Ryan A. Rossi

Senior Staff Research Scientist

✉ ryan.rossi.ai@gmail.com
☎ +1 (415) 819 8500
Erdős Number: 3

Personalized AI, Agentic Systems, Reasoning, Recommendation, Model Routing, Multimodal & Vision-Language Models, Video & Image Generation, Efficient Algorithms, Scalable Inference, Training, Small Models, RAG, Graph Neural Networks, Trust & Privacy, Knowledge Graphs

🎓 *Google Scholar* · ***200+** publications* · ***100+** patents* · ***20000+** citations* · ***57** h-index*

Summary

Product-oriented AI research leader focused on building and shipping agentic AI systems, multimodal & LLM generative, and foundation-model applications across Adobe products (from Adobe Analytics/CJA to Firefly). Led large cross-functional AI initiatives from product requirements and architecture through implementation, evaluation, IP, and customer-facing release. Drove technology-transfer and research efforts producing over 15+ tech-transfers into products, 200+ publications, 100+ patents, and 20,000+ citations.

Technical Strategy, Research & Systems Impact

AI Personalization: Lead personalization of AI systems at Adobe. Work spans personalized text generation, personalized graph-based retrieval, and multimodal personalization.

Simulation: Leading research on LLM-based conversational user simulation, introducing formal taxonomies and methods for high-fidelity synthetic user generation. This work enables scalable evaluation and optimization of conversational AI.

Multimodal and Vision-Language Models: Driving research on image and video generation and editing, post-training, judge and reward modeling, agentic systems for video and image editing, post-training techniques, and multimodal retrieval-augmented generation.

Agentic AI: Leading LLM-based agents beyond predefined actions, reward-weighted conversation optimization, and preference-guided code generation at Adobe Research.

Graph Representation Learning: Established foundational contributions in graph neural networks, knowledge graphs, recommendation, persona/role discovery, temporal graph embeddings, influence maximization, and network analysis.

Scientific Impact: Author of **200+ publications** (with **20,000+** citations, **57** h-index, **235** i10-index) in top venues including NeurIPS, ICML, ICLR, ACL, EMNLP, AAAI, KDD, ICDM, WWW, UAI, AISTATS, COLT, among many others.

Innovation and IP: Inventor on over **100+ patents**, spanning AI personalization, agents, RAG, AI fairness/bias, AI-as-a-judge, simulation, graph learning, knowledge graphs, provably-efficient algorithms, scalable AI, recommendation, AI interfaces, vis generation.

Product and Technology Transfer: Led **15+ technology transfers** of AI research into Adobe products across Experience Cloud, Creative Cloud, Document Cloud, and Commerce spanning Analytics, CJA, Experience Platform, Photoshop, Firefly, Real-Time CDP, Campaign, Acrobat Assistant, among others.

Education

2009–2015 **Ph.D., Computer Science.**, *Purdue University*, USA.
Title: "Improving Relational Machine Learning by Modeling Temporal Dependencies"
Recipient of Four Ph.D. Fellowships:
– National Science Foundation Graduate Fellowship (NSF GRFP)
– DoD: National Defense Science and Engineering Graduate Fellowship (NDSEG)
– Bilsland Dissertation Fellowship Awarded to Outstanding Ph.D. candidates
– Purdue University Fredrick N. Andrews Doctoral Fellowship

2013 **MS in Computer Science.**, *Purdue University*, USA.
Concentrate in Machine Learning

2005–2009 **Bachelor of Science in Computer Science.**, *Coastal Carolina University (CCU)*, USA.
* *Valedictorian class of 2009.* GPA: **4.0**., Summa Cum Laude
Advisor: Jean-Louis Lassez (Retired IBM T.J. Watson Research Center)

Summer 11–12 **LLNL Scholar** , *Lawrence Livermore National Laboratory*, USA.

Summer 2010 **NREIP**, *Naval Research Laboratory (AI Center)*, USA.

2009 **NASA Fellow**, *California Institute of Technology, JPL*, USA.

2009 **USRP Fellow**, *Jet Propulsion Laboratory*, USA.

Summer 2008 **NSF REU Fellow**, *University of Massachusetts at Amherst*, USA.

Summer 2007 **Research Fellow**, *New Mexico Institute of Technology*, USA.

Experience

2017–present **Senior Staff Research Scientist**, Adobe Research.

- Leading technical strategy and architecture for large-scale AI systems that translate frontier models into scalable product capabilities, spanning agentic AI, production multimodal and LLM platforms, model orchestration, personalization, recommendation, evaluation, and scalable inference.
- Led a large-scale agentic data intelligence platform for Adobe Analytics, Customer Journey Analytics, and Adobe Experience Platform. Lead 30+ cross-functional contributors across research, engineering, product, UI, legal, and business stakeholders.
- Led and contributed to Firefly multimodal generative AI systems, including agentic systems and judge models for image and video generation and editing.
- Architected and delivered production-grade multimodal & LLM data-agent capabilities.
- Developed and adapted architectures, LLMs, multimodal foundation models, fine-tuned models, retrieval-augmented systems, structured-output pipelines, judge models, and task-specialized model workflows to solve Adobe-specific product problems and support customer-facing releases.
- Improved product quality, latency, scalability, and cost through production-aligned model design, efficient data structures, parallelization, and rigorous offline and continuous evaluation pipelines.
- Led *15+ technology transfers* of AI research into Adobe products across Experience Cloud, Creative Cloud, Document Cloud, and Commerce spanning Analytics, CJA, Experience Platform, Photoshop, Firefly, Real-Time CDP, Campaign, Acrobat Assistant, among others.

2015–2017 **Member of Research Staff**, *Palo Alto Research Center (PARC, a Xerox company)*, Machine Learning group .

2013–2015 **Visiting Researcher**, *Palo Alto Research Center (PARC)*, Palo Alto, CA USA.
Research focused on theory, algorithms & applications of relational (graph-based) machine learning

Summer 2013 **Research Intern**, *Palo Alto Research Center (PARC)*, Palo Alto, CA USA.
Developed recommendation systems via collective matrix-tensor factorization

2009–2015 **Research Assistant**, *Purdue University*, USA.
Research: Machine Learning, Statistical Relational Learning
Proposed methods for *role discovery in large dynamic graphs* and *dynamic relational classification*.

Summer 2011–2012 **Research Assistant**, *Lawrence Livermore National Laboratory (ISCR)*, USA.
Research focused on developing ML algorithms to characterize and model user behavior for detecting malicious intent/intrusions in real-time. Invited back for second year.
Resulted in two papers on modeling dynamic roles in large networks.

Summer 2010 **Research Assistant**, *Naval Research Laboratory (Artificial Intelligence Center)*, USA.
Advisor: David Aha, Co-advisor: Luke McDowell (U.S. Naval Academy), NREIP
Resulted in the JAIR paper "Transformation of Graph Data for Statistical Relational Learning"

Summer 2009 **Research Assistant**, *California Institute of Technology (NASA JPL)*, USA.
Advisor: Mark W. Powell, Summer Research Fellowship (returned to continue my research).

Spring 2009 **Research Assistant**, *NASA Jet Propulsion Laboratory*, USA.
Advisor: Mark W. Powell, Spring USRP Fellowship.

Summer 2008 **Research Assistant**, *University of Massachusetts at Amherst*, USA.
Advisor: David Jensen, Co-advisor: Brian Taylor. *REU NSF Fellowship*.
"Experimental Methods for Improving the Design of Participatory Sensing Systems"

Summer 2007 **Research Assistant**, *New Mexico Institute of Technology, ICASA*, USA.
Advisor: Srinivas Mukkamala, Senior Research Scientist, ICASA

2005–2009 **Research Assistant**, *Coastal Carolina University*, USA.
Advisor: Jean-Louis Lassez, Retired IBM T.J. Watson Research Center

Journal Publications

[J32] Zhehao Zhang, **Ryan A. Rossi**, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, Ruiyi Zhang, Jiuxiang Gu, Tyler Derr, Hongjie Chen, Junda Wu, Xiang Chen, Zichao Wang, Subrata Mitra, Nedim Lipka, Nesreen K. Ahmed, Yu Wang, *Personalization of Large Language Models: A Survey*, Transactions on Machine Learning Research (TMLR), 2025.

[J31] Reuben Luera, **Ryan A. Rossi**, Alexa F. Siu, Franck Dernoncourt, Tong Yu, Sungchul Kim, Ruiyi Zhang, and Xiang Chen, *Survey on User Interface Design and Interactions for Generative AI Applications*, Foundations and Trends in Human-Computer Interaction, 2025.

[J30] Hui Wei, Dong Yoon Lee, Shubham Rohal, Zhizhang Hu, **Ryan A. Rossi**, Shiwei Fang, and Shijia Pan, *A Survey of Foundation Models for IoT: Taxonomy and Criteria-based Analysis*, CCF Transactions on Pervasive Computing and Interaction, 2026.

[J29] Hongjie Chen, **Ryan A. Rossi**, Nesreen K. Ahmed, Namyong Park, Yu Wang, Tyler Derr, and Hoda Eldardiry, *Edges Matter: An Analysis of Graph Time-Series Representations for Temporal Networks*, IEEE Transactions on Network Science and Engineering, 2025.

[J28] Mustafa Abdallah, **Ryan A. Rossi**, Kanak Mahadik, Sungchul Kim, Handong Zhao, and Saurabh Bagchi, *Evaluation-Free Time-Series Forecasting Model Selection via Meta-Learning*, ACM Transactions on Knowledge Discovery from Data (TKDD), 2025.

[J27] Kewei Cheng, Nesreen K. Ahmed, **Ryan A. Rossi**, Theodore Willke, and Yizhou Sun, *Neural-Symbolic Methods for Knowledge Graph Reasoning: A Survey*, ACM Transactions on Knowledge Discovery from Data (TKDD), 2025.

[J26] Adrian Chen, **Ryan A. Rossi**, Namyong Park, Puja Trivedi, Yu Wang, Tong Yu, Sungchul Kim, Franck Dernoncourt, and Nesreen K. Ahmed, *Fairness-aware Graph Neural Networks: A Survey*, ACM Transactions on Knowledge Discovery from Data (TKDD), 2024.

[J25] Isabel O. Gallegos, **Ryan A. Rossi**, Joe Barrow, Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed, *Bias and Fairness in Large Language Models: A Survey*, Computational Linguistics, 2024.

[J24] Chenhan Yuan, **Ryan A. Rossi**, Andrew Katz, and Hoda Eldardiry, *A Reinforcement Learning Framework for N-Ary Document-Level Relation Extraction*, IEEE Transactions on Big Data, 2024.

[J23] Rashmi Ranjan Bhuyan, Adel Javanmard, Sungchul Kim, Gourab Mukherjee, **Ryan A. Rossi**, Tong Yu, and Handong Zhao, *Structured Dynamic Pricing: Optimal Regret in a Global Shrinkage Model*, Journal of Machine Learning Research (JMLR), 2024.

[J22] Hongjie Chen, **Ryan A. Rossi**, Kanak Mahadik, Sungchul Kim and Hoda Eldardiry, *Graph Deep Factors for Probabilistic Time-series Forecasting*, Transactions on Knowledge Discovery from Data (TKDD), 2022.

[J21] Nesreen K. Ahmed, Nick Duffield, **Ryan A. Rossi**, *Online Sampling of Temporal Networks*, Transactions on Knowledge Discovery from Data (TKDD), 1–43, 2021.

[J20] Hoda Eldardiry, Jennifer Neville, **Ryan A. Rossi**, *Ensemble Learning for Relational Data*, Journal of Machine Learning Research (JMLR), 2020.

[J19] **Ryan A. Rossi**, Nesreen K. Ahmed, Aldo Carranza, David Arbour, Anup Rao, Sungchul Kim, and Eunyee Koh, *Heterogeneous Graphlets*, Transactions on Knowledge Discovery from Data (TKDD), 1–43, 2020.

[J18] Xin Qian, **Ryan A. Rossi**, Fan Du, Sungchul Kim, Eunyee Koh, Sana Malik, Tak Yeon Lee, Nesreen K. Ahmed, *Personalized Visualization Recommendation*, ACM Transactions on the Web (TWEB), 1–43, 2021.

[J17] **Ryan A. Rossi**, Di Jin, Sungchul Kim, Nesreen K. Ahmed, Danai Koutra, and John Boaz Lee, *On Proximity and Structural Role-based Embeddings in Networks: Misconceptions, Techniques, and Applications*, Transactions on Knowledge Discovery from Data (TKDD), 1–32, 2020.

[J16] John Boaz Lee, Giang Nguyen, **Ryan A. Rossi**, Nesreen K. Ahmed, Eunyee Koh, and Sungchul Kim, *Dynamic Node Embeddings from Edge Streams*, IEEE Transactions on Emerging Topics in Computational Intelligence, 1–15, 2020.

[J15] John Boaz Lee, **Ryan A. Rossi**, Sungchul Kim, Nesreen K. Ahmed, and Eunyee Koh, *Attention Models in Graphs: A Survey*, Transactions on Knowledge Discovery from Data (TKDD), 1–19, 2019.

[J14] Nesreen K. Ahmed, **Ryan A. Rossi**, Rong Zhou, John Boaz Lee, Xiangnan Kong, Theodore L. Willke, Hoda Eldardiry, *Role-based Graph Embeddings*, IEEE Transactions on Knowledge and Data Engineering (TKDE), 2020.

[J13] **Ryan A. Rossi** and Nesreen K. Ahmed, *Complex networks are structurally distinguishable by domain*, Social Network Analysis and Mining (SNAM), Vol. 9, No. 1, 51 pages, 2020.

[J12] **Ryan A. Rossi**, Rong Zhou, and Nesreen K. Ahmed, *Deep Inductive Graph Representation Learning*, IEEE Transactions on Knowledge and Data Engineering (TKDE), 1–14, 2018.

[J11] **Ryan A. Rossi**, Rong Zhou, and Nesreen K. Ahmed, *Estimation of Graphlet Counts in Massive Networks*, IEEE Transactions on Neural Networks and Learning Systems (TNNLS), 14 pages, 2018.

[J10] **Ryan A. Rossi**, Nesreen K. Ahmed, and Rong Zhou, *Interactive Visual Graph Mining and Learning*, ACM Transactions on Intelligent Systems and Technology (TIST), 1–24, 2017.

[J9] **Ryan A. Rossi** and Rong Zhou, *GraphZIP: A Clique-based Sparse Graph Compression Method*, Journal of Big Data, 2018.

[J8] Nesreen K. Ahmed, Jennifer Neville, **Ryan A. Rossi**, Nick Duffield, Theodore L. Willke, *Graphlet Decomposition: Framework, Algorithms, and Applications*, Knowledge and Information Systems (KAIS), 689–722, 2016 **Invited paper to KAIS Journal Special Issue (ICDM Best papers)*.

[J7] **Ryan Rossi**, *Relational Time Series Learning*, Knowledge Engineering Review (KER), Cambridge University Press, 1–15, 2018.

[J6] **Ryan Rossi** and Rong Zhou, *Parallel Collective Factorization for Modeling Large Heterogeneous Networks*, Social Network Analysis and Mining (SNAM), 2016.

[J5] **Ryan Rossi**, David F. Gleich, and Assefaw H. Gebremedhin, *Parallel Maximum Clique Algorithms with Applications to Network Analysis*, SIAM Journal on Scientific Computing (SISC), 37(5), C589–C616 (28 pages), 2015.

[J4] **Ryan Rossi** and Nesreen K. Ahmed, *Role Discovery in Networks*, IEEE Transactions on Knowledge and Data Engineering (TKDE), 1112–1131, 2014.

[J3] **Ryan Rossi** and Nesreen K. Ahmed, *Coloring Large Complex Networks*, Social Network Analysis and Mining (SNAM), Vol. 4, No. 1-228, 37 pages, 2014.

[J2] David F. Gleich, **Ryan A. Rossi**, *A Dynamical System for PageRank with Time-Dependent Teleportation*, Internet Mathematics, 188–217, 2014 .

[J1] **Ryan Rossi**, Luke McDowell, David Aha, and Jennifer Neville, *Transforming Graph Representations for Statistical Relational Learning*, Journal of Artificial Intelligence Research (JAIR), pages 363–441, 2012. *Invited for presentation at IJCAI 2013 journal track.

Other Peer-reviewed Publications

[C233] Mohsen Fayyaz, Ali Modarressi, Hanieh Deilamsalehy, Franck Dernoncourt, **Ryan A. Rossi**, Trung Bui, Hinrich Schütze, and Nanyun Peng, *Steering MoE LLMs via Expert (De)Activation*, International Conference on Learning Representations (ICLR), 2026.

[C232] Shawn Li, **Ryan A. Rossi**, Sungchul Kim, Sunav Choudhary, Franck Dernoncourt, Puneet Mathur, Zhengzhong Tu, and Yue Zhao, *Charts Are Not Images: On the Challenges of Scientific Chart Editing*, International Conference on Learning Representations (ICLR), 2026.

[C231] Chien Van Nguyen, Ruiyi Zhang, Hanieh Deilamsalehy, Puneet Mathur, Viet Dac Lai, Haoliang Wang, Jayakumar Subramanian, **Ryan A. Rossi**, Trung Bui, Nikos Vlassis, Franck Dernoncourt, and Thien Huu Nguyen, *Lizard: An Efficient Linearization Framework for LLMs*, Annual Meeting of the Association for Computational Linguistics (ACL), 2026.

[C230] Chien Van Nguyen, **Ryan A. Rossi**, Linh Ngo Van, Franck Dernoncourt, Thien Huu Nguyen, *Octopus: Gated Selective Attention for Memory-Bounded Long-Context Inference*, Annual Meeting of the Association for Computational Linguistics (ACL), 2026.

[C229] Manan Suri, Puneet Mathur, Nedim Lipka, Franck Dernoncourt, **Ryan A. Rossi**, and Dinesh Manocha, *Structured Uncertainty Guided Clarification for LLM Agents*, Findings of the Association for Computational Linguistics: ACL, 2026.

[C228] Brian Zhang, Deepti Guntur, Zhiyang Zuo, Abhinav Sharma, Shreyas Chaudhari, Wenlong Zhao, Franck Dernoncourt, Puneet Mathur, **Ryan A. Rossi**, and Nedim Lipka, *Test-Time Strategies for More Efficient and Accurate Agentic RAG*, ACL Student Research Workshop (ACL SRW), 2026.

[C227] Anish Pahilajani, Devasha Trivedi, Jincen Shuai, Khin S. Yone, Samyak Rajesh Jain, Namyong Park, **Ryan A. Rossi**, Nesreen K. Ahmed, Franck Dernoncourt, and Yu Wang, *Reasoning Graph-Structured Question Answering*, International Conference on Language Resources and Evaluation (LREC), 2026.

[C226] Li Li, Peilin Cai, **Ryan A. Rossi**, Franck Dernoncourt, Branislav Kveton, Junda Wu, Tong Yu, Linxin Song, Tiankai Yang, Yuehan Qin, Nesreen K. Ahmed, Samyadeep Basu, Subhojyoti Mukherjee, Ruiyi Zhang, Zhengmian Hu, Bo Ni, Yuxiao Zhou, Zichao Wang, Yue Huang, Yu Wang, Xiangliang Zhang, Philip S. Yu, Xiyang Hu, and Yue Zhao, *A Personalized Conversational Benchmark*, NeurIPS 2025 Workshop on Multi-Turn Interactions in LLMs (MTI-LLM), 2025 (Spotlight, top 5%).

[C225] Ho Yin Sam Ng, Ting-Yao Hsu, Aashish Anantha Ramakrishnan, Branislav Kveton, Nedim Lipka, Franck Dernoncourt, Dongwon Lee, Tong Yu, Sungchul Kim, **Ryan A. Rossi**, and Ting-Hao Kenneth Huang, *LaMP-Cap: Personalized Figure Caption Generation with Multimodal Figure Profiles*, Findings of the Association for Computational Linguistics: EMNLP, 2025.

[C224] Ali Modarressi, Hanieh Deilamsalehy, Franck Dernoncourt, Trung Bui, **Ryan A. Rossi**, Seunghyun Yoon, and Hinrich Schütze, *NoLiMa: Long-Context Evaluation Beyond Literal Matching*, International Conference on Machine Learning (ICML), 2025.

[C223] Yu Wang, **Ryan A. Rossi**, Namyong Park, Nesreen K. Ahmed, Danai Koutra, Franck Dernoncourt, and Tyler Derr, *Demystifying the Power of Large Language Models in Graph Generation*, Findings of the Association for Computational Linguistics: NAACL, 2025.

[C222] Yu Xia, Sungchul Kim, Tong Yu, **Ryan A. Rossi**, and Julian McAuley, *Multi-Agent Collaborative Filtering: Orchestrating Users and Items for Agentic Recommendations*, The ACM Web Conference (WWW), 2026.

[C221] Zhehao Zhang, **Ryan A. Rossi**, Tong Yu, Franck Dernoncourt, Ruiyi Zhang, Jiuxiang Gu, Sungchul Kim, Xiang Chen, Zichao Wang, and Nedim Lipka, *VipAct: Visual-Perception Enhancement via Specialized VLM Agent Collaboration and Tool-use*, AAAI Conference on Artificial Intelligence (AAAI), 2026.

[C220] Bo Ni, Yu Wang, Leyao Wang, Branislav Kveton, Franck Dernoncourt, Yu Xia, Hongjie Chen, Reuben Luera, Samyadeep Basu, Subhojyoti Mukherjee, Puneet Mathur, Nesreen K. Ahmed, Junda Wu, Li Li, Huixin Zhang, Ruiyi Zhang, Tong Yu, Sungchul Kim, Jiuxiang Gu, Zhengzhong Tu, Alexa F. Siu, Zichao Wang, Seunghyun Yoon, Nedim Lipka, Namyong Park, Zihao Lin, Trung Bui, Yue Zhao, Tyler Derr, and **Ryan A. Rossi**, *A Survey on LLM-based Conversational User Simulation*, European Chapter of the Association for Computational Linguistics (EACL), 2026.

[C219] Sriram Balasubramanian, Samyadeep Basu, Koustava Goswami, **Ryan A. Rossi**, Varun Manjunatha, Roshan Santhosh, Ruiyi Zhang, Soheil Feizi, and Nedim Lipka, *Decomposition-Enhanced Training for Post-Hoc Attributions in Language Models*, European Chapter of the Association for Computational Linguistics (EACL), 2026.

[C218] Utkarsh Sahu, Zhisheng Qi, Mahantesh Halappanavar, Nedim Lipka, **Ryan A. Rossi**, Franck Dernoncourt, Yu Zhang, Yao Ma, and Yu Wang, *Knowledge Homophily in Large Language Models*, ACM International Conference on Web Search and Data Mining (WSDM), 2026.

[C217] Koustava Goswami, Puneet Mathur, **Ryan A. Rossi**, Franck Dernoncourt, Vivek Gupta, and Dinesh Manocha, *ChartEval: LLM-Driven Chart Generation Evaluation Using Scene Graph Parsing*, International Joint Conference on Natural Language Processing (IJCNLP), 2025.

[C216] Manas Bumb, Aditya Vemulapalli, Sri Harsha V. P. Jella, Anvesh Gupta, Anurag La, **Ryan A. Rossi**, and Hongjie Chen, *Forecasting Time Series with LLMs via Patch-based Prompting and Decomposition*, Pacific Asia Conference on Language, Information and Computation (PACLIC), 2025.

[C215] Steven Au, Cameron J. Dimacali, Ojasmitha Pedirappagari, Namyong Park, Franck Dernoncourt, Yu Wang, Nesreen K. Ahmed, **Ryan A. Rossi**, Nedim Lipka, and Tyler Derr, *Personalized Graph-based Retrieval for Large Language Models*, Pacific Asia Conference on Language, Information and Computation (PACLIC), 2025.

[C214] Mehrab Tanjim, Yeonjun In, Xiang Chen, Victor Bursztyn, **Ryan A. Rossi**, Sungchul Kim, Guang-Jie Ren, Vaishnavi Muppala, Shun Jiang, Yujie Zhong, and Chien Van Nguyen, *Disambiguation in Conversational Question Answering in the Era of LLMs and Agents: A Survey*, Empirical Methods in Natural Language Processing (EMNLP), 2025.

[C213] Yu Xia, Yangshijie Shen, Junda Wu, Tong Yu, Sungchul Kim, **Ryan A. Rossi**, Lina Yao, and Julian McAuley, *SAND: Boosting LLM Agents with Self-Taught Action Deliberation*, Empirical Methods in Natural Language Processing (EMNLP), 2025.

[C212] Manan Suri, Puneet Mathur, Nedim Lipka, Franck Dernoncourt, **Ryan A. Rossi**, Vivek Gupta, and Dinesh Manocha, *Follow the Flow: Fine-grained Flowchart Attribution with Neurosymbolic Agents*, Empirical Methods in Natural Language Processing (EMNLP), 2025.

[C211] Dongliang Guo, Hang Zhao, **Ryan A. Rossi**, Sungchul Kim, Nedim Lipka, Tong Yu, and Sheng Li, *Few-shot Fine-grained Image Classification with Interpretable Prompt Learning through Distribution Alignment*, International Conference on Multimodal Interaction (ICMI), 2025.

[C210] Jy-Yong Choi, Sung-Gyum Kim, Jeongwon Jeong, **Ryan A. Rossi**, Jin-Hwa Kil, and Tak Yeon Lee, *Gaze2Prompt: Turning Eye-Tracking Data into Visual Prompts for Multimodal LLMs*, ACM International Joint Conference on Pervasive and Ubiquitous Computing Companion (UbiComp), 2025.

[C209] Kourosh Banihashem, Xi Chen, MohammadTaghi Hajiaghayi, Sungchul Kim, Kanak Mahadik, **Ryan A. Rossi**, and Tong Yu, *Fully Dynamic Embedding into ℓ_p Spaces*, International Conference on Machine Learning (ICML), 2025.

[C208] Haitao Xu, Zheng Yao, Yushun Dong, Zichao Wang, **Ryan A. Rossi**, Mengnan Li, and Yue Zhao, *Few-shot Graph Out-of-distribution Detection with LLMs*, European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD), 2025.

[C207] Jiahe Fang, Chia-Tung Liu, Jaewook Kim, Yashas Bhedaru, Elvis Liu, Neel Singh, Nedim Lipka, Puneet Mathur, **Ryan A. Rossi**, Franck Dernoncourt, Vivek Gupta, and Dinesh Manocha, *Multi-LLM Text Summarization*, Recent Advances in Natural Language Processing (RANLP), 2025.

[C206] Chien Van Nguyen, Xuan Shen, Ryan Aponte, Yu Xia, Samyadeep Basu, Zhengzhong Hu, Jian Chen, Mihir Parmar, **Ryan A. Rossi**, and others, *A Survey on Small Language Models*, Recent Advances in Natural Language Processing (RANLP), 2025.

[C205] Deonna M. Owens, **Ryan A. Rossi**, Sungchul Kim, Tong Yu, Franck Dernoncourt, Xiang Chen, Ruiyi Zhang, Jiuxiang Gu, Hanieh Deilamsalehy, and Nedim Lipka, *Multi-LLM Debiasing Framework*, Recent Advances in Natural Language Processing (RANLP), 2025.

[C204] Ryan Aponte, **Ryan A. Rossi**, Shunan Guo, Franck Dernoncourt, Tong Yu, Xiang Chen, Saayan Mitra, and Georgios Theocharous, *A Framework for Fine-tuning LLMs Using Heterogeneous Feedback*, Recent Advances in Natural Language Processing (RANLP), 2025.

[C203] Ashish Singh, Prateek Agarwal, Zixuan Huang, Arpita Singh, Tong Yu, Sungchul Kim, Victor Bursztyn, Nikos Vlassis, and **Ryan A. Rossi**, *FigCaps-HF: A Figure-to-Caption Generative Framework and Benchmark with Human Feedback*, Recent Advances in Natural Language Processing (RANLP), 2025.

[C202] Hongjie Chen, **Ryan A. Rossi**, Sungchul Kim, Kanak Mahadik, and Hoda Eldardiry, *Probabilistic Hypergraph Recurrent Neural Networks for Time-Series Forecasting*, ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2025.

[C201] Hieu Man, Nghia Trung Ngo, Viet Dac Lai, **Ryan A. Rossi**, Franck Dernoncourt, and Thien Huu Nguyen, *LUSIFER: Language Universal Space Integration for Enhanced Representation in Multilingual Text Embedding Models*, ACM SIGIR Conference on Research and Development in Information Retrieval, 2025.

[C200] Junda Wu, Yu Xia, Tong Yu, Xiang Chen, Sai Sree Harsha, Akash V. Maharaj, Ruiyi Zhang, Victor Bursztyn, and **Ryan A. Rossi**, *Doc-ReAct: Multi-page Heterogeneous Document Question-Answering*, Annual Meeting of the Association for Computational Linguistics (ACL), 2025.

[C199] Manan Suri, Puneet Mathur, Nedim Lipka, Franck Dernoncourt, **Ryan A. Rossi**, and Dinesh Manocha, *ChartLens: Fine-grained Visual Attribution in Charts*, Annual Meeting of the Association for Computational Linguistics (ACL), 2025.

[C198] Vipula Rawte, **Ryan A. Rossi**, Franck Dernoncourt, and Nedim Lipka, *Document Attribution: Examining Citation Relationships using Large Language Models*, Workshop on Scholarly Document Processing (SDP) at ACL, 2025.

[C197] Yongjia Lei, Haoyu Han, **Ryan A. Rossi**, Franck Dernoncourt, Nedim Lipka, Mahantesh M. Halappanavar, and Yu Wang, *Mixture of Structural-and-Textual Retrieval over Text-rich Graph Knowledge Bases*, Findings of the Association for Computational Linguistics: ACL, 2025.

[C196] Yu Xia, Subhojyoti Mukherjee, Zhouhang Xie, Junda Wu, Xintong Li, Ryan Aponte, Hanjia Lyu, Joe Barrow, Hongjie Chen, and **Ryan A. Rossi**, *From Selection to Generation: A Survey of LLM-based Active Learning*, Annual Meeting of the Association for Computational Linguistics (ACL), 2025.

[C195] Dang Nguyen, Jian Chen, Yu Wang, Gang Wu, Namyong Park, Zhengzhong Hu, Hanjia Lyu, Junda Wu, Ryan Aponte, and **Ryan A. Rossi**, *GUI Agents: A Survey*, Findings of the Association for Computational Linguistics: ACL, 2025.

[C194] Tiankai Yang, Yi Nian, Li Li, Ruiyao Xu, Yuangang Li, Jiaqi Li, Zhuo Xiao, Xiyang Hu, **Ryan A. Rossi**, Kaize Ding, and Xia Hu, *AD-LLM: Benchmarking Large Language Models for Anomaly Detection*, Findings of the Association for Computational Linguistics: ACL, 2025.

[C193] Zhixin Lin, Zhiyang Wang, Yanan Pan, Varun Manjunatha, **Ryan A. Rossi**, Adrian Lau, Lijuan Huang, and Tong Sun, *Persona-SQ: A Personalized Suggested Question Generation Framework For Real-world Documents*, Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL), 2025.

[C192] Manan Suri, Puneet Mathur, Franck Dernoncourt, Koustava Goswami, **Ryan A. Rossi**, and Dinesh Manocha, *VisDoM: Multi-document QA with Visually Rich Elements Using Multimodal Retrieval-Augmented Generation*, Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL), 2025.

[C191] Yu Xia, Junda Wu, Sungchul Kim, Tong Yu, **Ryan A. Rossi**, Haoliang Wang, and Julian McAuley, *Knowledge-Aware Query Expansion with Large Language Models for Textual and Relational Retrieval*, Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL), 2025.

[C190] Yeonjun In, Sungchul Kim, **Ryan A. Rossi**, Mehrab Tanjim, Tong Yu, Ritwik Sinha, and Chanyoung Park, *Diversify-Verify-Adapt: Efficient and Robust Retrieval-Augmented Ambiguous Question Answering*, Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL), 2025.

[C189] Isabel O. Gallegos, Ryan Aponte, **Ryan A. Rossi**, Joe Barrow, Mehrab Tanjim, Tong Yu, Hanieh Deilamsalehy, Ruiyi Zhang, Sungchul Kim, Franck Dernoncourt, Nedim Lipka, Seunghyun Yoon, Jiuxiang Gu, Xiang Chen, Apoorv Saxena, and Thien Huu Nguyen, *Self-Debiasing Large Language Models: Zero-Shot Recognition and Reduction of Stereotypes*, Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL), 2025.

[C188] Xuan Shen, Zhenyu Song, Yufan Zhou, Bo Chen, Jiaying Liu, Ruiyi Zhang, **Ryan A. Rossi**, Hao Tan, and Tong Yu, *Numerical Pruning for Efficient Autoregressive Models*, AAAI Conference on Artificial Intelligence (AAAI), 2025.

[C187] Koustava Goswami, Puneet Mathur, **Ryan A. Rossi**, and Franck Dernoncourt, *PlotEdit: Natural Language-Driven Accessible Chart Editing in PDFs via Multimodal LLM Agents*, European Conference on Information Retrieval (ECIR), 2025.

[C186] Jian Chen, Ruiyi Zhang, Yufan Zhou, Tong Yu, Franck Dernoncourt, Jiuxiang Gu, **Ryan A. Rossi**, Changyou Chen, and Tong Sun, *SV-RAG: LoRA-Contextualizing Adaptation of MLLMs for Long Document Understanding*, International Conference on Learning Representations (ICLR), 2025.

[C185] Dang Nguyen, Viet Dac Lai, Seunghyun Yoon, **Ryan A. Rossi**, Handong Zhao, Ruiyi Zhang, and Puneet Mathur, *DynaSaur: Large Language Agents Beyond Predefined Actions*, International Conference on Learning Representations (ICLR), 2025.

[C184] Yiwei Wang, **Ryan A. Rossi**, Namyong Park, Hongjie Chen, Nesreen K. Ahmed, Puja Trivedi, Franck Dernoncourt, Danai Koutra, and Tyler Derr, *A Large-scale Training Paradigm for Graph Generative Models*, International Conference on Learning Representations (ICLR), 2025.

[C183] Songwen Hu, **Ryan A. Rossi**, Tong Yu, Junda Wu, Handong Zhao, Sungchul Kim, and Sheng Li, *Interactive Visualization Recommendation with Hier-SUCB*, The ACM Web Conference (WWW), 2025.

[C182] Reuben Luera, **Ryan A. Rossi**, Franck Dernoncourt, Alexa F. Siu, Sungchul Kim, Tong Yu, Ruiyi Zhang, and Xiang Chen, *Personalizing Data Delivery: Investigating User Characteristics and Enhancing LLM Predictions*, Companion Proceedings of the ACM Web Conference (WWW Companion), 2025.

[C181] Koustava Goswami, Puneet Mathur, **Ryan A. Rossi**, and Franck Dernoncourt, *ChartCitor: Answer Citations for ChartQA via Multi-Agent LLM Retrieval*, Companion Proceedings of the ACM Web Conference (WWW Companion), 2025.

[C180] Koustava Goswami, Puneet Mathur, **Ryan A. Rossi**, and Franck Dernoncourt, *PlotGen: Multi-agent LLM-based Scientific Data Visualization via Multimodal Retrieval Feedback*, Companion Proceedings of the ACM Web Conference (WWW Companion), 2025.

[C179] Sung-Gyum Kim, Jy-Yong Choi, **Ryan A. Rossi**, Eunyee Koh, and Tak Yeon Lee, *Chart-to-Experience: Benchmarking Multimodal LLMs for Predicting Experiential Impact of Charts*, IEEE Pacific Visualization Conference (PacificVis), 2025.

[C178] Yoonsuh Lee, Nedim Lipka, Zichao Wang, **Ryan A. Rossi**, Puneet Mathur, Tong Sun, and Alexa F. Siu, *DocVoyager: Anticipating User's Information Needs and Guiding Document Reading through Question Answering*, CHI Conference on Human Factors in Computing Systems Extended Abstracts, 2025.

[C177] Xinyu Shi, Yinghou Wang, **Ryan A. Rossi**, and Jian Zhao, *Brickify: Enabling Expressive Design Intent Specification through Direct Manipulation on Design Tokens*, CHI Conference on Human Factors in Computing Systems (CHI), 2025.

[C176] Moayed Haji Ali, Lilin Zhao, Xiaohang Shen, Somdeb Mukherjee, Yufan Zhou, Gang Wu, Viet Dac Lai, and **Ryan A. Rossi**, *A Survey on Long-Video Storytelling Generation: Architectures, Consistency, and Cinematic Quality*, IEEE International Conference on Computer Vision (ICCV), 2025.

[C175] Sung-Gyum Kim, Jy-Yong Choi, Pyo-Yeon Lee, Jiwon Chung, **Ryan A. Rossi**, Jin-Hwa Kil, Eunyee Koh, and Tak Yeon Lee, *Grounded Generation of Embellished Bar Chart Ensuring Chart Integrity*, IEEE Visualization and Visual Analytics (VIS) Short Papers, 2025.

[C174] Hongjie Chen, Ayoub Dahmane Beachnau, Puranjay Thomas, Pranav Maneriker, John Kimball, and **Ryan A. Rossi**, *Live-ITS: LSH-based Interactive Visualization Explorer for Large-scale Incomplete Time Series*, IEEE International Conference on Big Data (BigData), 2024.

[C173] Hongjie Chen, **Ryan A. Rossi**, Sungchul Kim, Kanak Mahadik, and Hoda Eldardiry, *A Study of Foundation Models for Large-scale Time-series Forecasting*, IEEE International Conference on Big Data (BigData), 2024.

[C172] Shravika Mittal, Kanak Mahadik, **Ryan A. Rossi**, Sungchul Kim, and Handong Zhao, *GPURank: A Cloud GPU Instance Recommendation System*, IEEE International Conference on Big Data (BigData), 2024.

[C171] Mihir Parmar, Hanieh Deilamsalehy, Franck Dernoncourt, Seunghyun Yoon, **Ryan A. Rossi**, and Trung Bui, *Towards Enhancing Coherence in Extractive Summarization: Dataset and Experiments with LLMs*, Empirical Methods in Natural Language Processing (EMNLP), 2024.

[C170] Jon Saad-Falcon, Joe Barrow, Alexa F. Siu, Ani Nenkova, Seunghyun Yoon, **Ryan A. Rossi**, and Franck Dernoncourt, *PDFTriage: Question Answering over Long, Structured Documents*, Empirical Methods in Natural Language Processing (EMNLP) Industry Track, 2024.

[C169] Yu Wang, Nedim Lipka, Ruiyi Zhang, Alexa Siu, Yuying Zhao, Bo Ni, Xin Wang, **Ryan A. Rossi**, and Tyler Derr, *Topology-aware Retrieval Augmentation for Text Generation*, ACM International Conference on Information and Knowledge Management (CIKM), 2024.

[C168] Yu Wang, Kaize Ding, Xiaorui Liu, Jian Kang, **Ryan A. Rossi**, and Tyler Derr, *Data Quality-aware Graph Machine Learning*, ACM International Conference on Information and Knowledge Management (CIKM), 2024.

[C167] Yu Fu, Shunan Guo, Jane Hoffswell, Victor S. Bursztyn, **Ryan A. Rossi**, and John Stasko, *"The Data Says Otherwise": Towards Automated Fact-checking and Communication of Data Claims*, ACM Symposium on User Interface Software and Technology (UIST), 2024.

[C166] Junda Wu, Tong Yu, Xiang Chen, Haoliang Wang, **Ryan A. Rossi**, Sungchul Kim, Anup Rao, and Julian McAuley, *DeCoT: Debiasing Chain-of-Thought for Knowledge-Intensive Tasks in Large Language Models via Causal Intervention*, Annual Meeting of the Association for Computational Linguistics (ACL), 2024.

[C165] Swetha Eppalapally, Daksh Dangi, Chaithra Bhat, Ankita Gupta, Ruiyi Zhang, Shubham Agarwal, Karishma Bagga, and **Ryan A. Rossi**, *KAPQA: Knowledge-Augmented Product Question-Answering*, Workshop on Knowledge Augmented Methods for NLP at ACL, 2024.

[C164] Puja Trivedi, **Ryan A. Rossi**, David Arbour, Tong Yu, Franck Dernoncourt, Sungchul Kim, Nedim Lipka, Namyong Park, Nesreen K. Ahmed, and Danai Koutra, *Editing Partially Observable Networks via Graph Diffusion Models*, International Conference on Machine Learning (ICML), 2024.

[C163] Samyadeep Basu, Keivan Rezaei, Priyatham Kattakinda, Vlad I. Morariu, Nanxuan Zhao, **Ryan A. Rossi**, Varun Manjunatha, and Soheil Feizi, *On Mechanistic Knowledge Localization in Text-to-Image Generative Models*, International Conference on Machine Learning (ICML), 2024.

[C162] Zhengzhong Hu, Yi Zhang, **Ryan A. Rossi**, Tong Yu, Sungchul Kim, and Shijia Pan, *Are Large Language Models Capable of Causal Reasoning for Sensing Data Analysis?*, Workshop on Edge and Mobile Foundation Models, 2024.

[C161] Rui Wang, Tong Yu, Ruiyi Zhang, Sungchul Kim, **Ryan A. Rossi**, Handong Zhao, Junda Wu, Saayan Mitra, Lina Yao, and Ricardo Henao, *Personalized Federated Learning for Text Classification with Gradient-Free Prompt Tuning*, Findings of the Association for Computational Linguistics: NAACL, 2024.

[C160] Yu Xia, Xu Liu, Tong Yu, Sungchul Kim, **Ryan A. Rossi**, Anup Rao, Tung Mai, and Shuai Li, *Hallucination Diversity-Aware Active Learning for Text Summarization*, Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL), 2024.

[C159] Mehrnoosh Mirtaheri, **Ryan A. Rossi**, Sungchul Kim, Kanak Mahadik, Tong Yu, Xiang Chen, and Mohammad Rostami, *Tackling Long-Tail Entities for Temporal Knowledge Graph Completion*, Companion Proceedings of the ACM Web Conference (WWW Companion), 2024.

[C158] Yu Xia, Fang Kong, Tong Yu, Liya Guo, **Ryan A. Rossi**, Sungchul Kim, and Shuai Li, *Which LLM to Play? Convergence-aware Online Model Selection with Time-Increasing Bandits*, The ACM Web Conference (WWW), 2024.

[C157] Xinyu Shi, Minyue Liu, Ziqi Zhou, Ali Neshati, **Ryan A. Rossi**, and Jian Zhao, *Exploring Interactive Color Palettes for Abstraction-Driven Exploratory Image Colorization*, CHI Conference on Human Factors in Computing Systems (CHI), 2024.

[C156] Fatima Nouraei, Alexa F. Siu, **Ryan A. Rossi**, and Nedim Lipka, *Thinking Outside the Box: Non-Designer Perspectives and Recommendations for Template-Based Graphic Design Tools*, CHI Conference on Human Factors in Computing Systems Extended Abstracts, 2024.

[C155] Ting-Yao Hsu, Chieh-Yang Huang, Shih-Hong Huang, **Ryan A. Rossi**, Sungchul Kim, Tong Yu, C. Lee Giles, and Ting-Hao Kenneth Huang, *SciCapenter: Supporting Caption Composition for Scientific Figures with Machine-Generated Captions and Ratings*, CHI Conference on Human Factors in Computing Systems Extended Abstracts, 2024.

[C154] Thuat Nguyen, Chien Van Nguyen, Viet Dac Lai, Hieu Man, Nghia Trung Ngo, Franck Dernoncourt, **Ryan A. Rossi**, and Thien Huu Nguyen, *CulturaX: A Cleaned, Enormous, and Multilingual Dataset for Large Language Models in 167 Languages*, Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING), 2024.

[C153] Hongjie Chen, **Ryan A. Rossi**, Sungchul Kim, Kanak Mahadik, and Hoda Eldardiry, *Evolving Super Graph Neural Networks for Large-scale Time-Series Forecasting*, Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), 2024.

[C152] Gaurav Shazan Ahmad, Sharmin Agarwal, Saayan Mitra, **Ryan A. Rossi**, Malay K. Doshi, Vibhor Porwal, and S. M. K. Paila, *ScaleViz: Scaling Visualization Recommendation Models on Large Data*, Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), 2024.

[C151] Namyong Park, Xing Wang, Antoine Simoulin, Shuai Yang, Grey Yang, **Ryan A. Rossi**, Puja Trivedi, and Nesreen K. Ahmed, *Forward Learning of Graph Neural Networks*, International Conference on Learning Representations (ICLR), 2024.

[C150] Yu Wang, Nedim Lipka, **Ryan A. Rossi**, Alexa Siu, Ruiyi Zhang, and Tyler Derr, *Knowledge Graph Prompting for Multi-Document Question Answering*, AAAI Conference on Artificial Intelligence (AAAI), 2024.

[C149] Jian Chen, Ruiyi Zhang, Yufan Zhou, Rajiv Jain, Zhiqiang Xu, **Ryan A. Rossi**, and Changyou Chen, *Towards Aligned Layout Generation via Diffusion Model with Aesthetic Constraints*, International Conference on Learning Representations (ICLR), 2024.

[C148] **Ryan A. Rossi**, Somdeb Sarkhel, and Nesreen K. Ahmed, *Efficient Estimation of Local Causal Effects in Graphs via Neighborhood Pooling*, IEEE International Conference on Big Data (BigData), 2023.

[C147] Hongjie Chen, **Ryan A. Rossi**, Kanak Mahadik, Sungchul Kim, and Hoda Eldardiry, *Hypergraph Neural Networks for Time-Series Forecasting*, IEEE International Conference on Big Data (BigData), 2023.

[C146] Jianjun Luo, **Ryan A. Rossi**, Xiangnan Kong, and Yanhua Li, *Multi-item Continuous Influence Maximization*, IEEE International Conference on Big Data (BigData), 2023.

[C145] Brandon Stigall, **Ryan A. Rossi**, Jane Hoffswell, Xiang Chen, Shunan Guo, Fan Du, Eunyee Koh, and Kelly Caine, *On Chatbots for Visual Exploratory Data Analysis*, IEEE International Conference on Big Data (BigData), 2023.

[C144] Gaurav Gupta, Anup Rao, Tung Mai, **Ryan A. Rossi**, Xiang Chen, Saayan Mitra, and Anshumali Shrivastava, *Near Neighbor Search for Constraint Queries*, IEEE International Conference on Big Data (BigData), 2023.

[C143] Jaeho Bang, Sungchul Kim, **Ryan A. Rossi**, Tong Yu, and Handong Zhao, *Interpretable Unsupervised Log Anomaly Detection*, IEEE International Conference on Big Data (BigData), 2023.

[C142] Camille Harris, **Ryan A. Rossi**, Sana Malik, Jane Hoffswell, Fan Du, Tak Yeon Lee, Eunyee Koh, and Handong Zhao, *Visual Insight Recommendation: From Ranking Insight Visualizations to Insight Types*, IEEE International Conference on Big Data (BigData), 2023.

[C141] Arpit Narechania, Shubham Chakraborty, Shubham Agarwal, Atanu R. Sinha, **Ryan A. Rossi**, Fan Du, and Eunyee Koh, *DataCockpit: A Toolkit for Data Lake Navigation and Monitoring Utilizing Quality and Usage Information*, IEEE International Conference on Big Data (BigData), 2023.

[C140] Pattara Sukprasert, Gromit Yeuk-Yin Chan, **Ryan A. Rossi**, Fan Du, and Eunyee Koh, *Discovery and Matching Numerical Attributes in Data Lakes*, IEEE International Conference on Big Data (BigData), 2023.

[C139] Alex Tang, Gromit Chan, **Ryan A. Rossi**, Chang Xiao, and Eunyee Koh, *Tabular Data to Image Generation: Benchmark Data, Approaches, and Evaluation*, IEEE International Conference on Big Data (BigData), 2023.

[C138] Gautam Choudhary, Iftikhar Ahamath Burhanuddin, Eunyee Koh, Fan Du, and **Ryan A. Rossi**, *PersonaSAGE: A Multi-Persona Graph Neural Network*, IEEE International Conference on Big Data (BigData), 2023.

[C137] Namyong Park, **Ryan A. Rossi**, Xing Wang, Antoine Simoulin, Nesreen K. Ahmed, and Christos Faloutsos, *GLEMOS: Benchmark for Instantaneous Graph Learning Model Selection*, Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track, 2023.

[C136] Shanyun Gao, Raghavendra Addanki, Tong Yu, **Ryan A. Rossi**, and Murat Kocaoglu, *Causal Discovery in Semi-Stationary Time Series*, Advances in Neural Information Processing Systems (NeurIPS), 2023.

[C135] Sudhanshu Chanpuriya, **Ryan A. Rossi**, Anup B. Rao, Tung Mai, Nedim Lipka, Zhao Song, and Cameron Musco, *Exact Representation of Sparse Networks with Symmetric Nonnegative Embeddings*, Advances in Neural Information Processing Systems (NeurIPS), 2023.

[C134] Yang Li, Gang Wu, Stefano Petrangeli, Haoliang Wang, **Ryan A. Rossi**, and Viswanathan Swaminathan, *Active Context Modeling for Efficient Image and Burst Compression*, IEEE International Symposium on Multimedia (ISM), 2023.

[C133] Ting-Yao Hsu, Chieh-Yang Huang, **Ryan A. Rossi**, Sungchul Kim, C. Lee Giles, and Ting-Hao Kenneth Huang, *GPT-4 as an Effective Zero-Shot Evaluator for Scientific Figure Captions*, Findings of the Association for Computational Linguistics: EMNLP, 2023.

[C132] Viet Lai, Chien Van Nguyen, Nghia Trung Ngo, Thuat Nguyen, Franck Dernoncourt, **Ryan A. Rossi**, and Thien Huu Nguyen, *Okapi: Instruction-tuned Large Language Models in Multiple Languages with Reinforcement Learning from Human Feedback*, Empirical Methods in Natural Language Processing (EMNLP) Demonstration Track, 2023.

[C131] Guande Wu, Shunan Guo, Jane Hoffswell, Gromit Yeuk-Yin Chan, **Ryan A. Rossi**, and Eunyee Koh, *Socrates: Data Story Generation via Adaptive Machine-Guided Elicitation of User Feedback*, IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS), 2023.

[C130] Jean-Peic Chou, Alexa F. Siu, Nedim Lipka, **Ryan A. Rossi**, Franck Dernoncourt, and Maneesh Agrawala, *TaleStream: Supporting Story Ideation with Trope Knowledge*, ACM Symposium on User Interface Software and Technology (UIST), 2023.

[C126] Hyeok Kim, **Ryan A. Rossi**, Jessica Hullman, and Jane Hoffswell, *DuPo: A Mixed-Initiative Authoring Tool for Responsive Visualization*, IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS), 2023.

[C125] Harsh Chopra, Atanu R. Sinha, Shubham Choudhary, **Ryan A. Rossi**, Praneet Indela, Venkata Parwatala, and others, *Delivery Optimized Discovery in Behavioral User Segmentation under Budget Constraint*, ACM International Conference on Information and Knowledge Management (CIKM), 2023.

[C124] Charles Chen, Jane Hoffswell, Shunan Guo, **Ryan A. Rossi**, Yeuk-Yin Chan, Fan Du, Eunyee Koh, and Zhicheng Liu, *WhatsNext: Guidance-enriched Exploratory Data Analysis with Interactive, Low-code Notebooks*, IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC), 2023.

[C123] Chieh-Yang Huang, Ting-Yao Hsu, **Ryan A. Rossi**, Ani Nenkova, Sungchul Kim, Gromit Yeuk-Yin Chan, Eunyee Koh, C. Lee Giles, and Ting-Hao Kenneth Huang, *Summaries as Captions: Generating Figure Captions for Scientific Documents with Automated Text Summarization*, International Natural Language Generation Conference (INLG), 2023.

[C122] Yu Xia, Junda Wu, Tong Yu, Sungchul Kim, **Ryan A. Rossi**, and Shuai Li, *User-Regulation Deconfounded Conversational Recommender System with Bandit Feedback*, ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2023.

[C121] Mohammad Mehrabi and **Ryan A. Rossi**, *A Model-Free Closeness-of-Influence Test for Features in Supervised Learning*, International Conference on Machine Learning (ICML), 2023.

[C120] Camille Harris, **Ryan A. Rossi**, Sana Malik, Jane Hoffswell, Fan Du, Tak Yeon Lee, Eunyee Koh, and Handong Zhao, *SpotLight: Visual Insight Recommendation*, Companion Proceedings of the ACM Web Conference (WWW Companion), 2023.

[C119] Md Main Uddin Rony, Fan Du, **Ryan A. Rossi**, Jane Hoffswell, Niyati Chhaya, Iftikhar Burhanuddin, and Eunyee Koh, *Augmenting Visualizations with Predictive and Investigative Insights to Facilitate Decision Making*, Companion Proceedings of the ACM Web Conference (WWW Companion), 2023.

[C118] **Ryan A. Rossi**, Nesreen K. Ahmed, and Namyong Park, *On Graph Time-Series Representations for Temporal Networks*, Companion Proceedings of the ACM Web Conference (WWW Companion), 2023.

[C117] Ryan Aponte, **Ryan A. Rossi**, Shunan Guo, Jane Hoffswell, Nedim Lipka, Chang Xiao, Gromit Yeuk-Yin Chan, Eunyee Koh, and Fan Du, *A ML-based Approach for HTML-based Style Recommendation*, Companion Proceedings of the ACM Web Conference (WWW Companion), 2023.

[C116] Enyu Cai, **Ryan A. Rossi**, and Chang Xiao, *Improving Learning-Based Camera Pose Estimation for Image-Based Augmented Reality Applications*, CHI Conference on Human Factors in Computing Systems Extended Abstracts, 2023.

[C115] Jianna So, Nedim Lipka, Alexa F. Siu, **Ryan A. Rossi** and Franck Dernoncourt, *How Small Businesses Transform PDF Agreements into Action*, CSCW, 2023.

[C114] Mohammad Mehrabi and **Ryan A. Rossi**, *A model-free closeness-of-influence test for features in supervised learning*, International Conference on Machine Learning (ICML), 2023.

[C113] Sudhanshu Chanpuriya, **Ryan A. Rossi**, Sungchul Kim, Tong Yu, Jane Hoffswell, Nedim Lipka, Shunan Guo and Cameron N. Musco, *Direct Embedding of Temporal Network Edges via Time-Decayed Line Graphs*, International Conference on Learning Representations (ICLR), 2023.

[C112] Namyong Park, **Ryan A. Rossi**, Nesreen K. Ahmed, and Christos Faloutsos, *MetaGL: Evaluation-Free Selection of Graph Learning Models via Meta-Learning*, International Conference on Learning Representations (ICLR), 2023.

[C111] Arpit Narechania, Fan Du, Atanu R. Sinha, **Ryan A. Rossi**, Jane Hoffswell, Shunan Guo, Eunyee Koh, Shamkant B Navathe and Alex Endert, *DataPilot: Utilizing Quality and Usage Information for Subset Selection during Visual Data Preparation*, CHI, 2023.

[C110] Xinyu Shi, Ziqi Zhou, Jing Wen Zhang, Ali Neshati, Anjul Kumar Tyagi, **Ryan A. Rossi**, Shunan Guo, Fan Du and Jian Zhao, *De-Stijl: Facilitating Graphics Design with Interactive 2D Color Palette Recommendation*, CHI, 2023.

[C109] Gaurav Verma, **Ryan A. Rossi**, Christopher Tensmeyer, Jiuxiang Gu, and Ani Nenkova, *Learning the Visualness of Text Using Large Vision-Language Models*, AAAI CreativeAI Workshop, 2023.

[C108] Shashank Goel, Hritik Bansal, Sumit Bhatia, **Ryan A. Rossi**, Vishwa Vinay, and Aditya Grover, *CyCLIP: Cyclic Contrastive Language-Image Pretraining*, Advances in Neural Information Processing Systems (NeurIPS), 2022.

[C107] Enyu Cai, **Ryan A. Rossi** and Chang Xiao, *Improving Learning-based Camera Pose Estimation for Image-based Augmented Reality Applications*, CHI LBW, 2023.

[C106] Nesreen K. Ahmed, **Ryan A. Rossi**, Rong Zhou, John Boaz Lee, Xiangnan Kong, Theodore L. Willke, Hoda Eldardiry, *Role-based Graph Embeddings*, IEEE International Conference on Data Engineering (ICDE), 2023.

[C105] Sejoon Oh, Sungchul Kim, **Ryan A. Rossi**, and Srijan Kumar, *Influence-guided Data Augmentation for Neural Tensor Completion*, Proceedings of the 30th ACM International Conference on Information & Knowledge Management (CIKM), 2021.

[C104] **Ryan A. Rossi**, Anup Rao, Sungchul Kim, Eunyee Koh, Nesreen K. Ahmed and Gang Wu, *From Closing Triangles to Higher-Order Motif Closures for Better Unsupervised Online Link Prediction*, Proceedings of the 30th ACM International Conference on Information & Knowledge Management (CIKM), 2021.

[C103] Yue Zhao, **Ryan A. Rossi**, and Leman Akoglu, *Automatic Unsupervised Outlier Model Selection*, Advances in Neural Information Processing Systems (NeurIPS), 2021.

[C102] Saed Rezayi, Nedim Lipka, Vishwa Vinay, **Ryan A. Rossi**, Franck Dernoncourt, Tracy H. King and Sheng Li, *A Framework for Knowledge-Derived Query Suggestions*, IEEE BigData, 2021.

[C101] Aravind Reddy, **Ryan A. Rossi**, Zhao Song, Anup Rao, Tung Mai, Nedim Lipka, Gang Wu, Eunyee Koh, Nesreen Ahmed, *One-Pass Algorithms for MAP Inference of Nonsymmetric Determinantal Point Processes*, Proceedings of the 39th International Conference on Machine Learning (ICML), 2022.

[C100] Gaurav Verma, Vishwa Vinay, **Ryan A. Rossi**, Srijan Kumar, *Robustness of Fusion-based Multimodal Classifiers to Cross-Modal Content Dilutions*, EMNLP, 2022.

[C99] Namyong Park, **Ryan A. Rossi**, Eunyee Koh, Iftikhar Ahamath Burhanuddin, Sungchul Kim, Fan Du, Nesreen Ahmed, Christos Faloutsos, *CGC: Contrastive Graph Clustering for Community Detection and Tracking*, Proceedings of the ACM Web Conference (WWW), 2022.

[C98] Mustafa Abdallah, **Ryan A. Rossi**, Kanak Mahadik, Sungchul Kim, Handong Zhao, Saurabh Bagchi, *AutoForecast: Automatic Time-Series Forecasting Model Selection*, Proceedings of the 31st ACM International Conference on Information & Knowledge Management (CIKM), 2022.

[C97] Duc Hoang, Haotao Wang, Handong Zhao, **Ryan A. Rossi**, Sungchul Kim, Kanak Mahadik, Zhangyang Wang, *AutoMARS: Searching to Compress Multi-Modality Recommendation Systems*, Proceedings of the 31st ACM International Conference on Information & Knowledge Management (CIKM), 2022.

[C96] Chenhan Yuan, **Ryan A. Rossi**, Andrew Katz, Hoda Eldardiry, *Clustering-based Unsupervised Generative Relation Extraction*, IEEE BigData, 2022.

[C95] Chang Xiao, **Ryan A. Rossi**, Eunyee Koh, *iMarker: Instant and True-to-scale AR with Invisible Markers*, Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST), 2022.

[C94] Weixin Jiang, Gang Wu, Vishy Swaminathan, Stefano Petrangeli, Haoliang Wang, **Ryan A. Rossi**, Nedim Lipka, *Task-Oriented Near-Lossless Burst Compression*, IEEE International Symposium on Multimedia (ISM), 2022.

[C93] Hongjie Chen, **Ryan A. Rossi**, Kanak Mahadik, Hoda Eldardiry, *Context Integrated Relational Spatio-Temporal Resource Forecasting*, IEEE International Conference on Big Data (IEEE BigData), 2022.

[C92] Can Qin, Sungchul Kim, Handong Zhao, Tong Yu, **Ryan A. Rossi**, Yun Fu, *External Knowledge Infusion for Tabular Pre-training Models with Dual-adapters*, Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2022.

[C91] Sejoon Oh, Ankur Bhardwaj, Jongseok Han, Sungchul Kim, **Ryan A. Rossi**, Srijan Kumar, *Implicit Session Contexts for Next-Item Recommendations*, Proceedings of the 31st ACM International Conference on Information & Knowledge Management, 2022.

[C90] Swapneel Mehta, Somdeb Sarkhel, Xiang Chen, Saayan Mitra, Viswanathan Swaminathan, **Ryan A. Rossi**, Ali Aminian, Han Guo, Kshitiz Garg, *Open-Domain Trending Hashtag Recommendation for Videos*, IEEE International Symposium on Multimedia (ISM), 2022.

[C89] Abhraneel Sarma, Shunan Guo, Jane Hoffswell, **Ryan A. Rossi**, Fan Du, Eunyee Koh, Matthew Kay, *Evaluating the Use of Uncertainty Visualisations for Imputations of Data Missing At Random in Scatterplots*, IEEE VIS, 2022.

[C88] Bingjie Xu, Shunan Guo, Eunyee Koh, Jane Hoffswell, **Ryan A. Rossi**, Fan Du, *ARShopping: In-Store Shopping Decision Support Through Augmented Reality and Immersive Visualization*, IEEE VIS, 2022.

[C87] Beleicia Bullock, Shunan Guo, Eunyee Koh, **Ryan A. Rossi**, Fan Du, Jane Hoffswell, *Let's Get Personal: Exploring the Design of Personalized Visualizations*, IEEE VIS, 2022.

[C86] Fayokemi Ojo, **Ryan A. Rossi**, Jane Hoffswell, Shunan Guo, Fan Du, Sungchul Kim, Chang Xiao, Eunyee Koh, *VisGNN: Personalized Visualization Recommendation via Graph Neural Networks*, Proceedings of the ACM Web Conference (WWW), 2022.

[C85] Hyeok Kim, **Ryan A. Rossi**, Fan Du, Eunyee Koh, Shunan Guo, Jessica Hullman, Jane Hoffswell, *Cicero: A Declarative Grammar for Responsive Visualization*, ACM Human Factors in Computing Systems (CHI), 2022.

[C84] Zihao Zhou, Xingyi Yang, **Ryan A. Rossi**, Handong Zhao, Rose Yu, *Neural Point Process for Learning Spatiotemporal Event Dynamics*, Learning for Dynamics and Control Conference (L4DC), 2022.

[C83] Zhuohao Zhang, Sana Malik, Shunan Guo, Jane Hoffswell, **Ryan A. Rossi**, Fan Du, and Eunyee Koh, *Understanding Business Analysts' Needs for Data Report Authoring*, EuroVis Workshop on Visual Analytics (EuroVA), 2022.

[C82] Zhuohao Zhang, Sana Malik, Shunan Guo, Jane Hoffswell, **Ryan A. Rossi**, Fan Du, and Eunyee Koh, *CODAS: Integrating Business Analytics and Report Authoring*, EuroVis Workshop on Visual Analytics (EuroVA), 2022.

[C81] Zening Qu, Fan Du, **Ryan A. Rossi**, and Bill Howe, *Aunt Lily Can Say Her Visualizations: Directing Analysis, Design, and Storytelling in Natural Language*, NL VIZ: Workshop on Exploring Opportunities and Challenges for Natural Language Techniques to Support Visual Analysis, 2021.

[C80] Di Jin, Sungchul Kim, **Ryan A. Rossi**, and Danai Koutra, *On Generalizing Static Node Embedding to Dynamic Settings*, The Fifteenth International Conference on Web Search and Data Mining (WSDM), 2022.

[C79] Galen Weld, Peter West, Maria Glenski, David Arbour, **Ryan A. Rossi**, and Tim Althoff, *Adjusting for Confounders with Text: Challenges and an Empirical Evaluation Framework for Causal Inference*, AAAI Conference on Web and Social Media (ICWSM), 2022.

[C78] Hyeok Kim, **Ryan A. Rossi**, Abhraneel Sarma, Dominik Moritz, and Jessica Hullman, *An Automated Approach to Reasoning About Task-Oriented Insights in Responsive Visualization*, IEEE VIS, 2021.

[C77] Enayat Ullah, Tung Mai, Anup Rao, **Ryan A. Rossi**, and Raman Arora, *Machine unlearning via algorithmic stability*, Conference on Learning Theory (COLT), 2021.

[C76] Xin Qian, **Ryan A. Rossi**, Fan Du, Sungchul Kim, Eunyee Koh, Sana Malik, Tak Yeon Lee, Joel Chan, *Learning to Recommend Visualizations from Data*, ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2021.

[C75] Hongjie Chen, **Ryan A. Rossi**, Kanak Mahadik, Sungchul Kim, Hoda Eldardiry, *Graph Deep Factor Model for Cloud Utilization Forecasting*, ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2021.

[C74] Gromit Yeuk-Yin Chan, Tung Mai, Anup Rao, **Ryan A. Rossi**, Fan Du, Cláudio T. Silva, Juliana Freire, *Interactive Audience Expansion On Large Scale Online Visitor Data*, ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2021.

[C73] Yikun Xian, Handong Zhao, Tak Yeon Lee, Sungchul Kim, **Ryan A. Rossi**, Zuohui Fu, Gerard de Melo, S. Muthukrishnan, *EXACTA: Explainable Column Annotation*, ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2021.

[C72] Mohammad Mehrabi, Adel Javanmard, **Ryan A. Rossi**, Anup Rao and Tung Mai, *Fundamental Tradeoffs in Distributionally Adversarial Training*, International Conference on Machine Learning (ICML), 2021.

[C71] Mojtaba Sahraee-Ardakan, Tung Mai, Anup Rao, **Ryan A. Rossi**, Sundeep Rangan and Alyson Fletcher, *Asymptotics of Ridge Regression in Convolutional Models*, International Conference on Machine Learning (ICML), 2021.

[C70] Jun Yan, Mrigank Raman, Aaron Chan, Tianyu Zhang, **Ryan A. Rossi**, Handong Zhao, Sungchul Kim, Nedim Lipka and Xiang Ren, *Learning Contextualized Knowledge Structures for Commonsense Reasoning*, ACL, 2021.

[C69] Mrigank Raman, Hansen Wang, PeiFeng Wang, Siddhant Agarwal, Sungchul Kim, **Ryan A. Rossi**, Handong Zhao, Nedim Lipka, Xiang Ren, *Learning to Deceive Knowledge Graph Augmented Models via Targeted Perturbation*, International Conference on Learning Representations (ICLR), 2021, Vienna, Austria.

[C68] Jiong Zhu, **Ryan A. Rossi**, Anup Rao, Tung Mai, Nedim Lipka and Danai Koutra, *Graph Neural Networks with Heterophily*, AAAI, 2021.

[C67] Enayat Ullah, Tung Mai, Anup Rao, **Ryan A. Rossi**, and Raman Arora, *Machine unlearning via algorithmic stability*, Symposium on Foundations of Responsible Computing (FORC).

[C66] Hongchang Gao, Gang Wu, **Ryan A. Rossi**, *Provable Distributed Stochastic Gradient Descent with Delayed Updates*, SDM, 2021.

[C65] Xin Qian, Eunyee Koh, Fan Du, Sungchul Kim, Joel Chan, **Ryan A. Rossi**, Sana Malik and Tak Yeon Lee, *Generating Accurate Caption Units for Figure Captioning*, WWW, 2021.

[C64] Saed Rezayi, Handong Zhao, Sungchul Kim, **Ryan A. Rossi**, Nedim Lipka, Sheng Li, *Edge: Enriching Knowledge Graph Embeddings with External Text*, NAACL, 2021.

[C63] Jun Yan, Mrigank Raman, Tianyu Zhang, **Ryan A. Rossi**, Handong Zhao, Sungchul Kim, Nedim Lipka, Xiang Ren, *Learning Contextualized Knowledge Structures for Commonsense Reasoning*, NeurIPS KR2ML Workshop, 2021.

[C62] Nesreen Ahmed, Richard Alo, Catherine Amelink, Young Yun Baek, Aashish Chudhary, Kristy Collins, Albert Esterline, Edward Fox, Geoffrey Fox, Aric Hagberg, Ron Kenyon, Chris Kuhlman, Jure Leskovec, Dustin Machi, Madhav Marathe, Nataragan Meghanathan, Yasuo Miyasaki, Judy Qiu, Naren Ramakrishnan, S. S. Ravi, **Ryan A. Rossi**, Roc Sosic and Gregor von Laszewski, *net.science: A Cyberinfrastructure for Sustained Innovation in Network Science and Engineering*, Gateways, 2020.

[C61] Youngsuk Park, **Ryan A. Rossi**, Zheng Wen, Gang Wu and Handong Zhao, *Structured Policy Iteration for Linear Quadratic Regulator*, International Conference on Machine Learning (ICML), 2020, Vienna, Austria.

[C60] Aldo Carranza, **Ryan A. Rossi**, Anup Rao and Eunyee Koh, *Higher-order Clustering in Complex Heterogeneous Networks*, KDD, 2020, San Diego, California USA.

[C59] **Ryan A. Rossi**, Nesreen K. Ahmed, Eunyee Koh, Sungchul Kim, Anup Rao, Yasin Abbasi-Yadkori, *A Structural Graph Representation Learning Framework*, Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM), 2020, Houston, Texas USA.

[C58] Gromit Yeuk-Yin Chan, Fan Du, **Ryan A. Rossi**, Anup Rao, Eunyee Koh, Cláudio T. Silva, Juliana Freire, *Real-Time Clustering for Large Sparse Online Visitor Data*, Proceedings of The Web Conference (WWW), 2020, Taipei, Taiwan.

[C57] **Ryan A. Rossi**, Anup Rao, Tung Mai, Nesreen K Ahmed, *Fast and Accurate Estimation of Typed Graphlets*, Proceedings of The Web Conference (WWW), 2020, Taipei, Taiwan.

[C56] **Ryan A. Rossi**, Anup Rao, Sungchul Kim, Eunyee Koh, Nesreen K Ahmed, *From Closing Triangles to Closing Higher-Order Motifs*, Proceedings of The Web Conference (WWW), 2020, Taipei, Taiwan.

[C55] **Ryan A. Rossi**, Nesreen K. Ahmed, Eunyee Koh, Sungchul Kim, *Fast Hierarchical Graph Clustering in Linear-Time*, Proceedings of The Web Conference (WWW), 2020, Taipei, Taiwan.

[C54] John Boaz Lee, **Ryan A. Rossi**, Xiangnan Kong, Sungchul Kim, Eunyee Koh and Anup Rao, *Graph Convolutional Networks with Motif-based Attention*, Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM), 2019, Beijing, China.

[C53] Alireza Farhadi, MohammadTaghi Hajiaghayi, Tung Mai, Anup Rao, and **Ryan A. Rossi**, *Approximate Maximum Matching in Random Streams*, Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2020.

[C52] **Ryan A. Rossi**, Somdeb Sarkhel, Nesreen K. Ahmed, *Inferring Individual Level Causal Models from Graph-based Relational Time Series*, AAAI StarAI, 2020, New York, New York USA.

[C51] Tung Mai, Anup Rao, Matt Kapilevich, **Ryan A. Rossi**, Yasin Abbasi-Yadkori, and Ritwik Sinha, *On Densification for Minwise Hashing*, UAI, 2019, Tel Aviv, Israel.

[C50] Di Jin, **Ryan A. Rossi**, Eunyee Koh, Sungchul Kim, Anup Rao, and Danai Koutra, *Latent Network Summarization: Bridging Network Embedding and Summarization*, KDD, 2019, Anchorage, Alaska USA.

[C49] Di Jin, Mark Heimann, **Ryan A. Rossi**, and Danai Koutra, *node2bits: Compact Time- and Attribute-aware Node Representations for User Stitching*, ECML/PKDD, 2019, Würzburg, Germany.

[C48] Nesreen K. Ahmed, **Ryan A. Rossi**, John Boaz Lee, Theodore L. Willke, Rong Zhou, Xiangnan Kong, Hoda Eldardiry, *role2vec: Role-based Network Embeddings*, DLG KDD, 2019, Anchorage, Alaska USA.

[C47] Charles Chen, Ruiyi Zhang, Eunyee Koh, Sungchul Kim, Scott Cohen, **Ryan A. Rossi**, *FigCaps: Relational Reasoning in Figure Captioning*, IEEE Winter Conference on Applications of Computer Vision (WACV), 2020.

[C46] **Ryan A. Rossi**, Nesreen K. Ahmed, Aldo Carranza, David Arbour, Anup Rao, Sungchul Kim and Eunyee Koh, *Heterogeneous Graphlets*, MLG KDD, 2019, Anchorage, Alaska USA.

[C45] Charles Chen, Eunyee Koh, Sungchul Kim, **Ryan A. Rossi**, Scott Cohen, *Neural caption generation over figures*, UbiComp/ISWC, 2019.

[C44] Di Jin, **Ryan A. Rossi**, Danai Koutra, Eunyee Koh, Sungchul Kim and Anup Rao, *Latent Network Summarization*, MLG KDD, 2019, Anchorage, Alaska USA.

[C43] Youngsuk Park, Kanak Mahadik, **Ryan A. Rossi**, Gang Wu, Handong Zhao, *Linear Quadratic Regulator for Resource-Efficient Cloud Services*, ACM Symposium on Cloud Computing (SoCC'19), 2019.

[C42] Donghyun Kim, Sungchul Kim, Handong Zhao, **Ryan A. Rossi**, Sheng Li and Eunyee Koh, *Domain Switch-Aware Holistic Recurrent Neural Network for Modeling Multi-Domain User Behavior*, WSDM, 2019, Melbourne, Australia.

[C41] Giang Hoang Nguyen, John Boaz Lee, **Ryan A. Rossi**, Nesreen K. Ahmed, Eunyee Koh, and Sungchul Kim, *Dynamic Network Embeddings: From Random Walks to Temporal Random Walks*, BigData, 2018, Seattle, Washington.

[C40] John Boaz Lee, **Ryan A. Rossi**, Xiangnan Kong, *Graph Classification using Structural Attention*, KDD, 2018, London, UK.

[C39] **Ryan A. Rossi**, Rong Zhou, Nesreen K. Ahmed, and Hoda Eldardiry, *Relational Similarity Machines (RSM): A Similarity-based Learning Framework for Graphs*, BigData, 2018, Seattle, Washington.

[C38] Charles Chen, Sungchul Kim, Hung Bui, **Ryan A. Rossi**, Branislav Kveton, Eunyee Koh, and Razvan Bunescu, *Predictive Analysis by Leveraging Temporal User Behavior and User Embeddings*, CIKM, 2018, Turin, Italy.

[C37] **Ryan A. Rossi**, Nesreen K. Ahmed, and Eunyee Koh, *Higher-order Network Representation Learning*, WWW, 2018, Lyon, France.

[C36] **Ryan A. Rossi**, Rong Zhou, and Nesreen K. Ahmed, *Deep Inductive Network Representation Learning*, 3rd International Workshop on Learning Representations for Big Networks (WWW BigNet), 2018, Lyon, France.

[C35] Giang Hoang Nguyen, John Boaz Lee, **Ryan A. Rossi**, Nesreen K. Ahmed, Eunyee Koh, and Sungchul Kim, *Continuous-Time Dynamic Network Embeddings*, 3rd International Workshop on Learning Representations for Big Networks (WWW BigNet), 2018, Lyon, France.

[C34] **Ryan A. Rossi**, Nesreen K. Ahmed, Hoda Eldardiry, and Rong Zhou, *Similarity-based Multi-label Learning*, International Joint Conference on Neural Networks (IJCNN), 2018, Rio de Janeiro, Brazil.

[C33] Nesreen K. Ahmed, **Ryan A. Rossi**, Rong Zhou, John Boaz Lee, Xiangnan Kong, Theodore L. Willke, Hoda Eldardiry, *Learning Role-based Graph Embeddings*, StarAI IJCAI, 2018, Stockholm, Sweden.

[C32] Nesreen K. Ahmed, **Ryan A. Rossi**, Rong Zhou, John Boaz Lee, Xiangnan Kong, Theodore L. Willke, Hoda Eldardiry, *Representation Learning in Large Attributed Graphs*, WiML NIPS, 2017, Long Beach, CA USA.

[C31] James P. Canning, Emma E. Ingram, Sammantha Nowak-Wolff, Adriana M. Ortiz, Nesreen K. Ahmed, **Ryan A. Rossi**, Karl R. B. Schmitt, Sucheta Soundarajan, *Network Classification and Categorization*, International Conference on Complex Networks (CompleNet), 2018, Boston, MA USA.

[C30] Nesreen K. Ahmed, Nick Duffield, Theodore L. Willke, **Ryan A. Rossi**, *On Sampling from Massive Graph Streams*, VLDB, 1430–1441, 2017, Munich, Germany.

[C29] Nesreen K. Ahmed, **Ryan A. Rossi**, Theodore L. Willke, and Rong Zhou, *Edge Role Discovery via Higher-order Structures*, PAKDD, 2017, Jeju, Korea.

[C28] Nesreen K. Ahmed, **Ryan A. Rossi**, Theodore L. Willke, and Rong Zhou, *A Higher-order Latent Space Network Model*, Proceedings of the AAAI PAIR (Plan, Activity, and Intent Recognition) Workshop, 2017, San Francisco, CA USA.

[C27] **Ryan A. Rossi** and Rong Zhou, *Leveraging Multiple GPUs and CPUs for Graphlet Counting in Large Networks*, ACM International Conference on Information and Knowledge Management (CIKM), 2016, Indianapolis, IN USA.

[C26] Nesreen K. Ahmed, **Ryan A. Rossi**, Theodore L. Willke, and Rong Zhou, *Estimation of Local Subgraph Counts*, Proceedings of the IEEE International Conference on Big Data, 2016, Washington D.C., USA.

[C25] **Ryan A. Rossi**, Rong Zhou, and Nesreen K. Ahmed, *Relational Similarity Machines*, Proceedings of the 12th International Workshop on Mining and Learning with Graphs (KDD MLG), Pages 8, 2016, San Francisco, CA USA.

[C24] Nesreen Ahmed, Theodore L. Willke, and **Ryan A. Rossi**, *Exact and Estimation of Local Edge-centric Graphlet Counts*, KDD BigMine, 2016, San Francisco, CA USA.

[C23] **Ryan A. Rossi** and Nesreen K. Ahmed, *An Interactive Data Repository with Visual Analytics*, SIGKDD Explorations, 2016.

[C22] **Ryan Rossi** and Rong Zhou, *Towards Interactive Relational Learning*, Twenty-Ninth Conference on Artificial Intelligence (AAAI), 2016, Phoenix, Arizona USA.

[C21] Nesreen K. Ahmed, Jennifer Neville, **Ryan A. Rossi**, Nick Duffield, *Efficient Graphlet Counting for Large Networks*, IEEE International Conference on Data Mining (ICDM), pages 10, 2015. **Invited to KAIS Journal Special Issue (ICDM Best papers)*, Atlantic City, NJ USA.

[C20] **Ryan A. Rossi** and Rong Zhou, *Scalable Relational Learning for Large Heterogeneous Networks*, IEEE International Conference on Data Science and Advanced Analytics (DSAA), 10 pages, 2015, Paris, France.

[C19] Nesreen K. Ahmed and **Ryan A. Rossi**, *Interactive Visual Graph Analytics on the Web*, Proceedings of the 9th International AAAI Conference on Web and Social Media, pages 566–569, 2015, Oxford, UK.

[C18] **Ryan Rossi** and Nesreen K. Ahmed, *The Network Data Repository with Interactive Graph Analytics and Visualization*, Twenty-Ninth Conference on Artificial Intelligence (AAAI) DT, pages 4292–4293, 2015, Austin, Texas USA.

[C17] **Ryan Rossi**, *Fast Triangle Core Decomposition for Mining Large Graphs*, Advances in Knowledge Discovery and Data Mining, 310–322, 2014, Tainan, Taiwan.

[C16] **Ryan A. Rossi**, David F. Gleich, Assefaw H. Gebremedhin, Md. Mostofa Ali Patwary, *A Parallel Maximum Clique Algorithm for Large Graphs*, Proceedings of the 23rd ACM International Conference Companion on World Wide Web, 2014, Seoul, Korea.

[C15] **Ryan Rossi**, Sonia Fahmy, and Nilothpal Talukder, *A Multi-Level Approach for Evaluating Internet Topology Generators*, Networking, 1–9, 2013, New York, USA.

[C14] **Ryan Rossi**, Brian Gallagher, Jennifer Neville, and Keith Henderson, *Modeling Dynamic Behavior in Large Evolving Graphs*, In Proceedings of the Sixth ACM International Conference on Web Search and Data Mining (WSDM), pp. 667–676, 2013, Rome, Italy.

[C13] **Ryan A. Rossi**, David F. Gleich, Assefaw H. Gebremedhin, *Triangle Core Decomposition and Maximum Cliques*, SIAM Workshop on Network Science, 1–2, 2013.

[C12] **Ryan Rossi** and David Gleich, *Dynamic PageRank using Evolving Teleportation*, Algorithms and Models for the Web Graph, volume 7323 of Lecture Notes in Computer Science, pages 126–137. Springer, 2012, Nova Scotia, Canada.

[C11] **Ryan Rossi**, Brian Gallagher, Jennifer Neville, and Keith Henderson, *Role-Dynamics: Fast Mining of Large Dynamic Networks*, Proceedings of the 21st ACM International Conference Companion on World Wide Web (WWW), pages 997–1006, 2012, Lyon, France.

[C10] **Ryan Rossi** and Jennifer Neville, *Time-Evolving Relational Classification and Ensemble Methods*, In Proceedings of the Pacific-Asia International Conference on Knowledge Discovery and Data Mining (PAKDD), pages 1–13, 2012, Kuala Lumpur, Malaysia.

[C9] **Ryan Rossi** and Jennifer Neville, *Modeling the Evolution of Discussion Topics and Communication to Improve Relational Classification*, In Proceedings of the 1st SOMA Workshop, 16th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 89–97, 2010, Washington D.C., USA.

[C8] Khawaja S. Shams, Mark W. Powell, Tom M. Crockett, Jeffrey S. Norris, **Ryan Rossi**, Tom Soderstrom, *Polyphony: A Workflow Orchestration Framework for Cloud Computing*, 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing (CCGrid), pages 606–611, 2010., Melbourne, Australia.

[C7] **Ryan A. Rossi**, *Latent Semantic Analysis of the Languages of Life, Communications in Computer and Information Science*, ISICA, CCIS 51, 128–137, 2009.

[C6] Mark W. Powell, **Ryan A. Rossi**, and Khawaja S. Shams, *A Scalable Image Processing Framework for Gigapixel Mars and Other Celestial Body Images*, IEEE Aerospace, 1–11, 2009, Big Sky, Montana.

[C5] John Stamey, **Ryan A. Rossi**, *Automatically Identifying Relations in Privacy Policies*, Proceedings of the 27th ACM International Conference on Design of Communication, 233–238, 2009, Bloomington, Indiana.

[C4] Jean-Louis Lassez, **Ryan A. Rossi**, Kumar Jeev, *Ranking Links on the Web: Search and Surf Engines*, New Frontiers in Applied Artificial Intelligence (IEA/AIE), volume 5027 of Lecture Notes of Artificial Intelligence, 199–208 (2008), Wroclaw, Poland.

[C3] Jean-Louis Lassez, **Ryan A. Rossi**, Stephen Sheel, Srinivas Mukkamala, *Signature Based Intrusion Detection System using Latent Semantic Analysis*, IEEE World Congress on Computational Intelligence, International Joint Conference of Neural Networks, IJCNN, 1068–1074 (2008), Hong Kong.

[C2] John Stamey, Jean-Louis Lassez, **Ryan Rossi**, Daniel Boorn, *Client-Side Dynamic Metadata in Web 2.0*, Proceedings of the 25th ACM International Conference on Design of Communication, 155–161 (2007), El Paso, Texas.

[C1] Jean-Louis Lassez, **Ryan A. Rossi**, Axel E. Bernal, *Crick's Hypothesis Revisited: The Existence of a Universal Coding Frame, IEEE International Conference on Bioinformatics and Life Science Computing*, AINA/BLSC, 745–751 (2007), Niagara Falls, Canada.

Patents

[P108] **Ryan A. Rossi** et al., *Graph-Based Context Retrieval for Large Language Models, Adobe Research*, 2025.

[P107] **Ryan A. Rossi** et al., *Generating Coherent Extractive Summaries by Training LLMs on Feedback Annotations From Initially Generated Summaries, Adobe Research*, 2025.

[P106] **Ryan A. Rossi** et al., *Knowledge-Aware Query Expansion With Large Language Models for Textual and Relational Retrieval, Adobe Research*, 2025.

[P105] **Ryan A. Rossi** et al., *Systems and Methods for Debiasing Language Models, Adobe Research*, 2025.

[P104] **Ryan A. Rossi** et al., *Document-Related Query Answering Using a Language Generation Model, Adobe Research*, 2025.

[P103] **Ryan A. Rossi** et al., *Retrieval-Augmented Question Answering, Adobe Research*, 2025.

[P102] **Ryan A. Rossi** et al., *Multi-Agent Visual Perception, Adobe Research*, 2025.

[P101] **Ryan A. Rossi** et al., *Boosting LLM Code Generation via Preference-Guided Refinement, Adobe Research*, 2025.

[P100] **Ryan A. Rossi** et al., *Utilizing Multimodal Large Language Models to Identify Grounding Text Content Within Digital Images, Adobe Research*, 2025.

[P99] **Ryan A. Rossi** et al., *Suggested Question Generation Using Machine-Learning, Adobe Research*, 2025.

[P98] **Ryan A. Rossi** et al., *Systems and Methods for Generating Attribution Information, Adobe Research*, 2025.

[P97] **Ryan A. Rossi** et al., *Text Generation Using Selective Attention, Adobe Research*, 2025.

[P96] **Ryan A. Rossi** et al., *Task Performance Using a Machine Learning Agent-Generated Action, Adobe Research*, 2025.

[P95] **Ryan A. Rossi** et al., *Dialogue-Oriented Document Navigation, Adobe Research*, 2025.

[P94] **Ryan A. Rossi** et al., *Iterative Retrieval and Query of a Multi-Page Multi-Modal Document, Adobe Research*, 2025.

[P93] **Ryan A. Rossi** et al., *Generating Comprehensive Analytics for Evaluating AI Agents, Adobe Research*, 2025.

[P92] **Ryan A. Rossi** et al., *Expert System Selection Using Limited Feedback, Adobe Research*, 2025.

[P91] **Ryan A. Rossi** et al., *Reasoning-Based Personalized Generation for Users With Sparse Data, Adobe Research*, 2025.

[P90] **Ryan A. Rossi** et al., *Lizard: An Efficient Linearization Framework for Infinite-Context Large Language Models, Adobe Research*, 2025.

[P89] **Ryan A. Rossi** et al., *Steering Generation of Prompt Response Behavior via Adjusting Activation Statuses of Feed-Forward Networks, Adobe Research*, 2025.

[P88] **Ryan A. Rossi** et al., *Data Question Answering With Auxiliary Recommendations, Adobe Research*, 2025.

[P87] **Ryan A. Rossi** et al., *Using Reinforcement Learning to Recommend Data Visualizations, Adobe Research*, 2025.

[P86] Sunav Choudhary, Atanu R. Sinha, Harshita Chopra, **Ryan A. Rossi**, Veda Pranav Parwatala, Paavan Kumar Indela, Srinjayee Paul, Shunan Guo, *Segment Discovery and Channel Delivery, Adobe Research*, Patent application filed 12/18/2023, App.#18/543,666, 2023.

[P85] Haoliang Wang, Xiang Chen, Tong Yu, Sungchul Kim, **Ryan A. Rossi**, Junda Wu, Anup Bandigadi Rao, *Chain-of-Thought Machine-Learning Model Debiasing, Adobe Research*, Patent application filed 5/24/2024, Pub.#20250363409, 2024.

[P84] Younghun Lee, Xiang Chen, Tong Yu, Sungchul Kim, **Ryan A. Rossi**, *Fine-Tuning a Language Model to Reduce Representations of Structured Data, Adobe Research*, Patent application filed 1/8/2024, Pub.#20/250,225,373, 2024.

[P83] **Ryan A. Rossi**, Aravind Reddy Talla, Zhao Song, Anup Rao, Tung Mai, Nedim Lipka, Gang Wu, Eunyee Koh, *Online Inference and Learning for Nonsymmetric Determinantal Point Processes, Adobe Research*, Patent application filed 5/12/2022, US Patent No. 12,288,237, 2022.

[P82] **Ryan A. Rossi**, Xiang Chen, Tong Yu, Sungchul Kim, Subrata Mitra, Shunan Guo, Ryan Alexander Aponte, Nedim Lipka, Franck Dernoncourt, *Generation and Utilization of Unified Fine-Tuning Datasets, Adobe Research*, Patent application filed 2/1/2024, Pub.#20250252268, 2024.

[P81] William Brandon George, Wei Zhang, Tyler Rasmussen, Tung Mai, Tong Yu, Sungchul Kim, Shunan Guo, Samuel Nephi Grigg, Said Kobeissi, **Ryan A. Rossi**, Ritwik Sinha, Eunyee Koh, Prithvi Bhutani, Jordan Henson Walker, Abhisek Trivedi, *Learning-Based Method for Generating Visualizations via Natural Language, Adobe Research*, Patent application filed 8/31/2023, US Patent No. 12,265,557, 2023.

[P80] Chen Chen, Jane Elizabeth Hoffswell, Shunan Guo, Fan Du, Nathan Carl Ross, **Ryan A. Rossi**, Yeuk Yin Chan, Eunyee Koh, *Generating Data Insights, Adobe Research*, Patent application filed 6/5/2023, US Patent No. 12,216,677, 2023.

[P79] **Ryan A. Rossi**, Ryan Aponte, Shunan Guo, Jane Elizabeth Hoffswell, Nedim Lipka, Chang Xiao, Yeuk-yin Chan, Eunyee Koh, *Hypergraph Representation Learning, Adobe Research*, Patent application filed 7/25/2023, Pub.#20250036936, 2023.

[P78] **Ryan A. Rossi**, Ryan Aponte, Shunan Guo, Nedim Lipka, Jane Hoffswell, Chang Xiao, Eunyee Koh, Yeuk-yin Chan, *Performing Machine Learning Techniques for Hypertext Markup Language-Based Style Recommendations, Adobe Research*, Patent application filed 7/25/2023, Pub.#20250036858, 2023.

[P77] **Ryan A. Rossi**, Iftikhar Ahamath Burhanuddin, Gautam Choudhary, Fan Du, Eunyee Koh, *Generating Node Embeddings for Multiple Roles, Adobe Research*, Patent application filed 6/22/2022, US Patent No. 12,475,376, 2022.

[P76] Md Main Uddin Rony, Fan Du, Iftikhar Ahamath Burhanuddin, **Ryan A. Rossi**, Niyati Himanshu Chhaya, Eunyee Koh, *Facilitating Generation and Presentation of Advanced Insights, Adobe Research*, Patent application filed 10/11/2023, US Patent No. 12,182,493, 2023.

[P75] Namyong Park, **Ryan A. Rossi**, Eunyee Koh, Iftikhar Ahamath Burhanuddin, Sungchul Kim, Fan Du, *Systems and Methods for Contrastive Graph Learning, Adobe Research*, Patent application filed 11/3/2022, App.#18/052,463, 2022.

[P74] Can Qin, Sungchul Kim, Handong Zhao, Tong Yu, **Ryan A. Rossi**, *Tabular Data Machine-Learning Models, Adobe Research*, Patent application filed 11/3/2022, App.#17/979,843, 2022.

[P73] Sudhanshu Chanpuriya, **Ryan A. Rossi**, Nedim Lipka, Anup Rao, Tung Mai, Zhao Song, *Generative Graph Modeling Framework, Adobe Research*, Patent application filed 10/31/2022, App.#18/051,364, 2022.

[P72] **Ryan A. Rossi**, Kanak Mahadik, Mustafa Abdallah, Sungchul Kim, Handong Zhao, *Automatic Forecasting Using Meta-Learning, Adobe Research*, Patent application filed 10/28/2022, App.#18/050,607, 2022.

[P71] Nathan Ng, Tung Mai, Thomas Greger, Kelly Nicholes, Antonio Cuevas, Saayan Mitra, Somdeb Sarkhel, Anup Rao, **Ryan A. Rossi**, Viswanathan Swaminathan, Shivakumar Vaithyanathan, *Dynamic User Profile Projection, Adobe Research*, Patent application filed 10/24/2022, App.#18/049,069, 2022.

[P70] **Ryan A. Rossi**, *Fast Graph Model Selection via Meta-Learning, Adobe Research*, Patent application filed 09/28/2022, App.#17/936,099, 2022.

[P68] **Ryan A. Rossi**, *Feature-Based Network Embedding, Adobe Research*, Patent application filed 09/2/2022, App.#17/902,378, 2022.

[P67] Swapneel Mehta, Somdeb Sarkhel, Xiang Chen, **Ryan A. Rossi**, Saayan Mitra, Vishy Swaminathan, Ali Aminian, Han Guo, Kshitiz Garg, *Open-Domain Trending Hashtag Recommendations, Adobe Research*, Patent application filed 07/29/2022, App.#17/877,469, 2022.

[P65] **Ryan A. Rossi**, Tung Mai, Anup Rao, *System And Methods for Streaming String Similarity and Pattern Matching, Adobe Research*, Patent application filed 06/8/2022, App.#17/835,791, 2022.

[P64] Chang Xiao, **Ryan A. Rossi**, Eunyee Koh, *System And Methods for Providing Invisible Augmented Reality Markers, Adobe Research*, Patent application filed 05/25/2022, App.#17/664,972, 2022.

[P63] Gang Wu, Yang Li, Stefano Petrangeli, Viswanathan Swaminathan, Haoliang Wang, **Ryan A. Rossi**, Zhao Song, *Actively Learned Context Modeling for Image Compression, Adobe Research*, Patent application filed 05/20/2022, App.#17/749,846, 2022.

[P61] Fan Du, Yeuk-Yin Chan, Anup Rao, **Ryan A. Rossi**, Eunyee Koh, Charles Menguy, Margarita Savova, *Segmenting Users with Sparse Data Utilizing Hash Partitions, Adobe Research*, Patent application filed 4/22/2022, App.#17/660,328, 2022.

[P60] Shunan Guo, **Ryan A. Rossi**, Jane Hoffswell, Fan Du, Eunyee Koh, Bingjie Xu, *Augmented Reality Systems for Comparing Physical Objects, Adobe Research*, Patent application filed 04/20/2022, App.#17/724,686, 2022.

[P59] Nathan Ng, Tung Mai, Thomas Greger, Kelly Nicholes, Antonio Cuevas, Saayan Mitra, Somdeb Sarkhel, Anup Rao, **Ryan A. Rossi**, Viswanathan Swaminathan, Shivakumar Vaithyanathan, *Dynamic User Profile Management, Adobe Research*, Patent application filed 10/24/2022, App.#18/049,083, 2022.

[P58] **Ryan A. Rossi**, Fayokemi Ojo, Jane Hoffswell, Shunan Guo, Sungchul Kim, Fan Du, Chang Xiao, Eunyee Koh, *Utilizing A Graph Neural Network to Generate Visualization and Attribute Recommendations, Adobe Research*, Patent application filed 03/15/2022, App.#17/654,933, 2022.

[P57] Arpit Ajay Narechania, Fan Du, Atanu Sinha, **Ryan A. Rossi**, Jane Hoffswell, Shunan Guo, Eunyee Koh, John Anderson, Sonali Surange, Saurabh Mahapatra, Vasanthi Holtcamp, *Dashboard for Monitoring Current and Historical Consumption and Quality Metrics for Attributes and Records of a Dataset, Adobe Research*, Patent application filed 03/14/2022, App.#17/693,811, 2022.

[P56] Arpit Ajay Narechania, Fan Du, Atanu Sinha, **Ryan A. Rossi**, Jane Hoffswell, Shunan Guo, Eunyee Koh, John Anderson, Sonali Surange, Saurabh Mahapatra, Vasanthi Holtcamp, *Data Selection Based on Consumption and Quality Metrics for Attributes and Records of a Dataset, Adobe Research*, Patent application filed 03/14/2022, App.#17/693,799, 2022.

[P55] Arpit Ajay Narechania, Fan Du, Atanu Sinha, **Ryan A. Rossi**, Jane Hoffswell, Shunan Guo, Eunyee Koh, John Anderson, Sonali Surange, Saurabh Mahapatra, Vasanthi Holtcamp, *Interactive Tree Representing Attribute Quality or Consumption Metrics for Data Ingestion and Other Applications, Adobe Research*, Patent application filed 03/14/2022, App.#17/693,778, 2022.

[P54] Hyeok Kim, Jane Hoffswell, **Ryan A. Rossi**, Fan Du, Eunyee Koh, Shunan Guo, *Design Recommender for Responsive Visualization Transformations, Adobe Research*, Patent application filed 02/23/2022, App.#17/678,346, 2022.

[P53] Gromit Yeuk-Yin Chan, Tung Mai, Anup Rao, **Ryan A. Rossi**, Fan Du, Moumita Sinha, Charles Menguy, Margarita Savova, Matt Kapilevich, *Trait Expansion Techniques in Binary Matrix Datasets, Adobe Research*, Patent application filed 02/22/2022, App.#17/677,323, 2022.

[P52] Sungchul Kim, Sejoon Oh, **Ryan A. Rossi**, *Enhancing Neural-Based Prediction of Multi-Dimensional Data Via Influence and Data Augmentation, Adobe Research*, Patent application filed 02/03/2022, App.#17/592,186, 2022.

[P51] Anup Rao, **Ryan A. Rossi**, Tung Mai, Enayat Ullah, *Machine Unlearning and Retraining of a Machine Learning Model based on a Modified Training Dataset, Adobe Research*, Patent application filed 10/18/2021, App.#17/451,260, 2022.

[P50] Tung Mai, Saayan Mitra, **Ryan A. Rossi**, Gaurav Gupta, Anup Rao, Xiang Chen, *Sparse Embedding Index for Search, Adobe Research*, Patent application filed 11/15/2021, App.#17/527,001, 2021.

[P49] Fan Du, **Ryan A. Rossi**, Eunyee Koh, Sungchul Kim, Handong Zhao, Keshav Vadrevu, Saurabh Mahapatra, Vasanthi Holtcamp, *An Automatic, Personalized, And Explainable Approach for Measuring, Monitoring, And Improving Data Efficacy, Adobe Research*, Patent application filed 10/28/2021, App.#17/513,571, 2021.

[P48] Sana Lee, Zhuohao Zhang, Zhicheng Liu, Tak Yeon Lee, Shunan Guo, **Ryan A. Rossi**, Fan Du, Eunyee Koh, *Systems for Generating Interactive Reports, Adobe Research*, Patent application filed 09/14/2021, App.#17/474,188, 2021.

[P47] Fan Du, Zening Qu, **Ryan A. Rossi**, Sungchul Kim, Sumit Shekhar, Eunyee Koh, Tak Yeon Lee, Sana Lee, Saurabh Mahapatra, Vasanthi Holtcamp, Nikhil Belsare, Andrew Thomson, *Natural Language Interface and Mixed-initiative Recommendation Model for Creating Data Stories, Adobe Research*, Patent application filed 06/3/2021, App.#17/YYY,YYY, 2021.

[P46] Eunyee Koh, Tak Yeon Lee, Andrew Thomson, Vasanthi Holtcamp, **Ryan A. Rossi**, Fan Du, Caroline Kim, Tong Yu, Shunan Guo, Nedim Lipka, and Shriram Revankar, *Content Go: Helping Marketers Create Better Email Contents with Smart Suggestions, Adobe Research*, Patent application filed 07/23/2021, App. #17/YYY,YYY, 2021.

[P45] Tung Mai, Kirankumar Shiragur, Anup Rao, **Ryan A. Rossi**, Georgios Theocharous, Michele Saad, *Item Transfer Systems, Adobe Research*, Patent application filed 07/15/2021, App. #17/YYY,YYY, 2021.

[P44] **Ryan A. Rossi**, Jiong Zhu, Tung Mai, Anup Rao, Nedim Lipka, Viswanathan Swaminathan, *Graph Neural Networks For Datasets With Heterophily, Adobe Research*, Patent application filed 3/23/2021, App. #17/210,157, 2021.

[P43] **Ryan A. Rossi**, Xin Qian, Eunyee Koh, Fan Du, Tak Yeon Lee, Sana Lee, Sungchul Kim, *Machine Learning Techniques for Generating Visualization Recommendations, Adobe Research*, Patent application filed 3/22/2021, App. #17/207,959, 2021.

[P42] Tak Yeon Lee, Qisheng Li, Eunyee Koh, Fan Du, **Ryan A. Rossi**, Sana Lee, *Systems for Suggesting Content Components, Adobe Research*, Patent application filed 2/23/2021, App. #17/183,055, 2021. US Patent No. 11,157,680. Awarded 10/26/2021.

[P41] Nedim Lipka, Seyedsaed Rezayidemne, Vishwa Vinay, **Ryan A. Rossi**, Franck Dernoncourt, Tracy King, *Knowledge-derived Search Suggestion, Adobe Research*, Patent application filed 2/8/2021, App. #17/170,520, 2021.

[P40] Sana Lee, **Ryan A. Rossi**, Camille Harris, Zening Qu, Fan Du, Eunyee Koh, Tak Yeon Lee, Sungchul Kim, Handong Zhao, Sumit Shekhar, *Configuration of User Interface for Intuitive Selection of Insight Visualizations, Adobe Research*, Patent application filed 1/29/2021, App. #17/161,770, 2021.

[P39] Eunyee Koh, Shenyu Xu, **Ryan A. Rossi**, Tak Yeon Lee, Fan Du, Sana Lee, *Generating Visual Data Stories, Adobe Research*, Patent application filed 1/28/2021, App. #17/161,406, 2021.

[P38] **Ryan A. Rossi**, *Selection of Outlier-detection Programs Specific to Dataset Meta-features, Adobe Research*, Patent application filed 1/15/2021, App. #17/150,890, 2021.

[P37] Sungchul Kim, Di Jin, **Ryan A. Rossi**, Eunyee Koh, *Temporal-Based Network Embedding and Prediction, Adobe Research*, Patent application filed 11/11/2020, App. #17/095,070, 2020.

[P36] Eunyee Koh, Xin Qian, Tak Yeon Lee, Sana Lee, **Ryan A. Rossi**, Fan Du, Duy-Trung Dinh, *Automated Caption Generation from a Dataset, Adobe Research*, Patent application filed 11/10/2020, App. #17/094,435, 2020.

[P35] **Ryan A. Rossi**, Vasanthi Holtcamp, Nathan Ross, John Anderson, Eunyee Koh, Sungchul Kim, Fan Du, Sana Lee, Tak Yeon Lee, *Personalized Visualization Recommendation System, Adobe Research*, Patent application filed 11/6/2020, App. #17/091,941, 2020.

[P34] **Ryan A. Rossi**, Hongjie Chen, Kanak Mahadik, Sungchul Kim, *Systems for Forecasting Computing Metric Values, Adobe Research*, Patent application filed 11/4/2020, App. #17/089,157, 2020.

[P33] Manoj A. Kilaru, Vishwa Vinay, Vidit Jain, Shaurya Goel, **Ryan A. Rossi**, Pratyush Garg, Nedim Lipka, Harkanwar Singh, *Generating Occurrence Contexts for Objects in Digital Content Collections, Adobe Research*, Patent application filed 10/26/2020, App. #17/079,945, 2020.

[P32] **Ryan A. Rossi**, Vasanthi Holtcamp, Nathan Ross, John Anderson, Eunyee Koh, Fan Du, Sana Lee, Tak Yeon Lee, *Graph-based Configuration of User Interface for Selection of Features in Visualization Applications, Adobe Research*, Patent application filed 9/9/2020, App. #17/015,495, 2020. US Patent No. 11,288,541. Awarded 03/22/2022.

[P31] **Ryan A. Rossi**, Anup Rao, Tung Mai, *System and Methods for Estimating Typed Graphlets in Large Data, Adobe Research*, Patent application filed 8/31/2020, App. #17/008,339, 2020. US Patent No. 11,343,325. Awarded 5/24/2022.

[P30] Gang Wu, Hongchang Gao, **Ryan A. Rossi**, Viswanathan Swaminathan, *Deep Relational Factorization Machine Techniques for Content Usage Prediction via Multiple Interaction Types, Adobe Research*, Patent application filed 7/27/2020, App. #16/939,661, 2020.

[P29] Kanak Mahadik, **Ryan A. Rossi**, Sana Malik, Georgios Theocharous, Handong Zhao, Gang Wu, *Predictive Resource Scaling for Efficient Resource Management, Adobe Research*, Patent application filed 5/5/2020, App. #16/867,104, 2020. US Patent No. 11,487,579. Awarded 11/1/2022.

[P28] Fan Du, Yeuk-Yin Chan, **Ryan A. Rossi**, Anup Rao, Eunyee Koh, Charles Menguy, Margarita Savova, *Dynamic Clustering of Sparse Data Utilizing Hash Partitions, Adobe Research*, Patent application filed 4/17/2020, App. #16/852,110, 2020. US Patent No. 11,328,002. Awarded 5/10/2022.

[P27] Yikun Xian, Handong Zhao, Tak Yeon Lee, **Ryan A. Rossi**, Sungchul Kim, *Generating Explanatory Paths for Predicted Column Annotations, Adobe Research*, Patent application filed 2/20/2020, App. #16/796,681, 2020.

[P26] Yikun Xian, Handong Zhao, Tak Yeon Lee, Sungchul Kim, **Ryan A. Rossi**, *Dynamically Determining Schema Labels Using a Hybrid Neural Network Encoder, Adobe Research*, Patent application filed 1/24/2020, App. #16/751,755, 2020.

[P25] Alireza Farhadi, **Ryan A. Rossi**, Anup Rao, Tung Mai, *Single-Pass Matching in Large Data Streams, Adobe Research*, Patent application filed 11/19/2019, App. #16/688,700, 2019. US Patent No. 11,526,907. Awarded 12/13/2022.

[P24] **Ryan A. Rossi**, *Multi-Item Influence Maximization, Adobe Research*, Patent application filed 11/7/2019, App. #16/677,007, 2019.

[P23] Di Jin, Sungchul Kim, **Ryan A. Rossi**, Eunyee Koh, *A System and Method for Enriching Data via Temporal Embeddings and Models, Adobe Research*, Patent application filed 7/22/2020, App. #XX, 2020.

[P22] **Ryan A. Rossi**, Sungchul Kim, Eunyee Koh, Anup Rao, Russell Stringham, *Analytics System Entity Resolution, Adobe Research*, Patent application filed 9/12/2019, App. #16/569,484, 2019.

[P21] **Ryan A. Rossi**, *Feature-based Temporal Walks for Inductive Temporal Network Representation Learning in Graph Streams, Adobe Research*, Patent application filed 7/10/2019, App. #16/507,204, 2019. US Patent No. 11,483,408. Awarded 10/25/2022.

[P20] **Ryan A. Rossi**, Aldo Gael Carranza, David Arbour, Anup Rao, Sungchul Kim, Eunyee Koh, *System for Identifying Typed Graphlets, Adobe Research*, Patent application filed 6/25/2019, App. #16/451,956, 2019. US Patent No. 11,170,048. Awarded 11/9/2021.

[P19] **Ryan A. Rossi**, Aldo Carranza, Anup Rao, Eunyee Koh, *Higher-order Network Clustering and Embedding, Adobe Research*, Patent application filed 4/29/2019, App. #16/397,839, 2019. US Patent No. 11,163,803. Awarded 11/02/2021.

[P18] Anup Rao, Tung Mai, Yasin Abbasi-Yadkori, **Ryan A. Rossi**, Ritwik Sinha, Matt Kapilevich, *Utilizing One Hash Permutation And Populated-value-Slot-Based Densification For Generating Audience Segment Trait Recommendations, Adobe Research*, Patent application filed 3/28/2019, App. #16/367,628, 2019. US Patent No. 11,109,085. Awarded 08/31/2021.

[P17] John Boaz Lee, **Ryan A. Rossi**, Sungchul Kim, Eunyee Koh and Anup Rao, *Graph Convolutional Networks with Motif-based Attention, Adobe Research*, Patent application filed 3/8/2019, App. #16/297,024, 2019.

[P16] Charles Chen, Eunyee Koh, Sungchul Kim, **Ryan A. Rossi**, Scott Cohen, *Figure Captioning System and Related Methods, Adobe Research*, Patent application filed 3/7/2019, App. #16/296,076, 2019. US Patent No. 11,461,638. Awarded 10/4/2022.

[P15] **Ryan A. Rossi**, Di Jin, Sungchul Kim, Anup Rao, Eunyee Koh, *Latent Network Summarization, Adobe Research*, Application filed 1/18/2019, App. #16/252,169, 2019. US Patent No. 11,113,293. Awarded 09/7/2021.

[P14] **Ryan A. Rossi**, Sungchul Kim, Eunyee Koh, Anup Rao, *Higher-Order Network Embedding, Adobe Research*, Patent application filed 11/29/2018, App. #16/204,616, 2018. US Patent No. 10,728,105. Awarded 07/28/2020.

[P13] **Ryan A. Rossi**, Sungchul Kim, Eunyee Koh, *Time-Dependent Network Embedding, Adobe Research*, Patent application filed 11/15/2018, App. #16/192,313, 2018. US Patent No. 10,728,104. Awarded 07/28/2020.

[P12] **Ryan A. Rossi**, Hoda Eldardiry, *Similarity-based Multi-label Learning, Palo Alto Research Center*, Patent application filed 12/31/2018, App. #16/237,439, 2018.

[P11] Jungho Park, Ajay Raghavan, **Ryan A. Rossi**, Yosuke Tajika, Akira Minegishi, and Tetsuyoshi Ogura, *System and Method for Anomaly Characterization Based on Joint Historical and Time-series Analysis, Palo Alto Research Center (Xerox PARC)*, Patent application filed, App. #16/170815, 2018.

[P10] Ajay Raghavan, **Ryan A. Rossi**, and Jungho Park, *System and Method for Binned Inter-Quartile Range Analysis in Anomaly Detection of Time-series Data, Palo Alto Research Center (Xerox PARC)*, Patent application filed, 2018.

[P9] **Ryan A. Rossi**, Ajay Raghavan, and Jungho Park, *One-class Similarity Machines for Anomaly Detection, Palo Alto Research Center (Xerox PARC)*, Patent application filed, 2018.

[P8] **Ryan A. Rossi**, Rong Zhou, *Deep Graph Representation Learning, Palo Alto Research Center*, US Patent No. 10482375, 2019.

[P7] **Ryan A. Rossi**, Rong Zhou, *A Graph Search Engine, Palo Alto Research Center*, Patent application filed, 2017.

[P6] **Ryan A. Rossi**, Rong Zhou, *System and Method for Hybrid Task Management Across CPU and GPU for Efficient Data Mining, Palo Alto Research Center*, US Patent No. 10235182, 2019.

[P5] **Ryan A. Rossi**, Rong Zhou, *Localized Visual Graph Filters For Complex Graph Queries, Palo Alto Research Center*, Patent application filed, 2016.

[P4] **Ryan A. Rossi**, Rong Zhou, *Fast and Accurate Unbiased Graphlet Estimation, Palo Alto Research Center*, Patent application filed, 2015.

[P3] **Ryan A. Rossi**, Rong Zhou, *A System and Method for Compressing Graphs via Cliques to Speedup Graph Algorithms and Reduce Storage Requirements, Palo Alto Research Center*, Patent application filed, 2015.

[P2] **Ryan A. Rossi**, Rong Zhou, *Relational Time Series Classification using Similarity, Palo Alto Research Center*, Patent application filed, 2015.

[P1] **Ryan A. Rossi**, Rong Zhou, *Parallel Collective Matrix Factorization Framework for Big Data, Palo Alto Research Center*, US Patent No. 10235403, 2019.

Research Awards & Funding

2018-2023 NSF, *CINES: A Scalable Cyberinfrastructure for Sustained Innovation in Network Engineering and Science*, PIs: Madhav Marathe (VTech), Jure Leskovec (Stanford), Geoffrey Fox (IU), Total: **$4,000,000**

2017 Xerox Research: DocuShare.
A Deep Similarity-based Graph Model for Learning Relevant Tags from Unstructured Data.
PI: *Ryan Rossi*, Co-PI: Hoda Eldardiry. Awarded **$450,000** total

2016-2018 XIG Explore Research: Self Machine Learning Program.
Relational Deep Learning.
PI: *Ryan Rossi*. **$300,000** awarded over 3 years, **$900,000** total
Deep learning in graphs including inductive network representation learning, graph embeddings, and learning representations from attributed graphs

2017 Xerox Research: Horizon 3 Research Program.
Intelligent Workflow Automation using AI Planning, Learning, and Conversation Agents.
PI: Rong Zhou, Co-PI: *Ryan Rossi*. Awarded **$720,000** total

2016 NVIDIA Hardware Research Gift.
Deep Graph Learning using Higher-order Functions based on Network Motifs.
PI: *Ryan Rossi*.

2015-2017 Xerox Foundation: University Research Collaboration.
Learning from an Expert in Noisy, Structured Domains: Adapting to Healthcare Problems.
co-PI & Liaison Manager: *Ryan Rossi* (subcontract with Indiana University), **$90,000**

Selected Talks & Outreach

2026 WSDM, *Tutorial on "Rigorizing Retrieval-augmented Generation with Structured Knowledge"*.

2026 AACL, *Co-Organizer of "RARA: The Second Workshop on Grounding Documents with Reasoning, Agents, Retrieval, and Attribution"*.

2025 COLM, *Co-Organizer of the "Third Scientific Figure Captioning Challenge"*.

2025 ICDM, *Co-Organizer of "RARA: Workshop on Grounding Documents with Reasoning, Agents, Retrieval, and Attribution"*.

2024 CIKM, *Tutorial on "Data Quality-aware Graph Machine Learning"*.

2024 IJCAI, *Co-Organizer of the "Second Scientific Figure Captioning (SciCap) Challenge"*.

2024 Adobe, *Personalization in AI: From LLMs to MLLMs*.

2023 ICCV, *Co-Organizer of the "First Scientific Figure Captioning (SciCap) Challenge"*.

2018–present GEM, *Successfully mentored over 20 GEM interns, resulting in publications at top-tier conferences and journals and significant business impact. Many of these students subsequently joined top universities and industrial research labs.*.

2018–present Lead Adobe's Industry Mentorship Programs, *Lead Adobe's industry mentorship programs where small groups of masters students work with myself, and other researchers on critical and high-impact problems. Nearly all mentorships have resulted in publications at top-tier conferences and journals, and many students have subsequently joined top universities and top industry positions and research labs.*.

Selected Honors and Awards

- 2025 **Best Paper**, *"Chart-to-Experience: Benchmarking Multimodal LLMs for Predicting Experiential Impact of Charts"*, IEEE PacificVis.
- 2025 **Best Poster**, *"Gaze2Prompt: Turning Eye-Tracking Data into Visual Prompts for Multimodal LLMs"*, UbiComp.
- 2023 **Best Paper** , *"Knowledge Graph Prompting for Multi-Document Question Answering"*, NeurIPS New Frontiers in Graph Learning Workshop.
- 2023 **Best Paper**, *"Summaries as Captions: Generating Figure Captions for Scientific Documents with Automated Text Summarization"*, International Natural Language Generation Conference (INLG).
- 2022 **Best Paper Honorable Mention**, *"Evaluating the Use of Uncertainty Visualisations for Imputations of Data Missing At Random in Scatterplots"*, IEEE VIS.
- 2022 **Best Short Paper Honorable Mention** , *"ARShopping: In-Store Shopping Decision Support Through Augmented Reality and Immersive Visualization"*, IEEE VIS.

2019–present Adobe Distinguished Inventor.

2020 Elevation to IEEE Senior Member.

2017 Research Impact Award (PARC).

2016 PARC Research Award (Nominated by peers) .

2012–2014 National Science Foundation (NSF) Graduate Fellow (GRFP).

2009–2012 Department of Defense: National Defense Science & Engineering PhD (NDSEG) Fellow.

2014–2015 Bilsland Dissertation Fellowship Awarded to Outstanding Ph.D. candidates.

2009 Purdue University Fredrick N. Andrews Doctoral Fellow.

2011–2012 LLNL Scholar (Cyber Defenders).

2010 Naval Research Laboratory (NREIP) Fellow: Center for Applied Research in Artificial Intelligence, Office of Naval Research (DoN).

Summer 2009 NASA USRP Fellow, Jet Propulsion Laboratory, California Institute of Technology.

2009 National Aeronautics and Space Administration SC Space Grant.

Spring 2009 NASA USRP Fellow, Jet Propulsion Laboratory, California Institute of Technology.

– Erdős Number 3, Rossi-Duffield-Alon-Erdős.

2003–2004 Obtained 11 Software Engineering and Information Technology Certifications at age 16-17: Microsoft Certified Solution Developer (MCSD), MCAD, CIW-A, CCNA, Linux+, Project+, i-Net+, Server+, Security+, Network+, A+, and MCP.

Open Source Software & Datasets

*Stats were last updated on April 2026.

2023–present **Fair LLM Benchmark**
A comprehensive benchmark compiling over 20 publicly available bias and fairness evaluation datasets for LLMs, enabling standardized evaluation across diverse dimensions.
1,700+ citations · 160 stars · 16 forks

2026–present **Orthrus**
A dual-architecture framework that unifies the exact generation fidelity of autoregressive large language models with the high-speed parallel token generation of diffusion models, sharing one key-value cache across both views for lossless, memory-efficient inference.
353 stars · 13 forks

2024–present **CulturaX**
A cleaned, multilingual dataset of 6.3 trillion tokens spanning 167 languages for large language model development, fully released on Hugging Face.
596 likes · 26K+ downloads/month

2025–present **DynaSaur**
An LLM agent framework that generates and executes Python programs as a universal action representation, creating and accumulating reusable actions for open-ended tasks.
358 stars · 28 forks

2025–present **NoLiMa**
A long-context evaluation benchmark that minimizes lexical overlap between questions and target content, requiring models to infer latent assoc. rather than rely on literal matching.
195 stars · 17 forks

2025–present **PGraphRAG**
Framework & benchmark for personalized graph-based retrieval-augmented generation that integrates user-centric knowledge graphs to enrich personalization under sparse user history.
34 stars · 3 forks

2024–present **LongLaMP**
A benchmark for personalized long-form text generation, providing a diverse evaluation framework across long-text tasks such as personalized email writing, review generation, and topic writing.
9 stars · 2 forks

2025–present **AD-LLM**
The first benchmark evaluating large language models for anomaly detection, spanning zero-shot detection, data augmentation, and model selection.
44 stars · 9 forks

2025–present **VisDoM**
A benchmark and multimodal retrieval-augmented generation approach for multi-document question answering over visually rich elements such as tables, charts, and slides.
43 stars · 6 forks

2025–present **LGGM**
A large-scale training paradigm for graph generative models, pretrained over thousands of graphs, with zero-shot, fine-tuning, and text-to-graph generation capabilities.
29 stars · 10 forks

2024–present **KGP**
A knowledge graph prompting method for multi-document question answering that pairs graph construction over documents with an LLM-guided graph traversal agent.
324 stars · 33 forks

2023–present **GLEMOS**
A comprehensive benchmark for instantaneous graph learning model selection, providing extensive performance records, evaluation testbeds, and meta-graph features.
6 stars

2023–present **FigCaps-HF**
A figure-to-caption generative framework and benchmark that incorporates human feedback to optimize generated captions for reader preferences.
3 stars · 1 fork

2023–present **Okapi**
A framework and resources for instruction tuning large language models in 26 languages using reinforcement learning from human feedback.
96 stars · 3 forks

2022–present **CyCLIP**
A contrastive vision-language pretraining framework that enforces geometric consistency in the image and text representation spaces.
125 stars · 16 forks

2021–present **CPGNN**
A graph neural network framework that learns an interpretable compatibility matrix, generalizing message passing to graphs with either homophily or heterophily.
31 stars · 3 forks

2019–present **node2bits**
An efficient framework for user stitching that encodes multi-dimensional node context from feature-based temporal walks into compact binary hashcodes.
7 stars · 6 forks

2019–present **role2vec**
A widely used community implementation of our popular role2vec role-based network embedding method.
169 stars · 33 forks

2015–present **Parameterized Graphlet Decomposition Library (PGD)**
A fast parallel high-performance parameterized graphlet decomposition library for massive networks. Code at .
90 stars · 31 forks

2012–present **Network Repository (NR)**
The first interactive data repository that integrates visualization with state-of-the-art statistical methods and analytic techniques to support discovery and exploration of data in real-time. NR is the largest network data repository, with over 6,000 donations across 30+ collections and growing.
4,000+ citations · 250 million downloads

2014–present **GraphVIS**
Interactive visual graph mining and machine learning on the web. Visualize and explore network data easily. GraphVIS is the result of years of research in relational machine learning and graph mining. A free demo version is available at ***http://networkrepository.com/graphvis***

2013 **Parallel Maximum Clique Library (PMC)**
A parallel high-performance library for solving the maximum clique problem on dense graphs and large sparse networks.
120 stars · 51 forks

2013–present **MLVis**
An interactive data repository that makes it easy to find, explore, and understand machine learning data, providing researchers with open, persistent, and accessible data alongside web-based visual analytic tools.

2012 **Dynamic PageRank**
A package for modeling the importance and influence of nodes in dynamic networks with external interest and attributes.
16 stars · 4 forks

Ph.D. Dissertation Committees

2024–present **Wang Wei**, *Virginia Tech*, Ph.D. Candidate,
Dissertation: "LLM/MLLM-based Model Routing".

2019–2024 **Yu Wang**, *Vanderbilt University*,
Dissertation: "Data-quality-aware Graph Machine Learning"
– 2025 ACM SIGKDD Best Dissertation Award Honorable Mention.

2020–2024 **Hongjie Chen**, *Virginia Tech*.

2019–2022 **George Panagopoulos**, *École Polytechnique*, Ph.D. Candidate,
Dissertation: "Learning Influence Representations: Methods and Applications".

2018–2021 **Di Jin**, *University of Michigan*,
Dissertation: "Network Summarization and Embedding".

2016-2020 **John Boaz Lee**, *Worcester Polytechnic Institute (WPI)*,
Dissertation: "Deep Learning on Graph-structured Data".

Students Supervised

2024–2026 **Bo Ni**, *Vanderbilt University*, Ph.D. Student.

2024–2026 **Wang Wei**, *Virginia Tech*, Ph.D. Candidate.

2024–2026 **Reuben Luera**, *UC Berkeley*, GEM Fellow, M.S. Student.

2024–2026 **Deonna M. Owens**, *Stanford University*, GEM Fellow, M.S. Student.

2024–2025 **Yu Xia**, *UC San Diego*, Ph.D. Student.

2024–2025 **Dang Nguyen**, *University of Oregon*, Ph.D. Student.

2024–2025 **Yongjia Lei**, Ph.D. Student.

2024–2025 **Tiankai Yang**, *Texas A&M University*, Ph.D. Student.

2024–2025 **Zhehao Zhang**, *Dartmouth College*, M.S. Student.

2024–2025 **Yeonjun In**, *KAIST*, Ph.D. Student.

2024–2025 **Ashish Singh**, *UMass*, Ph.D. Student.

2024–2025 **Songwen Hu**, *University of Georgia*, Ph.D. Student.

2023–2024 **Isabel O. Gallegos**, *Stanford University*, Ph.D. Student.

2023–2024 **Shanyun Gao**, *Purdue University*, Ph.D. Student.

2023–2024 **Xinyu Shi**, *University of Waterloo*, Ph.D. Student.

2023–2024 **Yoonjoo Lee**, *KAIST*, Ph.D. Student.

2023–2024 **Puja Trivedi**, *University of Michigan*, Ph.D. Student.

2023–2024 **Mehrnoosh Mirtaheri**, *USC*, Ph.D. Student.

2023–2024 **Jian Chen**, *University at Buffalo*, Ph.D. Student.

2022–2023 **Fayokemi Ojo**, Ph.D. Student.

2022–2023 **Enyu Cai**, Ph.D. Student.

2020-2022 **Namyong Park**, *CMU*, Meta Research.

2020-2022 **Sudhanshu Chanpuriya**, *UMass Amherst*, Ph.D. Student.

2020-2022 **Hyeok Kim**, *Northwestern*, Ph.D. Student.

2021-2022 **Arpit Narechania**, *Georgia Tech*, Ph.D. Student.

2021-2022 **Mustafa Abdallah**, *Purdue University*, Assistant Professor at IUPUI.

2021-2022 **Gaurav Verma**, *Georgia Tech*, Ph.D. Student.

2021-2022 **Mohammad Mehrabi**, *USC*, Assistant Professor at IUPUI.

2022 **Shravika Mittal**, *Georgia Tech*, Ph.D. Student.

2022 **April Chen**, *Harvard University*, Ph.D. Student.

2022 **Xinyu Shi**, *Univ. of Waterloo*, Ph.D. Student.

2022 **Tingyao Hsu**, *Penn. State*, Ph.D. Student.

2022 **Yuhang Yao**, *CMU*, Ph.D. Student.

2022 **Carol (Xinyi) Zheng**, *CMU*, Ph.D. Student.

2022 **Shanyun Gao**, *Purdue University*, Ph.D. Student.

2022 **Rashmi Ranjan Bhuyan**, *USC*, Ph.D. Student.

2022 **Luke Snyder**, *University of Washington*, Ph.D. Student.

2022 **Guande Wu**, *NYU*, Ph.D. Student.

2022 **Ryan Aponte**, *CMU*, Ph.D. Student.

2022 **Chen Chen**, *UMD*, Ph.D. Student.

2022 **Alex Tang**, *Northwestern*, Ph.D. Student.

2022 **Jean-Peic Chou**, *Stanford University*, Ph.D. Student.

2022 **Melanie Bancilhon**, *Washington University in St. Louis*, Ph.D. Student.

2022 **Jianna Audrey So**, *Stanford University*, Ph.D. Student.

2022 **Jaeho Bang**, *Georgia Tech*, Ph.D. Student.

2022 **Dario Garigliotti**, *Aalborg University*, Ph.D. Student.

2022 **Princess Sampson**, *University of Pennsylvania*, Ph.D. Student.

2022 **Enyu Cai**, *Purdue University*, Ph.D. Student.

2022 **Dongliang Guo**, *University of Virginia*, Ph.D. Student.

2022 **Songwen Hu**, *Shanghai Jiao Tong University*, Ph.D. Student.

2021 **Gaurav Gupta**, *Rice University*, Ph.D. Student.

2021 **Yang Li**, *UNC*, Ph.D. Student.

2021 **Nathan Ng**, *UMass Amherst*, Ph.D. Student.

2021 **Bingjie (Jenny) Xu**, *Northwestern*, Ph.D. Student.

2021 **Beleicia Bullock**, *Stanford University*, Ph.D. Student.

2021 **Md Main Uddin Rony**, *UMD*, Ph.D. Student.

2021 **Fayokemi Ojo**, *John Hopkins University*, Ph.D. Student.

2021 **Can Qin**, *Northeastern*, Ph.D. Student.

2021 **Shivam Srivastava**, *UMass Amherst*, Ph.D. Student.

2021 **Abhraneel Sarma**, *Northwestern*, Ph.D. Student.

2021 **Weixin Jiang**, *Northwestern*, Ph.D. Student.

2021 **Duc Hoang**, *University of Texas at Austin*, Ph.D. Student.

2021 **Benjamin Coleman**, *Rice University*, Ph.D. Student.

2021 **Pattara Sukprasert**, *Northwestern*, Ph.D. Student.

2020 **Jiong Zhu**, *University of Michigan*, Ph.D. Student.

2019-2021 **Xin Qian**, *University of Maryland*, Ph.D. Student.

2020 **Yue Zhao**, *CMU*, Ph.D. Student.

2020-2021 **Sejoon Oh**, *Georgia Tech*, Ph.D. Student.

2020 **Mojtaba Sahraee-Ardakan**, *USC*, Ph.D. Student.

2020 **Enayat Ullah**, *John Hopkins University*, Ph.D. Student.

2019-2020 **Jun Yan**, *USC*, Ph.D. Student.

2020 **Shenyu Xu**, *Georgia Tech*, Ph.D. Student.

2020 **Zening Qu**, *University of Washington*, Ph.D. Student.

2020 **Ihudiya Finda Ogbonnaya-Ogburu**, *University of Michigan*, Ph.D. Student.

2020 **Zhuohao Zhang**, *University of Illinois at Urbana-Champaign*, Ph.D. Student.

2020 **Galen Weld**, *University of Washington*, Ph.D. Student.

2020-2021 **Chenhan Yuan**, *Virginia Tech*, Ph.D. Student.

2020 **Camille Harris**, *Georgia Tech*, Ph.D. Student.

2020 **Mrigank Raman**, *USC*, Ph.D. Student.

2020 **Zihao Zhou**, *UCSD*, Ph.D. Student.

2019-2022 **Alireza Farhadi**, *University of Maryland (UMD)*, Ph.D. Candidate.

2019-2022 **Hongjie Chen**, *Virginia Tech*, Dolby Labs.

2019 **Youngsuk Park**, *Stanford University*, Ph.D. Candidate.

2019 **Yikun Xian**, *Rutgers University*, Ph.D. Candidate.

2019-2020 **Saed Rezayi**, *University of Georgia (UGA)*, Ph.D. Candidate.

2019-2020 **Gromit Yeuk-Yin Chan**, *New York University (NYU)*, Ph.D. Candidate.

2019 **Kirankumar Shiragur**, *Stanford University*, Ph.D. Candidate.

2019 **Hongchang Gao**, *University of Pittsburgh*, Ph.D. Candidate.

2019 **He Jia**, *Georgia Institute of Technology*, Ph.D. Candidate.

2018-2021 **Di Jin**, *University of Michigan*, Ph.D. Candidate,
Research led to a paper entitled "Deep Graph Attention Models" published at KDD.
Patent application filed on the research.
Serving on his Ph.D. committee .

2018-2019 **Mark Heimann**, *University of Michigan*, Ph.D. Candidate.

2018 **Aldo Carranza**, *Stanford University*, Ph.D. Student.

2018 **Zahra Shakeri**, *Rutgers University*, Ph.D. Candidate.

2018 **Donghyun Kim**, *POSTECH*, Graduated @ Oath/Yahoo Research.

2018-2019 **Jianjun Luo**, *Worcester Polytechnic Institute (WPI)*, Ph.D. Candidate.

2018-2019 **Charles Chen**, *Ohio State University*, Ph.D. Candidate.

2018 **Tung Mai**, *Georgia Institute of Technology*, Ph.D. Candidate.

2017 **Jungho Park**, *Seoul National University*, Ph.D Candidate,
Filed 3 patent applications related to anomaly detection in large multi-variate time series.

2017-2021 **John Boaz Lee**, *Worcester Polytechnic Institute (WPI)*, Ph.D. Candidate,
Research led to a paper entitled "Deep Graph Attention Models" published at KDD.
Patent application filed on the research.
Served on his Ph.D. committee "Deep Learning on Graph-structured Data". Successfully defended and now at Facebook Research. .

2017-2018 **Giang Hoang Nguyen**, *Worcester Polytechnic Institute (WPI)*, Masters Student.

Professional Service & Leadership

Invited SPC/PC Member/Organizer:

2021-2026 Associate Editor for Frontiers in Big Data (Machine Learning and Artificial Intelligence)

2020-2026 International Conference on Learning Representations (ICLR)

2019-2026 Neural Information Processing Systems (NeurIPS)

2019-2026 International Conference on Machine Learning (ICML)

2019-2026 European Conference on Machine Learning (ECML/PKDD)

2018-2025 SIAM International Conference on Data Mining (SDM)

2016-2026 International Joint Conferences on Artificial Intelligence (IJCAI)

2016-2026 World Wide Web Conference (WWW)

2015-2026 AAAI Conference on Artificial Intelligence

2019-2025 International Conference on Information and Knowledge Management (CIKM)

2021-2026 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)

2021-2025 SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)

2019-2021 International Conference on Artificial Neural Networks (ICANN)

2019-2021 ACM CHI Conference on Human Factors in Computing Systems

2016-2021 IEEE Big Data

Journal Reviewing

2019-2025 Journal of Machine Learning Research

2016-2025 IEEE Transactions on Neural Networks and Learning Systems (TNNLS)

2020-2026 ACM Computing Surveys

2018-present Review Editor, Machine Learning and Artificial Intelligence, Frontiers in Big Data

2012-2026 Transactions on Knowledge and Data Engineering (TKDE)

2015-2025 Artificial Intelligence

2012-13, 2017-20 Data Mining and Knowledge Discovery (DMKD)

2018-20 Theoretical Computer Science (TSC)

2019-20 Information Systems

2019-20 Network Science

2020 IEEE Transactions on Intelligent Transportation Systems

2020 Journal of Statistical Mechanics: Theory and Experiment (JSTAT)

2020 IEEE Transactions on Network Science and Engineering

2019-20 IEEE Intelligent Systems

2018-19 ACM Transactions on Intelligent Systems and Technology (TIST)

2018-19 Social Network Analysis and Mining

2019 Information Processing and Management

2019 Applied Network Science

2019 Future Internet

2015-2017 Internet Mathematics

2015-2019 Knowledge and Information Systems (KAIS)

2018 IEEE Internet of Things Journal

2018 IEEE Transactions on Emerging Topics in Computational Intelligence

2013-2014 The Annals of Applied Statistics

2016 Physica A: Statistical Mechanics and its Applications

2016-2017 ACM Transactions on Information Systems (TOIS)

2015-2018 Journal of Big Data

2016 Optimization Methods and Software

Proposal Reviewing

2019 Natural Sciences and Engineering Research Council of Canada (NSERC)

2018 Swiss National Science Foundation (SNSF)

2017 NSF Information & Intelligent Systems Panel