Cold-starting collaborative filtering for music recommendations using social tags

In music streaming services, users' listening activity represents an abundant source of data (commonly referred to as implicit feedback). Collaborative filtering methods can be successfully applied to implicit feedback [1], but they suffer from the cold-start problem. We show how to mitigate the cold-start problem for new artists using a hybrid model that combines listening activity and social tags into a joint matrix factorization scheme. In contrast to previous models, we address the weak labeling nature of social tags by weighting them according to their relative importance in the dataset, similarly as it is done with implicit feedback [2]. Our model is able to mitigate the cold-start problem for new artists, based on experiments on real Last.fm data [3].

Andreu Vall
Department of Computational Perception
Johannes Kepler University, Linz, Austria
andreu.vall@jku.at

» Hybrid collaborative filtering with social tags

The proposed model combines user-artist playcounts with user-tag counts and artist-tag counts (i.e., how often a user has applied a tag and how often an artist has been labeled with a tag). Latent factors for users, artists and tags are estimated in a joint matrix factorization scheme to approximate binary versions of the data matrices.

The cost function (equipped with a regularization term) is minimized through Alternating Least Squares and hyperparameters are determined by grid search.

Once the model is trained, the expected preferences for each user-artist pair are computed on the basis of user and artist factors only.

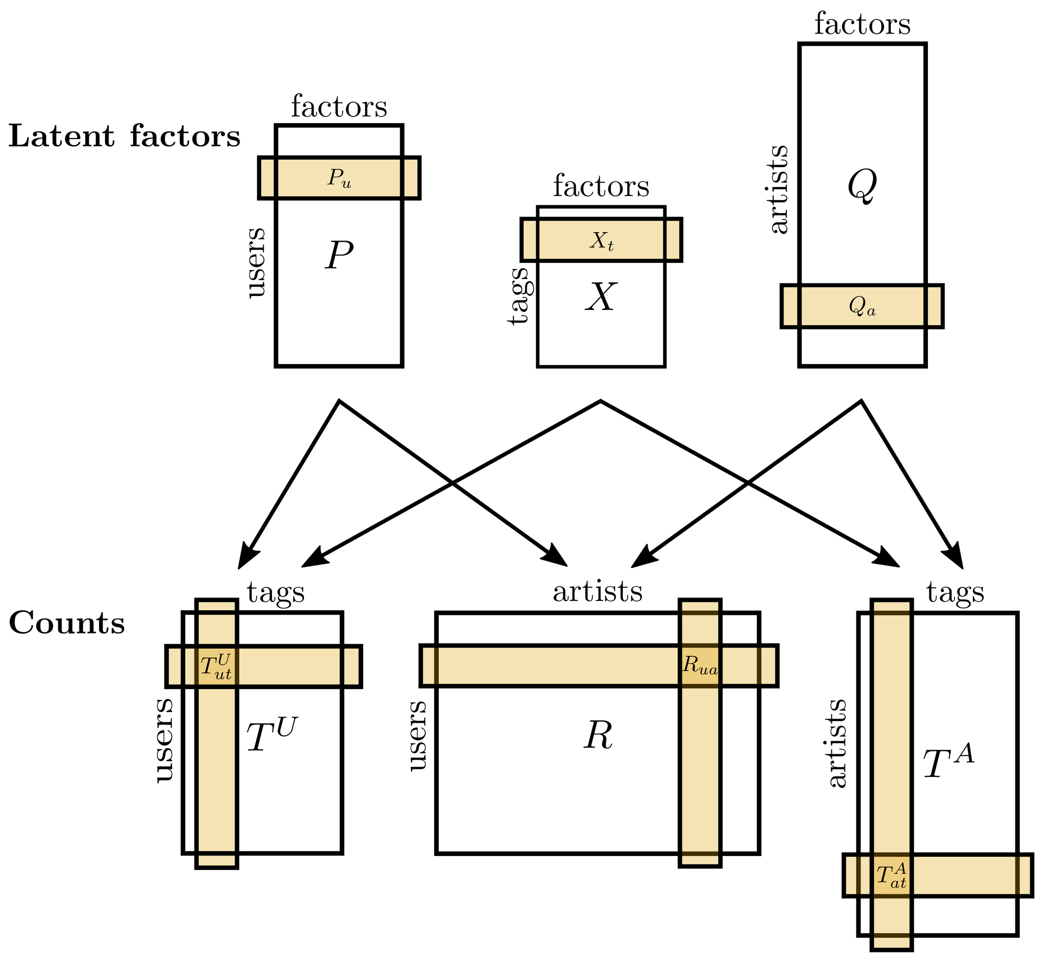

Cost function

$$J(P,Q,X) = \sum_{ua \in R} w(\alpha, R_{ua}) \left(\widetilde{R}_{ua} - P_u Q_a^T\right)^2 + \mu_1 \sum_{ut \in T^U} w(\beta, T^U_{ut}) \left(\widetilde{T}^U_{ut} - P_u X_t^T\right)^2 + \mu_2 \sum_{at \in T^A} w(\gamma, T^A_{at}) \left(\widetilde{T}^A_{at} - Q_a X_t^T\right)^2 + \lambda\left(\|P\|_F^2 + \|Q\|_F^2 + \|X\|_F^2\right)$$

Predicted preference

$$Z = PQ^T$$

» Numerical study

We compared our model combining implicit feedback and tags (i.e., *Hybrid*) with a plain collaborative filtering (*CF*) model for implicit feedback. We used a dataset of Last.fm listening events, top tags used by users and top tags applied to artists, collected through the Last.fm API [4]. For each observed user-artist pair in the dataset, we measured the ability of each model to identify the artist as relevant for the user. The user-artist pairs were assigned to subsets defined by the artist frequency in the dataset. We report the expected percentile rank computed by 5-fold cross validation for each model and each subset.

Dataset

- 2,902 users
- 71,223 artists
- 687,833 user-artist pairs
- 630 unique user tags
- 12,902 unique artist tags

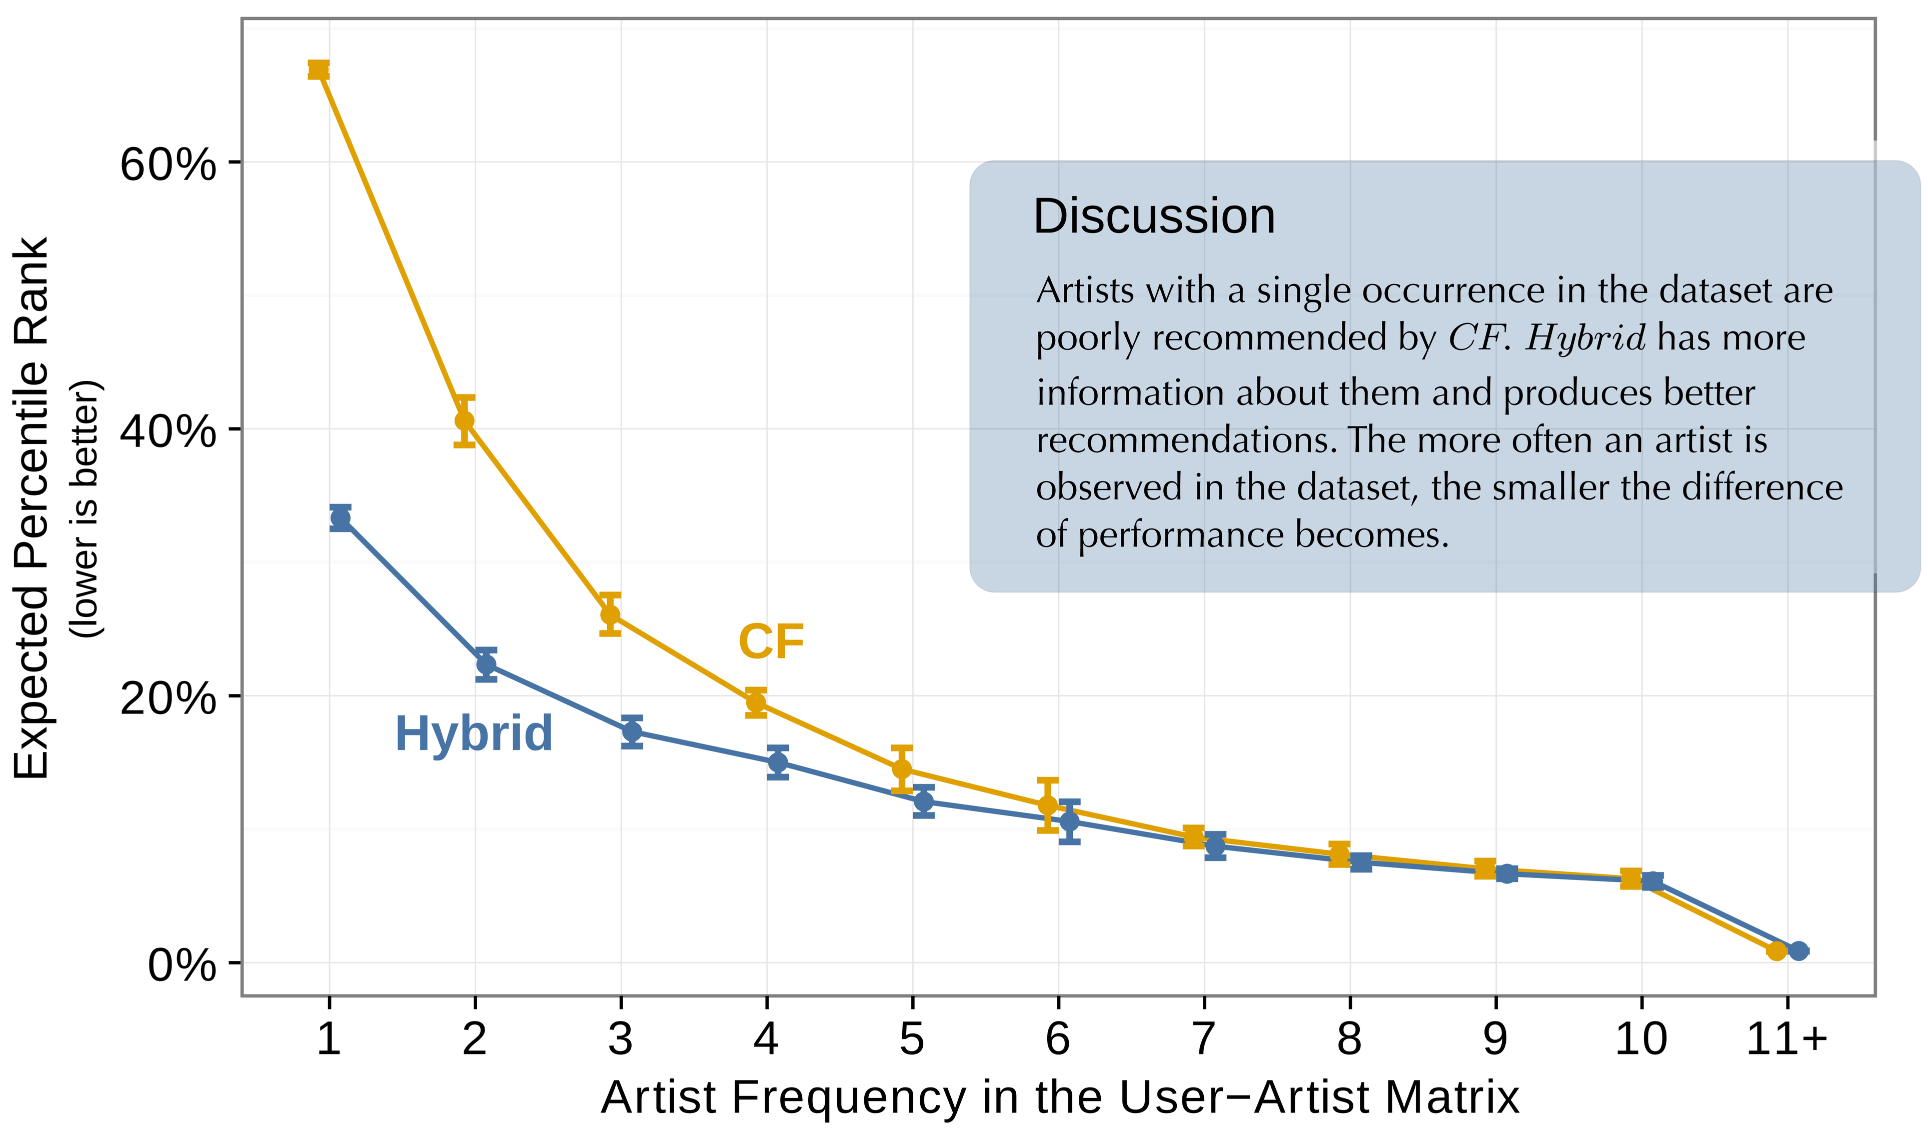

Discussion

Artists with a single occurrence in the dataset are poorly recommended by *CF*. *Hybrid* has more information about them and produces better recommendations. The more often an artist is observed in the dataset, the smaller the difference of performance becomes.

» Qualitative example

We trained *CF* and *Hybrid* using 80% of the data. The table shows the predicted expected preference Z_{ua} of a selected user u for four different artists she listened to, but belong to the 20% of data withheld for test. *CF* is not able to identify that the first and the fourth artists were interesting for the user, while *Hybrid* is able to do so. The intuition behind the model is depicted in the figure.

Artist Name	Z_{ua}^{CF}	Z_{ua}^{Hybrid}
Feliu Ventura	0.00	0.61
Joan Colomo	0.35	0.75
Manos de Topo	0.23	0.64
Mazoni	0.00	0.69

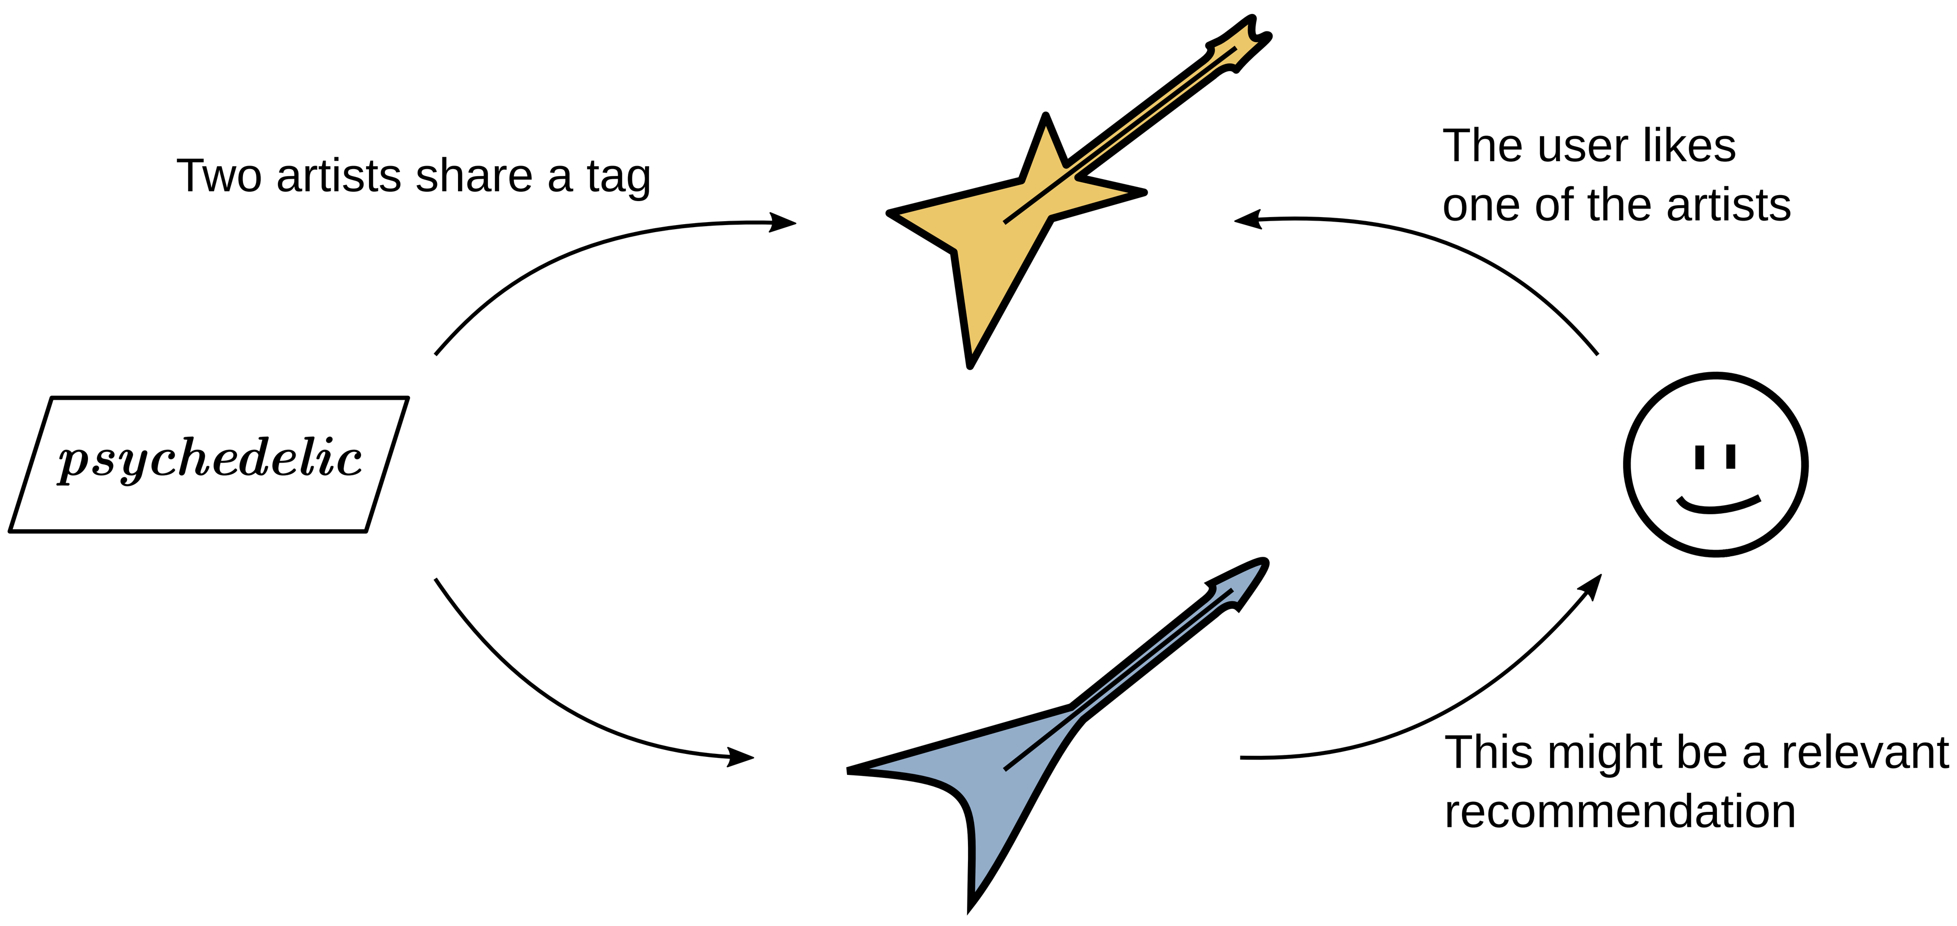

References

[1] Y. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In Proc. ICDM, pages 263–272. IEEE, 2008.
[2] A. Vall, M. Skowron, P. Knees, and M. Schedl. Improving music recommendations with a weighted factorization of the tagging activity. In Proc. ISMIR, 2015.
[3] A. Vall. Listener-inspired automated music playlist generation. In Proc. RecSys, 2015.
[4] www.last.fm/api

This research is supported by the Austrian Science Fund (FWF): P25655.

JOHANNES KEPLER UNIVERSITÄT LINZ
Department of Computational Perception