DLRS 2018 - Third 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
Intuit
Tel-Aviv, Israel
oren.sarshalom@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

Flavian Vasile
Criteo
Paris, France
f.vasile@criteo.com

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

ABSTRACT

Deep learning is now an integral part of recommender systems, but the research is still in its early phase. New research topics pop up frequently and established topics are extended in new, interesting directions. DLRS 2018 is a venue for pioneering work in the intersection of deep learning and recommender systems research.

KEYWORDS

deep learning; neural networks; recommender systems

ACM Reference Format:
Balázs Hidasi, Alexandros Karatzoglou, Oren Sar-Shalom, Bracha Shapira, Domonkos Tikk, Flavian Vasile, and Sander Dieleman. 2018. DLRS 2018 - Third Workshop on Deep Learning for Recommender Systems. In *Twelfth ACM Conference on Recommender Systems (RecSys '18), October 2–7, 2018, Vancouver, BC, Canada.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3240323.3240333

1 INTRODUCTION

DLRS 2018 is the third instance in the Deep Learning for Recommender Systems workshop series [8, 11]. As the acceptance of deep learning changes within the recommender systems community, so does the role of the DLRS workshop. The goal of the first DLRS was to popularize the idea of using deep learning technology in recommender systems, as research in this topic was few and far between before 2016. The second workshop's goal was to further strengthen the acceptance of this idea. While in 2017 deep learning was already part of the main conference, DLRS 2017 gave the additional needed exposure to the topic, including novel research directions and domains within the field.

RecSys '18, October 2–7, 2018, Vancouver, BC, Canada

ACM ISBN 978-1-4503-5901-6/18/10.
https://doi.org/10.1145/3240323.3240333

The main conference in 2018 received a record number of deep learning related submissions. This resulted in DLRS not breaking the record of 2017 in the number of submissions received, as it is now similar to what it was in 2016, which still implies massive interest. The topic is getting more and more integrated into the main conference, which is the original goal of the DLRS series.

By 2018, using deep learning in recommender systems became well accepted. Even then, we are still in the early days of the research, since the resurgence of using neural networks to solve recommendation tasks is only a few years old. There are many unexplored directions and untapped potential in deep learning for the RecSys community. This includes both (1) revisiting domains that were considered hard before deep learning was available as a tool and (2) looking at recent advances in deep learning research and using them for making recommender systems better. Therefore the main role of DLRS 2018 is to provide a venue and exposure to pioneering work as well as for papers in recently established research areas within the field.

DLRS 2018 builds upon the positively received traits of previous DLRS workshops. DLRS 2018 is a fast paced workshop with a focus on high quality paper presentations and keynotes. 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 PROGRESS IN THE FIELD

Last year we identified four main research topics that have been established within the field during 2015-2017. These topics are *learning item representations* [4, 18]; *feature extraction* from heterogeneous data including audio [17], images [5] and text [1]; *deep collaborative filtering* [19, 21]; *session-based recommendations* with recurrent neural networks (RNNs) [6, 7]; and of course the application of the aforementioned in live recommenders e.g. [2].

Since then, these research topics have been extended in many different directions. E.g. session-based RNN recommenders were extended to be context-aware by [16]; to handle information across user sessions [15]; to predict intent from browsing data using seq2seq models and attention [13]. Some papers transcend the categorization and combine elements from different research lines. E.g. [14] extends item representation learning by also using feature extraction techniques.

New research topics have also started emerging recently:

Deep generative models, such as Variational Autoencoders (VAE) [10] and Generative Adversarial Networks (GAN) [3] have been one of the most promising additions to deep learning in recent years. Lately, recommender systems researchers also started using these models. This research has been pioneered by IRGAN [20] and variational autoencoders for collaborative filtering (CF-VAE) [12].

IRGAN unifies of the generative and discriminative approach in informartion retrieval by setting up a GAN like framework. The generator selects relevant documents from a predefined set, conditioned on relevance information and the query. No new items are generated, which is a key difference, compared to the original GAN. The discriminator distinguishes generated and real documents. By competing, the generator acts as an adaptive negative sampler for the discriminative model. The framework can be used for different tasks, including recommendation.

CF-VAE replaces traditional (deterministic) autoencoders in recommender models with variational autoencoders. Thus the encoder captures the distribution of the code and the decoder remaps a sample from this distribution. To adapt to the domain, the decoder uses multinomial likelihood. Variational autoencoders can be also used for other tasks within the recommender systems domain, such as slate recommendation. [9]

Deep reinforcement learning is a promising direction for online recommender systems. It is hard to infer user goals and intent from static snapshots of data. While session-based recommendation are adaptive to a degree as they change their predictions when the user's session changes, they are also limited by focusing on a single aspect of the sessions (e.g. next click prediction), even though different user goals would require different approaches (e.g. next click prediction can be appropriate if the user browses, but not so much if he is looking for a specific item). Deep reinforcement learning can learn highly adaptive policies that can solve this problem. Reinforcement learning also enables optimization for long-term online KPIs (e.g. revenue) and sparse rewards (e.g. purchase events). There has been some research in this direction, but at the moment proposed methods haven't been evaluated in a proper way. Evaluation seems to be one of the most significant problems to be solved, because due to their nature reinforcement learning algorithms can not be evaluated in an offline manner.

3 SUMMARY

Deep learning is now an integral part of the toolbox of researchers of recommender systems and algorithms, and it is widely accepted by the RecSys community. The field has a few already established research lines and many other up and coming interesting directions, such as using generative models or reinforcement learning. Yet, the overall research is still in its early years, so rapid progress is to be expected as researchers work on finding the best techniques suitable for RecSys problems. DLRS 2018 aims to give exposure to pioneering work in the intersection of deep learning and recommender systems and to bring 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] 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.

[3] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.

[4] 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.

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

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

[7] 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).

[8] Balázs Hidasi, Alexandros Karatzoglou, Oren Sar-Shalom, Sander Dieleman, Bracha Shapira, and Domonkos Tikk. 2017. DLRS 2017: Second Workshop on Deep Learning for Recommender Systems. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17)*. ACM, New York, NY, USA, 370–371. https://doi.org/10.1145/3109859.3109953

[9] Ray Jiang, Sven Gowal, Timothy A Mann, and Danilo J Rezende. 2018. Optimizing Slate Recommendations via Slate-CVAE. *arXiv preprint arXiv:1803.01682* (2018).

[10] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082* (2014).

[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] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational Autoencoders for Collaborative Filtering. *arXiv preprint arXiv:1802.05814* (2018).

[13] Pablo Loyola, Chen Liu, and Yu Hirate. 2017. Modeling User Session and Intent with an Attention-based Encoder-Decoder Architecture. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17)*. ACM, New York, NY, USA, 147–151. https://doi.org/10.1145/3109859.3109917

[14] Thomas Nedelec, Elena Smirnova, and Flavian Vasile. 2017. Specializing Joint Representations for the task of Product Recommendation. In *Proceedings of the 2nd Workshop on Deep Learning for Recommender Systems*. ACM, 10–18.

[15] 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).

[16] Elena Smirnova and Flavian Vasile. 2017. Contextual Sequence Modeling for Recommendation with Recurrent Neural Networks. In *Proceedings of the 2nd Workshop on Deep Learning for Recommender Systems (DLRS 2017)*. 2–9. https://doi.org/10.1145/3125486.3125488

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

[18] 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.

[19] 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.

[20] Jun Wang, Lantao Yu, Weinan Zhang, Yu Gong, Yinghui Xu, Benyou Wang, Peng Zhang, and Dell Zhang. 2017. Irgan: A minimax game for unifying generative and discriminative information retrieval models. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*. ACM, 515–524.

[21] 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.