Who do you think I am? Interactive User Modelling with Item Metadata

Joey De Pauw
joey.depauw@uantwerpen.be
University of Antwerp
Antwerp, Belgium

Koen Ruymbeek
koen.ruymbeek@uantwerpen.be
University of Antwerp
Antwerp, Belgium

Bart Goethals
bart.goethals@uantwerpen.be
University of Antwerp
Antwerp, Belgium
Monash University
Melbourne, Australia

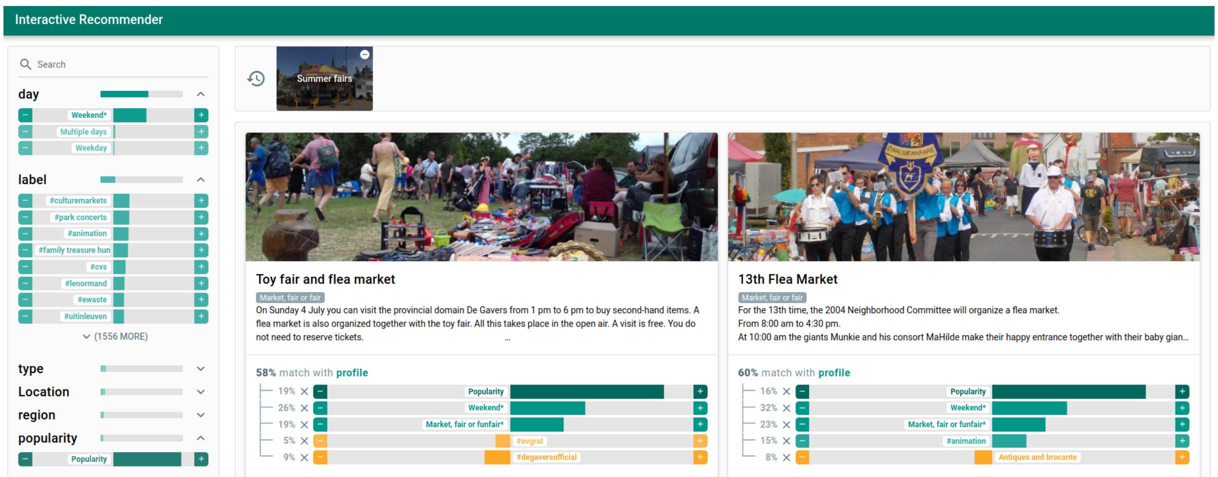

Figure 1: Screenshot of the interactive recommender system interface hosted on https://tease-rec.uantwerpen.be.

ABSTRACT

Recommender systems are used in many different applications and contexts, however their main goal can always be summarised as “connecting relevant content to interested users”. Explanations have been found to help recommender systems achieve this goal by giving users a look under the hood that helps them understand *why* they are recommended certain items. Furthermore, explanations can be considered to be the first step towards interacting with the system. Indeed, for a user to give feedback and guide the system towards better understanding her preferences, it helps if the user has a better idea of what the system has already learned.

To this end, we propose a linear collaborative filtering recommendation model that builds user profiles within the domain of item metadata. Our method is hence inherently transparent and explainable. Moreover, since recommendations are computed as a linear function of item metadata and the interpretable user profile, our method seamlessly supports interactive recommendation. In other words, users can directly tweak the weights of the learned profile for more fine-grained browsing and discovery of content based on their current interests. We demonstrate the interactive aspect of this model in an online application for discovering cultural events in Belgium.

CCS CONCEPTS

• **Information systems** → **Search interfaces**; *Collaborative search*; *Personalization*; *Learning to rank*.

KEYWORDS

interactive recommendation, personalization, transparency, explainability

ACM Reference Format:
Joey De Pauw, Koen Ruymbeek, and Bart Goethals. 2022. Who do you think I am? Interactive User Modelling with Item Metadata. In *Sixteenth ACM Conference on Recommender Systems (RecSys ’22), September 18–23, 2022, Seattle, WA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3523227.3551470

RecSys ’22, September 18–23, 2022, Seattle, WA, USA

ACM ISBN 978-1-4503-9278-5/22/09.
https://doi.org/10.1145/3523227.3551470

1 INTRODUCTION

Whichever recommendation method is used, it is found that a user’s trust in the system can be improved with truthful and relevant explanations, as explanations provide both context for the recommendations and insight into the system [8]. Furthermore, explanation also help the users to accept the given recommendation, to find relevant content faster and to increase the overall ease of use of the system [8, 22, 25, 26]. However, not all explanations and explanation types are equally informative and their usefulness also depends on the recommendation scenario and the current goals of the user [22, 23].

For example, explanations can be based on the computed similarities between users and/or items [22] or based on item metadata for feature based explanations. The latter are found to be more interpretable due to their typically smaller domain that is easier to understand [18]. Alternatively, explanation methods can be designed as post-processing steps [14, 16, 17]. This allows them to be used in combination with any recommendation algorithm, however, with the caveat that the *fidelity* is high enough. There is no guarantee that explanations computed this way actually reflect what the algorithm has learned and hence they are of limited use for gaining insight into the underlying model.

To overcome these limitations of explanation methods, we propose a hybrid recommender model [4] called TEASER for "Transparent & Explainable Aspect Space Embedding Recommender". It uses implicit feedback interaction data to learn the similarities between aspects of items. In other words, our method infers from the history of items that users consume, which aspects they might like, and subsequently recommends new items based on the learned profiles. As a result we combine the benefits of both collaborative and content based filtering explanations to achieve a fully transparent recommendation algorithm. To the best of our knowledge, no hybrid linear recommendation model with a comparable degree of explainability, transparency and interactiveness exists in the literature, though numerous similar avenues are explored in related works [1–3, 5, 7, 9, 11, 12, 15, 19, 24]. An online application with the TEASER model that showcases its strengths is discussed in Section 2.

2 APPLICATION

In the demo application hosted on https://tease-rec.uantwerpen.be, users can interact with the TEASER model to browse for events or other things to do in Belgium. Figure 1 shows a screenshot of the application and a short video demonstration is available at https://www.youtube.com/watch?v=gkmasN7FmPs. To help new users get acquainted with our interface, an interactive tour automatically starts on the first visit. We summarise this tour here for completeness.

On the left-hand side is a list of *categories* and the respective *tags* they contain. For example the category 'day' has three tags: 'Weekend', 'Multiple days' and 'Weekday'. For each tag the model computes an affinity score based on the user history. This collection of tags and scores is what we call the *user profile*. Higher affinity scores are shaded more teal and the progress bar is filled more to the right, whereas for negative scores the bar is filled to the left and shaded in orange. Additionally, we can see the estimated impact that each category has on the recommendations with the positive-only progress bars next to them. These sum up to one across all features.

On the top-right then, we find an ordered list of recently viewed events in the *user history*. These are what make up the profile on the left and can help the user make the connection between the two. It is also possible for the user to remove any of the history items and receive recommendations based on the remaining ones.

Finally on the remainder of the page, we show personalized *recommendations* for the user. Alongside the basic information of each event, we also include detailed information about why the event was selected specifically for the user. These *item explanations* are computed as mentioned in Section 3.2 and only the top-5 explanations are selected to be displayed if they have a contribution of 5% or more on the final score in absolute value. One can see that this also allows *negative explanations* to be given, indicating that the model recommended an event to you, despite knowing that you might not like it for specific reasons. This of course is a design choice and can be disabled if negative explanations are undesirable.

To dive deeper into the explanations, we first dissect the overall %-match score into its normalized sum of aspects. Here the percentages show how much each aspect contributed to the final score. The aspects themselves are the same as in the user profile. We hence provide a way for the user to both learn more about the event *and* about why the model thinks she will like it, effectively ranking the most interesting information about the event for the user.

The capstone feature of TEASER and the demo is its ability to interact with the user. As can be seen, each tag of the user profile also has a plus and a minus button. These behave as expected and increase respectively decrease the weight of the tag, which allows the user to manually indicate her preferences on top of what the model has learned. After making changes, the top recommendations are recomputed to reflect the updated preferences.

3 RECOMMENDATION MODEL

TEASER is a hybrid recommender model for implicit feedback data. It is similar in conception to the well-known linear regression model EASE [20], and in fact it is based on one of its variants, namely EDLAE [21]. We adapt the training objective of EDLAE by replacing the decoder matrix such that the user embeddings "take on the meaning" of the item metadata features. This choice was motivated by previous work that demonstrates how using metadata for explanation can greatly increase the understandability for end users [18]. Indeed, though EASE and EDLAE are also linear models, one can argue that they are still not understandable due to the intractable scale of their features [10, 13]. Condensing the learned information down to less and more interpretable weights is the main design choice behind TEASER.

3.1 Definition

Let $X \in \{0, 1\}^{m \times n}$ be the binary interaction matrix with m the number of users and n the number of items, then the TEASER model is defined as

$$\hat{E} = \underset{E}{\arg\min} \left\| X - X(ES^\top - \mathrm{diagM}(\mathrm{diag}(ES^\top))) \right\|_F^2 + \lambda_1 \left\| ES^\top - \mathrm{diagM}(\mathrm{diag}(ES^\top)) \right\|_F^2 + \lambda_2 \left\| E \right\|_F^2 . \quad (1)$$

with diagM(.) the diagonal matrix with given diagonal, $E, S \in \mathbb{R}^{n \times t}$ the encoder respectively the decoder, and rank t smaller than $\min(m, n)$. Notice that in this definition it is not possible for the model to learn anything from the diagonal elements. By eliminating the diagonal of $ES^\top$ from the training objective, we prevent the

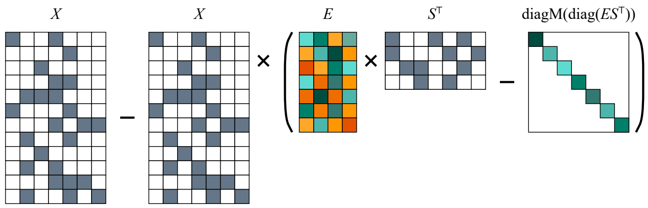

Figure 2: Graphical illustration of TEASER. "×" represents matrix multiplication, teal colours indicate positive values, white and grey are zero and one respectively, and shades of orange are used for negative values.

model from overfitting towards the identity as it can no longer learn to predict items based on their own presence in the history [21].

TEASER modifies EDLAE to use the fixed one-hot encoded item metadata matrix (S) as item embeddings instead. Since users and items are embedded in a common latent space, this effectively makes it so the dimensions of this space correspond to the item features. As a result, the encoder E now connects the interaction matrix X with the metadata matrix S, which means that our model is a hybrid of collaborative filtering (via X) and content-based filtering (via S).

See Figure 2 for a graphical illustration of the model. More details and an in depth discussion of the benefits of the model, including experimental results, can be found in the following technical report [6]. Additionally, the source code is available at https://github.com/JoeyDP/TEASER.

3.2 Benefits

The training objective of TEASER is quite restrictive: recommendation scores are no longer computed from pairwise item-item weights, but rather from tag-item weights. This seemingly simple change compared to EDLAE however has several benefits for *explainability*, *transparency* and *interactiveness*, as explained in the next subsections.

3.2.1 Explainability. The first and foremost benefit of using item metadata, is that explanations of recommendations can take on the form "Recommendation based for {37%} on your affinity score of {0.4} with {outdoors activities}" (or any other aspect of the item). These kinds of explanations can trivially be computed from the model, and are even the actual weights that lead to the recommendation. Namely, an item's score is calculated by summing over the aspects of that item in the "user profile". Given user history $\vec{x}$, an item i's score is computed as $\langle \vec{x}E, S_i \rangle$. From this the contribution of each aspect can be computed as in the example.

Secondly, alongside item explanations, our method can also explain the *profile* it has learned of the user. The aforementioned affinity scores with certain aspects are simply the user embedding ($\vec{x}E$) in our model. The higher the score, the more the system expects the user has interest in this aspect. Each row E_i reflects which aspects a user is expected to find interesting (or not) in other items, if she *consumed* item i. For example when recommending events, if someone likes an event in Brussels, they can also be given a higher affinity with the neighbouring cities and villages or even with other big cities in Belgium, if this signal is present in the interaction data.

3.2.2 Transparency. Furthermore, the matrix E encodes affinity scores of users for aspects on a per-item basis. For example, the row E_i reflects which aspects a user is expected to find interesting (or not) in other items, if she *consumed* item i. This kind of scrutability of the model itself can come in handy for data scientists or practitioners that maintains the system. Either to gain insight into the interaction data or to debug the model by identifying missing or unexpected relations.

3.2.3 Interactiveness. The final benefit related to the transparency of our model lies in the fact that the explanations are exactly the weights of the model, and not the result of a post-processing step to approximate what the model really computed. This means that, when a user gives feedback on the explanations, the feedback can seamlessly be integrated into the model and be used for interactive recommendation.

4 CONCLUSIONS

A new highly transparent, explainable and interactive hybrid recommender, TEASER, is presented and demonstrated in an online application. The main benefit of TEASER lies in the use of item metadata to build a profile of the user, which is then used to compute recommendations and their accompanying explanations. As such, the model is fully transparent and explainable with item metadata. This domain both scales well and is intuitive for the end user. Furthermore, TEASER seamlessly enables interactive recommendation, where the user can provide feedback on the explanations and on the learned profile for the model to incorporate.

ACKNOWLEDGMENTS

This work was supported by the Research Foundation — Flanders (FWO) [11E5921N to J. De Pauw] and the Flemish Government under the "Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen" (https://airesearchflanders.be/) programme. Special thanks to Noah Daniëls for designing and developing the web application.

REFERENCES

[1] Ivana Andjelkovic, Denis Parra, and John O'Donovan. 2019. Moodplay: Interactive music recommendation based on Artists' mood similarity. *International Journal of Human-Computer Studies* 121 (2019), 142–159. https://doi.org/10.1016/j.ijhcs.2018.04.004 Advances in Computer-Human Interaction for Recommender Systems.
[2] Diego Antognini and Boi Faltings. 2021. Fast Multi-Step Critiquing for VAE-based Recommender Systems. In *Fifteenth ACM Conference on Recommender Systems*. 209–219.
[3] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: a visual interactive hybrid recommender system. In *Proceedings of the sixth ACM conference on Recommender systems*. 35–42.
[4] Robin Burke. 2002. Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction* 12, 4 (2002), 331–370.
[5] Robin D Burke, Kristian J Hammond, and Benjamin C Young. 1996. Knowledge-based navigation of complex information spaces. In *Proceedings of the national conference on artificial intelligence*, Vol. 462. 468.
[6] Joey De Pauw, Koen Ruymbeek, and Bart Goethals. 2022. Modelling Users with Item Metadata for Explainable and Interactive Recommendation. *arXiv preprint arXiv:2207.00350* (2022).
[7] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications* 56 (2016), 9–27.
[8] 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.

[9] Hanze Li, Scott Sanner, Kai Luo, and Ga Wu. 2020. A Ranking Optimization Approach to Latent Linear Critiquing for Conversational Recommender Systems. In *Fourteenth ACM Conference on Recommender Systems*. 13–22.

[10] 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.

[11] Benedikt Loepp, Tim Donkers, Timm Kleemann, and Jürgen Ziegler. 2019. Interactive recommending with tag-enhanced matrix factorization (TagMF). *International Journal of Human-Computer Studies* 121 (2019), 21–41.

[12] Kai Luo, Scott Sanner, Ga Wu, Hanze Li, and Hojin Yang. 2020. Latent linear critiquing for conversational recommender systems. In *Proceedings of The Web Conference 2020*. 2535–2541.

[13] Christoph Molnar. 2020. *Interpretable machine learning*. Lulu. com.

[14] Cataldo Musto, Fedelucio Narducci, Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. 2019. Linked open data-based explanations for transparent recommender systems. *International Journal of Human-Computer Studies* 121 (2019), 93–107.

[15] John O'Donovan, Barry Smyth, Brynjar Gretarsson, Svetlin Bostandjiev, and Tobias Höllerer. 2008. PeerChooser: Visual Interactive Recommendation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) *(CHI '08)*. Association for Computing Machinery, New York, NY, USA, 1085–1088. https://doi.org/10.1145/1357054.1357222

[16] Alexis Papadimitriou, Panagiotis Symeonidis, and Yannis Manolopoulos. 2012. A generalized taxonomy of explanations styles for traditional and social recommender systems. *Data Mining and Knowledge Discovery* 24, 3 (2012), 555–583.

[17] Georgina Peake and Jun Wang. 2018. Explanation mining: Post hoc interpretability of latent factor models for recommendation systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2060–2069.

[18] Yanou Ramon, David Martens, Theodoros Evgeniou, and Stiene Praet. 2020. Metafeatures-based Rule-Extraction for Classifiers on Behavioral and Textual Data. *arXiv preprint arXiv:2003.04792* (2020).

[19] Shilad Sen, Jesse Vig, and John Riedl. 2009. Tagommenders: connecting users to items through tags. In *Proceedings of the 18th international conference on World wide web*. 671–680.

[20] Harald Steck. 2019. Embarrassingly shallow autoencoders for sparse data. In *The World Wide Web Conference*. 3251–3257.

[21] Harald Steck. 2020. Autoencoders that don't overfit towards the Identity. *Advances in Neural Information Processing Systems* 33 (2020).

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

[23] Thi Ngoc Trang Tran, Viet Man Le, Müslüm Atas, Alexander Felfernig, Martin Stettinger, and Andrei Popescu. 2021. Do Users Appreciate Explanations of Recommendations? An Analysis in the Movie Domain. In *Fifteenth ACM Conference on Recommender Systems*. 645–650.

[24] 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.

[25] Jingjing Zhang and Shawn P Curley. 2018. Exploring explanation effects on consumers' trust in online recommender agents. *International Journal of Human–Computer Interaction* 34, 5 (2018), 421–432.

[26] Yongfeng Zhang and Xu Chen. 2018. Explainable recommendation: A survey and new perspectives. *arXiv preprint arXiv:1804.11192* (2018).