RecSys'16 Workshop on Deep Learning for Recommender Systems (DLRS)

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

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

Domonkos Tikk
Gravity R&D
Budapest, Hungary
domonkos.tikk@gravityrd.com

Oren Sar-Shalom
IBM Research
Haifa, Israel
orensr@il.ibm.com

Haggai Roitman
IBM Research
Haifa, Israel
haggai@il.ibm.com

Bracha Shapira
Ben-Gurion University
Beersheba, Israel
bshapira@bgu.ac.il

ABSTRACT

We believe that 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 tasks such as computer vision, natural language processing and speech recognition. Despite this, only little work has been published on Deep Learning methods for Recommender Systems. Notable recent application areas are music recommendation, news recommendation, and session-based recommendation. The aim of the workshop is to encourage the application of Deep Learning techniques in Recommender Systems, to promote research in deep learning methods for Recommender Systems, and to bring together researchers from the Recommender Systems and Deep Learning communities.

Keywords

deep learning; neural networks; recommender systems

1. INTRODUCTION

The past few years have seen the tremendous success of deep neural networks in a number of complex tasks such as computer vision, natural language processing and speech recognition. Therefore it is one of the hottest topics in the machine learning community recently. Deep learning is a class of algorithms where data is passed through a cascade of nonlinear processing layers organized in complex architectures where each layer learns latent representations of the data.

Despite this yet relatively limited uptake of the technology in the recsys community, we believe that deep learning is one of the next big things in recommendation systems technology due to several reasons:

(1) Deep learning is capable of capturing complex dependencies and nonlinearities in large sets of data that other algorithms can't, thus it can bring a new and deeper modelling of user preferences encoded into the algorithmic layer, without the necessity to understand the exact nature of patterns contributing to the prediction.

(2) Deep learning techniques can integrate additional data sources (the content itself, text, images, music, video, etc.) to enrich the preference models of recommender algorithms by extracting high quality feature representations. Previously recommender systems have focused on utilizing the transactions and/or the descriptors of the content (i.e. metadata). However, the visual information users see about the product (i.e. an image of the product) makes a large contribution to user decisions (e.g. if they will click or not on the recommendations). Deep learning methods allow for a better and more accurate modeling of the content features of the item, particularly visual features such as images, textual features such as the description of the item or the audio signal of a song.

(3) Deep learning has the potential to allow for direct modeling of user preferences, with enough data on user preferences it is conceivable that user models of product preferences could be derived based on the content features of the items alone. For example one could potentially fine-tune a generic convolutional image classification model to recognize items that a user will like e.g. a particular pair of shoes or sportswear or even a potential partner (reciprocal recommendation). This will allow for user preference models that do not depend on collaborative filtering and thus the collection and modeling of large amounts of preference data on a single server.

(4) Finally deep learning methods provide a rich toolkit for key recommendation problems e.g. session-based recommendation with recurrent neural networks. The uptake of machine learning methods in the recommendation community is relatively limited to a small set of methods e.g. matrix factorization, LDA. Deep learning methods can and have been used in collaborative filtering models e.g. denoising autoencoders, Restricted Boltzmann Machines etc. We expect to see more research on how to use deep learning methods in a collaborative filtering setting with models that take into account content and context. Deep learning methods have the potential to expand upon this set of methods and to enhance existing ones.

The aim of DLRS 2016 is to help introduce and speed-

RecSys '16 September 15-19, 2016, Boston , MA, USA

ACM ISBN 978-1-4503-4035-9/16/09.

DOI: http://dx.doi.org/10.1145/2959100.2959202

up the use of deep learning in the RecSys Community by bringing together researchers from the Recommender Systems and Deep Learning communities.

2. DEEP LEARNING IN RECSYS

Most of the work on deep models and recommendations focus on the classical collaborative filtering (CF) user–item setting. Restricted Boltzmann Machines (RBM) were one of the first neural networks to be used for classical CF and recommender systems [5]. More recently denoising autoencoders have been used to perform CF in a similar manner [8]. Deep networks have also been used in cross-domain recommendation whereby items are mapped to a joint latent space using deep neural networks [1].

Including automatically generated content features in the recommender is one of the most exciting imminent research paths with deep learning. Prior work typically used these features together with more conventional collaborative filtering models. Convolutional deep networks have been used to extract features from music files that are then used in a factor model [6]. More recently [7] introduced a more generic approach whereby a deep network is used to extract generic content-features from items, these features are then incorporated in a standard CF model to enhance the recommendation performance. This approach seems to be particularly useful in settings where there is not sufficient user–item interaction information. Image features that have been extracted using convolutional networks have been used in classical matrix factorization-type CF in [2, 4] to enhance the quality of recommendations. Deep learning can also address novel recommendation problems. For example [3] used a custom recurrent neural network to model click-session data to provide session-based recommendations.

3. THE POTENTIAL OF DEEP LEARNING

We outline four major broad research topics, which have huge potential to significantly improve recommender algorithms.

Content modeling: Incorporating unstructured data sources such as text, audio, video or image into recommendation algorithms is one of the more trivial ways of applying deep learning. As we argued earlier, the inclusion of the content instead of its metadata makes more sense and expected to give better recommendations. True content based methods can help recommenders with long standing issues like the item cold-start, while simultaneously recommending relevant content.

User and session modeling: Traditional collaborative filtering is largely focused on creating static user models, which are good for domains where tastes are expected to change slowly and item relevance is preference-based rather than necessity based (e.g. movies as opposed to household appliances). Some domains require more dynamic behavior modeling. In others, it is important to find the balance between using long-term and short-term user histories. Modeling intent is also important for knowing what the user really wants and recommending accordingly. Last, but not least, using sessions rather than user histories is important in situations where the permanent user cold-start problem is present. All of the aforementioned problems can be potentially addressed by the proper use of deep learning's sequence modeling (e.g. recurrent neural networks).

Situation modeling: Surrounding ourselves with more and more sensors allows recommenders to get more data about the situation that we are in and recommend accordingly. Moving towards this situation based (or context-driven) recommendation will be one of the main shifts in the next few years. Understanding the situation from a large amount of sensory data requires powerful pattern recognition methods, such as deep neural networks.

Adapting deep learning: Even though deep neural networks have great potential, their application in the field of recommender systems should be done with care. Practical considerations – e.g. training time, response time and scalability – must be taken into account when developing novel methods. Therefore neural models should be adapted to fit the recommendation task.

4. SUMMARY

Deep learning achieved tremendous success on tasks, considered to be hard for computers. We believe it is time to apply them in the RecSys domain as well. Deep learning has the potential to be the core of the next generation of recommender systems. DLRS 2016 aims to speed up the spread and acceptance of these methods in the RecSys Community.

Acknowledgments

The organizers are supported in part by European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 610594 known as the CrowdRec project http://crowdrec.eu/.

5. ADDITIONAL AUTHORS

Additional authors: Lior Rokach (Ben-Gurion University, email: `liorrk@bgu.ac.il`).

6. REFERENCES

[1] A. M. Elkahky, Y. Song, and X. He. A multi-view deep learning approach for cross domain user modeling in recommendation systems. In *WWW'15: 24th Int. Conf. on World Wide Web*, pages 278–288, 2015.

[2] R. He and J. McAuley. VBPR: Visual Bayesian Personalized Ranking from implicit feedback. *CoRR*, 1510.01784, 2015.

[3] B. Hidasi, A. Karatzoglou, L. Baltrunas, and D. Tikk. Session-based recommendations with recurrent neural networks. *International Conference on Learning Representations*, 2016.

[4] J. McAuley, C. Targett, Q. Shi, and A. van den Hengel. Image-based recommendations on styles and substitutes. In *SIGIR'15: 38th Int. ACM SIGIR Conf. on Research and Development in Information Retrieval*, pages 43–52, 2015.

[5] R. Salakhutdinov, A. Mnih, and G. Hinton. Restricted boltzmann machines for collaborative filtering. In *ICML'07: 24th Int. Conf. on Machine Learning*, pages 791–798, 2007.

[6] A. Van den Oord, S. Dieleman, and B. Schrauwen. Deep content-based music recommendation. In *Advances in Neural Information Processing Systems*, pages 2643–2651, 2013.

[7] H. Wang, N. Wang, and D.-Y. Yeung. Collaborative deep learning for recommender systems. In *KDD'15: 21th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, pages 1235–1244, 2015.

[8] Y. Wu, C. DuBois, A. X. Zheng, and M. Ester. Collaborative denoising auto-encoders for top-N recommender systems. In *WSDM'16: 9th ACM Int. Conf. on Web Search and Data Mining*, pages 153–162, 2016.