Machine learning tasks and representations for heterogeneous information networks

Fang, Y.

Publication date
2023

Bibliography

[1] A. Ahmed, N. Shervashidze, S. M. Narayanamurthy, V. Josifovski, and A. J. Smola. Distributed large-scale natural graph factorization. In *WWW*, pages 37–48, 2013. (Cited on pages 12 and 41.)

[2] L. Akoglu, H. Tong, and D. Koutra. Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29(3):626–688, 2015. (Cited on page 31.)

[3] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *NeurIPS*, pages 585–591, 2001. (Cited on pages 12 and 41.)

[4] R. Bian, Y. S. Koh, G. Dobbie, and A. Divoli. Network embedding and change modeling in dynamic heterogeneous networks. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, pages 861–864, 2019. (Cited on pages 13 and 26.)

[5] P. Bonacich. Power and centrality: A family of measures. *American Journal of Sociology*, 92(5): 1170–1182, 1987. (Cited on pages 44, 68, and 89.)

[6] A. J. Bose, A. Jain, P. Molino, and W. L. Hamilton. Meta-Graph: Few shot link prediction via meta learning. *CoRR*, abs/1912.09867, 2019. (Cited on pages 61, 62, 65, and 75.)

[7] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. In *ICLR*, 2014. (Cited on pages 42 and 64.)

[8] H. Cai, V. W. Zheng, and K. C. Chang. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Trans. Knowl. Data Eng.*, 30(9):1616–1637, 2018. (Cited on page 1.)

[9] S. Cao, W. Lu, and Q. Xu. Grarep: Learning graph representations with global structural information. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM 2015, Melbourne, VIC, Australia, October 19 - 23, 2015*, pages 891–900, 2015. (Cited on page 12.)

[10] P. Cheng, Y. Li, X. Zhang, L. Chen, D. E. Carlson, and L. Carin. Dynamic embedding on textual networks via a gaussian process. In *AAAI, New York, NY, USA, February 7-12, 2020*, pages 7562–7569. AAAI Press, 2020. (Cited on pages 88 and 93.)

[11] M. A. A. Cox and T. F. Cox. *Multidimensional Scaling*. Springer Berlin Heidelberg, 2008. (Cited on pages 12 and 41.)

[12] Y. Cui, H. Sun, Y. Zhao, H. Yin, and K. Zheng. Sequential-knowledge-aware next POI recommendation: A meta-learning approach. *ACM Trans. Inf. Syst.*, 40(2):23:1–23:22, 2022. (Cited on page 65.)

[13] A. M. Dai and Q. V. Le. Semi-supervised sequence learning. In *NeurIPS*, pages 3079–3087, 2015. (Cited on page 39.)

[14] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *NeurIPS*, pages 3837–3845, 2016. (Cited on pages 42 and 64.)

[15] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2019. (Cited on pages 40, 46, and 47.)

[16] K. Ding, J. Li, and H. Liu. Interactive anomaly detection on attributed networks. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, VIC, Australia, February 11-15, 2019*, pages 357–365. ACM, 2019. (Cited on page 78.)

[17] K. Ding, Q. Zhou, H. Tong, and H. Liu. Few-shot network anomaly detection via cross-network meta-learning. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, pages 2448–2456. ACM / IW3C2, 2021. (Cited on pages 61, 62, 66, and 75.)

[18] Y. Dong, N. V. Chawla, and A. Swami. metapath2vec: Scalable representation learning for heterogeneous networks. In *KDD*, pages 135–144, 2017. (Cited on pages 1, 10, 12, 25, 41, and 49.)

[19] Y. Dong, Z. Hu, K. Wang, Y. Sun, and J. Tang. Heterogeneous network representation learning. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 4861–4867. ijcai.org, 2020. (Cited on page 41.)

[20] M. Elseidy, E. Abdelhamid, and et al. GRAMI: frequent subgraph and pattern mining in a single large graph. *PVLDB*, 7(7):517–528, 2014. (Cited on page 16.)

[21] Y. Fan, Y. Ye, Q. Peng, J. Zhang, Y. Zhang, X. Xiao, C. Shi, Q. Xiong, F. Shao, and L. Zhao. Metagraph aggregated heterogeneous graph neural network for illicit traded product identification in underground market. In *20th IEEE International Conference on Data Mining, ICDM 2020, Sorrento, Italy, November 17-20, 2020*, pages 132–141. IEEE, 2020. (Cited on page 12.)

[22] Y. Fang, X. Zhao, Y. Chen, W. Xiao, and M. de Rijke. PF-HIN: Pre-training for heterogeneous information networks. *IEEE Trans. Knowl. Data Eng.*, 35(8):8372–8385.

[23] Y. Fang, X. Zhao, P. Huang, W. Xiao, and M. de Rijke. M-HIN: complex embeddings for heterogeneous

information networks via metagraphs. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, pages 913–916, 2019. (Cited on page 10.)

[24] Y. Fang, X. Zhao, P. Huang, W. Xiao, and M. de Rijke. Scalable representation learning for dynamic heterogeneous information networks via metagraphs. *ACM Trans. Inf. Syst.*, 40(4):64:1–64:27, 2022.

[25] Y. Fang, X. Zhao, W. Xiao, and M. de Rijke. Few-shot learning for heterogeneous information networks. *ACM Trans. Inf. Syst.*, Under review.

[26] Y. Fang, X. Zhao, W. Xiao, and M. de Rijke. Prompt learning for textual heterogeneous information networks. *IEEE Trans. Knowl. Data Eng.*, Under review.

[27] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR, 2017. (Cited on pages 62 and 75.)

[28] T. Fu, W. Lee, and Z. Lei. Hin2vec: Explore meta-paths in heterogeneous information networks for representation learning. In *CIKM*, pages 1797–1806, 2017. (Cited on pages 1, 10, 12, 26, 41, and 49.)

[29] H. Gao, Z. Wang, and S. Ji. Large-scale learnable graph convolutional networks. In *KDD*, pages 1416–1424, 2018. (Cited on pages 12, 42, and 64.)

[30] P. Goyal, N. Kamra, X. He, and Y. Liu. Dyngem: Deep embedding method for dynamic graphs. *CoRR*, abs/1805.11273, 2018. (Cited on pages 13 and 26.)

[31] P. Goyal, S. R. Chhetri, and A. Canedo. dyngraph2vec: Capturing network dynamics using dynamic graph representation learning. *Knowledge Based Systems*, 187, 2020. (Cited on pages 13 and 26.)

[32] M. Grootendorst. KeyBERT: Minimal keyword extraction with BERT. `https://doi.org/10.5281/zenodo.4461265`, 2020. (Cited on page 95.)

[33] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *KDD*, pages 855–864, 2016. (Cited on pages 1, 9, 12, 25, 41, 49, and 50.)

[34] Q. Guo, X. Zhao, Y. Fang, S. Yang, X. Lin, and D. Ouyang. Learning hypersphere for few-shot anomaly detection on attributed networks. In *CIKM 2022: 31st ACM International Conference on Information and Knowledge Management*, pages 635–645. ACM, 2022.

[35] Z. Guo, C. Zhang, W. Yu, J. Herr, O. Wiest, M. Jiang, and N. V. Chawla. Few-shot graph learning for molecular property prediction. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, pages 2559–2567. ACM / IW3C2, 2021. (Cited on pages 61, 62, 66, and 75.)

[36] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, pages 1735–1742. IEEE Computer Society, 2006. (Cited on page 64.)

[37] H. Hafidi, M. Ghogho, P. Ciblat, and A. Swami. GraphCL: Contrastive self-supervised learning of graph representations. *CoRR*, abs/2007.08025, 2020. (Cited on page 65.)

[38] W. L. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 1024–1034, 2017. (Cited on pages 42, 49, 64, and 75.)

[39] K. Hassani and A. H. K. Ahmadi. Contrastive multi-view representation learning on graphs. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 4116–4126. PMLR, 2020. (Cited on page 65.)

[40] K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735. IEEE, 2020. (Cited on page 64.)

[41] M. Henaff, J. Bruna, and Y. LeCun. Deep convolutional networks on graph-structured data. *CoRR*, abs/1506.05163, 2015. (Cited on pages 42 and 64.)

[42] R. D. Hjelm, A. Fedorov, S. Lavoie-Marchildon, K. Grewal, P. Bachman, A. Trischler, and Y. Bengio. Learning deep representations by mutual information estimation and maximization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. (Cited on page 71.)

[43] J. Howard and S. Ruder. Universal language model fine-tuning for text classification. In *ACL*, pages 328–339, 2018. (Cited on page 39.)

[44] W. Hu, B. Liu, J. Gomes, M. Zitnik, P. Liang, V. S. Pande, and J. Leskovec. Strategies for pre-training graph neural networks. In *ICLR*, 2020. (Cited on page 42.)

[45] Z. Hu, C. Fan, T. Chen, K. Chang, and Y. Sun. Pre-training graph neural networks for generic structural feature extraction. *CoRR*, abs/1905.13728, 2019. (Cited on page 43.)

[46] Z. Hu, Y. Dong, K. Wang, K. Chang, and Y. Sun. GPT-GNN: generative pre-training of graph neural

networks. In *KDD*, pages 1857–1867, 2020. (Cited on pages 39, 43, and 50.)
[47] Z. Hu, Y. Dong, K. Wang, and Y. Sun. Heterogeneous graph transformer. In *WWW*, pages 2704–2710, 2020. (Cited on pages 42, 43, 49, and 75.)
[48] K. Huang and M. Zitnik. Graph meta learning via local subgraphs. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. (Cited on pages 61, 65, 67, and 75.)
[49] P. Huang, X. Zhao, M. Hu, Y. Fang, X. Li, and W. Xiao. Extract-select: A span selection framework for nested named entity recognition with generative adversarial training. In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 85–96. Association for Computational Linguistics, 2022.
[50] Z. Huang and N. Mamoulis. Heterogeneous information network embedding for meta path based proximity. *CoRR*, abs/1701.05291, 2017. (Cited on pages 1, 10, 12, 41, and 49.)
[51] X. Jiang, T. Jia, Y. Fang, C. Shi, Z. Lin, and H. Wang. Pre-training on large-scale heterogeneous graph. In *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14-18, 2021*, pages 756–766. ACM, 2021. (Cited on pages 43 and 50.)
[52] Z. Jiang, F. F. Xu, J. Araki, and G. Neubig. How can we know what language models know. *Trans. Assoc. Comput. Linguistics*, 8:423–438, 2020. (Cited on page 88.)
[53] Z. Jiang, J. Gao, and X. Lv. Metap: Meta pattern learning for one-shot knowledge graph completion. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 2232–2236. ACM, 2021. (Cited on page 65.)
[54] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. (Cited on page 20.)
[55] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL `https://openreview.net/forum?id=SJU4ayYgl`. (Cited on pages 1, 12, 42, 49, 61, 64, and 75.)
[56] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut. ALBERT: A lite BERT for self-supervised learning of language representations. In *ICLR*, 2020. (Cited on page 47.)
[57] J. B. Lee, R. A. Rossi, and X. Kong. Graph classification using structural attention. In *KDD*, pages 1666–1674, 2018. (Cited on pages 42 and 64.)
[58] B. Lester, R. Al-Rfou, and N. Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 3045–3059. Association for Computational Linguistics, 2021. (Cited on page 88.)
[59] R. Levie, F. Monti, X. Bresson, and M. M. Bronstein. Cayleynets: Graph convolutional neural networks with complex rational spectral filters. *IEEE Trans. Signal Process.*, 67(1):97–109, 2019. (Cited on pages 42 and 64.)
[60] R. Li, S. Wang, F. Zhu, and J. Huang. Adaptive graph convolutional neural networks. In *AAAI*, pages 3546–3553, 2018. (Cited on pages 42 and 64.)
[61] X. Li, D. Ding, B. Kao, Y. Sun, and N. Mamoulis. Leveraging meta-path contexts for classification in heterogeneous information networks. In *37th IEEE International Conference on Data Engineering, ICDE 2021, Chania, Greece, April 19-22, 2021*, pages 912–923. IEEE, 2021. (Cited on page 61.)
[62] X. L. Li and P. Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 4582–4597. Association for Computational Linguistics, 2021. (Cited on page 88.)
[63] Z. Li, F. Nie, X. Chang, Y. Yang, C. Zhang, and N. Sebe. Dynamic affinity graph construction for spectral clustering using multiple features. *IEEE Trans. Neural Networks Learn. Syst.*, 29(12): 6323–6332, 2018. (Cited on page 13.)
[64] Z. Li, W. Zheng, X. Lin, Z. Zhao, Z. Wang, Y. Wang, X. Jian, L. Chen, Q. Yan, and T. Mao. Transn: Heterogeneous network representation learning by translating node embeddings. In *36th IEEE International Conference on Data Engineering, ICDE 2020, Dallas, TX, USA, April 20-24, 2020*, pages 589–600. IEEE, 2020. (Cited on page 61.)
[65] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *CoRR*, 2021. (Cited on

page 88.)

[66] Z. Liu, C. Chen, X. Yang, J. Zhou, X. Li, and L. Song. Heterogeneous graph neural networks for malicious account detection. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM 2018, Torino, Italy, October 22-26, 2018*, pages 2077–2085. ACM, 2018. (Cited on page 12.)

[67] Z. Liu, Y. Fang, C. Liu, and S. C. H. Hoi. Relative and absolute location embedding for few-shot node classification on graph. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 4267–4275. AAAI Press, 2021. (Cited on page 65.)

[68] M. Luo, F. Nie, X. Chang, Y. Yang, A. G. Hauptmann, and Q. Zheng. Adaptive unsupervised feature selection with structure regularization. *IEEE Trans. Neural Networks Learn. Syst.*, 29(4):944–956, 2018. (Cited on page 13.)

[69] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to information retrieval*. Cambridge University Press, 2008. (Cited on page 78.)

[70] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119, 2013. (Cited on page 64.)

[71] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *CVPR*, pages 5425–5434, 2017. (Cited on pages 42 and 64.)

[72] S. Mu, Y. Li, W. X. Zhao, S. Li, and J. Wen. Knowledge-guided disentangled representation learning for recommender systems. *ACM Trans. Inf. Syst.*, 40(1):6:1–6:26, 2022. (Cited on page 65.)

[73] M. Niepert, M. Ahmed, and K. Kutzkov. Learning convolutional neural networks for graphs. In *ICML*, pages 2014–2023, 2016. (Cited on pages 42 and 64.)

[74] G. Niu, Y. Li, C. Tang, R. Geng, J. Dai, Q. Liu, H. Wang, J. Sun, F. Huang, and L. Si. Relational learning with gated and attentive neighbor aggregator for few-shot knowledge graph completion. In *SIGIR '21*, pages 213–222. ACM, 2021. (Cited on page 65.)

[75] S. Nowozin, B. Cseke, and R. Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 271–279, 2016. (Cited on page 71.)

[76] Z. Peng, W. Huang, M. Luo, Q. Zheng, Y. Rong, T. Xu, and J. Huang. Graph representation learning via graphical mutual information maximization. In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, pages 259–270. ACM / IW3C2, 2020. (Cited on page 65.)

[77] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: online learning of social representations. In *KDD*, pages 701–710, 2014. (Cited on pages 1, 9, 12, 25, 41, and 49.)

[78] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *NAACL-HLT*, pages 2227–2237, 2018. (Cited on page 39.)

[79] J. Qiu, Q. Chen, Y. Dong, J. Zhang, H. Yang, M. Ding, K. Wang, and J. Tang. GCC: graph contrastive coding for graph neural network pre-training. In *KDD*, pages 1150–1160, 2020. (Cited on pages 39, 43, 50, and 65.)

[80] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021. (Cited on page 88.)

[81] M. Rahman, T. K. Saha, M. A. Hasan, K. S. Xu, and C. K. Reddy. Dylink2vec: Effective feature representation for link prediction in dynamic networks. *CoRR*, 2018. (Cited on page 13.)

[82] N. Reimers and I. Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3980–3990. Association for Computational Linguistics, 2019. (Cited on pages 86 and 88.)

[83] R. Reinanda, E. Meij, and M. de Rijke. Knowledge graphs: An information retrieval perspective. *Foundations and Trends in Information Retrieval*, 14(4):289–444, October 2020. (Cited on page 1.)

[84] T. Rooker. Review of neurocomputing: Foundations of research. *AI Mag.*, 10(4):64–66, 1989. (Cited on page 20.)

[85] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000. (Cited on pages 12 and 41.)

[86] A. Sankar, X. Zhang, and K. C. Chang. Meta-gnn: metagraph neural network for semi-supervised

learning in attributed heterogeneous information networks. In *ASONAM '19: International Conference on Advances in Social Networks Analysis and Mining, Vancouver, British Columbia, Canada, 27-30 August, 2019*, pages 137–144. ACM, 2019. (Cited on page 12.)

[87] J. Shang, M. Qu, J. Liu, L. M. Kaplan, J. Han, and J. Peng. Meta-path guided embedding for similarity search in large-scale heterogeneous information networks. *CoRR*, 2016. (Cited on page 12.)

[88] D. Shen, X. Zhang, R. Henao, and L. Carin. Improved semantic-aware network embedding with fine-grained word alignment. In *EMNLP, Brussels, Belgium, October 31 - November 4, 2018*, pages 1829–1838. Association for Computational Linguistics, 2018. (Cited on pages 88 and 93.)

[89] J. Snell, K. Swersky, and R. S. Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4077–4087, 2017. (Cited on page 75.)

[90] F. Sun, J. Hoffmann, V. Verma, and J. Tang. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. In *ICLR*, 2020. (Cited on pages 42 and 65.)

[91] M. Sun, K. Zhou, X. He, Y. Wang, and X. Wang. GPPT: graph pre-training and prompt tuning to generalize graph neural networks. In A. Zhang and H. Rangwala, editors, *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022*, pages 1717–1727. ACM, 2022. doi: 10.1145/3534678.3539249. URL `https://doi.org/10.1145/3534678.3539249`. (Cited on pages 88 and 93.)

[92] X. Sun, J. Guo, X. Ding, and T. Liu. A general framework for content-enhanced network representation learning. *CoRR*, abs/1610.02906, 2016. (Cited on pages 88 and 93.)

[93] Y. Sun and J. Han. *Mining Heterogeneous Information Networks: Principles and Methodologies*. Synthesis Lectures on Data Mining and Knowledge Discovery. Morgan & Claypool Publishers, 2012. (Cited on page 1.)

[94] J. Tang, M. Qu, and Q. Mei. PTE: predictive text embedding through large-scale heterogeneous text networks. In *KDD*, pages 1165–1174, 2015. (Cited on pages 12 and 41.)

[95] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei. LINE: large-scale information network embedding. In *WWW*, pages 1067–1077, 2015. (Cited on pages 1, 9, 12, 41, and 49.)

[96] Y. Tian, D. Krishnan, and P. Isola. Contrastive multiview coding. In *ECCV*, volume 12356 of *Lecture Notes in Computer Science*, pages 776–794. Springer, 2020. (Cited on page 64.)

[97] E. Triantafillou, T. Zhu, V. Dumoulin, P. Lamblin, U. Evci, K. Xu, R. Goroshin, C. Gelada, K. Swersky, P. Manzagol, and H. Larochelle. Meta-dataset: A dataset of datasets for learning to learn from few examples. In *ICLR*, 2020. (Cited on page 75.)

[98] T. Trouillon, C. R. Dance, and et al. Knowledge graph completion via complex tensor factorization. *JMLR*, 18:1–38, 2017. (Cited on pages 10 and 15.)

[99] C. Tu, H. Liu, Z. Liu, and M. Sun. CANE: context-aware network embedding for relation modeling. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1722–1731, 2017. (Cited on pages 85, 88, and 93.)

[100] A. van den Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748, 2018. (Cited on pages 65 and 72.)

[101] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *ICLR*, 2018. (Cited on pages 1, 42, 49, and 64.)

[102] P. Velickovic, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm. Deep graph infomax. In *ICLR*, 2019. (Cited on page 65.)

[103] D. Wang, P. Cui, and W. Zhu. Structural deep network embedding. In *KDD*, pages 1225–1234, 2016. (Cited on pages 12, 41, 86, and 89.)

[104] N. Wang, M. Luo, K. Ding, L. Zhang, J. Li, and Q. Zheng. Graph few-shot learning with attribute matching. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, pages 1545–1554. ACM, 2020. (Cited on page 65.)

[105] X. Wang, H. Ji, C. Shi, B. Wang, Y. Ye, P. Cui, and P. S. Yu. Heterogeneous graph attention network. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 2022–2032. ACM, 2019. (Cited on pages 12 and 75.)

[106] X. Wang, N. Liu, H. Han, and C. Shi. Self-supervised heterogeneous graph neural network with co-contrastive learning. In *KDD '21*, pages 1726–1736. ACM, 2021. (Cited on page 42.)

[107] D. J. Watts and S. H. Strogatz. Collective dynamics of —[lsquo]—small-world—[rsquo]— networks.

Nature, 393:440–442, 1998. (Cited on page 10.)

[108] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, pages 3733–3742. IEEE Computer Society, 2018. (Cited on page 65.)

[109] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *ICLR*, 2019. (Cited on pages 42, 49, and 64.)

[110] Z. Xu, Q. Su, X. Quan, and W. Zhang. A deep neural information fusion architecture for textual network embeddings. In *EMNLP-IJCNLP, Hong Kong, China, November 3-7, 2019*, pages 4697–4705. Association for Computational Linguistics, 2019. (Cited on pages 85, 88, and 93.)

[111] C. Yang, Z. Liu, D. Zhao, M. Sun, and E. Y. Chang. Network representation learning with rich text information. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2111–2117, 2015. (Cited on pages 85, 88, and 93.)

[112] C. Yang, Y. Xiao, Y. Zhang, Y. Sun, and J. Han. Heterogeneous network representation learning: Survey, benchmark, evaluation, and beyond. *CoRR*, 2020. (Cited on page 41.)

[113] H. Yao, X. Wu, Z. Tao, Y. Li, B. Ding, R. Li, and Z. Li. Automated relational meta-learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020. (Cited on page 61.)

[114] H. Yao, C. Zhang, Y. Wei, M. Jiang, S. Wang, J. Huang, N. V. Chawla, and Z. Li. Graph few-shot learning via knowledge transfer. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 6656–6663. AAAI Press, 2020. (Cited on page 66.)

[115] C. Zhang, D. Song, C. Huang, A. Swami, and N. V. Chawla. Heterogeneous graph neural network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, pages 793–803. ACM, 2019. (Cited on pages 1, 12, 40, 42, 44, 49, and 75.)

[116] D. Zhang, J. Yin, and et al. Metagraph2vec: Complex semantic path augmented heterogeneous network embedding. In *PAKDD*, pages 196–208, 2018. (Cited on page 12.)

[117] J. Zhang, X. Shi, J. Xie, H. Ma, I. King, and D. Yeung. Gaan: Gated attention networks for learning on large and spatiotemporal graphs. In *UAI*, pages 339–349, 2018. (Cited on pages 42 and 64.)

[118] M. Zhang and Y. Chen. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 5171–5181, 2018. (Cited on page 76.)

[119] Y. Zhang, Q. Yao, W. Dai, and L. Chen. Autosf: Searching scoring functions for knowledge graph embedding. In *36th IEEE International Conference on Data Engineering, ICDE 2020, Dallas, TX, USA, April 20-24, 2020*, pages 433–444. IEEE, 2020. (Cited on page 61.)

[120] Z. Zhang, P. Cui, J. Pei, X. Wang, and W. Zhu. TIMERS: error-bounded SVD restart on dynamic networks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 224–231, 2018. (Cited on page 13.)

[121] Z. Zhong, D. Friedman, and D. Chen. Factual probing is [MASK]: learning vs. learning to recall. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 5017–5033. Association for Computational Linguistics, 2021. (Cited on page 88.)

[122] F. Zhou, C. Cao, K. Zhang, G. Trajcevski, T. Zhong, and J. Geng. Meta-GNN: On few-shot node classification in graph meta-learning. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, pages 2357–2360. ACM, 2019. (Cited on pages 61, 62, 65, and 75.)

[123] K. Zhou, J. Yang, C. C. Loy, and Z. Liu. Learning to prompt for vision-language models. *CoRR*, abs/2109.01134, 2021. URL `https://arxiv.org/abs/2109.01134`. (Cited on page 88.)

[124] K. Zhou, J. Yang, C. C. Loy, and Z. Liu. Learning to prompt for vision-language models. *Int. J. Comput. Vis.*, 130(9):2337–2348, 2022. (Cited on page 87.)

[125] L. Zhou, Y. Yang, X. Ren, F. Wu, and Y. Zhuang. Dynamic network embedding by modeling triadic closure process. In *AAAI*, pages 571–578, 2018. (Cited on pages 10, 13, 17, 26, and 28.)

[126] L. Zhu, D. Guo, J. Yin, G. V. Steeg, and A. Galstyan. Scalable temporal latent space inference for link prediction in dynamic social networks. *IEEE Trans. Knowl. Data Eng.*, 28(10):2765–2777, 2016.

(Cited on page 13.)
[127] Z. Zhuang, X. Xiang, S. Huang, and D. Wang. Hinfshot: A challenge dataset for few-shot node classification in heterogeneous information network. In *ICMR '21: International Conference on Multimedia Retrieval, Taipei, Taiwan, August 21-24, 2021*, pages 429–436. ACM, 2021. (Cited on page 66.)

Summary

Heterogeneous information networks are ubiquitous. Social networks, knowledge graphs, and interactions between users and items in search and recommender systems can be modeled as networks with multiple types of nodes and edges. In order to mine the rich information captured by a heterogeneous information network, network representation learning embeds a network into a low-dimensional space. Network representation learning has drawn a significant amount of interest from the research community. However, traditional network representation learning models simply learn network embeddings based on a given network, ignoring various real-life scenarios and limitations.

The work in this thesis provides a series of algorithms that are able to deal with different heterogeneous information network scenarios and different downstream tasks. We first focus on dynamic heterogeneous information networks as real-life networks are always evolving. M-DHIN is proposed as a method to learn dynamic embeddings; it is also able to predict the future network.

After that, we study the pre-training problem in heterogeneous information networks. Since most networks are unlabeled, we propose a pre-training and fine-tuning framework PF-HIN that uses two self-supervised pre-training tasks, i.e., masked node modeling and adjacent node prediction. The pre-trained encoder can quickly be adapted to datasets of different domains and different downstream tasks.

We also investigate the few-shot problem for heterogeneous information networks as in practice, only a handful of nodes are labeled. We propose META-HIN, which uses a meta-learning framework that is able to handle the few-shot learning task in both a single-graph scenario and multiple-graph scenarios.

In the final research chapter, we study how to leverage meaningful textual information that may be contained in a heterogeneous information network. We propose a novel prompt learning framework P-HIN that is able to simultaneously employ the textual information and handle the few-shot problems. We align the text representation and node representation using a contrastive learning mechanism, so that textual information is incorporated in an effective manner.

Finally, we suggest a number of directions for future research. These include model compression to improve the model efficiency, the use of multi-modal HIN attributes, application-oriented HIN representation learning, and interpretability of HIN representation learning.

Samenvatting

Heterogene informatienetwerken zijn overal te vinden. Sociale netwerken, kennisgrafen, en interacties tussen gebruikers en items in zoek- en aanbevelingssystemen kunnen worden gemodelleerd als netwerken met meerdere soorten knopen en kanten. Om de rijke informatie te ontginnen die is vastgelegd met een heterogeen informatienetwerk, bedden we bij het leren van netwerkrepresentaties een netwerk in in een laagdimensionale ruimte. Het leren van netwerkrepresentaties trekt veel aandacht binnen de onderzoeksgemeenschap. Traditionele modellen voor het leren van netwerkrepresentaties leren netwerkinbeddingen terwijl ze voorbijgaan aan realistische scenario's en beperkingen.

Het werk in dit proefschrift biedt een reeks algoritmen die in staat zijn om om te gaan met verschillende heterogene informatienetwerkscenario's en verschillende *downstream*-taken. We richten ons eerst op dynamische heterogene informatienetwerken, aangezien real-life netwerken altijd in ontwikkeling zijn. We stellen M-DHIN voor als een methode om dynamische inbeddingen te leren; de methode is ook in staat om toekomstige netwerken te voorspellen.

Daarna bestuderen we het *pre-trainings*-probleem in heterogene informatienetwerken. Aangezien de meeste netwerken niet-gelabeld zijn, stellen we een *pre-training* en *fine-tuning* framework PF-HIN voor dat gebruik maakt van twee *pre-trainings*-taken met zelf-supervisie, namelijk het modelleren van gemaskeerde knopen en het voorspellen van aangrenzende knopen. De vooraf getrainde *encoder* kan snel worden aangepast aan datasets van verschillende domeinen en voor verschillende *downstream*-taken.

We onderzoeken ook het *few shot*-probleem voor heterogene informatienetwerken, aangezien in de praktijk slechts een handvol knopen is gelabeld. We stellen META-HIN voor, dat een meta-leerraamwerk gebruikt dat in staat is om leertaken aan de hand van een paar voorbeelden aan te kunnen in het scenario met één enkele graaf en in scenario's met meerdere grafen.

In het laatste onderzoekshoofdstuk bestuderen we hoe we zinvolle tekstuele informatie kunnen gebruiken die zich in een heterogeen informatienetwerk kan bevinden. We stellen een nieuw P-HIN-raamwerk voor voor het leren van zogenaamde *prompt* waarmee tegelijkertijd tekstuele informatie gebruikt kan worden en leerproblemen met weinig voorbeelden kunnen worden opgelost. We stemmen de tekstrepresentatie en representatie van knopen op elkaar af met behulp van een contrastief leermechanisme, zodat tekstuele informatie op een effectieve manier wordt opgenomen.

Tot slot stellen we een aantal richtingen voor toekomstig onderzoek voor, waaronder modelcompressie om de efficiëntie van modellen te verbeteren, het gebruik van multimodale HIN-attributen, toepassingsgericht leren van HIN-representaties, en interpreteerbaarheid van het leren van HIN-representaties.

COMPLEX INFORMATION OFTEN INVOLVES MULTIPLE TYPES OF OBJECTS AND RELATIONS. SUCH INFORMATION CAN BE REPRESENTED VIA HETEROGENEOUS INFORMATION NETWORKS (HINS). IN THIS THESIS, WE EXPLORE DIFFERENT LEARNING MECHANISMS FOR NETWORK REPRESENTATION LEARNING TO FIT INTO DIFFERENT SCENARIOS AND DIFFERENT MACHINE LEARNING TASKS OF HINS. WE FIRST STUDY REPRESENTATION LEARNING OF DYNAMIC HINS, AS REAL-LIFE NETWORKS ARE ALWAYS EVOLVING. THEN, WE PRE-TRAIN A HIN WITHOUT USING LABELED INFORMATION, AND A PRE-TRAINED GNN MODEL CAN EASILY BE ADAPTED TO DIFFERENT DATASETS AND DIFFERENT DOWNSTREAM TASKS. WE ALSO INVESTIGATE FEW-SHOT LEARNING OF HINS, WHERE ONLY A HANDFUL OF LABELS ARE GIVEN. AND FINALLY, WE MAKE USE OF AUXILIARY TEXTUAL INFORMATION TO FURTHER FACILITATE THE LEARNING OF HINS.