It Is Different When Items Are Older
Debiasing Recommendations When Selection Bias and User Preferences Are Dynamic

Jin Huang[1], Harrie Oosterhuis[2], Maarten de Rijke[1]

University of Amsterdam[1], Radboud University Nijmegen[2]

Introduction

User interactions with recommender systems (RSs) are affected by user selection bias, e.g., popularity bias and positivity bias. Methods exist for mitigating the effects of selection bias in user ratings on the evaluation and optimization of RSs. However, these methods treat selection bias as static, despite the fact that the popularity of an item may change drastically over time and the fact that user preferences may also change over time.

In this paper, we make a three-fold contribution:

- an analysis and estimation of dynamic selection bias and dynamic user preferences in the MovieLens-Latest-small dataset [1];
- DANCER: a general debiasing method that is adaptable for DebiAsing in the dyNamiC scEnaRio;
- time-aware matrix factorization (TMF)-DANCER: to our knowledge it is the first recommendation method that corrects for dynamic selection bias and models dynamic user preferences.

Estimation Ignoring Dynamic Bias

We consider a simple example $\mathcal{X}$ with one user u, one item i and two time periods t_1 and t_2. Due to dynamic user preferences and dynamic selection bias, the user ratings and observation probabilities are not constant over the different time periods: $y_{t_1} \neq y_{t_2}$, $p_{t_1} \neq p_{t_2}$. The ideal loss is the standard estimation of how well the predicted user preferences reflect the true user preferences based on all the ratings:

$$\mathcal{L}^{\mathcal{X}} = \frac{1}{2}\left(L(\hat{y}_{t_1}, y_{t_1}) + L(\hat{y}_{t_2}, y_{t_2})\right). \tag{1}$$

The widely-used debiasing method [2] uses inverse propensity scoring (IPS) to correct for the probability that the user rates the item at time t_1 or t_2: $p_{u,i} = p_{t_1} + (1 - p_{t_1})p_{t_2}$. If we consider the expected value of this *static* IPS estimator:

$$\mathbb{E}[\mathcal{L}^{\mathcal{X}}_{\text{staticIPS}}] = \frac{1}{2}\left(\frac{p_{t_1}}{p_{u,i}} L(\hat{y}_{t_1}, y_{t_1}) + \frac{p_{t_2}}{p_{u,i}} L(\hat{y}_{t_2}, y_{t_2})\right) \not\propto \mathcal{L}^{\mathcal{X}}. \tag{2}$$

It is not proportional to the true loss, because the *static* IPS estimation fails to address the problem that the user's rating at a time with a higher probability of being observed is more likely to be represented in logged data than at any other time.

DANCER: DebiAsing in the dyNamiC scEnaRio

With accurate propensities $p_{u,i,t}$, dynamic selection bias can be fully corrected by applying DANCER to inversely weight the evaluation of the predicted ratings:

$$\mathcal{L}_{\text{DANCER}} = \frac{1}{|\mathcal{U}| \cdot |\mathcal{I}| \cdot |\mathcal{T}|} \sum_{\substack{u,i,t:\\ o_{u,i,t}=1}} \frac{L(\hat{y}_{u,i,t}, y_{u,i,t})}{p_{u,i,t}}. \tag{3}$$

Unlike the *static* IPS approach with a *static* estimator $\mathcal{L}_{\text{staticIPS}}$, the proposed debiasing method $\mathcal{L}_{\text{DANCER}}$ is unbiased in the dynamic scenario:

$$\mathbb{E}[\mathcal{L}^{\mathcal{X}}_{\text{DANCER}}] = \frac{1}{2}\left(p_{t_1}\frac{L(\hat{y}_{t_1}, y_{t_1})}{p_{t_1}} + p_{t_2}\frac{L(\hat{y}_{t_2}, y_{t_2})}{p_{t_2}}\right) = \mathcal{L}^{\mathcal{X}}, \tag{4}$$

Because DANCER utilizes propensities that vary per time period t, it can correct for dynamic effects of bias that the existing static IPS estimators cannot.

TMF-DANCER:While DANCER is not model specific, we will apply it to a time-aware matrix factorization (TMF) model that accounts for temporal effects. Given an observed rating $y_{u,i,t}$ from user u on item i at time t, TMF computes the predicted rating $\hat{y}_{u,i,t}$ as: $\hat{y}_{u,i,t} = \boldsymbol{p}_u^T \boldsymbol{q}_i + b_u + b_i + b + b_t$, where the $\boldsymbol{p}_u \in \mathbb{R}^d$ and $\boldsymbol{q}_i \in \mathbb{R}^d$ are embedding vectors of user u and item i, and $b_u \in \mathbb{R}$, $b_i \in \mathbb{R}$, and $b \in \mathbb{R}$ are user, item and global offsets, respectively. Under this model, the proposed TMF-DANCER is optimized by minimizing the following loss:

$$\arg\min_{\boldsymbol{P},\boldsymbol{Q},\boldsymbol{B}} \left[\sum_{u,i,t:o_{u,i,t}=1} \frac{\delta(\hat{y}_{u,i,t}, y_{u,i,t})}{p_{u,i,t}} + \lambda\left(\|\boldsymbol{P}\|_F^2 + \|\boldsymbol{Q}\|_F^2 + \|\boldsymbol{B}\|_F^2\right)\right], \tag{5}$$

where $\boldsymbol{P}$, $\boldsymbol{Q}$ and $\boldsymbol{B}$ denote the embeddings of all users, all items and all the offset terms, respectively; δ is the mean square error (MSE) loss function.

RQ1: Is Selection Bias Dynamic? – Yes!

Table: Performance in observation prediction.

Method	Random NLL		Random PPL		Time-Based NLL		Time-Based PPL	
Constant	0.0973		1.1022		0.0337		1.0343	
Pop	0.0890		1.0931		0.0404		1.0412	
MF	0.0697	(0.0015)	1.0722	(0.0016)	0.0271	(0.0000)	1.0275	(0.0000)
T-Pop	0.1234		1.1314		0.0523		1.0537	
TMF	0.0658†	(0.0001)	1.0680†	(0.0001)	**0.0267**†	(0.0000)	**1.0271**†	(0.0000)
TTF++	0.0632†	(0.0002)	1.0653†	(0.0002)	0.0268†	(0.0001)	1.0271†	(0.0001)
TMTF	**0.0621**†	(0.0001)	**1.0641**†	(0.0001)	0.0268†	(0.0000)	1.0272†	(0.0000)

Time-aware methods can better predict selection bias than static methods.

RQ2: Are User Preferences Dynamic? – Yes!

Table: Performance comparison of different methods in predicting ratings.

Method	Observed MSE ↓	Observed MAE ↓	Observed ACC↑	Debiased SNIPS-MSE↓	Debiased SNIPS-MAE↓	Debiased SNIPS-ACC↑
Avg	0.9535	0.7540	0.2241	1.1436	0.8360	0.2048
MF	0.7551	0.6679	0.2515	1.2911	0.8985	0.1829
T-Avg	1.0850	0.7974	0.2181	1.3105	0.8865	0.1955
TMF	0.7505	0.6656	0.2525	1.1210†	0.8383†	0.1944†
TTF	1.1515	0.8187	0.2120	1.8834	1.0879	0.1504
TTF++	0.7526	0.6645†	**0.2552**†	1.0839†	0.8067†	**0.2134**†
TMTF	**0.7503**†	**0.6637**†	0.2533†	**1.0727**†	**0.8026**†	0.2127†

Most time-aware methods outperform time-unaware methods, and therefore, we answer **RQ2** in the affirmative: item-age has a significant effect on user preferences.

RQ3: Can TMF-DANCER Better Mitigate Dynamic Selection Bias? – Yes!

Table: Performance of TMF-DANCER compared with different methods.

Method	MSE↓		MAE↓		ACC↑	
Avg	0.3155		0.4321		0.3623	
T-Avg	0.3280		0.4326		0.3614	
MF	0.1811	(0.0030)	0.3314	(0.0028)	0.4680	(0.0040)
TMF	0.1338	(0.0019)	0.2818	(0.0022)	0.5396	(0.0038)
MF-StaticIPS	0.1879	(0.0035)	0.3377	(0.0032)	0.4598	(0.0044)
TMF-StaticIPS	0.1086	(0.0021)	0.2491	(0.0027)	0.6065	(0.0057)
MF-DANCER	0.1533	(0.0016)	0.3032	(0.0017)	0.5074	(0.0023)
TMF-DANCER	**0.1045**†	(0.0014)	**0.2444**†	(0.0018)	**0.6151**†	(0.0039)

We observe that:

- The time-based methods outperform their static counterparts: TMF $\succ$ MF, TMF-StaticIPS $\succ$ MF-StaticIPS, and TMF-DANCER $\succ$ MF-DANCER.
- The debiased methods increase performance: MF-DANCER $\succ$ MF and TMF-DANCER $\succ$ TMF-StaticIPS $\succ$ TMF.
- The best performing method is TMF-DANCER.

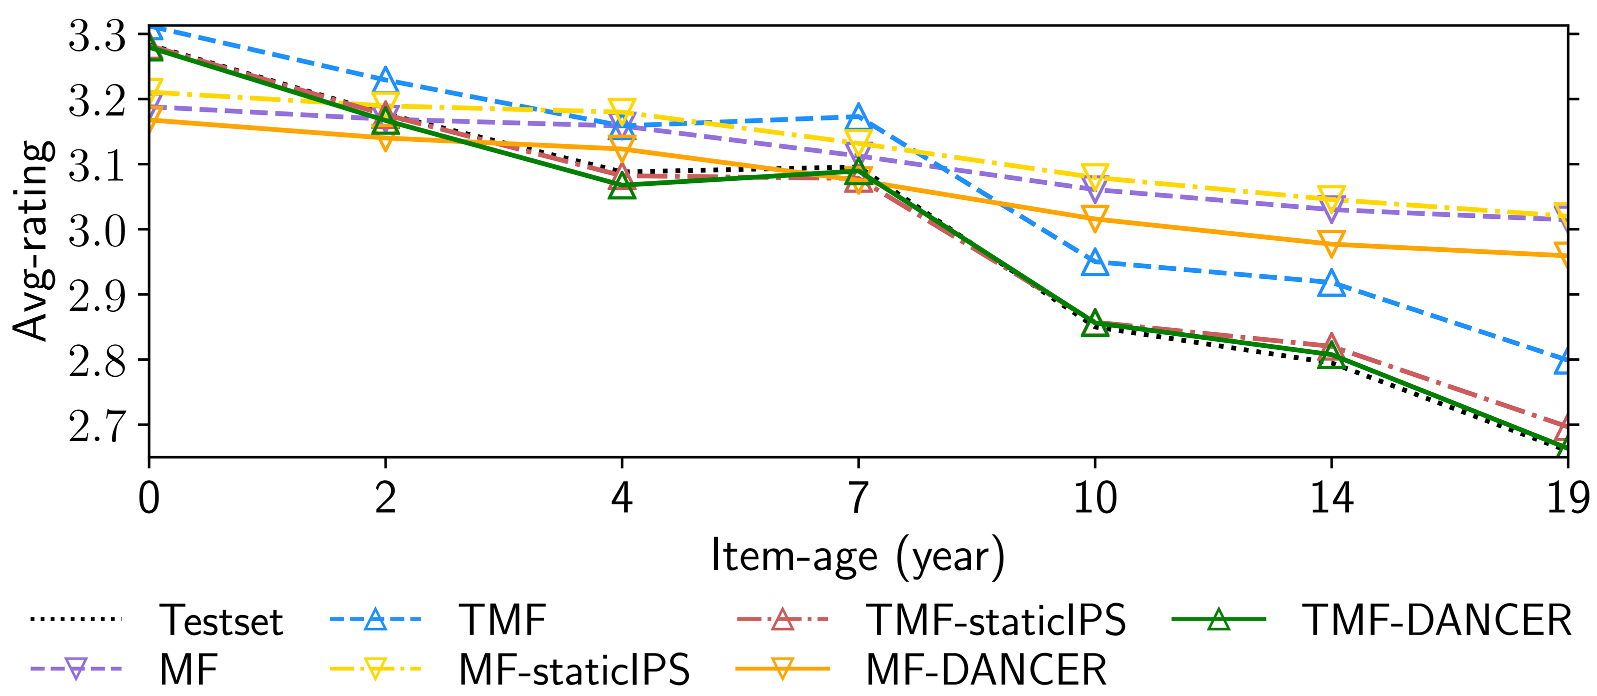

Figure: Average rating on items predicted by different models over the item-age.

We observe:

- The MF methods are unable to model changes in ratings as items get older.
- The TMF methods better capture the overall trend.
- TMF consistently overestimates ratings; TMF-staticIPS reduces overestimation by correcting for static bias; the overestimation becomes worse for older items.
- TMF-DANCER approximates the actual average rating at each item-age; its accuracy is quite consistent over time.

Conclusion

In real-world MovieLens dataset [1], the dynamic scenario is real: selection bias and user preferences are dynamic. The proposed DANCER debiasing method that takes into account the dynamic aspects of bias and user preferences is unbiased in the dynamic scenario. The dynamic scenario have implications for state-of-the-art recommendation methods, as they are strongly affected by dynamic selection bias. With the DANCER debiasing method, RSs can now be expanded to deal with dynamic scenarios.

References

[1] F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, 5(4):1–19, 2015.

[2] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *ICML*, pages 1670–1679, 2016.

University of Amsterdam