

The content providers (e.g. artists and publishers) can also gain insights in the taste and trend of different user groups.

The Session-based Heterogeneous graph Embedding (SHE) learns the feature representation for each session, which may be applicable to inferring account attributes, predicting session life-cycle, and detecting abnormal accounts, in addition to user identification. Moreover, the proposed SHE-UI can be applied to any activity and event sequences, allowing for incorporation of metadata. For example, one may study activity traces from smart devices and apply SHE-UI for activity recognition.

ACKNOWLEDGMENT

We would like to thank the anonymous reviewers for their helpful comments. The work was partially supported by NIH U01HG008488, NIH R01GM115833, NIH U54GM114833, and NSF IIS-1313606. This work was also supported by Ministry of Science and Technology (MOST) Taiwan: MOST 107-2636-E-006-002 and MOST 106-3114-E-006-002, and by Academia Sinica AS-107-TP-A05.

REFERENCES

- [1] Michal Aharon, Eshcar Hillel, Amit Kagian, Ronny Lempel, Hayim Makabee, and Raz Nissim. 2015. Watch-It-Next: A Contextual TV Recommendation System. In *Proceedings of European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD '15)*. 180–195.
- [2] Leman Akoglu, Rishi Chandy, and Christos Faloutsos. 2013. Opinion Fraud Detection in Online Reviews by Network Effects. *ICWSM 13* (2013), 2–11.
- [3] Yuichiro Anzai. 2012. *Pattern recognition and machine learning*. Elsevier.
- [4] David Arthur and Sergei Vassilvitskii. 2007. k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1027–1035.
- [5] Payal Bajaj and Sumit Shekhar. 2016. Experience Individualization on Online TV Platforms Through Persona-based Account Decomposition. In *Proceedings of ACM International Conference on Multimedia (MM '16)*. 252–256.
- [6] Jesus Bobadilla, Fernando Javier Ortega, Antonio Hernandez, and Abraham Gutierrez. 2013. Recommender Systems Survey. *Knowledge-Based Systems* 46 (2013), 109–132.
- [7] Léon Bottou. 1991. Stochastic gradient learning in neural networks. *Proceedings of Neuro-Nimes* 91, 8 (1991).
- [8] Qiang Cao, Xiaowei Yang, Jieqi Yu, and Christopher Palow. 2014. Uncovering large groups of active malicious accounts in online social networks. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 477–488.
- [9] Oscar Celma. 2010. *Music Recommendation and Discovery in the Long Tail*. Springer.
- [10] Shiyu Chang, Wei Han, Jiliang Tang, Guo-Jun Qi, Charu C Aggarwal, and Thomas S Huang. 2015. Heterogeneous network embedding via deep architectures. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 119–128.
- [11] Ting Chen and Yizhou Sun. 2017. Task-guided and path-augmented heterogeneous network embedding for author identification. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, 295–304.
- [12] W Bruce Croft, Donald Metzler, and Trevor Strohmann. 2010. *Search engines: Information Retrieval in Practice*. Pearson Education.
- [13] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 135–144.
- [14] Brian S Everitt. 2006. *The Cambridge dictionary of statistics*. Cambridge University Press.
- [15] Brendan J Frey and Delbert Dueck. 2007. Clustering by passing messages between data points. *science* 315, 5814 (2007), 972–976.
- [16] Prem Gopalan, Francisco J Ruiz, Rajesh Ranganath, and David M Blei. 2014. Bayesian Nonparametric Poisson Factorization for Recommendation Systems. In *AISTATS*. 275–283.
- [17] Prem K Gopalan, Laurent Charlin, and David Blei. 2014. Content-based recommendations with poisson factorization. In *NIPS*. 3176–3184.
- [18] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. 855–864.
- [19] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. (2016).
- [20] Balazs Hidasi and Domonkos Tikk. 2016. General Factorization Framework for Context-aware Recommendations. *Data Mining and Knowledge Discovery (DMKD)* 30, 2 (March 2016), 342–371.
- [21] Bryan Hooi, Neil Shah, Alex Beutel, Stephan Günnemann, Leman Akoglu, Mohit Kumar, Disha Makhija, and Christos Faloutsos. 2016. Birdnest: Bayesian inference for ratings-fraud detection. In *Proceedings of the 2016 SIAM International Conference on Data Mining*. SIAM, 495–503.
- [22] Rob J Hyndman and Anne B Koehler. 2006. Another look at measures of forecast accuracy. *International journal of forecasting* 22, 4 (2006), 679–688.
- [23] Nitin Jindal, Bing Liu, and Ee-Peng Lim. 2010. Finding unusual review patterns using unexpected rules. In *Proceedings of the 19th ACM international conference on Information and knowledge management*. ACM, 1549–1552.
- [24] Niels Landwehr. 2008. Modeling interleaved hidden processes. In *Proceedings of the 25th international conference on Machine learning*. ACM, 520–527.
- [25] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proceedings of Workshop at International Conference on Learning Representations*.
- [26] 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.
- [27] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of Empirical Methods in Natural Language Processing (EMNLP '14)*, Vol. 14. 1532–1543.
- [28] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '14)*. 701–710.
- [29] David Martin Powers. 2011. Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation. (2011).
- [30] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*. AUAI Press, 452–461.
- [31] Guy Shani, David Heckerman, and Ronen I. Brafman. 2005. An MDP-Based Recommender System. *Journal of Machine Learning Research (JMLR)* 6 (2005), 1265–1295.
- [32] Yue Shi, Alexandros Karatzoglou, Linas Baltrunas, Martha Larson, Nuria Oliver, and Alan Hanjalic. 2012. CLIMF: learning to maximize reciprocal rank with collaborative less-is-more filtering. In *RecSys*. ACM, 139–146.
- [33] Adish Singla, Ryan W. White, Ahmed Hassan, and Eric Horvitz. 2014. Enhancing Personalization via Search Activity Attribution. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '14)*. 1063–1066.
- [34] Alexander Strehl and Joydeep Ghosh. 2002. Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *Journal of machine learning research* 3, Dec (2002), 583–617.
- [35] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15)*. 1067–1077.
- [36] Bartłomiej Twardowski. 2016. Modelling Contextual Information in Session-Aware Recommender Systems with Neural Networks. In *Proceedings of the 10th ACM International Conference on Recommender Systems (RecSys '16)*. 273–276.
- [37] Koen Verstrepen and Bart Goethals. 2015. Top-N Recommendation for Shared Accounts. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys '15)*. 59–66.
- [38] Zhijian Wang, Yan Yang, Liang He, and Junzhong Gu. 2014. User Identification within a Shared Account: Improving IP-TV Recommender Performance. In *Proceedings of the 18th East European Conference on Advances in Databases and Information Systems (ADBIS '14)*. 219–233.
- [39] Markus Weimer, Alexandros Karatzoglou, and Alex Smola. 2008. Improving maximum margin matrix factorization. *Machine Learning* 72, 3 (2008), 263–276.
- [40] Ryan W. White and Ahmed Hassan Awadallah. 2015. Personalizing Search on Shared Devices. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '15)*. 523–532.
- [41] Ryan W. White, Ahmed Hassan, Adish Singla, and Eric Horvitz. 2014. From Devices to People: Attribution of Search Activity in Multi-user Settings. In *Proceedings of the 23rd International Conference on World Wide Web (WWW '14)*. 431–442.
- [42] Yan Yang, Qinmin Hu, Liang He, Minjie Ni, and Zhijian Wang. 2015. Adaptive Temporal Model for IPTV Recommendation. In *Proceedings of the 16th International Conference on Web-Age Information Management (WAIM '15)*. 260–271.
- [43] Amy Zhang, Nadia Fawaz, Stratis Ioannidis, and Andrea Montanari. 2012. Guess Who Rated This Movie: Identifying Users Through Subspace Clustering. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI'12)*. 944–953.
- [44] Yafeng Zhao, Jian Cao, and Yudong Tan. 2016. Passenger Prediction in Shared Accounts for Flight Service Recommendation. In *Proceedings of the 10th Asia-Pacific Services Computing Conference on Advances in Services Computing (APSCC '16)*. 159–172.