Learning to Rank for Personalized News Recommendation

Pavel Shashkin
National Research University Higher School of Economics
25/12 B. Pecherskaya str.
Nizhny Novgorod, Russia
p-sh@live.ru

Nikolay Karpov
National Research University Higher School of Economics
25/12 B. Pecherskaya str.
Nizhny Novgorod, Russia
nkarpov@hse.ru

ABSTRACT

Improving user experience through personalized recommendations is crucial to organizing the abundance of data on news websites. Modeling user preferences based on implicit feedback has recently gained lots of attention, partly due to growing volume of web generated click stream data. Matrix factorization learned with stochastic gradient descent has successfully been adopted to approximate various ranking objectives. The aim of this paper is to test the performance of learning to rank approaches on the real-world dataset and apply some simple heuristics to consider temporal dynamics present in news domain. Our model is based on WARP loss with changes to classic factorization model.

ACM Reference format:
Pavel Shashkin and Nikolay Karpov. 2017. Learning to Rank for Personalized News Recommendation. In *Proceedings of WI '17, Leipzig, Germany, August 23-26, 2017,* 3 pages.
https://doi.org/10.1145/3106426.3109432

1 INTRODUCTION

The news website is a complex dynamic graph structure simultaneously navigated by millions of users. During his visit, a user wants to find relevant and interesting articles among thousands of published ones after his last session. In order to save time that a user spends manually searching for content without even being sure of its existence, we produce a fixed size list of personalized recommendations. Ideally, we want to maximize the average number of articles viewed during user session.

We consider a variant of an implicit feedback recommender algorithm which was evaluated on news domain and on the MovieLens dataset. Our main contribution is the following:

(1) Make available specific news dataset;
(2) Successfully apply heuristic to factorization model to consider temporal dynamics present in news domain;
(3) Implement WARP algorithm for loss and sampling procedure to our factorization model;
(4) Implement AdaGrad learning algorithm with Hogwild parallelization to learning procedure.

WI '17, August 23-26, 2017, Leipzig, Germany

ACM ISBN 978-1-4503-4951-2/17/08…$15.00
https://doi.org/10.1145/3106426.3109432

The structure of the paper is as follows. Section 2 explore related approaches. Section 3 details an experiment methodology and our datasets to evaluate algorithms. Section 4 defines a notation that we use and our proposed approach. Section 5 describes experiments on data, and Section 6 concludes.

2 RELATED WORK

Two different approaches are commonly used in recommender systems: content-based recommendation and collaborative filtering. Several content based approaches have been successfully applied to on-line news recommendation [4].

Other approaches rely on collaborative filtering [1, 5, 10] where content of items is not considered, but opinions of peer users are taken into account to generate recommendations. More complex systems use hybrid models [2, 6, 9] that leverages both: information about user-article interactions and article metadata. Recently reinforcement learning approaches have gained attention [7, 8].

Recommendation task can be naturally posed as a ranking problem where for each user we are trying to rank positive items higher than negative ones. In implicit setting pairwise learning methods like Bayesian Personalized Ranking (BPR) [13] assume that all unobserved user-item interactions are negative and minimize ranking loss by triplets (user, positive item, negative item). To reduce computational complexity negative items are sampled uniformly.

It has been shown that more sophisticated adaptive sampling procedures[12, 14] perform more informative parameter updates and improve calculable performance metrics. In the WARP algorithm the loss function is similar to hinge loss weighted by the rank of positive item. Only triplets that violate the desired order are used to update model parameters.

3 EXPERIMENT METHODOLOGY

We present a simple experiment that illustrates the efficiency of proposed approach with the help of two types of data. Train-test split is made in such a way that all test interactions occur after train interactions.

Our goal is to produce a high quality ranked list of recomendations for each user it test set. Recommendation quality is measured by precision ($p@k$) and recall ($r@k$) at top k items averaged over all users.

3.1 News Dataset

The data is collected from popular Russian news website[1]. In addition to big (30 millions users by 140 thousands articles) and sparse (440 millions nonzero values) binary matrix the dataset also contains time for each interaction. Many users appear only once and

[1]https://life.ru

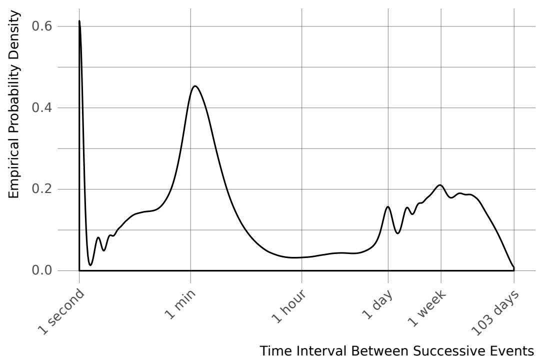

Figure 1: Distribution of time gaps between article interactions averaged over all users.

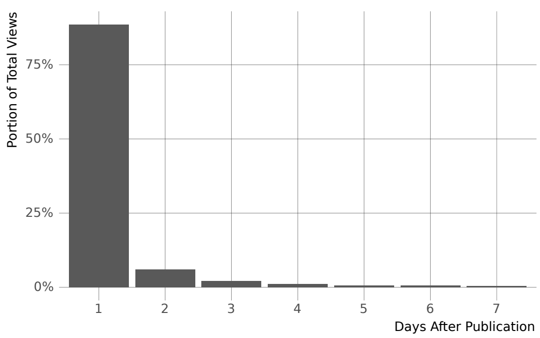

Figure 2: Distribution of views averaged over all articles.

even without taking them into account 49.7% of remaining users interact with the website only within 24 hours interval. We share this dataset[2] to make our results reproducible.

Figure 1 illustrates previously used notion of session as well as seasonality. We assume that the first two spikes (around one second and one minute marks) indicate whether user skipped content or interacted with it. The rightmost density increase demonstrates the time intervals between sessions (note the daily and weekly visitors). Unfortunately, the value used to identify the user is not very reliable - for the majority of users, it is not same across different devices or browsers.

News articles naturally tend to lose popularity as time passes and new stories become available (see Figure 2). There are some rare examples where new event or article alters the relevance of previously released one and instead of declining the popularity raises.

[2]goo.gl/dSkgGM

3.2 MovieLens Dataset

We apply our approach to a MovieLens 10M dataset because it is well known and widely used in education, research, and industry. It is a collaborative filtering dataset that contains 5-star movie ratings and time when user rates items. In our implicit experiment we don't use the rating values. If a user rates an item with any rating we believe that the item is interesting to the user in some sense.

4 PROPOSED APPROACH

We propose a factorization model that maps users, discrete time slots and items into a latent factor space $\mathbb{R}^m$ by the functions $\Phi_U(x)$, $\Phi_T(t)$ and $\Phi_I(y)$ respectively. The time of user-item interaction is discretized to take a value from a finite set. $\Phi_T(t)$ is used as an additive component to consider temporal dynamics present in news domain. For instance it can model temporal popularity of specific topics like elections, oscar etc. To model temporal popularity and availability of specific item we add a scalar item-time bias $\Psi(y, t)$ to inner products of item-time and user-item factors.

The proposed approach uses the following factorization model to assign scores to (user x, item y, time slot t) triplets:

$$f_y(x, t) = (\Phi_U(x) + \Phi_T(t))^\intercal \times \Phi_I(y) + \Psi(y, t) \tag{1}$$

Algorithm 1: Online Loss Optimization

Input: (user, item, time) triplets
Initialize model parameters uniformly (from $-1/2m$ to $1/2m$).
repeat
  Pick a random positive interaction (x, y, t)
  Let $f_y(x, t) = (\Phi_U(x) + \Phi_T(t))^\intercal \times \Phi_I(y) + \Psi(y, t)$
  Set $N = 0$. **repeat**
    Pick a random negative item $\bar{y} \notin D(x)$.
    Let $f_{\bar{y}}(x) = (\Phi_U(x) + \Phi_T(t))^\intercal \times \Phi_I(\bar{y}) + \Psi(\bar{y}, t)$
    $N = N + 1$.
  until $f_{\bar{y}}(x_i) > f_y(x) - 1$ *or* $N \geqslant Y - 1$;
  if $f_{\bar{y}}(x) > f_y(x) - 1$ **then**
    Make a gradient step to minimize:
    $L(\lfloor \frac{Y-1}{N} \rfloor) \left|1 - f_{\bar{y}}(x, t) + f_y(x, t)\right|_+$
    Project weights to enforce constraints.
  end
until *validation error does not improve.*;

To make the proposed method clear, we specify its optimization procedure in Algorithm 1. The loss and sampling processes are the same as in original WARP algorithm. Y is the maximum number of sampling iterations, $L()$ is the weighting function (natural logarithm in our experiments) for approximated rank and $D()$ is the item set viewed by user. For the stochastic gradient descent we use AdaGrad [3] learning schedule with Hogwild [11] parallelization.

5 EXPERIMENTAL RESULTS

We apply three approaches mentioned above: BPR, WARP and our proposed one. The accuracy is measured in terms of precision and recall at top 10 and 100 ranked items.

We use the following configuration for all experiments: 20 components in a vector of the latent factor, 0.1 initial learning rate, maximum 30 sampled items to find violating example. For news

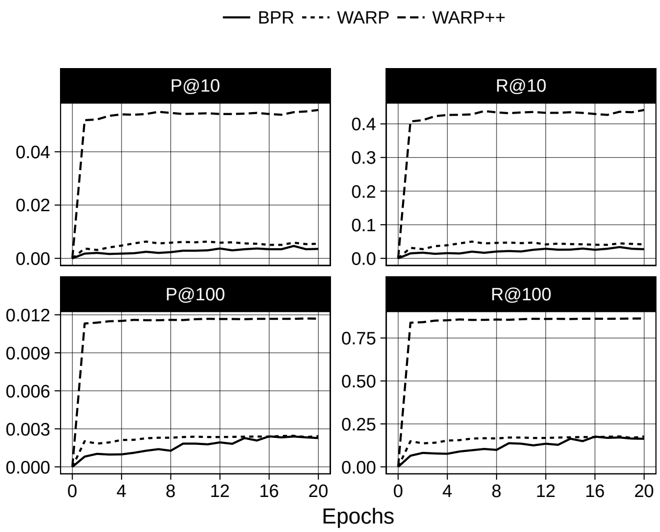

Figure 3: Precision and Recall for Top 10 and 100 Ranked Items from News Dataset

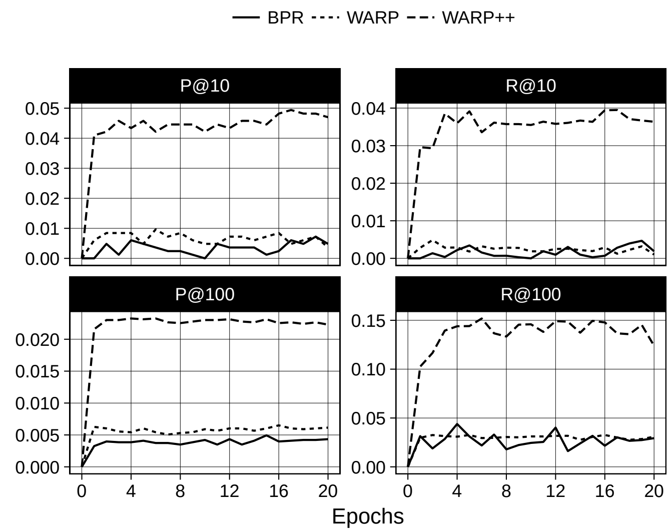

Figure 4: Precision and Recall for Top 10 and 100 Ranked Items from MovieLens 10M Dataset

data, we use 15 minutes time slot length, for movie data - one week. For each dataset, we test model on the last 1 000 observations.

Figure 3 shows precision and recall for the news dataset after *n* optimization rounds. As we can see that the proposed variant WARP++ helps to increase both precision and recall.

The precision and recall for the MovieLens data are shown in Figure 4. Our model outperforms others despite the fact that the content (movies), the rating data (5 stars), typical consumption patterns and consumption context are different. We can assume that relative improvement on movies dataset over matrix factorization model is smaller because there are less time-related effects in user-movie preferences.

6 CONCLUSIONS

We have introduced a model for implicit feedback recommender system. To consider temporal dynamics present in news domain we successfully apply a heuristic to factorization model. To our factorization model, we implement WARP algorithm for loss and sampling procedure. We also implement AdaGrad learning algorithm with Hogwild parallelization to a learning procedure.

This model was evaluated on our specific news dataset, which we made available, and on the MovieLens 10M well-known data.

7 ACKNOWLEDGMENTS

The article was prepared within the framework of the Academic Fund Program at the National Research University Higher School of Economics (HSE) in 2017 (grant N17-05-0007) and by the Russian Academic Excellence Project "5-100".

REFERENCES

[1] A. S. Das, M. Datar, A. Garg, and S. Rajaram. Google news personalization: scalable online collaborative filtering. In *Proceedings of the 16th international conference on World Wide Web*, pages 271–280. ACM, 2007.
[2] D. Doychev, A. Lawlor, R. Rafter, and B. Smyth. An analysis of recommender algorithms for online news. Citeseer, 2014.
[3] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
[4] F. Garcin, C. Dimitrakakis, and B. Faltings. Personalized news recommendation with context trees. *CoRR*, abs/1303.0665, 2013.
[5] F. Garcin, K. Zhou, B. Faltings, and V. Schickel. Personalized news recommendation based on collaborative filtering. In *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology-Volume 01*, pages 437–441. IEEE Computer Society, 2012.
[6] I. Ilievski and S. Roy. Personalized news recommendation based on implicit feedback. In *Proceedings of the 2013 International News Recommender Systems Workshop and Challenge*, pages 10–15. ACM, 2013.
[7] L. Li, W. Chu, and J. Langford. An unbiased, data-driven, offline evaluation method of contextual bandit algorithms. *CoRR*, abs/1003.5956, 2010.
[8] L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. *CoRR*, abs/1003.0146, 2010.
[9] L. Li, D. Wang, T. Li, D. Knox, and B. Padmanabhan. Scene: a scalable two-stage personalized news recommendation system. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, pages 125–134. ACM, 2011.
[10] Z. Lu, Z. Dou, J. Lian, X. Xie, and Q. Yang. Content-based collaborative filtering for news topic recommendation. Citeseer, 2015.
[11] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 693–701, 2011.
[12] S. Rendle and C. Freudenthaler. Improving pairwise learning for item recommendation from implicit feedback. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 273–282. ACM, 2014.
[13] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, pages 452–461. AUAI Press, 2009.
[14] J. Weston, H. Yee, and R. J. Weiss. Learning to rank recommendations with the k-order statistic loss. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 245–248. ACM, 2013.