The Performance of Total Variation Regularizer for Recommender Systems

1. Background

- Simple recommender systems are shown to perform better than state-of-the-art recommender systems on some metrics [3].
- Collaborative filtering is a simple, widely adopted technique for recommender systems [1].
- The collaborative filtering problem can be solved by interpolating missing user ratings over user similarity graphs.
- Total variation is a graph regularizer applied in the field of graph signal processing, that can interpolate missing values over graphs [2].

2. Research Question

"How does the total variation regularizer perform for user k-nearest neighbours collaborative filtering?"

3. User Graph Collaborative Filtering

- **Collaborative filtering** predicts user-movie ratings using the ratings of their most similar users [4].
- We use **Pearson Correlation coefficient** to calculate user similarities [4].
- User similarities are used to construct the **user similarity graph**, displayed in Figure 1.
- **Total variation** interpolates the missing ratings over the graph by solving equation (1) [6].

Equation (1): Total Variation Regularizer [6]

$$\hat{\mathbf{x}}_i = \min_{\mathbf{x}_i \in \mathbb{R}^U} \|\mathbf{y}_i - \mathbf{C}\mathbf{x}_i\|_2^2 + \mu \|\mathbf{x}_i - \mathbf{A}^{norm}\mathbf{x}_i\|_1$$

Where:

- $\hat{X}$ - Estimated ratings for movie *i*.
- U - Users.
- y_i - Observed (training) ratings for movie *i*.
- μ - Trade-off factor.
- A^{norm} - Normalized user similarity matrix.
- x_i - Rating predictions for movie *i*.
- C - Binary selection matrix which masks the unknown values.

Contact Details

Karolis Mariunas: k.mariunas@student.tudelft.nl
Responsible professor: Elvin Isufi
Supervisors: Maosheng Yang, Bishwadeep Das

4. Methodology

Implement Collaborative Filtering with Total Variation.

- Construct user similarity graphs (Figure 1).
- Solve total variation equation (1) for each item.

Measure Performance on MovieLens 100k data set [5]

- **Accuracy**: Root Mean Squared Error (RMSE).
- **Top-n recommendations:** Recall@n (REC@n) and Precision@n (Prec@n).

Experiment and Compare

- Graph construction - directed vs undirected similarity graphs.
- Accuracy vs top-n recommendation performance.
- Compare to the performance of traditional collaborative filtering (UserKNN).

Figure 1: User Similarity Graphs

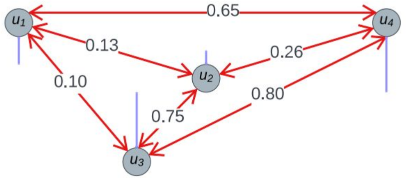

(a) Fully connected similarity graph.

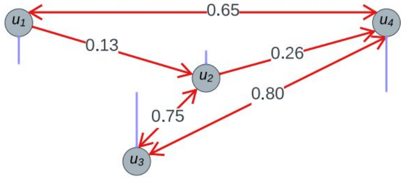

(b) Filtered k-nearest neighbours directed similarity graph.

Where:

- **Graph nodes** are users.
- **Edges** contain the connected **user similarity** and represent user neighbourhood.
- The nodes' blue lines indicate user **ratings** (signals).

5. Results

	RMSE	PREC@5	REC@5	PREC@10	REC@10	PREC@20	REC@20
UserMP	1.042	0.542	0.445	0.542	0.649	0.541	0.815
UserKNN^{RMSE}	1.005	0.679	0.511	0.631	0.709	**0.590**	**0.860**
UserKNN^{PREC}	1.018	**0.698**	**0.522**	**0.643**	**0.718**	**0.595**	**0.865**
UserTV$^{RMSE}_{undir}$	**0.960**	0.683	**0.518**	0.634	**0.713**	**0.591**	**0.862**
UserTV$^{RMSE}_{dir}$	**0.958**	0.681	0.516	0.633	**0.713**	**0.591**	**0.862**
UserTV$^{PREC}_{undir}$	**0.961**	0.683	**0.518**	0.634	**0.713**	**0.591**	**0.862**
UserTV$^{PREC}_{dir}$	**0.960**	0.681	**0.517**	0.632	**0.712**	**0.590**	**0.862**

- **Total variation** (UserTV) **outperforms** collaborative filtering (UserKNN) on **RMSE** by **4.68%.**
- **UserKNN** performs **better** for top-n recommendations, increasing **precision@5** and **precision@10** by **2.20%** and **1.42%,** respectively.

6. Conclusion

- Total variation performs **similarly** to collaborative filtering, with **improvements** in RMSE and a **decrease** in precision.
- However, the **performance** of total variation **increases** for bigger recommendation lists.
- Total variation indicates a **consistent balance** between RMSE, precision and recall.

7. Future Work

- Extensive research into user similarity calculations and their effects on total variation.
- Investigation into item-specific similarity graphs and their effects on total variation performance.
- Measurements of diversity and novelty of the recommendations provided by total variation.

References

[1] A. N. Nikolakopoulos, X. Ning, C. Desrosiers, and G. Karypis, Trust Your Neighbors: A Comprehensive Survey of Neighborhood-Based Methods for Recommender Systems, F. Ricci, L. Rokach, and B. Shapira, Eds. New York, NY: Springer US, 2022.
[2] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," IEEE signal processing magazine, vol. 30, no. 3, pp. 83–98, 2013.
[3] M. Ferrari Dacrema, P. Cremonesi, and D. Jannach, "Are we really making much progress? a worrying analysis of recent neural recommendation approaches," in Proceedings of the 13th ACM Conference on Recommender Systems, ser. RecSys '19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 101–109.
[4] C. C. Aggarwal, Recommender Systems, 1st ed. Springer, 2016.
[5] F. M. Harper and J. A. Konstan, "The movielens datasets: History and context," ACM Transactions on Interactive Intelligent Systems, vol. 5, no. 4, pp. 1–19, Dec. 2015.
[6] E. Isufi, B. Das, A. Natali, M. Yang, and M. Sabbaqi, "Graph filters for processing and learning from network data," Delft University of Technology, pp. 1–19, Sep. 2021.