Total Variation Regularisation for Item KNN Collaborative Filtering: Performance Analysis

Lars van Blokland
Supervisors: Elvin Isufi, Maosheng Yang
EEMCS, Delft University of Technology, The Netherlands

TU Delft

23 june 2022
L.W.vanBlokland@student.tudelft.nl

Research Objective

- The goal of this research is to evaluate the performance of the Total Variation regulariser when applied to item KNN collaborative filtering.
- Results produced using Total Variation are compared to results made by baseline recommender systems.

Background

- Recommender systems are algorithms that help users navigate the countless items available to them online.
- Ratings are used to find items similar to those rated highly by users.
- Predicting unknown ratings allows for the discovery of items that users might enjoy.
- Regularisers are a technique used in graph signal processing to reconstruct noisy data.
- By representing the data as a graph, the Total Variation regulariser is able to predict unknown ratings.

Method

- Ratings from the MovieLens 100K dataset are used to create a Normalized Adjacency matrix.
- This matrix represents a graph, with non-zero values indicating an edge between two nodes.
- The Total Variation function is minimized by an iterative solver, using the graph and ratings as input.

$$\min_{\mathbf{x}_i \varepsilon R^U} \|\mathbf{y}_i - \mathbf{x}_i\|_2^2 + \mu \|\mathbf{x}_i - \mathbf{A}_n \mathbf{x}_i\|_1$$

- The output of this process is a matrix of predicted ratings.
- The accuracy of these ratings is scored using metrics, and compared with the baselines.

Results

Metric scores for Total Variation and recommender baselines

	RMSE	Rec@5	Rec@10	Rec@20	Prec@5	Prec@10	Prec@20	NDCG@5	NDCG@10	NDCG@20
Item KNN CF	1.077	0.468	**0.677**	0.832	0.582	0.573	0.556	**0.841**	**0.872**	**0.900**
User KNN CF	1.035	**0.515**	**0.711**	**0.863**	**0.671**	**0.628**	**0.589**	**0.882**	**0.906**	**0.927**
Item Mean	1.034	**0.472**	0.673	**0.835**	0.582	0.569	0.558	0.805	0.845	0.880
User Mean	**1.011**	**0.522**	**0.716**	**0.865**	**0.683**	**0.634**	**0.591**	**0.887**	**0.910**	**0.930**
Item TV CF	1.025	0.468	0.676	0.832	0.582	0.580	0.561	0.823	0.867	0.891

Table 1: Performance metric scores achieved by the Item Total Variation Collaborative Filter (Item TV CF), the Item and User Collaborative Filters (Item KNN CF, User KNN CF) and the Item and User means. Both the TV and KNN Collaborative Filters had a K of 10. The TV Collaborative Filter had a μ of 1.0. Results in bold are better than those of the TV Collaborative Filter.

Behavior of the Total Variation method for different values of its input parameters K and μ.

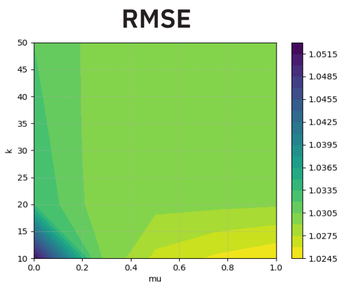

Figure 1

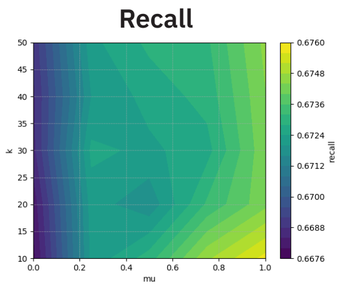

Figure 2

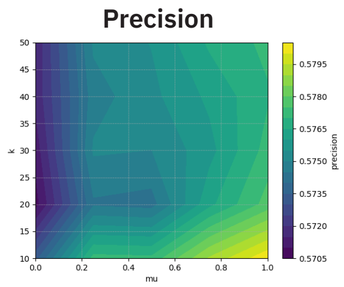

Figure 3

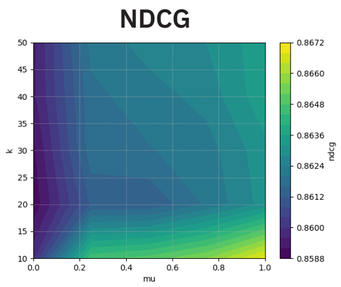

Figure 4

Performance of the Total Variation CF for RMSE (figure 1), Recall (2), Precision (figure 3) and NDCG (figure 3). Yellow colour hues indicate a better performance, whilst blue hues indicate a worse performance. K determines the number of neighbours each node has, whilst μ decides how much influence these neighbours have.

Conclusions

- For the MovieLens 100K dataset, Total Variation is outperformed by simple baselines.
- In particular, Total Variation has poor performance in metrics that measure the accuracy of important ratings.
- Scores of Item-based recommenders were worse than those of User-based methods, indicating a bias in the dataset, as shown by Figure 5.

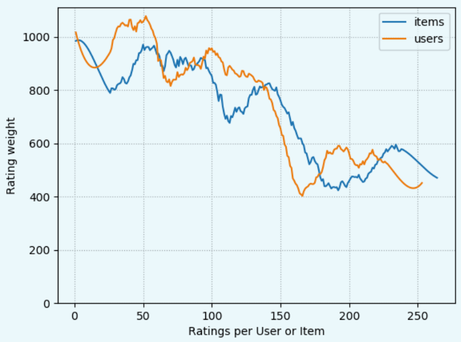

Figure 5: Comparison of the distribution of rating weight of users and items for the filtered ML100K dataset. Rating weight is described as the number of ratings given to a specific user or item, multiplied by how many users or items received that number of ratings.

Future Work

- Use multiple datasets to avoid potential bias introduced by inherent dataset rating distribution.
- Comparison with a wider variety of baseline recommenders.
- Experimentation with different setups for graph creation, such as other similarity measures.