Combining Rating and Review Data by Initializing Latent Factor Models with Topic Models for Top-N Recommendation

Francisco J. Peña, Diarmuid O'Reilly-Morgan, Elias Z. Tragos, Neil Hurley, Erika Duriakova, Barry Smyth, and Aonghus Lawlor
{first.last}@insight-centre.org
Insight Centre for Data Analytics,
University College Dublin
Dublin, Ireland

ABSTRACT

Nowadays we commonly have multiple sources of data associated with items. Users may provide numerical ratings, or implicit interactions, but may also provide textual reviews. Although many algorithms have been proposed to jointly learn a model over both interactions and textual data, there is room to improve the many factorization models that are proven to work well on interactions data, but are not designed to exploit textual information. Our focus in this work is to propose a simple, yet easily applicable and effective, method to incorporate review data into such factorization models. In particular, we propose to build the user and item embeddings within the topic space of a topic model learned from the review data. This has several advantages: we observe that initializing the user and item embeddings in topic space leads to faster convergence of the factorization algorithm to a model that out-performs models initialized randomly, or with other state-of-the-art initialization strategies. Moreover, constraining user and item factors to topic space allows for the learning of an interpretable model that users can visualise.

CCS CONCEPTS

• Computing methodologies → Topic modeling; **• Information systems → Recommender systems**.

ACM Reference Format:
Francisco J. Peña, Diarmuid O'Reilly-Morgan, Elias Z. Tragos, Neil Hurley, Erika Duriakova, Barry Smyth, and Aonghus Lawlor. 2020. Combining Rating and Review Data by Initializing Latent Factor Models with Topic Models for Top-N Recommendation. In *Fourteenth ACM Conference on Recommender Systems (RecSys '20), September 22–26, 2020, Virtual Event, Brazil.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3383313.3412207

1 INTRODUCTION

Recommender systems (RS) have presented themselves as powerful tools to help users make the right choice. The first RSs were trained on explicit rating data provided by users for items, and learn a model to predict the ratings of unrated items. Later, research focused on the Top-N recommendation problem, learning to predict the set of N items that are most likely to satisfy a user's need. Matrix factorization (MF) models have proven highly effective on this task. Such models [18, 19, 21, 22, 32] learn latent space embeddings of users and items either from explicit ratings, implicitly gathered interaction data or (recently) text data (i.e. user reviews).

Incorporating review data in RS models has proved to improve recommendation performance [7, 29]. In the area of text analytics, topic modelling algorithms are used to find structure in textual information, and then cluster the data into meaningful topics. Several approaches such as Collaborative Topic Regression (CTR) [33], Hidden Factors as Topics (HFT) [27], Ratings Meet Reviews (RMR) [25] and JMARS [10] strive to produce topic models and rating predictions by optimising a hybrid loss function reducing the rating error and maximising the corpus likelihood. The topic models learned in these works are probabilistic generative models. Our work differs from the above, as we focus on initialisation using topic models. Moreover, we forego joint learning, in favour of the flexibility of developing a methodology that can be applied to any existing latent factor RS model that learns user and item embeddings. Firstly we learn a topic model over the review data, extracting a topic space in which documents and words are embedded. Then, initialising the user and item factors of the MF problem within this topic space, we optimise these initial embeddings by minimising the loss function over the interaction data. Our proposed model outperforms a number of state-of-the-art initialisation strategies, yielding more accurate RS models as evaluated on standard datasets. The fact that user topics describe user preferences and item topics describe item qualities helps the algorithm to achieve a high prediction accuracy in a small number of iterations. The use of a data source different to the rating matrix during the initialisation process helps the algorithm to mitigate the problem of local minima and ultimately reach higher prediction accuracy at convergence.

Overall our main contributions are the following:

(1) We use the results of topic modelling to initialise the latent factors of three well-known RS algorithms.
(2) We show that our model provides better performance against a number of state-of-the-art methods that initialise the latent factors with other strategies.
(3) We show how we can obtain an interpretable model that users can visualise, at a cost of sacrificing prediction accuracy, but still in most cases with better performance than a randomly initialised model.

RecSys '20, September 22–26, 2020, Virtual Event, Brazil

ACM ISBN 978-1-4503-7583-2/20/09…$15.00
https://doi.org/10.1145/3383313.3412207

2 BACKGROUND

2.1 Initialization In Latent Factor Models

Latent factor models have proven to be very successful for both predicting user ratings and proposing Top-N recommendations. One of the first models designed specifically to produce Top-N recommendations is Weighted Regularized Matrix Facrtorization (WRMF) [18], which converts explicit feedback ratings into implicit feedback by binarising the ratings into a preference ρ and then assigns a confidence value c to that preference. Preferences are predicted by a dot multiplication between the user latent vector $\mathbf{p}_u$ and the item latent vector $\mathbf{q}_i$. Bayesian Personalized Ranking (BPR) [32] focuses on making a pairwise ordering between items in which seen items should always be ranked above unseen items for every user. Rank-SGD (Algorithm 1 in [19]) includes the actual scores in the pairwise loss function. These approaches have proven to work well for Top-N recommendations, however they do not consider user reviews.

MF models are traditionally initialised with random values [14, 23]. However, their performance can be improved if more sophisticated initialisation strategies are used. The two common goals of initialisation are: (i) achieve faster convergence and (ii) reach better performance. In this paper, we focus more on the latter objective, but we can also achieve better performance in fewer iterations than other algorithms. One well-known initialisation approach is NNDSVD [4] that is used for Non-negative Matrix Factorization (NMF). NNDSVD uses two Singular Value Decomposition (SVD) processes that are deterministic in order to find initialisation values for the latent factor matrices. Hidasi and Tikk present SimFactor, an initialisation method for Alternating Least Squares (ALS) that works with implicit feedback datasets [13, 14] and is based on the similarity between users and items. More recently, Nasiri and Minaei presented an initialisation method that completes the missing entries from the sparse rating matrix using user and item averages, followed by factorising the rating matrix with SVD [28], we call this method Average SVD. Our initialisation strategy is different from the above that use the rating matrix both for initialisation and model training, because for initialisation we use topic models extracted from reviews, and we exploit the rating matrix only at the algorithm training step. The only exception being [14] that uses tags and contextual information in order to build similarity matrices.

2.2 Topic Modelling

Topic modelling is an information retrieval technique that aims to find a latent semantic structure between terms based on their co-occurrence within documents without relying on any form of labelled data [15]. In topic models, terms are grouped together into topics that typically represent a concept or a theme and topics are grouped into documents. Topic modelling algorithms use a document-term frequency matrix in order to create topic models. Well known algorithms include Probabilistic Latent Semantic Analysis (pLSA) [15] and Latent Dirichlet Allocation (LDA) [3] which are probabilistic generative models. NMF can also be used to decompose the document-term matrix and produce topic models as it is done in [24] and an ensemble of NMF models is presented in [2].

Topic based recommenders have been a popular approach to mix textual reviews with ratings. Some approaches derive from LDA [3] and jointly learn the topic model and the latent factors matrices for rating prediction using a probabilistic generative model [10, 25, 27, 33]. McAuley and Leskovec present the HFT [27] model that learns alternating between minimising the prediction error in a step and then maximising the log likelihood of the corpus in the next step. HFT uses a transformation function to relate the latent factors with the topics. RMR is presented in [25] and uses a mixture of Gaussians to model the ratings assuming that the mixture proportion has the same distribution as the topic distribution. In this way, the need for a transformation function is also avoided. Diao et al. present JMARS [10] an unsupervised model that mines aspects from movie reviews using topic modelling and integrates the mined aspects into the recommendation engine. In [8, 9] aspect-aware model that correlates the user and item embeddings on a set of aspects obtained from the reviews is presented. Hou et al. introduces a model called AMF [16] that is built on top of the ALFM model [31], but differently to ALFM, AMF pre-trains the aspects matrix by using LDA, and once the topic model has been created it uses it as part of the model to learn the latent factor matrices that serve to predict ratings.

In recommender systems in general, there seems to be a gap: models that are used for Top-N recommendation do not incorporate information from reviews [18, 19, 32] and models that include reviews are mostly designed for rating prediction [5, 6, 10, 25–27, 33, 35], Joint Representation Learning (JRL) [34] being one of the few exceptions. We want to address this gap by improving the existing Top-N recommenders and adding reviews information into those models.

3 APPROACH

Our approach involves the following steps: (i) learn a topic model from review data; and (ii) initialise the user and item factors of a latent space recommender model in the topic space learned at step (i). Then, using the rating data, we optimise the user and item factors, by running an SGD to minimise the top-N loss function. We detail each step below.

Learning the topic model. Our datasets include for each item a set of user reviews, which constitute the "documents" over which the topic model is run. First, we mine features from the reviews. Similarly to [11, 29, 30], we create a bag-of-words using a vocabulary consisting only of the nouns in the reviews, obtaining a TF-IDF matrix $\mathbf{T}$ of size $|W| \times |R|$, where $|W|$ is the size of the vocabulary and $|R|$ is the number of reviews. For example, in a dataset of hotel reviews, the nouns might be words such as *swimming pool*, *bedroom*, *cleanliness* that capture a different aspect of a hotel. Then, using a topic modelling algorithm (i.e. LDA [3], NMF [24] or topic ensembling [2]) we obtain a $|W| \times k$ matrix, $\mathbf{H}$, representing an embedding of terms into a k-dimensional topic space, where k is the dimensional representation of each term in the vocabulary. To map users into topic space, we group all of the reviews written by a user into a single document and thus generate the term frequency matrix $\mathbf{T}_U$ of size $|U| \times |W|$, where $|U|$ is the number of users. $\mathbf{T}_U$ is a TF-IDF matrix if NMF or topic ensembling is used and a word-count matrix if LDA is used. Then, the user documents are "folded" into the topic model by applying a projection to $\mathbf{T}_U$ [2], $\mathbf{A} = \mathbf{T}_U \cdot \mathbf{H}$.

In a similar manner, all reviews associated with an item can be used to fold the item into topic space, using the $|I| \times |W|$ TF-IDF matrix $\mathbf{T}_I$, $\mathbf{B} = \mathbf{T}_I \cdot \mathbf{H}$, where $|I|$ is the number of items. Each row of $\mathbf{A}$ now corresponds to a user, with columns corresponding to the k topics. An entry $\mathbf{A}_{u,t}$ indicates the strength of association of user u to the topic t. The same applies to the item-topic matrix $\mathbf{B}$.

Topic Initialised Latent Factor Model. Given a latent space dimension f, the goal of an MF recommendation model is to find a vector $\mathbf{p}_u \in \mathbb{R}^f$ for each user u and a vector $\mathbf{q}_i \in \mathbb{R}^f$ for each item i such that a prediction $\hat{y}(u,i)$, for a given (u,i) pair can be obtained from the inner product $\mathbf{p}_u^\top \mathbf{q}_i$. Gathering user vectors and item vectors into the $|U| \times f$ matrix $\mathbf{P}$ and the $|I| \times f$ matrix $\mathbf{Q}$, respectively, we can associate topic space with the latent space of the ratings factorisation problem by setting $f = k$ and initialising $\mathbf{P} = \mathbf{A}$ or $\mathbf{Q} = \mathbf{B}$. From this initialisation in topic space, the MF model further optimises the factors, using the rating data. Our method can be applied to any model that employs latent factors. In fact, we will evaluate it on three such models, WRMF [18], BPR [32] and Rank-SGD [19].

4 EVALUATION

Datasets. We have selected four datasets from different domains to conduct our experiments[1 2]. For all datasets we executed a preprocessing step in which we removed reviews that were repeated, that were not in English language and that had missing or erroneous IDs. Subsequently, we performed part-of-speech tagging and lemmatisation for each one of the reviews. We perform lemmatisation in order to group together words that are syntactic variants of the same base word. In the final step of the preprocessing task, we built bag of words for the reviews. Given that we can not create a user-topics (or item-topics) matrix without any reviews, we remove from the test set users and items that are not present in the train set. Table 1 provides a summary of the datasets after the data preprocessing task has been executed.

Dataset	Records	Users	Items	Sparsity
Amazon Toys & Games	154,290	17,898	11,635	0.9993
Amazon Pet Supplies	147,385	18,645	8,395	0.9991
Amazon Health & Personal Care	323,553	36,432	17,996	0.9995
TripAdvisor Hotels	620,172	429,928	3,828	0.9996

Table 1: Description of the datasets.

Baselines. We compare our methodology against random, NNDSVD [4] and Average SVD [28] initialisations. We experiment initialising the following algorithms: (a) BPR[32], (b) Rank-SGD[19], and (c) WRMF[18]. Our experiments are focused on demonstrating the improvements that our initialisation method has on latent factor models, leaving as future work the integration of reviews and ratings in the optimisation function.

Evaluation Methodology. To evaluate our model we split the data three-ways in chronological order using a 80-10-10 ratio for training, validation and testing. The oldest 80% of the records is used to train the algorithms. The hyperparameter values are selected based on the Recall@10 performance on the validation set for each individual algorithm separately. The results show the performance of the algorithms on the test set, selecting the best run across all iterations. The topic model is trained on the train data and then is used as an input to our model. To train the model we use 10 negative samples for each positive one as described in [12, 17]. We follow the evaluation protocol used in [12, 17], where for each positive item associated with a target user in the test set, we randomly sample 50 negative items that have no interaction records with the user. We report Recall, Hit Ratio (HT), NDCG and Precision at rank N as the evaluation metrics for measuring the model's accuracy [8, 20]. We initialised each of the algorithms BPR, Rank-SGD and WRMF with topic models created using NMF. We chose this combination because all methods had slightly better performance with NMF than with LDA and topic ensembling.

Impact Of The Initialisation Strategy. In this section we compare Topic Initialized latent factors Model (TIM) against NNDSVD [4], one of most widely used methods for NMF initialisation, and the more recent Average SVD [28] method. We used `scikit-learn`'s version of NNDSVD and our own implementation of Average SVD. Figure 1 shows the Recall@10 performance improvement of BPR, Rank-SGD and WRMF using the initialisation methods compared to random initialisation across the four datasets. It is evident that, TIM consistently improves the performance for all algorithms achieving the best Recall@10 in all cases. These results suggest that our algorithm is model agnostic.

Convergence Analysis. We compare the four initialisation strategies used in combination with BPR, Rank-SGD and WRMF [33]. Table 2 displays the Recall@10 after 1, 10 and 100 iterations of each algorithm. We observe that NNDSVD outperforms TIM on the first epochs in a couple of cases, although not by much. This is expected since it is well known that NNDSVD has a fast convergence [1, 4, 23]. TIM performs the best with WRMF after 100 iterations across all of the datasets, the same happens with BPR and Rank-SGD after one iteration. To analyse the performance at convergence time we executed each algorithm for 500 iterations. Table 3 reports the best result over the 500 iterations. Notice that TIM has the best performance across all of the datasets. In summary TIM is able to achieve high accuracy both after a few and after many epochs. If one is looking for a quick solution then both TIM and NNDSVD perform well. However, for best performance which needs larger number of epochs, TIM demonstrated the ability to find the best solution.

Analysis Of The Influence Of The Number of Latent Factors On The Performance. In the following charts we can see how the number of latent factors has influence on the performance. In both datasets, the higher the number of topics the higher the recall, with 40 latent factors reaching the best performance. We remind the reader that the number of topics is also equal to the number of latent factors.

Interpretability. There are situations in which it is preferable to sacrifice prediction over interpretability, i.e. when explaining to a user why to stay in a certain hotel. NMF models are preferred over other models like SVD because its non-negativity allows to map the factor vectors to conceptual properties of the data [4, 23]. In the previous experiments, we initialised both user and item latent

[1] https://www.cs.cmu.edu/~jiweil/html/hotel-review.html
[2] http://jmcauley.ucsd.edu/data/amazon/

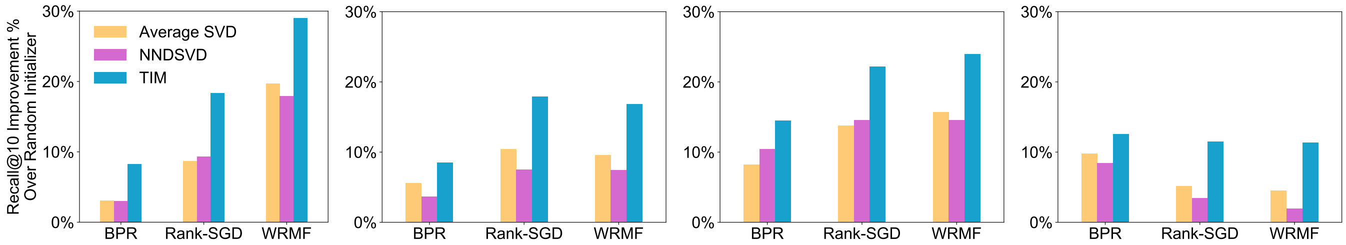

(a) Amazon Toys & Games **(b) Amazon Pet Supplies** **(c) Amazon Health** **(d) TripAdvisor Hotels**

Figure 1: Recall@10 percentage improvement of different initialisation models compared to random across multiple datasets

Algorithm	Initialization	Amazon Toys			Amazon Pets			Amazon Health			Tripadvisor Hotels		
		1	10	100	1	10	100	1	10	100	1	10	100
BPR	Average SVD	0.093	0.121	0.247	0.187	0.264	0.346	0.136	0.173	0.295	0.309	0.516	0.566
	NNDSVD	0.189	0.196	0.249	0.298	0.315	0.342	0.229	**0.243**	0.304	0.490	0.537	0.559
	Random	0.164	0.176	0.210	0.239	0.256	0.282	0.212	0.224	0.263	0.427	0.512	0.515
	TIM	**0.200**	**0.231**	**0.262**	**0.299**	**0.319**	**0.361**	**0.233**	0.204	**0.316**	**0.534**	**0.544**	**0.581**
Rank-SGD	Average SVD	0.103	0.160	0.232	0.214	0.296	0.327	0.149	0.229	0.280	0.233	0.521	0.496
	NNDSVD	0.191	0.224	0.232	0.299	0.320	0.311	0.231	**0.279**	0.282	0.494	0.528	0.495
	Random	0.113	0.133	0.195	0.153	0.210	0.265	0.146	0.193	0.244	0.478	0.511	0.512
	TIM	**0.211**	**0.231**	**0.247**	**0.321**	**0.344**	**0.339**	**0.238**	0.273	**0.300**	**0.536**	**0.565**	**0.546**
WRMF	Average SVD	0.099	0.144	0.258	0.198	0.270	0.347	0.148	0.201	0.311	0.164	0.534	0.513
	NNDSVD	**0.190**	**0.219**	0.254	0.298	0.303	0.340	0.235	**0.257**	0.302	0.499	0.508	0.492
	Random	0.142	0.189	0.207	0.158	0.266	0.302	0.148	0.229	0.257	0.242	0.507	0.517
	TIM	0.183	0.201	**0.281**	**0.310**	**0.309**	**0.362**	**0.242**	0.232	**0.330**	**0.521**	**0.536**	**0.576**

Table 2: Recall@10 after 1, 10 and 100 iterations when seeded by different initializations strategies. The highest recall is in bold.

Alg.	Initializ.	Amazon Toys				Amazon Pets				Amazon Health				Tripadvisor Hotels			
		Recall	HT	NDCG	Prec	Recall	HT	NDCG	Prec	Recall	HT	NDCG	Prec	Recall	HT	NDCG	Prec
BPR	Avg SVD	0.252	0.516	0.164	0.072	0.352	0.650	0.208	0.095	0.298	0.596	0.185	0.081	0.566	0.691	0.260	0.076
	NNDSVD	0.252	0.513	0.165	0.072	0.345	0.642	0.207	0.093	0.305	0.601	0.187	0.082	0.559	0.685	**0.261**	0.075
	Random	0.244	0.512	0.155	0.070	0.333	0.627	0.194	0.090	0.276	0.569	0.171	0.074	0.516	0.645	0.245	0.069
	TIM	**0.264**	**0.535**	**0.169**	**0.076**	**0.362**	**0.660**	**0.215**	**0.097**	**0.316**	**0.614**	**0.190**	**0.085**	**0.581**	**0.702**	0.260	**0.078**
Rank-SGD	Avg SVD	0.240	0.492	0.156	0.068	0.333	0.629	0.201	0.090	0.291	0.586	0.179	0.079	0.541	0.665	0.252	0.073
	NNDSVD	0.241	0.501	0.158	0.069	0.324	0.622	0.199	0.087	0.294	0.588	0.183	0.079	0.532	0.660	0.253	0.072
	Random	0.221	0.475	0.143	0.063	0.302	0.593	0.182	0.081	0.256	0.540	0.162	0.069	0.514	0.644	0.245	0.069
	TIM	**0.261**	**0.532**	**0.166**	**0.075**	**0.356**	**0.652**	**0.212**	**0.096**	**0.313**	**0.609**	**0.188**	**0.085**	**0.574**	**0.696**	**0.258**	**0.077**
WRMF	Avg SVD	0.262	0.536	0.164	0.075	0.349	0.653	0.208	0.094	0.312	0.620	0.189	0.084	0.544	0.671	0.256	0.073
	NNDSVD	0.258	0.527	0.164	0.074	0.342	0.645	0.208	0.092	0.309	0.615	0.189	0.084	0.530	0.657	0.251	0.071
	Random	0.219	0.467	0.143	0.063	0.319	0.607	0.191	0.086	0.270	0.560	0.169	0.073	0.520	0.650	0.246	0.070
	TIM	**0.282**	**0.566**	**0.173**	**0.081**	**0.372**	**0.674**	**0.217**	**0.100**	**0.335**	**0.644**	**0.197**	**0.090**	**0.579**	**0.705**	**0.261**	**0.078**

Table 3: Top-N performance @10 over various datasets. The best performance is highlighted in bold.

factors with topic models, allowing the RS algorithm to optimise both user/item latent factor matrices. Here, we fix one of the latent factor matrices and optimise only the other one. In this way, the model training minimises the ranking prediction error constrained to the topics on the other matrix. The practice of initialising one matrix and learning the remaining one is common among ALS algorithms [23]. Since the focus of this paper is on Stochastic Gradient Descent (SGD) we leave it to future work to explore ALS variants.

The advantage of relating topics directly to latent factors is that one can provide the users with visual explanations such as the one

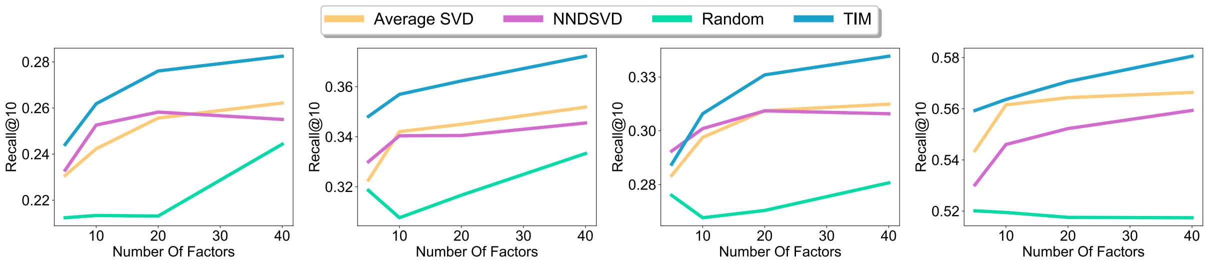

(a) Amazon Toys & Games (b) Amazon Pet Supplies (c) Amazon Health (d) Yelp RecSys

Figure 2: Influence of the number of latent factors on Recall@10 on multiple datasets

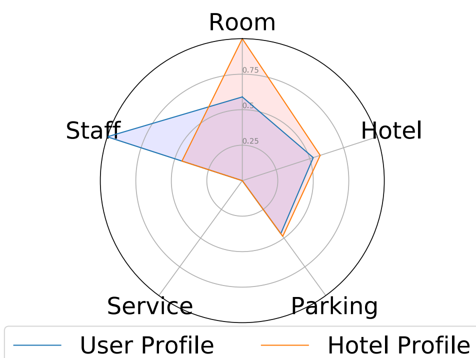

Figure 3: User-hotel profiles.

Alg.	Initializ.	Amazon Toys	Amazon Pets	Amazon Health	Tripadvisor Hotels
BPR	TIM-I	0.17	0.24	0.20	0.16
	TIM-U	0.22	0.32	0.24	0.54
	TIM	**0.26**	**0.36**	**0.31**	**0.58**
Rank-SGD	TIM-I	0.17	0.24	0.20	0.16
	TIM-U	0.22	0.32	0.24	0.54
	TIM	**0.26**	**0.35**	**0.31**	**0.57**
WRMF	TIM-I	0.17	0.24	0.20	0.16
	TIM-U	0.22	0.32	0.25	0.54
	TIM	**0.28**	**0.37**	**0.33**	**0.58**

Table 4: Recall@10 comparison between TIM-U, TIM-I and TIM initialisations.

presented in Figure 3 in which we have plotted the weight of each latent factor (that relates directly to a topic) for an example user and a recommended item. To plot Figure 3 we labelled each axis with the word with the highest weight for a topic, we then rescaled each latent feature vector so that its maximum weight had a value of 1.0. This graph was created using the TripAdvisor hotels dataset and a topic model with five topics (due to space reasons we are not including the best performing of 40 topics). Here, we indicate why we are suggesting a certain hotel to a user: we can see that the user and the hotel profiles are a close match.

As we can see in Table 4 interpretability comes at a cost. Here we have the TIM-U model, which we initialise with $\mathbf{P} = \mathbf{A}$ and $\mathbf{Q} = \mathbf{B}$. The parameters of $\mathbf{P}$ remain constant while the loss function is optimised by varying only the values of $\mathbf{Q}$ (in TIM-I $\mathbf{Q}$ is left fixed). Because only one matrix is to be learned, TIM-U has a fast convergence and in the end better performance, but unlike TIM, TIM-U does not outperform the random initialisation across all datasets (compared with the values in Table 3). We also see that there is a drop in the performance of TIM-U and TIM-I compared to TIM, which is expected since in TIM we don't fix the initialised latent factor matrices and we continue to optimise their values. We don't expect TIM-I to perform well since the grouped item reviews are written by different users forming a very heterogeneous document. On the other hand, grouped user reviews are written by the same users and therefore express preferences in a consistent way, resulting in cleaner preferences obtained by the topic model. Depending on the dataset and the situation, one might consider sacrificing prediction and choosing the TIM-U model to favour interpretability.

5 DISCUSSION AND CONCLUSIONS

We have presented TIM, a model that builds document topic matrices using them as a seed for the latent factor matrices in matrix factorization models for Top-N recommendations. The topics represent user preferences (or item qualities) and are used as ground truth to learn latent factors, allowing for interpretability. Our model can be integrated into several existing latent factor models such as BPR, Rank-SGD and WRMF among many others. We evaluated TIM on four datasets from different domains showing superior performance in terms of ranking prediction.

One particular difference between NNDSVD, Average SVD and TIM is that the former two exploit information from the rating matrix in order to build the initialisation matrices, while TIM uses the reviews transformed into topic models. Topic models (frequently used to rank documents) can be very helpful to rank items, which gives TIM an advantage for improved performance. Higher convergence speed is also achieved with topic models because they help gather user preferences and item qualities from the reviews, putting the algorithm in a better initial position for the optimisation.

Building a model that jointly learns from both reviews and ratings simultaneously is a problem with a higher complexity and a bigger search space than learning only ratings. In our case, the goal is to improve Top-N recommendations, therefore we discard the additional complexity given by the joint models and focus only on optimising the ranking prediction function. Nevertheless, if interpretability is the goal, we have provided the Topic-User Initialized latent factors Model (TIM-U) along with an example of a visual interpretation. However, as followup work we also plan to analyse and compare the approach of jointly learning our model with the two step approach adopted in this work.

REFERENCES

[1] Syed Muhammad Atif, Sameer Qazi, and Nicolas Gillis. 2019. Improved SVD-based initialization for nonnegative matrix factorization using low-rank correction. *Pattern Recognition Letters* 122 (2019), 53 – 59. https://doi.org/10.1016/j.patrec.2019.02.018

[2] Mark Belford, Brian Mac Namee, and Derek Greene. 2018. Stability of topic modeling via matrix factorization. *Expert Systems with Applications* 91 (2018), 159 – 169. https://doi.org/10.1016/j.eswa.2017.08.047

[3] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (March 2003), 993–1022.

[4] C. Boutsidis and E. Gallopoulos. 2008. SVD based initialization: A head start for nonnegative matrix factorization. *Pattern Recognition* 41, 4 (2008), 1350 – 1362. https://doi.org/10.1016/j.patcog.2007.09.010

[5] Rose Catherine and William Cohen. 2017. TransNets: Learning to Transform for Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems* (Como, Italy) *(RecSys '17)*. Association for Computing Machinery, New York, NY, USA, 288–296. https://doi.org/10.1145/3109859.3109878

[6] Chong Chen, Min Zhang, Yiqun Liu, and Shaoping Ma. 2018. Neural Attentional Rating Regression with Review-Level Explanations. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1583–1592. https://doi.org/10.1145/3178876.3186070

[7] Li Chen, Guanliang Chen, and Feng Wang. 2015. Recommender Systems Based on User Reviews: The State of the Art. *User Modeling and User-Adapted Interaction* 25, 2 (June 2015), 99–154. https://doi.org/10.1007/s11257-015-9155-5

[8] Zhiyong Cheng, Xiaojun Chang, Lei Zhu, Rose C. Kanjirathinkal, and Mohan Kankanhalli. 2019. MMALFM: Explainable Recommendation by Leveraging Reviews and Images. *ACM Trans. Inf. Syst.* 37, 2, Article 16 (Jan. 2019), 28 pages. https://doi.org/10.1145/3291060

[9] Zhiyong Cheng, Ying Ding, Lei Zhu, and Mohan Kankanhalli. 2018. Aspect-Aware Latent Factor Model: Rating Prediction with Ratings and Reviews. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 639–648. https://doi.org/10.1145/3178876.3186145

[10] Qiming Diao, Minghui Qiu, Chao-Yuan Wu, Alexander J. Smola, Jing Jiang, and Chong Wang. 2014. Jointly Modeling Aspects, Ratings and Sentiments for Movie Recommendation (JMARS). In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New York, New York, USA) *(KDD '14)*. ACM, New York, NY, USA, 193–202. https://doi.org/10.1145/2623330.2623758

[11] Ruihai Dong, Markus Schaal, Michael P. O'Mahony, Kevin McCarthy, and Barry Smyth. 2013. Opinionated Product Recommendation. In *Case-Based Reasoning Research and Development*, Sarah Jane Delany and Santiago Ontañón (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 44–58.

[12] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web* (Perth, Australia) *(WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 173–182. https://doi.org/10.1145/3038912.3052569

[13] Balázs Hidasi and Domonkos Tikk. 2012. Enhancing Matrix Factorization through Initialization for Implicit Feedback Databases. In *Proceedings of the 2nd Workshop on Context-Awareness in Retrieval and Recommendation* (Lisbon, Portugal) *(CaRR '12)*. Association for Computing Machinery, New York, NY, USA, 2–9. https://doi.org/10.1145/2162102.2162104

[14] Balazs Hidasi and Domonkos Tikk. 2013. Initializing Matrix Factorization Methods on Implicit Feedback Databases. 19, 12 (jun 2013), 1834–1853.

[15] Thomas Hofmann. 2001. Unsupervised Learning by Probabilistic Latent Semantic Analysis. *Mach. Learn.* 42, 1/2 (Jan. 2001), 177–196. https://doi.org/10.1023/A:1007617005950

[16] Yunfeng Hou, Ning Yang, Yi Wu, and Philip S. Yu. 2019. Explainable recommendation with fusion of aspect information. *World Wide Web* 22, 1 (01 Jan 2019), 221–240. https://doi.org/10.1007/s11280-018-0558-1

[17] Binbin Hu, Chuan Shi, Wayne Xin Zhao, and Philip S. Yu. 2018. Leveraging Meta-Path Based Context for Top- N Recommendation with A Neural Co-Attention Model. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (London, United Kingdom) *(KDD '18)*. Association for Computing Machinery, New York, NY, USA, 1531–1540. https://doi.org/10.1145/3219819.3219965

[18] Y. Hu, Y. Koren, and C. Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *2008 Eighth IEEE International Conference on Data Mining*. 263–272. https://doi.org/10.1109/ICDM.2008.22

[19] Michael Jahrer and Andreas Töscher. 2011. Collaborative Filtering Ensemble for Ranking. In *Proceedings of the 2011 International Conference on KDD Cup 2011 - Volume 18 (KDDCUP'11)*. JMLR.org, 153–167.

[20] Kalervo Järvelin and Jaana Kekäläinen. 2000. IR Evaluation Methods for Retrieving Highly Relevant Documents. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Athens, Greece) *(SIGIR '00)*. Association for Computing Machinery, New York, NY, USA, 41–48. https://doi.org/10.1145/345508.345545

[21] Yehuda Koren. 2008. Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Las Vegas, Nevada, USA) *(KDD '08)*. Association for Computing Machinery, New York, NY, USA, 426–434. https://doi.org/10.1145/1401890.1401944

[22] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (Aug. 2009), 30–37. https://doi.org/10.1109/MC.2009.263

[23] Amy N Langville, Carl D Meyer, Russell Albright, James Cox, and David Duling. 2014. Algorithms, initializations, and convergence for the nonnegative matrix factorization. *arXiv preprint arXiv:1407.7299* (2014).

[24] Daniel D Lee and H. Sebastian Seung. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401 (10 1999), 788–791. https://doi.org/10.1038/44565

[25] Guang Ling, Michael R. Lyu, and Irwin King. 2014. Ratings Meet Reviews, a Combined Approach to Recommend. In *Proceedings of the 8th ACM Conference on Recommender Systems* (Foster City, Silicon Valley, California, USA) *(RecSys '14)*. ACM, New York, NY, USA, 105–112. https://doi.org/10.1145/2645710.2645728

[26] Dionisis Margaris, Costas Vassilakis, and Dimitris Spiliotopoulos. 2020. What makes a review a reliable rating in recommender systems? *Information Processing & Management* 57, 6 (2020), 102304. https://doi.org/10.1016/j.ipm.2020.102304

[27] Julian McAuley and Jure Leskovec. 2013. Hidden Factors and Hidden Topics: Understanding Rating Dimensions with Review Text. In *Proceedings of the 7th ACM Conference on Recommender Systems* (Hong Kong, China) *(RecSys '13)*. ACM, New York, NY, USA, 165–172. https://doi.org/10.1145/2507157.2507163

[28] Mahdi Nasiri and Behrouz Minaei. 2016. Increasing Prediction Accuracy in Collaborative Filtering with Initialized Factor Matrices. *J. Supercomput.* 72, 6 (June 2016), 2157–2169. https://doi.org/10.1007/s11227-016-1717-8

[29] Francisco J. Peña. 2017. Unsupervised Context-Driven Recommendations Based On User Reviews. In *Proceedings of the Eleventh ACM Conference on Recommender Systems* (Como, Italy) *(RecSys '17)*. ACM, New York, NY, USA, 426–430. https://doi.org/10.1145/3109859.3109865

[30] Francisco J. Peña and Derek Bridge. 2017. Recommending from Experience. In *Proceedings of the Thirtieth Florida Artificial Intelligence Research Society Conference*. https://aaai.org/ocs/index.php/FLAIRS/FLAIRS17/paper/view/15439

[31] Lin Qiu, Sheng Gao, Wenlong Cheng, and Jun Guo. 2016. Aspect-based latent factor model by integrating ratings and reviews for recommender system. *Knowledge-Based Systems* 110 (2016), 233 – 243. https://doi.org/10.1016/j.knosys.2016.07.033

[32] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence* (Montreal, Quebec, Canada) *(UAI '09)*. AUAI Press, Arlington, Virginia, USA, 452–461.

[33] Chong Wang and David M. Blei. 2011. Collaborative Topic Modeling for Recommending Scientific Articles. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Diego, California, USA) *(KDD '11)*. ACM, New York, NY, USA, 448–456. https://doi.org/10.1145/2020408.2020480

[34] Yongfeng Zhang, Qingyao Ai, Xu Chen, and W. Bruce Croft. 2017. Joint Representation Learning for Top-N Recommendation with Heterogeneous Information Sources. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management* (Singapore, Singapore) *(CIKM '17)*. Association for Computing Machinery, New York, NY, USA, 1449–1458. https://doi.org/10.1145/3132847.3132892

[35] Lei Zheng, Vahid Noroozi, and Philip S. Yu. 2017. Joint Deep Modeling of Users and Items Using Reviews for Recommendation. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining* (Cambridge, United Kingdom) *(WSDM '17)*. Association for Computing Machinery, New York, NY, USA, 425–434. https://doi.org/10.1145/3018661.3018665