DLRS 2017 - Second Workshop on Deep Learning for Recommender Systems

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

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

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

Sander Dieleman
Google DeepMind
London, United Kingdom
sanderdieleman@gmail.com

Bracha Shapira
Ben Gurion University of the Negev
Beersheba, Israel
bshapira@bgu.ac.il

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

ABSTRACT

Deep learning methods became widely popular in the recommender systems community in 2016, in part thanks to the previous event of the DLRS workshop series. Now, deep learning has been embedded in the main conference as well and initial research directions have started forming, so the role of DLRS 2017 is to encourage starting new research directions, incentivize the application of very recent techniques from deep learning, and provide a venue for specialized discussion of this topic.

KEYWORDS

deep learning, neural networks, recommender systems

1 INTRODUCTION

Deep learning techniques achieved great success in complex domains, such as computer vision, natural language processing, machine translation, speech recognition and reinforcement learning. The uptake of deep learning by the recommender systems community was relatively slow, as the topic became popular only in 2016. The first Deep Learning for Recommender Systems workshop (DLRS 2016 [11]) – which was the workshop with the highest demand at RecSys 2016 – also took its share to propagate the topic. The popularity of this topic is still in its growing phase demonstrated – beside others – by the ~ 25% increase in the number of submissions received by DLRS compared to last year.

As deep learning methods gradually become part of the main conference, accordingly changes the role of DLRS. From being the primary venue for any RecSys related deep learning research, DLRS will gradually become a forum of discussing pioneering work that uses recent techniques from deep learning or proposes novel recommendation related problems that can be solved with deep learning. Fortunately, we see both among this year's submission.

DLRS 2017 builds upon the positively received traits of DLRS 2016. DLRS 2017 is a fast paced workshop with a focus on high quality paper presentations and keynote. We welcome original research using deep learning technology for solving recommender systems related problems. The workshop centers around the use of Deep Learning technology in Recommender Systems and algorithms.

2 STATE OF THE ART

While using neural networks for recommender systems is not a recent trend – e.g. see the application of Restricted Boltzmann Machines for collaborative filtering [15] from ten years ago – the recent boom of deep learning revitalized this topic in the last two years. With deep learning for recommender related tasks becoming more and more popular, we have seen distinct research directions forming within the area during last year. The topics form around few seminal papers with subsequent work extending or improving upon those. Here we give a very brief overview of these research directions that can be considered as the state-of-the-art of this area.

Embeddings & 2vec models: Several papers focus on finding good embeddings (also known as distributed representations, or latent features) for items. This direction is not foreign for the recommendation domain, e.g. matrix factorization methods are also latent models that find "embeddings" for items and users[1]. Recent methods focus on item embeddings without user identification and can be used as the basis of more advanced methods or as item-to-item recommenders. Most of the models use some variation of Word2Vec [13] – originally devised for word embeddings – on event data [6]. Subsequent papers extend the model (e.g. with side information [19]) or examine the latent space [20].

Extracting features from heterogenous data: Hybrid factorization algorithms in the past often relied on constraining item features on metadata to overcome the weaknesses of pure collaborative filtering, e.g. the cold-start. With deep learning, it became easier to reliably extract useful features directly from content and use them in CF. Depending on the domain the content and thus its processing can greatly differ. [18] used convolutional neural networks (CNNs) to extract features from music to be used in a factor model. More recently [1] used recurrent neural networks (RNNs) to include textual data in their end-to-end CF model. Using

RecSys'17, August 27–31, 2017, Como, Italy.
 ISBN 978-1-4503-4652-8/17/08.
DOI: http://dx.doi.org/10.1145/3109859.3109953

[1] MF can also be represented as a very simplistic neural network with one hidden layer without any nonlinearities. The input is the one-hot encoded user ID, the hidden layer is the user feature vector and the output is the user's preference over the items.

images is also a popular choice to enhance recommendations in domains like fashion [7, 12] or video recommendation [10].

Deep collaborative filtering: A natural application of deep models is collaborative filtering. Leveraging the versatility of deep models, multiple types of interaction and/or context is often integrated. [4] describes the deep learning algorithm behind YouTube's related video recommendation. [2] combines deep models with a wide logistic regression model to find a balance between generalization and memorization in the app recommendation domain. [5] proposes a model in which user and item features co-evolve and follow the conceptional drift in the interaction data.

Autoencoders for CF: A subclass of deep CF methods use denoising autoencoders as the central component of their model. This approach is popular with several different variants proposed [21, 23, 24]. These models can also be extended with additional information sources, such as text [22].

Session-based recommendations with RNNs: Session-based recommendation is a good example for a task that is very important in real life applications, yet was hardly addressed due to its complexity when using classic methods. However, [9] showed that gated RNN (such as the Gated Recurrent Unit, GRU [3]) can be adapted to be a perfect fit for this task. Subsequent papers building on this idea either improve the original algorithm by (1) proposing loss functions suiting the problem better [8], (2) data augmentation [16]; or other extend it e.g. by adding context [17] or personalization [14].

3 THE WAY FORWARD

This area of research is very young, thus there is much room for improvement in the aforementioned research directions. Besides, up to this point only a fraction of well established deep learning models have been used for recommendations. We expect the community to delve into exploring the use of more recent approaches (e.g. generative models) in the near future. We also expect that novel or previously discarded recommendation problems will come forward now that the toolset for tackling them is available; similarly to how session-based recommendation was revitalized by [8]. The submissions to DLRS 2017 show traits of all three types of research.

4 SUMMARY

After its success in other domains, deep learning became popular in the recommender systems community last year. As researchers started to use deep learning methods in the RecSys domain, different research directions revealed themselves, such as creating better item embeddings or session based recommendations with recurrent neural networks. The DLRS workshop series aims to further speed up the acceptance of deep learning methods in the RecSys community. Even though deep learning papers are now also present at the main conference, DLRS 2017 is the venue for discussing novel research directions, applications of recent techniques of deep learning and bringing together researchers from the deep learning and recommender systems communities.

REFERENCES

[1] Trapit Bansal, David Belanger, and Andrew McCallum. 2016. Ask the GRU: Multi-task Learning for Deep Text Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. 107–114.

[2] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishi Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, and others. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*. ACM, 7–10.

[3] Kyunghyun Cho, Bart van Merrienboer, Çaglar Gülçehre, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. *CoRR* 1406.1078 (2014). http://arxiv.org/abs/1406.1078

[4] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. 191–198.

[5] Hanjun Dai, Yichen Wang, Rakshit Trivedi, and Le Song. 2016. Recurrent coevolutionary latent feature processes for continuous-time recommendation. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*. ACM, 29–34.

[6] Mihajlo Grbovic, Vladan Radosavljevic, Nemanja Djuric, Narayan Bhamidipati, Jaikit Savla, Varun Bhagwan, and Doug Sharp. 2015. E-commerce in your inbox: Product recommendations at scale. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1809–1818.

[7] Ruining He and Julian McAuley. 2016. VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback.. In *AAAI*. 144–150.

[8] Balázs Hidasi and Alexandros Karatzoglou. 2017. Recurrent Neural Networks with Top-k Gains for Session-based Recommendations. *arXiv preprint arXiv:1706.03847* (2017).

[9] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939; ICLR 2016* (2015).

[10] 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 (RecSys '16)*. 241–248.

[11] Alexandros Karatzoglou, Balázs Hidasi, Domonkos Tikk, Oren Sar-Shalom, Haggai Roitman, Bracha Shapira, and Lior Rokach. 2016. RecSys'16 Workshop on Deep Learning for Recommender Systems (DLRS). In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. 415–416.

[12] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 43–52.

[13] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. *CoRR* 1310.4546 (2013). http://arxiv.org/abs/1310.4546

[14] Massimo Quadrana, Alexandros Karatzoglou, Balázs Hidasi, and Paolo Cremonesi. 2017. Personalizing Session-based Recommendations with Hierarchical Recurrent Neural Networks. *arXiv preprint arXiv:1706.04148* (2017).

[15] Ruslan Salakhutdinov, Andriy Mnih, and Geoffrey Hinton. 2007. Restricted Boltzmann machines for collaborative filtering. In *ICML'07: 24th Int. Conf. on Machine Learning*. 791–798.

[16] Yong Kiam Tan, Xinxing Xu, and Yong Liu. 2016. Improved Recurrent Neural Networks for Session-based Recommendations. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems (DLRS 2016)*. 17–22.

[17] Bartłomiej Twardowski. 2016. Modelling Contextual Information in Session-Aware Recommender Systems with Neural Networks. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. 273–276.

[18] Aaron Van den Oord, Sander Dieleman, and Benjamin Schrauwen. 2013. Deep content-based music recommendation. In *Advances in Neural Information Processing Systems*. 2643–2651.

[19] 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 (RecSys '16)*. 225–232.

[20] Jeroen B. P. 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 (DLRS 2016)*. 23–28.

[21] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative deep learning for recommender systems. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1235–1244.

[22] Hao Wang, SHI Xingjian, and Dit-Yan Yeung. 2016. Collaborative recurrent autoencoder: Recommend while learning to fill in the blanks. In *Advances in Neural Information Processing Systems*. 415–423.

[23] 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.

[24] Yin Zheng, Cailiang Liu, Bangsheng Tang, and Hanning Zhou. 2016. Neural Autoregressive Collaborative Filtering for Implicit Feedback. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*. ACM, 2–6.