Hybrid Session-based News Recommendation using Recurrent Neural Networks

Gabriel de Souza P. Moreira [1 2] **Dietmar Jannach** [3] **Adilson Marques da Cunha** [2]

Abstract

We describe a hybrid meta-architecture – the CHAMELEON – for session-based news recommendation that is able to leverage a variety of information types using Recurrent Neural Networks. We evaluated our approach on two public datasets, using a temporal evaluation protocol that simulates the dynamics of a news portal in a realistic way. Our results confirm the benefits of modeling the sequence of session clicks with RNNs and leveraging side information about users and articles, resulting in significantly higher recommendation accuracy and catalog coverage than other session-based algorithms.

1. Introduction

Recommender systems help users to deal with information overload by providing tailored item suggestions to them. One of the earliest application domains is the recommendation of online *news* (Karimi et al., 2018). News recommendation is sometimes considered as being particularly difficult, as it has a number of distinctive characteristics (Zheng et al., 2018). Among other challenges, news recommenders have to deal with a constant stream of news articles being published, which at the same time can become outdated very quickly. Another challenge is that the system often cannot rely on long-term user preference profiles. Typically, most users are not logged in and their short-term reading interests must be estimated from only a few logged interactions, leading to a *session-based recommendation problem* (Quadrana et al., 2018).

In recent years, we observed an increased interest in the problem of session-based recommendation, where the task is to recommend relevant items given an ongoing user session. Recurrent Neural Networks (RNN) represent a natural choice for sequence prediction tasks, as they can learn models from sequential data. *GRU4Rec* (Hidasi et al., 2016) was one of the first neural session-based recommendation techniques, and a number of other approaches were proposed in recent years that rely on deep learning architectures, as in (Liu et al., 2018; Li et al., 2017).

However, as shown in (Jannach & Ludewig, 2017; Ludewig & Jannach, 2018; Ludewig et al., 2019), neural approaches that only rely on logged item interactions have certain limitations and they can, depending on the experimental setting, be outperformed by simpler approaches based, e.g., nearest-neighbor techniques. Differently from previous works, we therefore leverage multiple types of side information with RNNs, including textual article embeddings, as well as the context of users and articles. Furthermore, we propose a meta-architecture to address the aforementioned challenges of recommending in the news domain.

2. Technical Contribution

Our approach is based upon CHAMELEON (Moreira, 2018; Moreira et al., 2018; 2019b;a), which is a Deep Learning Meta-Architecture for News Recommendation. It supports session-based news recommendation scenarios, modeling the sequence of user clicks using Recurrent Neural Networks. The resulting system is a hybrid recommender system, which addresses the permanent user and item cold-start problem in the news domain by leveraging the textual content of news articles, the article context (e.g., recent popularity and recency) and the user context (e.g., time, location, device, previous session clicks).

Figure 1 shows our instantiation of the *CHAMELEON* framework with its two main modules: the *ACR* module on the left creates distributed representations of articles' textual content. The *NAR* module on the right is responsible to generate next-click predictions. The *NAR* module is trained on a ranking loss function based on similarities, which is designed to recommend fresh articles without retraining. As proposed for the *DSSM* loss function (Huang et al., 2013), it is trained to maximize the likelihood of correctly predicting the next clicked article given a user session.

[1]NVIDIA - 2788 San Tomas Expressway, Santa Clara, CA, 95051. [2]Department of Electrical Engineering and Computing, Instituto Tecnológico de Aeronáutica (ITA) - São José dos Campos, São Paulo, Brazil. [3]Department of Applied Informatics, University of Klagenfurt, Austria. Correspondence to: Gabriel de Souza P. Moreira <gspmoreira@gmail.com>.

Proceedings of the LatinX in AI Research (LXAI) at ICML 2020, Vienna, Austria, 2020.

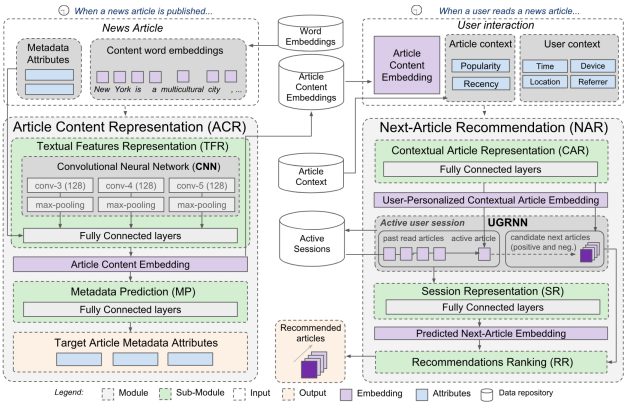

Figure 1. An architecture instantiation of *CHAMELEON*

3. Evaluation Protocol

The evaluation was performed as follows: (1) Recommenders are continuously trained on users' sessions ordered by time and grouped by hours. Each five hours, the recommenders are evaluated on sessions from the next hour; (2) For each session in the evaluation set, we incrementally revealed one click after the other to the recommender; and (3) For each click to be predicted, we created a set containing 50 negative samples articles (not clicked by the user in her session) and compute top-N metrics about accuracy, item coverage, and novelty.

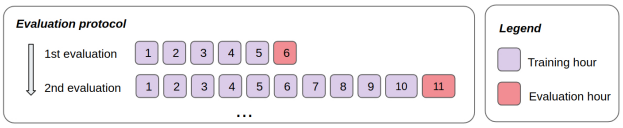

Figure 2. Illustration of the evaluation protocol. After training for 5 hours, we evaluate using the sessions of the next hour.

Experiments were performed with two public real-world datasets from the *G1* (Moreira et al., 2018) and *Adressa* (Gulla et al., 2017) news portals, described in Table 1.

We evaluated the following recommendation quality factors for the top-N ranked items: accuracy – Hit Rate (*HR@n*) and Mean Reciprocal Rank (MRR@n); item coverage – (*COV*) (i.e., the number of distinct articles that appeared in any top-N list divided by the number of recommendable articles); and novelty – *ESI-R*, which is based on item popularity, returning higher values when recommending long-tail items.

As baseline algorithms for session-based recommendation, we have used: two neural approaches (*GRU4Rec* (Hidasi et al., 2016) and *SR-GNN* (Wu et al., 2019)); association rules-based methods (*Co-Occurrence (CO)* and *Sequential Rules (SR)* (Ludewig & Jannach, 2018)); neighborhood-based methods (*Item-kNN* (Hidasi et al., 2016) and *Vector Multiplication Session-Based kNN (V-SkNN)* (Jannach & Ludewig, 2017)); and two other classical methods (*Recently Popular (RP)* (Ludmann, 2017) and *Content-Based (CB)*.

Table 1. Statistics of the datasets used for the experiments.

	Globo.com (G1)	*Adressa*
Language	Portuguese	Norwegian
Period (days)	16	16
# users	322,897	314,661
# sessions	1,048,594	982,210
# clicks	2,988,181	2,648,999
# articles	46,033	13,820
Avg. Sessions length (clicks)	2.84	2.70

4. Results

The evaluation results are presented in Table 2, as originally reported in (Moreira et al., 2019a). The best results for a metric are printed in bold face and marked with * if they are significantly different [1] from all other algorithms.

Table 2. Evaluation of recommendation quality factors

Recommender	*HR@10*	*MRR@10*	*COV@10*	*ESI-R@10*
G1 dataset				
CHAMELEON	**0.6738***	**0.3458***	0.6373	6.4177
SR	0.5900	0.2889	0.2763	5.9747
Item-kNN	0.5707	0.2801	0.3913	6.5909
CO	0.5689	0.2626	0.2499	5.5728
V-SkNN	0.5467	0.2494	0.1355	5.1760
SR-GNN	0.5144	0.2467	0.3196	5.4280
GRU4Rec	0.4669	0.2092	0.6333	5.2332
RP	0.4577	0.1993	0.0218	4.4904
CB	0.3643	0.1676	**0.6774**	**8.1531***
Adressa dataset				
CHAMELEON	**0.7018***	**0.3421***	0.7926	5.3410
SR	0.6288	0.3022	0.4604	5.4443
Item-kNN	0.6179	0.2819	0.5314	5.4675
CO	0.6131	0.2768	0.4220	5.0789
V-SkNN	0.6140	0.2723	0.1997	4.6018
SR-GNN	0.6122	0.2991	0.5197	5.1013
GRU4Rec	0.4958	0.2200	0.5143	5.0571
RP	0.5648	0.2481	0.0542	4.1465
CB	0.3307	0.1253	**0.8875***	**7.6715***

5. Conclusion

CHAMELEON was specifically designed to address news recommendation challenges such as (a) the short lifetime of the recommendable items and (b) the lack of longer-term preference profiles of the users. In the extensive experiments performed, CHAMELEON was able to provide recommendations with much higher accuracy than all other evaluated algorithms, and it led to the second best item coverage.

[1] As errors around the reported averages were normally distributed, we used paired Student's t-tests with Bonferroni correction at $p < 0.001$ for significance tests.

References

Gulla, J. A., Zhang, L., Liu, P., Özgöbek, Ö., and Su, X. The adressa dataset for news recommendation. In *Proceedings of the International Conference on Web Intelligence (WI'17)*, pp. 1042–1048, 2017.

Hidasi, B., Karatzoglou, A., Baltrunas, L., and Tikk, D. Session-based recommendations with recurrent neural networks. In *Proceedings of Fourth International Conference on Learning Representations (ICLR'16)*, 2016.

Huang, P.-S., He, X., Gao, J., Deng, L., Acero, A., and Heck, L. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM International Conference on Conference on Information & Knowledge Management*, pp. 2333–2338, 2013.

Jannach, D. and Ludewig, M. When recurrent neural networks meet the neighborhood for session-based recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys'17)*, pp. 306–310, 2017.

Karimi, M., Jannach, D., and Jugovac, M. News recommender systems–survey and roads ahead. *Information Processing & Management*, 54(6):1203–1227, 2018.

Li, J., Ren, P., Chen, Z., Ren, Z., Lian, T., and Ma, J. Neural attentive session-based recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, CIKM '17, pp. 1419–1428, 2017.

Liu, Q., Zeng, Y., Mokhosi, R., and Zhang, H. STAMP: short-term attention/memory priority model for session-based recommendation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD'18*, pp. 1831–1839, 2018.

Ludewig, M. and Jannach, D. Evaluation of session-based recommendation algorithms. *User-Modeling and User-Adapted Interaction*, 28(4–5):331–390, 2018.

Ludewig, M., Mauro, N., Latifi, S., and Jannach, D. Performance comparison of neural and non-neural approaches to session-based recommendation. In *Proceedings of the 2019 ACM Conference on Recommender Systems (RecSys 2019)*, 2019.

Ludmann, C. A. Recommending news articles in the clef news recommendation evaluation lab with the data stream management system odysseus. In *Working notes of the Conference and Labs of the Evaluation Forum (CLEF'17)*, 2017.

Moreira, G. d. S. P. Chameleon: A deep learning meta-architecture for news recommender systems. In *Proceedings of the Doctoral Symposium at 12th ACM Conference on Recommender Systems (RecSys'18)*, pp. 578–583, 2018.

Moreira, G. d. S. P., Ferreira, F., and Cunha, A. M. d. News session-based recommendations using deep neural networks. In *Proceedings of the 3rd Workshop on Deep Learning for Recommender Systems (DLRS) at ACM RecSys'18*, pp. 15–23, 2018.

Moreira, G. D. S. P., Jannach, D., and da Cunha, A. M. Contextual hybrid session-based news recommendation with recurrent neural networks. *IEEE Access*, 7:169185–169203, 12 2019a. ISSN 2169-3536. doi: 10.1109/ACCESS.2019.2954957.

Moreira, G. D. S. P., Jannach, D., and da Cunha, A. M. On the importance of news content representation in hybrid neural session-based recommender systems. In *Proceedings of the 7th International Workshop on News Recommendation and Analytics (INRA 2019) @ RecSys 2019*, volume Vol-2554. CEUR-WS, 09 2019b.

Quadrana, M., Cremonesi, P., and Jannach, D. Sequence-aware recommender systems. *ACM Computing Surveys (CSUR)*, 51(4):66, 2018.

Wu, S., Tang, Y., Zhu, Y., Wang, L., Xie, X., and Tan, T. Session-based recommendation with graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 346–353, 2019.

Zheng, G., Zhang, F., Zheng, Z., Xiang, Y., Yuan, N. J., Xie, X., and Li, Z. Drn: A deep reinforcement learning framework for news recommendation. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, pp. 167–176, 2018.