

## REFERENCES

- Addanki, R., Kumar, S., Chen, M., and Zhang, L. (2023) 'Caerus: Predictive Cache Replacement Meets Content Delivery,' *Proceedings of the 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI '23)*, Boston, MA, USA, pp. 1-16.
- Adomavicius, G. and Tuzhilin, A. (2005) 'Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions,' *IEEE Transactions on Knowledge and Data Engineering*, 17(6), pp. 734-749.
- Ali, W. and Shamsuddin, S. (2015) 'Intelligent Dynamic Aging Approaches in Web Proxy Cache Replacement,' *Journal of Intelligent Learning Systems and Applications*, 7(4), pp. 117-127.
- Ali, W., Shamsuddin, S.M., and Ismail, A.S. (2011) 'A Survey of Web Caching and Prefetching,' *International Journal of Advances in Soft Computing and its Applications*, 3(1), pp. 18-44.
- Antonio, O., Smith, S.L., Gu, A., Fernando, A., Gulcehre, C., Pascanu, R., and De, S. (2023) 'Resurrecting Recurrent Neural Networks for Long Sequences,' *arXiv preprint arXiv:2303.06349*, pp. 1-14.
- Arlot, S. & Celisse, A. (2010) 'A Survey of Cross-validation Procedures for Model Selection,' *Statistics Surveys*, 4, pp. 40-79.
- Bahmani, B., Chowdhury, A., and Goel, A. (2010) 'Fast Incremental and Personalized PageRank,' *Proceedings of the VLDB Endowment*, 4(3), pp. 173-184.
- Ban, Z. and Gu, Z. (2007) 'An Online PPM Prediction Model for Web Prefetching,' *Proceedings of the 9th Annual ACM International Workshop on Web Information and Data Management*, Lisbon, Portugal, pp. 89-96.
- Baskaran, K. and Chellan, K. (2016) 'Improved Performance by Combining Web Pre-fetching Using Clustering with Web Caching Based on SVM Learning Method,' *International Journal of Computers Communications & Control*, 11(2), pp. 67-78.
- Bergmeir, C. & Benítez, J.M. (2012) 'On the Use of Cross-validation for Time Series Predictor Evaluation,' *Information Sciences*, 191, pp. 192-213.
- Bin, W. and Kshemkalyani, A.D. (2004) 'Objective-Greedy Algorithms for Long-Term Web Prefetching,' *Proceedings of Third IEEE International Symposium on Network Computing and Applications (NCA 2004)*, Cambridge, MA, USA, pp. 89-96.
- Brand, M. (2002) 'Incremental Singular Value Decomposition of Uncertain Data with Missing Values,' *Computer Vision — ECCV 2002. Lecture Notes in Computer Science*, 2350, pp. 707-720. doi: [https://doi.org/10.1007/3-540-47969-4\\_47](https://doi.org/10.1007/3-540-47969-4_47).



- Burke, R. (2002) 'Hybrid Recommender Systems: Survey and Experiments,' *User Modeling and User-Adapted Interaction*, 12(4), pp. 331-370.
- Chen, X., Zhang, G., Wang, Y., Wu, Y., Su, S., Zhan, K., and Wang, B. (2023) 'Cache-Aware Reinforcement Learning in Large-Scale Recommender Systems,' *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, Macao, SAR, pp. 2156-2164.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014) 'Learning Phrase Representations Using RNN Encoder-Decoder for Statistical Machine Translation,' *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, pp. 1724-1734.
- on *Internet Technologies and Systems* (pp. 13-22). USENIX Association.
- Cohen, E., & Kaplan, H. (2000). Prefetching the means for document transfer: A new approach for reducing web latency. *Computer Networks*, 33(1-6), 259-274.
- Cremonesi, P., Koren, Y., and Turrin, R. (2010) 'Performance of Recommender Algorithms on Top-N Recommendation Tasks,' *Proceedings of the 4th ACM Conference on Recommender Systems (RecSys '10)*, Barcelona, Spain, pp. 39-46.
- Dacrema, M.F., Cremonesi, P., and Jannach, D. (2019) 'Are We Really Making Much Progress? A Worrying Analysis of Recent Neural Recommendation Approaches,' *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys '19)*, Copenhagen, Denmark, pp. 101-109.
- Domenech, J., Pont, A., Sahuquillo, J., and Gil, J.A. (2010) 'Evaluation, Analysis and Adaptation of Web Prefetching Techniques in Current Web,' *Web-based Support Systems*, London, UK: Springer, pp. 239-271.
- Fan, W., Ma, Y., Li, Q., He, Y., Zhao, E., Tang, J., and Yin, D. (2019) 'Graph Neural Networks for Social Recommendation,' *Proceedings of the World Wide Web Conference (WWW '19)*, San Francisco, CA, USA, pp. 417-426.
- Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., and Bouchachia, A. (2014) 'A Survey on Concept Drift Adaptation,' *ACM Computing Surveys*, 46(4), pp. 1-37.
- Ge, M., Delgado-Battenfeld, C., and Jannach, D. (2010) 'Beyond Accuracy: Evaluating Recommender Systems by Coverage and Serendipity,' *Proceedings of the 4th ACM Conference on Recommender Systems (RecSys '10)*, Barcelona, Spain, pp. 257-260.
- Gechev, M. (2018) 'Introducing Guess.js - A Toolkit for Enabling Data-driven User-experiences on the Web,' *Blog Post*. Available at: <https://blog.mgechev.com/2018/05/09/introducing-guess-js-data-driven-user-experiences-web/> (Accessed: 8 December 2024).



- Giannakas, T., Sermpezis, P., and Spyropoulos, T. (2021) 'Show Me the Cache: Optimizing Cache-Friendly Recommendations for Sequential Content Access,' *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, Vancouver, BC, Canada, pp. 1-10.
- Hamilton, W.L., Ying, R., and Leskovec, J. (2017) 'Inductive Representation Learning on Large Graphs,' *Advances in Neural Information Processing Systems (NIPS)*, Long Beach, CA, USA, pp. 1024-1034.
- Harper, F.M. & Konstan, J.A. (2015) 'The MovieLens Datasets: History and Context,' *ACM Transactions on Interactive Intelligent Systems*, 5(4), pp. 1-19.
- Haveliwala, T.H. (2002) 'Topic-Sensitive PageRank,' *Proceedings of the 11th International Conference on World Wide Web (WWW '02)*, Honolulu, HI, USA, pp. 517-526.
- He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., and Wang, M. (2020) 'LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation,' *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, Xi'an, China, pp. 639-648.
- Herlocker, J.L., Konstan, J.A., Terveen, L.G., and Riedl, J.T. (2004) 'Evaluating Collaborative Filtering Recommender Systems,' *ACM Transactions on Information Systems*, 22(1), pp. 5-53.
- Hu, Y., Koren, Y., and Volinsky, C. (2008) 'Collaborative Filtering for Implicit Feedback Datasets,' *Proceedings of the 8th IEEE International Conference on Data Mining (ICDM '08)*, Pisa, Italy, pp. 263-272.
- Huang, Y.F. and Hsu, J.M. (2008) 'Mining Web Logs to Improve Hit Ratios of Prefetching and Caching,' *Knowledge-Based Systems*, 21(1), pp. 62-69.
- Hyndman, R.J. & Athanasopoulos, G. (2018) 'Forecasting: Principles and Practice,' 3rd edition, Melbourne: OTexts.
- Ibrahim, T.I. and Xu, C.Z. (2004) 'A Keyword-Based Semantic Pre-fetching Approach in Internet News Services,' *IEEE Transactions on Knowledge and Data Engineering*, 16(5), pp. 601-611.
- Järvelin, K. and Kekäläinen, J. (2002) 'Cumulated Gain-based Evaluation of IR Techniques,' *ACM Transactions on Information Systems*, 20(4), pp. 422-446.
- Järvelin, K. and Kekäläinen, J. (2017) 'IR Evaluation Methods for Retrieving Highly Relevant Documents,' *ACM SIGIR Forum*, 51(2), pp. 243-250.
- Jin, B., Sihua, T., Lin, C., Ren, X., and Huang, Y. (2007) 'An Integrated Prefetching and Caching Scheme for Mobile Web Caching System,' *Eighth ACIS International Conference on Software Engineering, Artificial*



- Intelligence, Networking, and Parallel/Distributed Computing*, Qingdao, China, pp. 278-283.
- Joseph, V.R. (2022) 'Optimal Ratio for Data Splitting,' *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 15(4), pp. 531-538.
- Kazemi, S.M., Goel, R., Jain, K., Kobayez, I., Sethi, A., Forsyth, P., and Poupart, P. (2020) 'Representation Learning for Dynamic Graphs: A Survey,' *Journal of Machine Learning Research*, 21(70), pp. 1-73.
- Kim, J., Rossi, E., Rao, A., Kumar, S., and Prabhakar, B. (2023) 'Temporal Graph Networks for Deep Learning on Dynamic Graphs,' *Proceedings of the International Conference on Machine Learning (ICML)*, Honolulu, HI, USA, pp. 5315-5326.
- Kipf, T.N. and Welling, M. (2017) 'Semi-supervised Classification with Graph Convolutional Networks,' *Proceedings of the International Conference on Learning Representations (ICLR)*, Toulon, France, pp. 1-14.
- Koren, Y., Bell, R., and Volinsky, C. (2009) 'Matrix Factorization Techniques for Recommender Systems,' *Computer*, 42(8), pp. 30-37.
- Kroeger, T.M., Long, D.D.E., and Mogul, J.C. (1997) 'Exploring the Bounds of Web Latency Reduction from Caching and Prefetching,' *Proceedings of the USENIX Symposium on Internet Technology and Systems*, Monterey, CA, USA, pp. 13-22.
- Kumar, S., Zhang, X., and Leskovec, J. (2018) 'Learning Dynamic Embeddings from Temporal Interactions,' *arXiv preprint arXiv:1812.02289*, pp. 1-12.
- Li, J., Wang, Y., Zhang, Y., and Zhang, Y. (2023) 'Zebra: When Temporal Graph Neural Networks Meet Temporal Personalized PageRank,' *Proceedings of the VLDB Endowment*, 16(6), pp. 1332-1345.
- Li, L., Chu, W., Langford, J., and Schapire, R.E. (2010) 'A Contextual-Bandit Approach to Personalized News Article Recommendation,' *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*, Raleigh, NC, USA, pp. 661-670.
- Liben-Nowell, D. and Kleinberg, J. (2007) 'The Link-prediction Problem for Social Networks,' *Journal of the American Society for Information Science and Technology*, 58(7), pp. 1019-1031.
- Minwoo, J., Yeonoh, A., Heejun, R., and Wonjun, L. (2021) 'Predictive Prefetching Based on User Interaction for Web Applications,' *IEEE Communications Letters*, 25(3), pp. 821-824.
- Padmanabhan, V.N. and Mogul, J.C. (1996) 'Using Predictive Prefetching to Improve World Wide Web Latency,' *Computer Communication Review*, 26(3), pp. 22-36.



- Page, L., Brin, S., Motwani, R., and Winograd, T. (1999) 'The PageRank Citation Ranking: Bringing Order to the Web,' *Stanford InfoLab Technical Report*, 1999-66.
- Palpanas, T. and Mendelzon, A. (1999) 'Web Prefetching Using Partial Match Prediction,' *Proceedings of the 4th International Web Caching Workshop*, San Diego, CA, USA, pp. 47-59.
- Pallis, G., Vakali, A., and Pokorny, J. (2008) 'A Clustering-based Prefetching Scheme on a Web Cache Environment,' *Computers and Electrical Engineering*, 34(4), pp. 309-323.
- Pazzani, M.J. and Billsus, D. (2007) 'Content-based Recommendation Systems,' *The Adaptive Web*, Berlin: Springer, pp. 325-341.
- Peng, Z., Huang, W., Luo, M., Zheng, Q., Rong, Y., Xu, T., and Huang, J. (2022) 'SVD-GCN: A Simplified Graph Convolution Paradigm for Recommendation,' *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, Atlanta, GA, USA, pp. 1470-1479.
- Proposed System Implementation (2025) 'model.py: TimeEncode and MLP Mixer Classes,' *Internal Project Repository*.
- Proposed System Implementation (2025) 'SERVING\_ARCHITECTURE.md: Production Serving Architecture Benchmarks,' *Internal Project Repository*.
- Proposed System Implementation (2025) 'train\_hybrid\_recsys.py: EAGLE Hybrid Training Module,' *Internal Project Repository*.
- Proposed System Implementation (2025) 'util.py: Incremental TPPR Node Finder Class,' *Internal Project Repository*.
- Rangarajan, S.K., Phoha, V.V., Balagani, K., Selmic, R.R., and Iyengar, S.S. (2004) 'Web User Clustering and Its Application to Prefetching Using ART Neural Networks,' *IEEE Computer*, 37(6), pp. 38-46.
- Rendle, S., Freudenthaler, C., Gantner, Z., and Schmidt-Thieme, L. (2009) 'BPR: Bayesian Personalized Ranking from Implicit Feedback,' *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, Montreal, QC, Canada, pp. 452-461.
- Resnick, P. and Varian, H.R. (1997) 'Recommender Systems,' *Communications of the ACM*, 40(3), pp. 56-58.
- Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., and Bronstein, M. (2020) 'Temporal Graph Networks for Deep Learning on Dynamic Graphs,' *arXiv preprint arXiv:2006.10637*, pp. 1-18.
- Sajeev, G. and Sebastian, M. (2010) 'A Novel Content Classification Scheme for Web Caches,' *Evolving Systems*, 2(2), pp. 101-118.



- Sarwar, B., Karypis, G., Konstan, J., and Riedl, J. (2001) 'Item-based Collaborative Filtering Recommendation Algorithms,' *Proceedings of the 10th International Conference on World Wide Web (WWW '01)*, Hong Kong, China, pp. 285-295.
- Sasaki, Y. (2007) 'The Truth of the F-measure,' *Teaching, Tutoring Materials*, 1(5), pp. 1-5.
- Shani, G. and Gunawardana, A. (2011) 'Evaluating Recommendation Systems,' *Recommender Systems Handbook*, Boston: Springer, pp. 257-297.
- Schreiber, T., Griwodz, C., & Halvorsen, P. (2011). Reducing latency through piggybacking in mobile networks. In *IEEE International Conference on Communications (ICC)* (pp. 1-6). IEEE.
- Smyth, B. and McClave, P. (2001) 'Similarity vs. Diversity,' *Proceedings of the International Conference on Case-Based Reasoning (ICCBR '01)*, Vancouver, BC, Canada, pp. 347-361.
- Teng, W., Chang, C., and Chen, M. (2005) 'Integrating Web Caching and Web Prefetching in Client-Side Proxies,' *IEEE Transactions on Parallel and Distributed Systems*, 16(5), pp. 444-455.
- Tolstikhin, I.O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., Lucic, M., and Dosovitskiy, A. (2021) 'MLP-Mixer: An all-MLP Architecture for Vision,' *Advances in Neural Information Processing Systems (NeurIPS)*, 34, pp. 24261-24272.
- Trivedi, R., Farajtabar, M., Biswal, P., and Zha, H. (2019) 'DyRep: Learning Representations over Dynamic Graphs,' *Proceedings of the International Conference on Learning Representations (ICLR)*, New Orleans, LA, USA, pp. 1-12.
- Vargas, S. and Castells, P. (2011) 'Rank and Relevance in Novelty and Diversity Metrics for Recommender Systems,' *Proceedings of the 5th ACM Conference on Recommender Systems (RecSys '11)*, Chicago, IL, USA, pp. 109-116.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. (2018) 'Graph Attention Networks,' *Proceedings of the International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, pp. 1-12.
- Venkataramani, A., Yalagandula, P., and Kokku, R. (2002) 'The Potential Costs and Benefits of Long-term Prefetching for Content Distribution,' *Computer Communications*, 25(4), pp. 367-375.
- Voorhees, E.M. (1999) 'The TREC-8 Question Answering Track Report,' *Proceedings of the 8th Text REtrieval Conference (TREC-8)*, Gaithersburg, MD, USA, pp. 77-82.



- Wang, Q., Lee, D., Zhang, Y., and Vuković, M. (2023) 'Cacheus: Cache Replacement Through Deep Reinforcement Learning,' *Proceedings of the 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI '23)*, Boston, MA, USA, pp. 1-16.
- Wang, X., He, X., Cao, Y., Liu, M., and Chua, T.S. (2019) 'KGAT: Knowledge Graph Attention Network for Recommendation,' *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, Anchorage, AK, USA, pp. 950-958.
- Wang, X., He, X., Wang, M., Feng, F., and Chua, T.S. (2019) 'Neural Graph Collaborative Filtering,' *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, Paris, France, pp. 165-174.
- Weston, J., Bengio, S., and Usunier, N. (2011) 'WSABIE: Scaling Up to Large Vocabulary Image Annotation,' *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, Barcelona, Spain, pp. 2764-2770.
- Willmott, C.J. and Matsuura, K. (2005) 'Advantages of the Mean Absolute Error (MAE) over the Root Mean Square Error (RMSE) in Assessing Average Model Performance,' *Climate Research*, 30(1), pp. 79-82.
- Wong, A.K.Y. (2006) 'Web Cache Replacement Policies: A Pragmatic Approach,' *IEEE Network*, 20(1), pp. 28-34.
- Wu, S., Sun, F., Zhang, W., and Cui, B. (2021) 'Graph Neural Networks in Recommender Systems: A Survey,' *arXiv preprint arXiv:2011.02260*, pp. 1-22.
- Wu, S., Tang, Y., Zhu, Y., Wang, L., Xie, X., and Tan, T. (2022) 'Session-based Recommendation with Graph Neural Networks,' *Proceedings of the Web Conference (WWW '22)*, Lyon, France, pp. 168-178.
- Wu, S., Zhang, Y., Gao, C., Tang, Y., Xie, X., and Tan, T. (2021) 'Self-supervised Graph Learning for Recommendation,' *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Virtual Event, Canada, pp. 726-735.
- Xu, D., Ruan, C., Korpeoglu, E., Kumar, S., and Achan, K. (2020) 'Inductive Representation Learning on Temporal Graphs,' *Proceedings of the International Conference on Learning Representations (ICLR)*, Addis Ababa, Ethiopia.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. (2019) 'How Powerful are Graph Neural Networks?' *Proceedings of the International Conference on Learning Representations (ICLR)*, New Orleans, LA, USA, pp. 1-10.



- Yang, Q., Huang, J.Z., and Ng, M. (2008) 'A Data Cube Model for Prediction-based Web Prefetching,' *Journal of Intelligent Information Systems*, 20(1), pp. 11-30.
- Zhang, Y. (2022) 'An Answer to an Open Question in the Incremental SVD,' *arXiv preprint arXiv:2204.05398v2*, pp. 1-8.
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., and Sun, M. (2020) 'Graph Neural Networks: A Review of Methods and Applications,' *AI Open*, 1, pp. 57-81.
- Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., and Wang, L. (2022) 'Training-free Graph Neural Networks for Graph-level Prediction,' *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22)*, Washington, DC, USA, pp. 2275-2285.
- Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., and Wang, L. (2020) 'Deep Graph Contrastive Representation Learning,' *ICML Workshop on Graph Representation Learning and Beyond*, Virtual Event, pp. 1-6.