

Deep Learning for Recommender Systems

Alexandros Karatzoglou
Telefonica Research
Barcelona, Spain
alexk@tid.es

Balázs Hidasi
Gravity R&D
Budapest, Hungary
hidasi.balazs@gravityrd.com

ABSTRACT

Deep Learning is one of the next big things in Recommendation Systems technology. The past few years have seen the tremendous success of deep neural networks in a number of complex machine learning tasks such as computer vision, natural language processing and speech recognition. After its relatively slow uptake by the recommender systems community, deep learning for recommender systems became widely popular in 2016.

We believe that a tutorial on the topic of deep learning will do its share to further popularize the topic. Notable recent application areas are music recommendation, news recommendation, and session-based recommendation. The aim of the tutorial is to encourage the application of Deep Learning techniques in Recommender Systems, to further promote research in deep learning methods for Recommender Systems.

CCS CONCEPTS

• **Information systems** → **Information retrieval**; • **Computing methodologies** → **Artificial intelligence**; **Machine learning**;

KEYWORDS

Deep Learning, Recommender Systems

1 INTRODUCTION

The aim of the tutorial is dual, 1) we would like to introduce Deep Learning techniques that have been used in recommender systems such as Recurrent Neural Networks and Convolutional Networks etc. 2) we would like to present the current state-of-the-art methods that use deep learning techniques to provide recommendations.

Machine learning methods such as matrix factorization and tensor factorization have been frequently used in the area of recommender systems deep learning methods. These methods share in fact some similarities to deep learning such as e.g. the use of Stochastic Gradient Descent (SGD) for optimization. One can in fact cast matrix factorization as a neural network. In that respect the use of deep learning methods in the recommender systems community represents a natural step.

Deep Learning provides a new toolkit for recommender systems practitioners to extract features and to model user generated data and item data that has the potential to provide large improvements

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

RecSys'17, August 27–31, 2017, Como, Italy.

© 2017 Copyright held by the owner/author(s). ISBN 978-1-4503-4652-8/17/08.

DOI: <http://dx.doi.org/10.1145/3109859.3109903>

in the quality of the recommendations provided to users. Part of the power of deep learning techniques in recommender systems stems from the fact that deep learning methods allow for much better feature extraction from item characteristics such as image, video and audio compared to traditional techniques. This allows for a more accurate modeling of items based on their content, thus one can potentially expect hybrid and content-based methods to perform better compared to traditional content based techniques whenever deep learning is used. Another advantage deep learning methods provide is that they allow for different views of the data, standard collaborative filtering techniques such as matrix factorization have often treated user-item interaction as flat matrix structured data often ignoring the temporal structure and order in the data. Deep learning techniques like convolutions and recurrent neural networks allow us to model the temporal structure in this data leading to significant performance improvements.

2 DEEP LEARNING TECHNIQUES FOR RECOMMENDER SYSTEMS

Deep Learning has become popular after a 2006 paper by Hinton [7] described how to train deep neural networks using unsupervised layer-wise training (Deep Belief Networks). Since then there has been an immense amount of research in the area with impressive results in areas such as computer vision, speech recognition and natural language processing. The uptake of deep networks in the recommender systems community has been relatively slow but recent advances in the area of deep learning and recommender systems have encouraged researchers and practitioners to use these methods. Deep neural networks are now used in a range of recommender systems algorithms and form a considerable part of the program of many Information Retrieval conferences and indeed of the ACM RecSys conference.

Over the last 3-4 years there has been a number of notable publications in the area of deep learning for recommender systems. This deep learning work can be broadly grouped into the following categories:

- **Embedding methods.** This class of methods utilize embedding techniques inspired by deep learning methods such as word2vec [8] to embed user, item, [11], [13] or context profiles in latent spaces. These embeddings can then be used either directly to provide recommendations or as input to other (typically supervised) (deep learning) methods that provide the recommendations. While matrix factorization can be also viewed as an embedding technique, 2vec type embedding techniques are often somewhat more flexible than the matrix factorization models often employed in collaborative filtering.

- **Feedforward Networks and Autoencoders for Collaborative Filtering.** These methods use either Feedforward or Autoencoder [9], [14] networks directly on the user-item interactions in order to build collaborative filtering models that can then be used for recommendations. These methods can also be seen as a form of deep factorization methods. Deep Factorization Machines [2] can be seen as an example of this class of methods. These methods often outperform standard model-based collaborative filtering methods.
- **Deep Feature extracting methods.** This class of deep learning for recommenders methods focuses on using deep networks to perform feature extraction on the item features. These features are then either used in a more standard (hybrid) collaborative filtering methods or often the feature extractor is part of a larger deep architecture that also models other aspects of the data. Examples of this class of methods includes convolutional methods for image feature extraction for hybrid collaborative filtering methods [3], or convolutional networks for audio and music feature extraction [10] or even text-based recommendations with neural feature extractors (RNN's or convolutional networks).
- **Session-based Recommendation with Recurrent Neural Networks.**

These methods are based on the fact that often user interactions with content are in sessions, e.g. music listening, shopping. Moreover one often does not have a reliable user-identifier to model user preferences. In this case one can resort to model sessions as sequential data and one of the best performing models for sequential data are Recurrent Neural Networks. Typically GRU's are used in these models as they seem to perform equivalently to LSTM's and have slightly less memory usage. These models have been shown to outperform traditional methods by a large margin [5], [6], [4].

3 FUTURE

Deep Learning techniques are constantly progressing and several new techniques have been recently introduced such as e.g. Adversarial Training, Siamese Networks, one-shot learning and much more. While the application of these techniques to the recommender systems domain is not obvious at the moment we do believe that these and other techniques will eventually be adopted by the recommender systems community.

We also expect a significant increase in deployment of deep learning models with content-based methods becoming more popular due to increased quality of features that can be extracted with these methods. Another area where we believe that deep learning methods will become instrumental is the area of conversational recommender systems [1]. Deep models are used in conversational agents (bots) [12] and it is only natural that they will eventually be used in conversational recommender systems.

REFERENCES

- [1] Konstantina Christakopoulou, Filip Radlinski, and Katja Hofmann. 2016. Towards Conversational Recommender Systems. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. ACM, New York, NY, USA, 815–824. <https://doi.org/10.1145/2939672.2939746>
- [2] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. *CoRR* abs/1703.04247 (2017). <http://arxiv.org/abs/1703.04247>
- [3] Ruiming He and Julian McAuley. 2016. VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*. AAAI Press, 144–150. <http://dl.acm.org/citation.cfm?id=3015812.3015834>
- [4] Balázs Hidasi and Alexandros Karatzoglou. 2017. Recurrent Neural Networks with Top-k Gains for Session-based Recommendations. abs/1706.03847 (2017). <https://arxiv.org/abs/1706.03847>
- [5] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. *International Conference on Learning Representations (ICLR)* (2016).
- [6] Balázs Hidasi, Massimo Quadrana, Alexandros Karatzoglou, and Domonkos Tikk. 2016. Parallel recurrent neural network architectures for feature-rich session-based recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 241–248.
- [7] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural computation* 18, 7 (2006), 1527–1554.
- [8] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [9] Suvash Sedhain, Aditya Krishna Menon, Scott Sanner, and Lexing Xie. 2015. Autorec: Autoencoders meet collaborative filtering. In *Proceedings of the 24th International Conference on World Wide Web*. ACM, 111–112.
- [10] Aaron van den Oord, Sander Dieleman, and Benjamin Schrauwen. 2013. Deep content-based music recommendation. In *Advances in Neural Information Processing Systems 26*. C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.), Curran Associates, Inc., 2643–2651. <http://papers.nips.cc/paper/5004-deep-content-based-music-recommendation.pdf>
- [11] Flavian Vasile, Elena Smirnova, and Alexis Conneau. 2016. Meta-Prod2Vec: Product Embeddings Using Side-Information for Recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 225–232.
- [12] Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869* (2015).
- [13] Jeroen BP Vuurens, Martha Larson, and Arjen P de Vries. 2016. Exploring Deep Space: Learning Personalized Ranking in a Semantic Space. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*. ACM, 23–28.
- [14] Yao Wu, Christopher DuBois, Alice X Zheng, and Martin Ester. 2016. Collaborative denoising auto-encoders for top-n recommender systems. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*. ACM, 153–162.