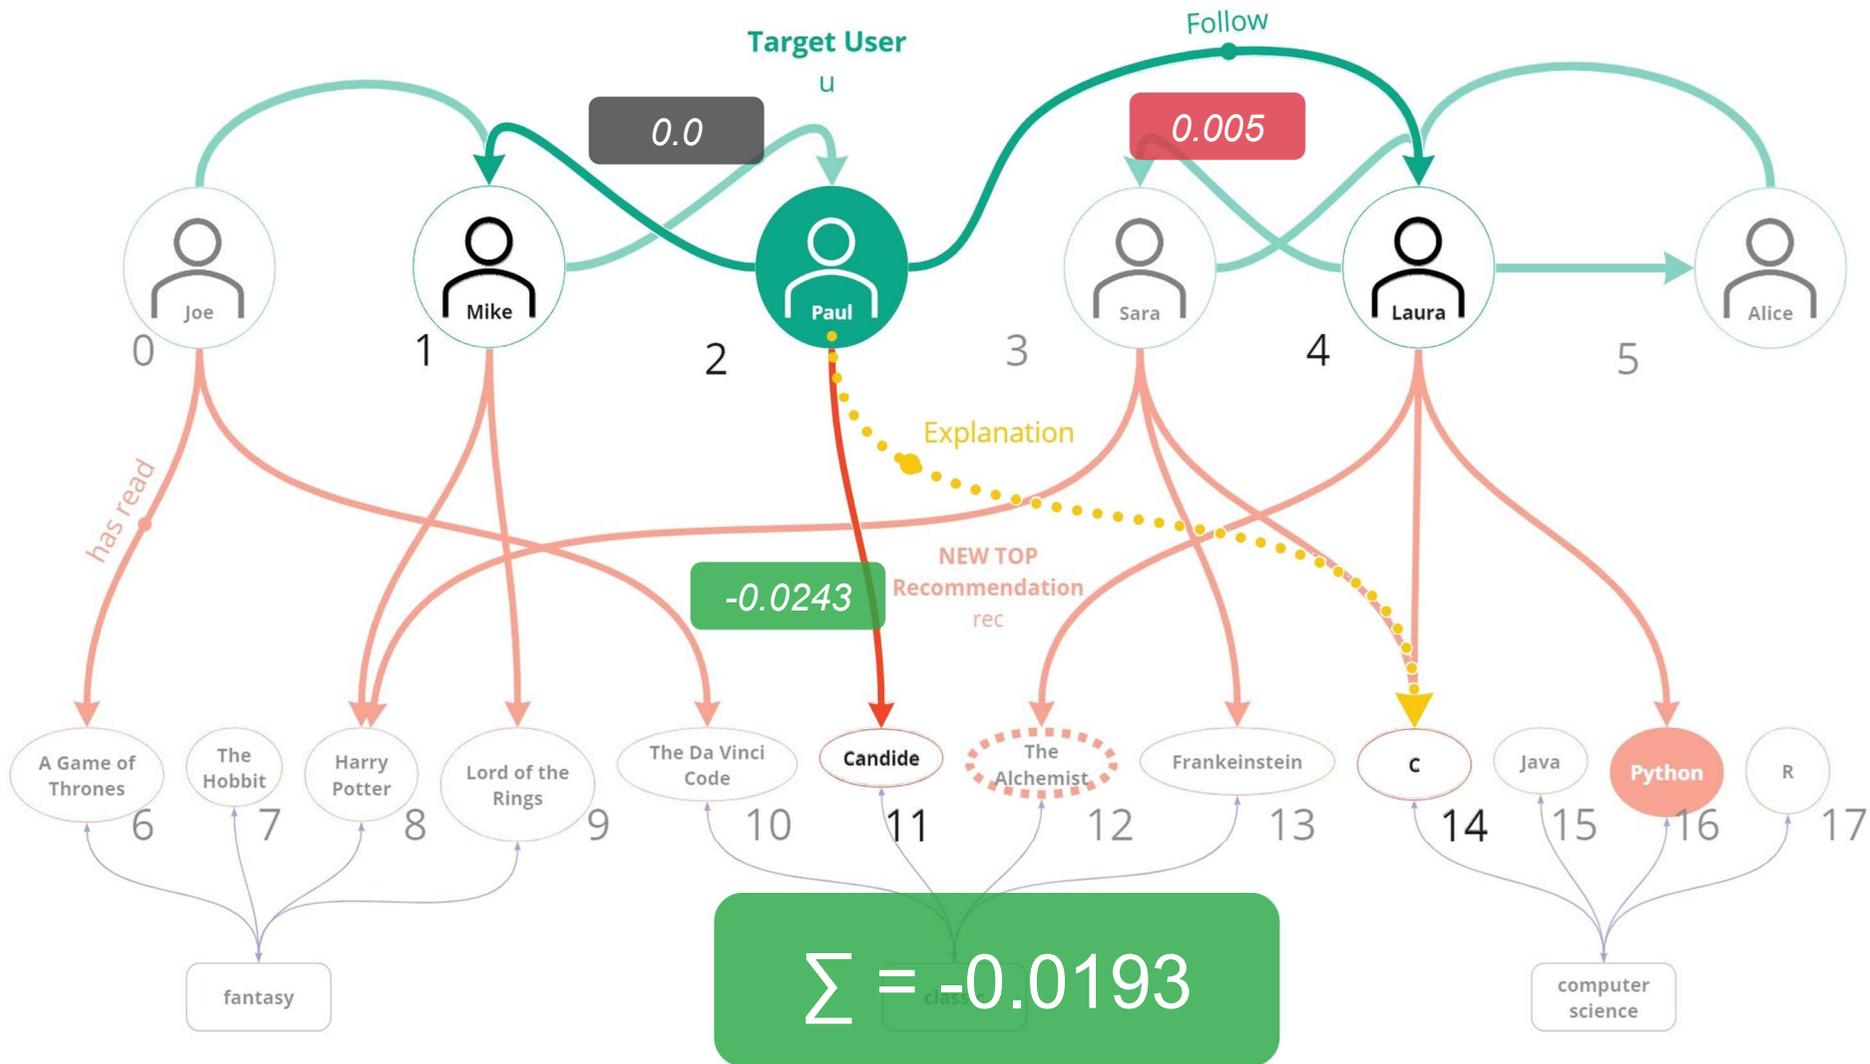




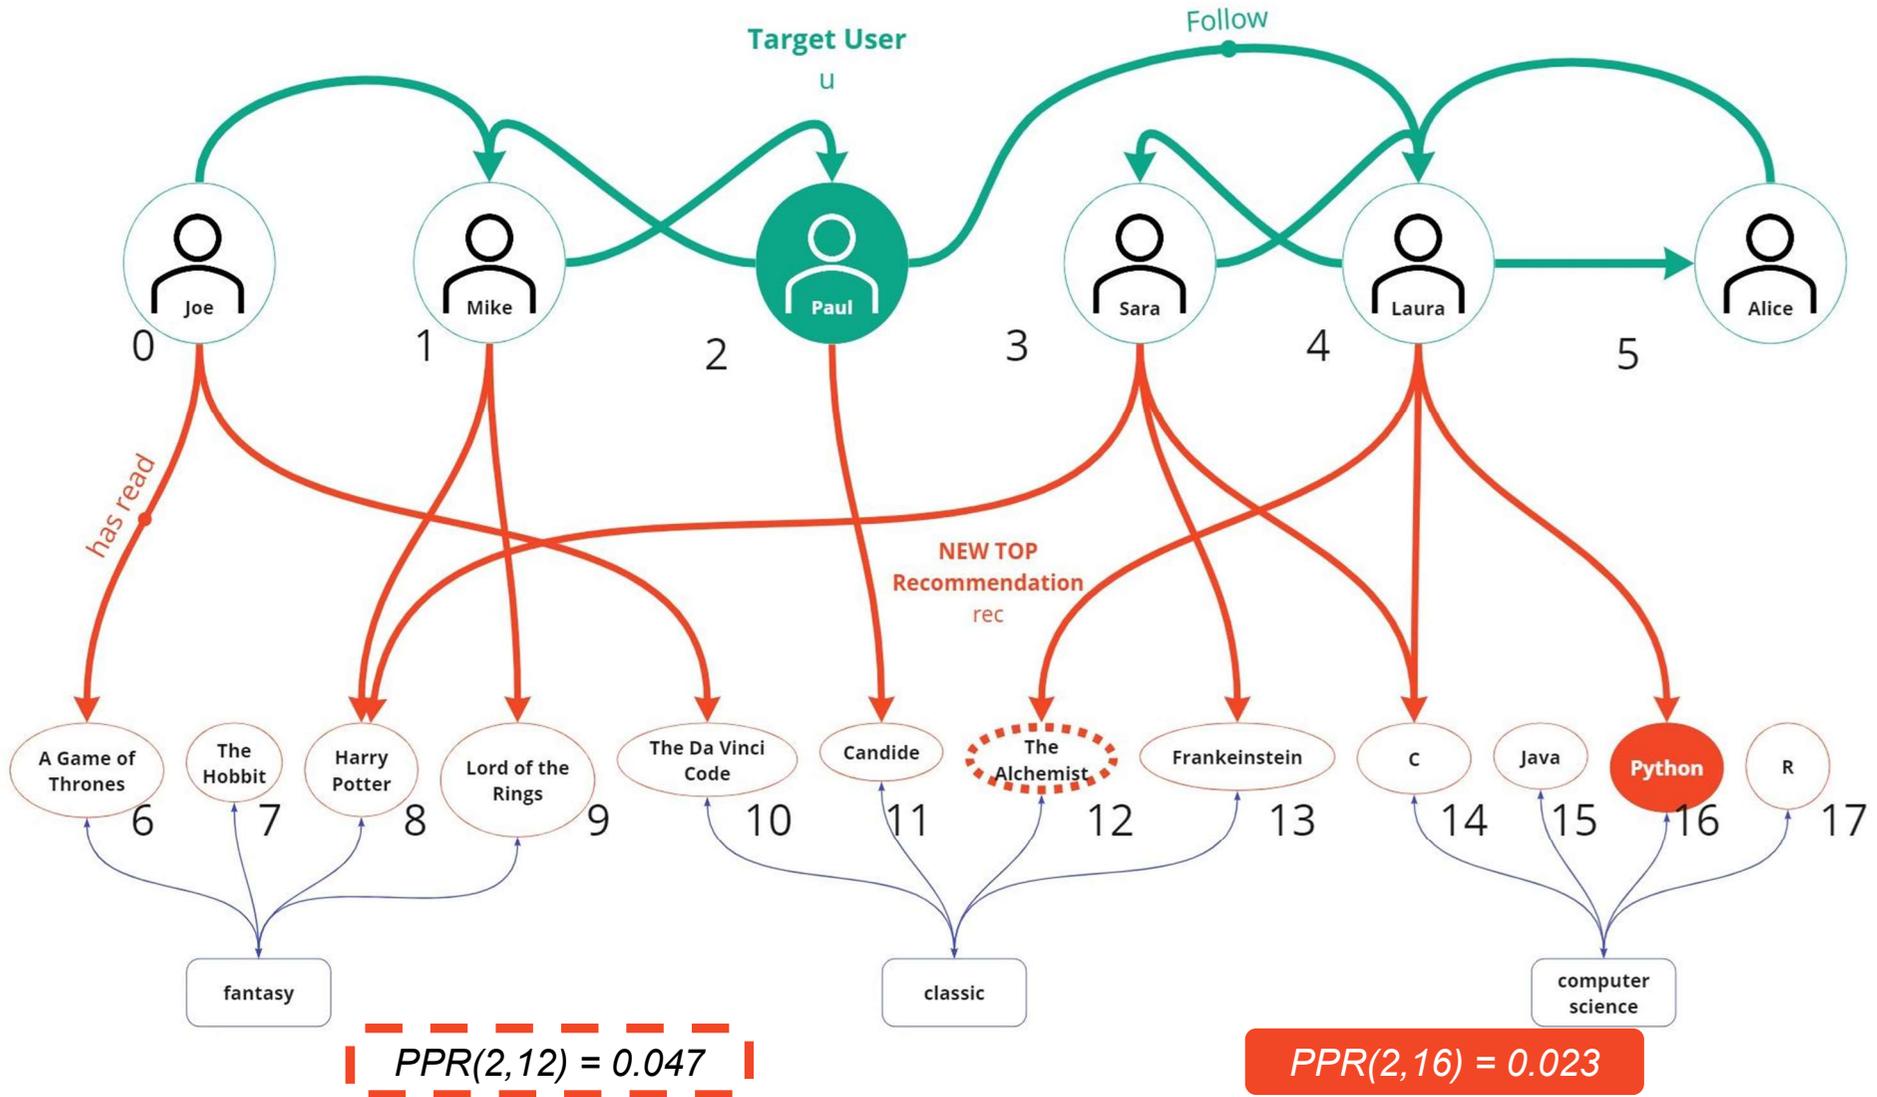


$$\text{contribution}_{rmv}(n_i) = W(u, n_i) \cdot (PPR(n_i, rec | A) - PPR(n_i, WNI | A))$$





After removing the explanation ...



Experimental Evaluation



Dataset Analysis

Why-Not scenario definition:

- **Target users:** 100 random “*normal*” users from an Amazon Dataset
- **Graph:** 4 radius subgraph around the target user node
- **Why-Not item:** items ranked 2 to 10 in the recommendation list

Dataset	#nodes	#edges	#users	#items	#reviews	#categories
Amazon	114 002	687 184	1 999	53 727	58 233	43

Edge Types	
type	nodes
“rated”	(user↔item)
“reviewed”	(user↔review)
“has-review”	(item↔review)
“belongs-to”	(item↔category)

Algorithms and Baselines

Algorithms:

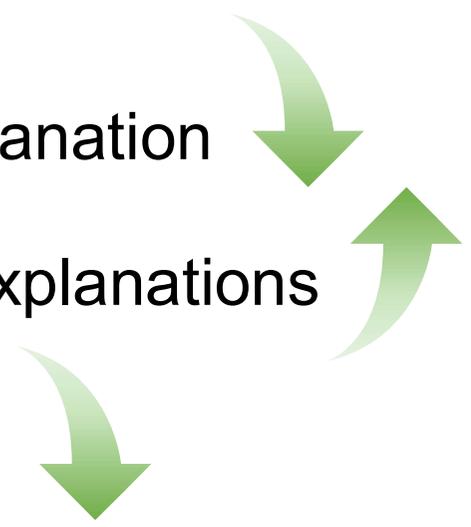
- add_Powerset, Powerset heuristic, Add Mode;
- add_Incremental, Incremental heuristic, Add Mode;
- add_ex Exhaustive Comparison heuristic, Add Mode;

- remove_Powerset Powerset heuristic, Remove Mode;
- remove_Incremental Incremental heuristic, Remove Mode;
- remove_ex Exhaustive Comparison heuristic, Remove Mode.

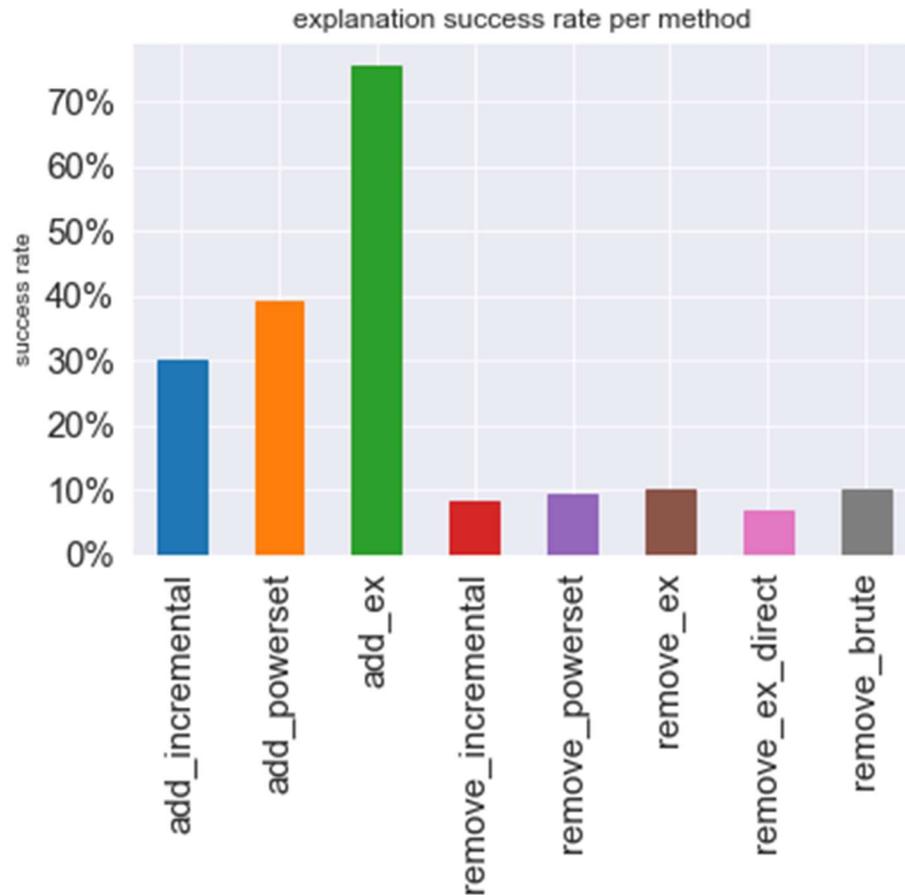
Baselines:

- Brute force (remove mode) : to evaluate the computation time and explanation size
- Exhaustive Comparison Direct : to evaluate the need for the check step

Evaluation metrics

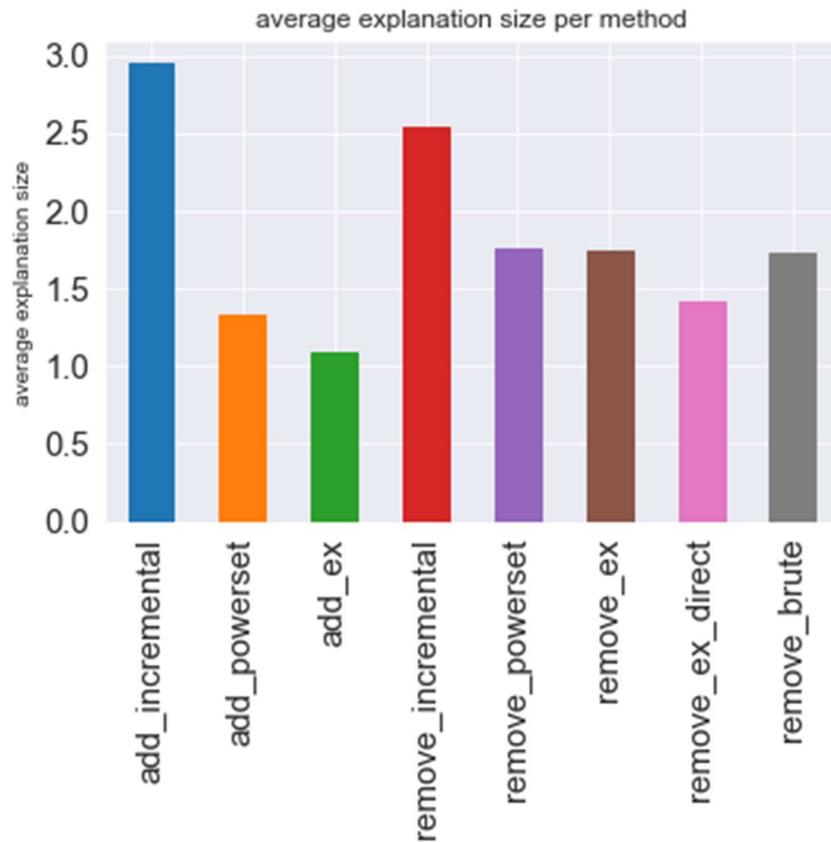
- **Explanation size** : a good explanation is a short explanation
 - **Success rate**: we need the algorithm to provide the explanations
 - **Time**: we need the algorithm to be as fast as possible
- 

Explanation success rate per method



- Low success rate in Remove Mode because of the limited candidate space.
- Add Mode is more successful than Remove Mode.
- Best heuristic: Exhaustive Method in Add Mode.

Average explanation size per method



- Low explanation size (smaller than 3 edges on average) for all heuristics.
- Remove Mode, best heuristics: Exhaustive and Powerset.
- Add Mode: low (or minimum) explanation size for all.

Average computation time

Method	(a)	(b)	(c)
<i>add_Incremental</i>	6,54	8,31	5,78
add_Powerset	57,55	133,96	8,19
add_ex	21618,32	23924,37	14646,56
<i>remove_Incremental</i>	9,07	8,20	9,15
remove_Powerset	287,91	15,32	315,31
remove_ex	173,44	24,48	190,13
remove_ex_direct	25,14	21,81	25,38
remove_brute	908,73	22,37	1008,07

Average runtime in seconds per method (a) in the general case, (b) when an explanation is found, and (c) when no explanation is found.

- The Exhaustive Comparison method in Remove Mode is comparable to brute force when successful.
- Incremental Methods give faster results, both in Add and Remove Modes.

Conclusion & Future Work



Conclusion & Future Work

Conclusions

1. We proposed counterfactual-like explanations for missing recommendations, in the form of a set of user-rooted, existing or potential edges.
2. We proposed various heuristics for computing why-not explanations, targeting small explanations or fast computation.
3. Our experiments on real-world data showed the feasibility of our solution.

Future Work

- Improve the success rate by extending the expressiveness of our explanations (considering edge weights, creating meta-explanations, etc).
- Further evaluate the quality of the explanations, e.g., by evaluating the user satisfaction property.

Why-Not Explainable Graph Recommender

Hervé-Madelein Attolou¹, Katerina Tzompanaki¹, Kostas Stefanidis², and Dimitris Kotzinos¹

¹ETIS, CY Cergy Paris University, ENSEA, CNRS, UMR8051

²Tampere University

Acknowledgments

This work is partially supported by the National French Agency, under the EXPIDA - ANR-22-CE23-0017 project.



For more, check:

<https://git.cyu.fr/hattolou/why-not-explainable-graph-recommender>



References

- [1] Behnoush Abdollahi and Olfa Nasraoui. Using Explainability for Constrained Matrix Factorization. In Proceedings of the Eleventh ACM Conference on Recommender Systems, RecSys '17, pages 79–83, New York, NY, USA, August 2017. Association for Computing Machinery.
- [2] Nicole Bidot, Melanie Herschel, and Katerina Tzompanaki. Query-Based WhyNot Provenance with NedExplain, 2014. type: dataset.
- [3] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yunhsuan Sung, Brian Strope, and Ray Kurzweil. Universal sentence encoder, 2018.
- [4] Adriane Chapman and H. V. Jagadish. Why not? In Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data, SIGMOD '09, page 523–534, New York, NY, USA, 2009. Association for Computing Machinery.
- [5] Azin Ghazimatin, Oana Balalau, Rishiraj Saha Roy, and Gerhard Weikum. PRINCE: Provider-side Interpretability with Counterfactual Explanations in Recommender Systems. Proceedings of the 13th International Conference on Web Search and Data Mining, pages 196–204, January 2020. arXiv: 1911.08378.
- [6] Md Islam. On answering why and why-not questions in databases. pages 298–301, 04 2013.
- [7] Yanhong Li, Wang Zhang, Changyin Luo, Xiaokun Du, and Jianjun Li. Answering why-not questions on top-k augmented spatial keyword queries. Knowledge Based Systems, 223:107047, 2021.
- [8] Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. Content-based Recommender Systems: State of the Art and Trends, pages 73–105. Springer US, Boston, MA, 2011.
- [9] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, page 4768–4777, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [10] Athanasios N. Nikolakopoulos and George Karypis. RecWalk: Nearly Uncoupled Random Walks for Top-N Recommendation. February 2019.
- [11] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should I trust you?": Explaining the predictions of any classifier. In Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi, editors, Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13–17, 2016, pages 1135–1144. ACM, 2016.
- [12] J. Ben Schafer, Joseph Konstan, and John Riedl. Recommender systems in ecommerce. In Proceedings of the 1st ACM conference on Electronic commerce, EC '99, pages 158–166, New York, NY, USA, November 1999. Association for Computing Machinery.
- [13] Maria Stratigi, Aikaterini Tzompanaki, and Kostas Stefanidis. Why-Not Questions & Explanations for Collaborative Filtering. In WISE, Amsterdam, Netherlands, October 2020.
- [14] Aikaterini Tzompanaki. Réponses manquantes : Débogage et Réparation de requêtes. (Query Debugging and Fixing to Recover Missing Query Results). PhD thesis, University of Paris-Saclay, France, 2015.
- [15] Mengting Wan and Julian J. McAuley. Item recommendation on monotonic behavior chains. In Sole Pera, Michael D. Ekstrand, Xavier Amatriain, and John O'Donovan, editors, Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018, pages 86–94. ACM, 2018.
- [16] Mengting Wan, Rishabh Misra, Ndapa Nakashole, and Julian J. McAuley. Fine grained spoiler detection from large-scale review corpora. In Anna Korhonen, David R. Traum, and Lluís Màrquez, editors, Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers, pages 2605–2610. Association for Computational Linguistics, 2019.
- [17] Minghua Xu and Shenghao Liu. Semantic-enhanced and context-aware hybrid collaborative filtering for event recommendation in event-based social networks. IEEE Access, 7:17493–17502, 2019.
- [18] Hongyang Zhang, Peter Lofgren, and Ashish Goel. Approximate personalized pagerank on dynamic graphs. In Proceedings of the 22nd ACM SIGKDD International Conference on knowledge discovery and data mining, pages 1315–1324, 2016.
- [19] Yongfeng Zhang and Xu Chen. Explainable Recommendation: A Survey and New Perspectives. INR, 14(1):1–101, March 2020. Number: 1 Publisher: Now Publishers, Inc.
- [20] M. Alshammari, O. Nasraoui and S. Sanders, "Mining Semantic Knowledge Graphs to Add Explainability to Black Box Recommender Systems," in IEEE Access, vol. 7, pp. 110563-110579, 2019, doi: 10.1109/ACCESS.2019.2934633.
- [21] Krisztián Balog, Filip Radlinski, and Shushan Arakelyan. 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, pages 265–274, 2019.
- [22] Zhang, Y., Lai, G., Zhang, M., Zhang, Y., Liu, Y., and Ma, S. (2014b). Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14, Gold Coast, QLD, Australia - July 06 - 11, 2014, pages 83–92.
- [23] GHAZIMATIN, Azin. Enhancing explainability and scrutability of recommender systems. 2021.
- [24] Wu, Y. and Ester, M. (2015). FLAME: A probabilistic model combining aspect based opinion mining and collaborative filtering. In Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM 2015, Shanghai, China, February 2-6, 2015, pages 199–208