Nesreen K. Ahmed, Ph.D.

n.kamel@gmail.com nkahmed (405)-255-1838
http://nesreenahmed.com
2292 Solace Loop, San Jose, CA 95122

Research Interests

My research is focused on large-scale machine learning with an emphasis on theory, algorithms, and applications of graph representation learning, LLMs, generative AI, as well as sampling and inference.

Education

2008 – 2015 **Ph.D. in Computer Science, Purdue University, CS Dept.**
Thesis title: *Scaling Up Network Analysis and Mining: Statistical Sampling, Estimation, and Pattern Discovery.*

2012 – 2014 **M.Sc. in Computer Science and Statistics, Purdue University, CS/Stat Departments.**

2003 – 2006 **M.Sc., Cairo University, Faculty of Computer Science and Information**
Thesis title: *A Proposed Intrusion Detection System for Encrypted Computer Networks.*

2001 **B.Sc., Cairo University, Faculty of Computer Science and Information** (Summa Cum Laude).
Thesis title: *Arabic Speech Recognition using Hidden Markov Models.*
Ranking 2nd among over 500 students

Experience

2015 – present **Senior Staff Research Scientist & Technical Lead,** Intel Labs, Santa Clara, CA.

2015 **Research Scientist,** Technicolor Research & Innovation Labs, Palo Alto, CA.

2008 – 2015 **PhD Candidate & Researcher,** Purdue University, West Lafayette, IN.

2014 **Research Intern,** Meta, Menlo Park, CA.

2013 **Research Intern,** Intel Corp, Folsom, CA.

2012 **Research Intern,** Adobe Research, San Jose, CA.

2005-2008 **Research Fellow,** Data Mining & Computer Modeling Center of Excellence in Cairo, United Nations Development Program.

2001-2008 **Researcher & Lecturer,** Computer Science, Cairo Univeristy, Egypt.

2006-2008 **Cisco Certified Instructor,** CCNA Networking Academy Program, Cairo University.

2008 **Visiting Researcher & Lecturer,** Computer Science, German Univeristy in Cairo, Media Engineering & Technology.

Research Publications

Journal Articles

1 A. Chen, R. A. Rossi, N. Park, P. Trivedi, Y. Wang, T. Yu, S. Kim, F. Dernoncourt, and **N. K. Ahmed**, "Fairness-aware graph neural networks: A survey," *ACM Transactions on Knowledge Discovery from Data*, vol. 18, no. 6, 2024, ISSN: 1556-4681. DOI: 10.1145/3649142.

2. I. O. Gallegos, R. A. Rossi, J. Barrow, M. M. Tanjim, S. Kim, F. Dernoncourt, T. Yu, R. Zhang, and **N. K. Ahmed**, "Bias and fairness in large language models: A survey," *Computational Linguistics (Accepted to appear) - arXiv:2309.00770*, 2024.

3. G. Ma, V. A. Vo, T. Willke, and **N. K. Ahmed**, "Memory-augmented graph neural networks: A brain-inspired review," *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 5, pp. 2011–2025, 2024. DOI: 10.1109/TAI.2023.3329454.

4. Y. Xiao, **N. K. Ahmed**, M. Capota, G. Ma, T. L. Willke, S. Nazarian, and P. Bogdan, "A graph-based learning framework for compiler loop auto-vectorization," *Intelligent Computing (to appear)*, 2024.

5. L. Cotta, B. Bevilacqua, **N. K. Ahmed**, and B. Ribeiro, "Causal lifting and link prediction," *Proceedings of the Royal Society A*, vol. 479, no. 2276, p. 20 230 121, 2023. DOI: 10.1098/rspa.2023.0121.

6. Y. Xiao, G. Ma, **N. K. Ahmed**, M. Capota, T. Willke, S. Nazarian, and P. Bogdan, "End-to-end programmable computing systems," *Nature Communications Engineering*, vol. 2, no. 1, p. 84, 2023. DOI: 10.1038/s44172-023-00127-7.

7. X. Qian, R. A. Rossi, F. Du, S. Kim, E. Koh, S. Malik, T. Y. Lee, and **N. K. Ahmed**, "Personalized visualization recommendation," *ACM Transactions on the Web (TWEB)*, vol. 16, no. 3, 2022, ISSN: 1559-1131. DOI: 10.1145/3538703.

8. **N. K. Ahmed**, N. Duffield, and R. A. Rossi, "Online sampling of temporal networks," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 15, no. 4, 2021, ISSN: 1556-4681. DOI: 10.1145/3442202.

9. J. B. Lee, G. Nguyen, R. A. Rossi, **N. K. Ahmed**, E. Koh, and S. Kim, "Dynamic node embeddings from edge streams," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 5, no. 6, pp. 931–946, 2021. DOI: 10.1109/TETCI.2020.3011432.

10. G. Ma, **N. K. Ahmed**, T. L. Willke, and S. Y. Philip, "Deep graph similarity learning: A survey," *Data Mining and Knowledge Discovery*, pp. 1–38, 2021.

11. G. Petri, S. Musslick, B. Dey, K. Özcimder, D. Turner, **N. K. Ahmed**, T. Willke, and J. D. Cohen, "Topological limits to the parallel processing capability of network architectures," *Nature Physics*, vol. 17, no. 5, pp. 646–651, 2021.

12. **N. K. Ahmed**, R. A. Rossi, J. B. Lee, T. L. Willke, R. Zhou, X. Kong, and H. Eldardiry, "Role-based graph embeddings," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 5, pp. 2401–2415, 2020.

13. R. A. Rossi, R. Zhou, and N. K. Ahmed, "Deep inductive graph representation learning," *IEEE Transactions on Knowledge & Data Engineering (TKDE)*, vol. 32, no. 03, pp. 438–452, 2020, ISSN: 1558-2191. DOI: 10.1109/TKDE.2018.2878247.

14. R. A. Rossi, D. Jin, S. Kim, **N. K. Ahmed**, D. Koutra, and J. B. Lee, "On proximity and structural role-based embeddings in networks: Misconceptions, techniques, and applications," *ACM Transactions on Knowledge Discovery from Data*, vol. 14, no. 5, 2020, ISSN: 1556-4681. DOI: 10.1145/3397191.

15. R. A. Rossi, **N. K. Ahmed**, A. Carranza, D. Arbour, A. Rao, S. Kim, and E. Koh, "Heterogeneous graphlets," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 15, no. 1, 2020, ISSN: 1556-4681. DOI: 10.1145/3418773.

16. J. B. Lee, R. A. Rossi, S. Kim, **N. K. Ahmed**, and E. Koh, "Attention models in graphs: A survey," *ACM Transactions on Knowledge Discovery from Data*, vol. 13, no. 6, 2019, ISSN: 1556-4681. DOI: 10.1145/3363574.

17. R. A. Rossi, R. Zhou, and **N. K. Ahmed**, "Estimation of graphlet counts in massive networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 1, pp. 44–57, 2019. DOI: 10.1109/TNNLS.2018.2826529.

18 R. A. Rossi, **N. K. Ahmed**, R. Zhou, and H. Eldardiry, "Interactive visual graph mining and learning," *ACM Transactions on Intelligent Systems and Technology ()*, vol. 9, no. 5, 2018, ISSN: 2157-6904. DOI: 10.1145/3200764.

19 **N. K. Ahmed**, J. Neville, R. A. Rossi, N. G. Duffield, and T. L. Willke, "Graphlet decomposition: Framework, algorithms, and applications," *Knowledge and Information Systems*, vol. 50, no. 3, 2017, ISSN: 0219-1377. DOI: 10.1007/s10115-016-0965-5.

20 J. P. Canning, E. E. Ingram, A. M. Ortiz, and K. R. Schmitt, "Network classification and categorization paper," 2017.

21 R. A. Rossi and **N. K. Ahmed**, "An interactive data repository with visual analytics," *ACM SIGKDD Explorations Newsletter*, vol. 17, no. 2, 2016, ISSN: 1931-0145. DOI: 10.1145/2897350.2897355.

22 R. A. Rossi and **N. K. Ahmed**, "Role discovery in networks," *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 27, no. 4, pp. 1112–1131, 2015, ISSN: 1041-4347. DOI: 10.1109/TKDE.2014.2349913.

23 R. A. Rossi and **N. K. Ahmed**, "Coloring large complex networks," *Social Network Analysis and Mining*, vol. 4, no. 1, pp. 1–37, 2014. DOI: 10.1007/s13278-014-0228-y.

24 **N. K. Ahmed**, J. Neville, and R. Kompella, "Network sampling: From static to streaming graphs," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 8, no. 2, 2013, ISSN: 1556-4681. DOI: 10.1145/2601438.

25 **N. K. Ahmed**, A. F. Atiya, N. E. Gayar, and H. El-Shishiny, "An empirical comparison of machine learning models for time series forecasting," *Econometric Reviews*, vol. 29, no. 5-6, pp. 594–621, 2010. DOI: 10.1080/07474938.2010.481556.

Conference Proceedings

1 L. Chen, A. Bhattacharjee, **N. K. Ahmed**, N. Hasabnis, G. Oren, V. Vo, and A. Jannesari, "OMPGPT: A generative pre-trained transformer model for openmp," in *Proc. of the International European Conference on Parallel and Distributed Computing (EuroPar)*, Aug. 2024.

2 N. Park, X. Wang, A. Simoulin, S. Yang, G. Yang, R. A. Rossi, P. Trivedi, and **N. K. Ahmed**, "Forward learning of graph neural networks," in *International Conference on Learning Representations (ICLR)*, May 2024.

3 N. Schneider, N. Hasabnis, V. A. Vo, T. Kadosh, N. Krien, M. Capotă, A. Wasay, G. Tamir, T. Willke, **N. K. Ahmed**, *et al.*, "MPIrigen: MPI code generation through domain-specific language models," in *AI4Sys Workshop at HPDC*, 2024.

4 P. Trivedi, R. Rossi, D. Arbour, T. Yu, F. Dernoncourt, S. Kim, N. Lipka, N. Park, **N. K. Ahmed**, and D. Koutra, "Leveraging graph diffusion models for network refinement tasks," in *Proc. of the 41st International Conference on Machine Learning (ICML*, Jul. 2024.

5 R. Aponte, R. Rossi, S. Guo, J. Hoffswell, N. Lipka, C. Xiao, G. Y.-Y. Chan, E. Koh, and **N. K. Ahmed**, "A ml-based approach for html-based style recommendation," in *Companion Proceedings of the ACM Web Conference*, Apr. 2023, pp. 9–13.

6 L. Chen, Q. I. Mahmud, H. Phan, **N. K. Ahmed**, and A. Jannesari, "Learning to parallelize with openmp by augmented heterogeneous ast representation," in *Proceedings of Machine Learning and Systems*, vol. 5, May 2023.

7 K. Cheng, **N. K. Ahmed**, and Y. Sun, "Neural compositional rule learning for knowledge graph reasoning," in *International Conference on Learning Representations (ICLR)*, May 2023.

8 G. Ma, V. A. Vo, T. L. Willke, and **N. K. Ahmed**, "Augmenting recurrent graph neural networks with a cache," in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Aug. 2023, pp. 1608–1619.

9. N. Park, R. Rossi, X. Wang, A. Simoulin, **N. K. Ahmed**, and C. Faloutsos, "Glemos: Benchmark for instantaneous graph learning model selection," in *Advances in Neural Information Processing Systems (Neurips)*, vol. 36, Dec. 2023.

10. N. Park, R. A. Rossi, **N. K. Ahmed**, and C. Faloutsos, "MetaGL: Evaluation-free selection of graph learning models via meta-learning," in *International Conference on Learning Representations (ICLR)*, May 2023.

11. R. Rossi, **N. K. Ahmed**, and N. Park, "On graph time-series representations for temporal networks," in *Companion Proceedings of the ACM Web Conference*, Apr. 2023, pp. 14–18.

12. R. Rossi, S. Sarkhel, and **N. K. Ahmed**, "Efficient estimation of local causal effects in graphs via neighborhood pooling," in *IEEE International Conference on Big Data (BigData)*, IEEE, Dec. 2023, pp. 542–547.

13. A. TehraniJamsaz, Q. I. Mahmud, L. Chen, **N. K. Ahmed**, and A. Jannesari, "Perfograph: A numerical aware program graph representation for performance optimization and program analysis," in *Advances in Neural Information Processing Systems (Neurips)*, vol. 36, Dec. 2023.

14. D. Cummings, A. Brahmaroutu, M. Nassar, and **N. K. Ahmed**, "Technology growth ranking using temporal graph representation learning," in *Proceedings of the Web Conference (WWW)*, Apr. 2022.

15. G. Ma, V. Vo, T. Willke, and **N. K. Ahmed**, "Cache-memory gated graph neural networks," in *NeurIPS workshop on Memory in Artificial and Real Intelligence (MemARI)*, 2022.

16. N. Park, R. Rossi, E. Koh, I. A. Burhanuddin, S. Kim, F. Du, **N. K. Ahmed**, and C. Faloutsos, "CGC: Contrastive graph clustering for community detection and tracking," in *Proceedings of the Web Conference (WWW)*, 2022.

17. A. Reddy, R. A. Rossi, Z. Song, A. Rao, T. Mai, N. Lipka, G. Wu, E. Koh, and **N. K. Ahmed**, "Online map inference and learning for nonsymmetric determinantal point processes," in *Proc. of the 39th International Conference on Machine Learning (ICML)*, Jul. 2022.

18. A. S. Tom, **N. K. Ahmed**, and G. Karypis, "Joint learning of hierarchical community structure and node representations: An unsupervised approach," in *European Conference on Machine Learning and Data Mining (ECML)*, 2022.

19. X. Zheng, R. A. Rossi, **N. K. Ahmed**, and D. Moritz, "Network report: A structured description for network datasets," in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (CIKM)*, 2022, ISBN: 9781450392365. DOI: 10.1145/3511808.3557115.

20. G. Ma, Y. Xiao, M. Capotă, T. L. Willke, S. Nazarian, P. Bogdan, and **N. K. Ahmed**, "Learning code representations using multifractal-based graph networks," in *2021 IEEE International Conference on Big Data (Big Data)*, IEEE, 2021, pp. 1858–1866.

21. G. Ma, Y. Xiao, T. Willke, **N. K. Ahmed**, S. Nazarian, and P. Bogdan, "A distributed graph-theoretic framework for automatic parallelization in multi-core systems," in *Proceedings of Machine Learning and Systems*, vol. 3, May 2021.

22. V. Md, S. Misra, G. Ma, R. Mohanty, E. Georganas, A. Heinecke, D. Kalamkar, **N. K. Ahmed**, and S. Avancha, "Distgnn: Scalable distributed training for large-scale graph neural networks," in *SC '21: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2021.

23. R. A. Rossi, A. Rao, S. Kim, E. Koh, **N. K. Ahmed**, and G. Wu, "From closing triangles to higher-order motif closures for better unsupervised online link prediction," in *Proc. of the 30th ACM International Conference on Information & Knowledge Management (CIKM)*, Oct. 2021, pp. 4085–4093.

24 Y. Xiao, G. Ma, **N. K. Ahmed**, T. L. Willke, S. Nazarian, and P. Bogdan, "Deep graph learning for program analysis and system optimization," in *Workshop on Graph Neural Networks and Systems - GNNSys 2021*, 2021.

25 J. Zhu, R. A. Rossi, A. Rao, T. Mai, N. Lipka, **N. K. Ahmed**, and D. Koutra, "Graph neural networks with heterophily," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.

26 **N. K. Ahmed**, R. Alo, C. Amelink, Y. Y. Baek, A. Chudhary, K. Collins, A. Esterline, E. Fox, G. Fox, A. Hagberg, *et al.*, "Net.science: A cyberinfrastructure for sustained innovation in network science and engineering," in *Gateways 2020 (Science Gateways Community Institute)*, 2020.

27 **N. K. Ahmed** and N. Duffield, "Adaptive shrinkage estimation for streaming graphs," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

28 A. Haj-Ali, **N. K. Ahmed**, T. Willke, S. Shao, K. Asanovic, and I. Stoica, "Neurovectorizer: End-to-end vectorization with deep reinforcement learning," in *CGO 2020: Proc. of the International Symposium on Code Generation and Optimization*, Feb. 2020.

29 J. B. Lee, X. Kong, C. M. Moore, and **N. K. Ahmed**, "Deep parametric model for discovering group-cohesive functional brain regions," in *Proceedings of the 2020 SIAM International Conference on Data Mining*, Society for Industrial and Applied Mathematics, 2020, pp. 631–639.

30 R. A. Rossi, **N. K. Ahmed**, E. Koh, and S. Kim, "Fast hierarchical graph clustering in linear-time," in *Proceedings of The Web Conference (WWW)*, 2020, pp. 10–12.

31 R. A. Rossi, **N. K. Ahmed**, E. Koh, S. Kim, A. Rao, and Y. Abbasi-Yadkori, "A structural graph representation learning framework," in *Proceedings of the 13th International Conference on Web Search and Data Mining*, Feb. 2020, pp. 483–491.

32 R. A. Rossi, A. Rao, S. Kim, E. Koh, and **N. K. Ahmed**, "From closing triangles to closing higher-order motifs," in *Companion Proceedings of the Web Conference*, Apr. 2020, pp. 42–43.

33 R. A. Rossi, S. Sarkhel, and **N. K. Ahmed**, "Inferring individual level causal models from graph-based relational time series," in *AAAI Workshop on Statistical Relational AI*, Feb. 2020.

34 A. Zweig, **N. K. Ahmed**, T. Willke, and G. Ma, "Neural algorithms for graph navigation," in *NeurIPS Workshop on Learning Meets Combinatorial Algorithms*, Dec. 2020.

35 **N. K. Ahmed**, R. A. Rossi, J. B. Lee, T. L. Willke, R. Zhou, X. Kong, and H. Eldardiry, "Role2vec: Role-based network embeddings," in *KDD Deep Learning Day 2019*, 2019, pp. 1–7.

36 A. Haj-Ali, **N. K. Ahmed**, T. Willke, S. Shao, K. Asanovic, and I. Stoica, "Learning to vectorize using deep reinforcement learning," in *Neurips Workshop on Machine Learning for Systems*, Dec. 2019.

37 G. Ma, **N. K. Ahmed**, T. L. Willke, D. Sengupta, M. W. Cole, N. B. Turk-Browne, and S. Y. Philip, "Deep graph similarity learning for brain data analysis," in *Proc. of the ACM International Conference on Information and Knowledge Management (CIKM)*, ACM, 2019, pp. 2743–2751.

38 **N. K. Ahmed**, N. Duffield, and L. Xia, "Sampling for approximate bipartite network projection," in *Proc. of the 27th Joint Conference on Artificial Intelligence (IJCAI-18)*, 2018.

39 **N. K. Ahmed**, R. Rossi, J. B. Lee, T. L. Willke, R. Zhou, X. Kong, and H. Eldardiry, "Learning role-based graph embeddings," in *IJCAI Workshop on Statistical Relational AI (StarAI)*, 2018.

40 G. H. Nguyen, J. B. Lee, R. A. Rossi, **N. K. Ahmed**, E. Koh, and S. Kim, "Dynamic network embeddings: From random walks to temporal random walks," in *IEEE Big Data*, 2018.

41 G. H. Nguyen, J. B. Lee, R. A. Rossi, **N. K. Ahmed**, E. Koh, and S. Kim, "Continuous-time dynamic network embeddings," in *Companion Proceedings of the The Web Conference*, 2018, pp. 969–976.

42 R. Rossi, **N. K. Ahmed**, and E. Koh, "Interactive higher-order network analysis," in *Demo at IEEE International Conference on Data Mining Workshops (ICDMW)*, 2018, pp. 1441–1446.

43 R. A. Rossi, **N. K. Ahmed**, H. Eldardiry, and R. Zhou, "Similarity-based multi-label learning," in *International Joint Conference on Neural Networks (IJCNN)*, 2018.

44 R. A. Rossi, **N. K. Ahmed**, and E. Koh, "Higher-order network representation learning," in *Companion Proceedings of the Web Conference*, 2018.

45 R. A. Rossi, R. Zhou, and **N. K. Ahmed**, "Deep inductive network representation learning," in *Companion Proceedings of the The Web Conference*, Apr. 2018, pp. 953–960.

46 R. A. Rossi, R. Zhou, **N. K. Ahmed**, and H. Eldardiry, "Relational similarity machines (rsm): A similarity-based learning framework for graphs," in *IEEE Big Data*, 2018.

47 **N. K. Ahmed**, N. Duffield, T. L. Willke, and R. A. Rossi, "On sampling from massive graph streams," in *Proceedings of the VLDB Endowment*, vol. 10, VLDB Endowment, 2017, pp. 1430–1441.

48 **N. K. Ahmed**, R. A. Rossi, T. Willke, and R. Zhou, "Edge role discovery via higher-order structures," in *Advances in Knowledge Discovery and Data Mining*, 2017.

49 V. S. Dave, **N. K. Ahmed**, and M. Al Hasan, "E-CLoG: Counting edge-centric local graphlets," in *IEEE Big Data*, 2017.

50 N. Duffield, Y. Xu, L. Xia, **N. K. Ahmed**, and M. Yu, "Stream aggregation through order sampling," in *Proc. of the 26th Conference on Information and Knowledge Management (CIKM)*, 2017.

51 K. Ozcimder, B. Dey, S. Musslick, G. Petri, **N. K. Ahmed**, T. L. Willke, and J. D. Cohen, "A formal approach to modeling the cost of cognitive control," in *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*, 2017.

52 S. Smith, X. Liu, **N. K. Ahmed**, A. S. Tom, F. Petrini, and G. Karypis, "Truss decomposition on shared-memory parallel systems," in *IEEE High Performance Extreme Computing Conference (HPEC)*, 2017.

53 A. S. Tom, N. Sundaram, **N. K. Ahmed**, S. Smith, S. Eyerman, M. Kodiyath, I. Hur, F. Petrini, and G. Karypis, "Exploring optimizations on shared-memory platforms for parallel triangle counting algorithms," in *IEEE High Performance Extreme Computing Conference (HPEC)*, 2017.

54 **N. K. Ahmed**, T. L. Willke, and R. A. Rossi, "Estimation of local subgraph counts," in *2016 IEEE International Conference on Big Data (Big Data)*, IEEE, 2016, pp. 586–595.

55 **N. K. Ahmed**, T. L. Willke, and R. A. Rossi, "Exact and estimation of local edge-centric graphlet counts," in *Workshop on Big Data, Streams and Heterogeneous Source Mining: Algorithms, Systems, Programming Models and Applications*, 2016, pp. 1–17.

56 **N. K. Ahmed**, J. Neville, R. A. Rossi, and N. Duffield, "Efficient graphlet counting for large networks," in *IEEE International Conference on Data Mining (ICDM)*, IEEE, 2015, pp. 1–10.

57 **N. K. Ahmed** and R. A. Rossi, "Interactive visual graph analytics on the web," in *Ninth International AAAI Conference on Web and Social Media (ICWSM)*, 2015, pp. 566–569.

58 R. Rossi and **N. K. Ahmed**, "The network data repository with interactive graph analytics and visualization," in *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

59 **N. K. Ahmed**, N. Duffield, J. Neville, and R. Kompella, "Graph sample and hold: A framework for big-graph analytics," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, ACM, 2014, pp. 1446–1455.

60 **N. K. Ahmed**, J. Neville, and R. Kompella, "Network sampling designs for relational classification," in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 6, 2012.

61 **N. K. Ahmed**, J. Neville, and R. Kompella, "Space-efficient sampling from social activity streams," in *Proceedings of the 1st international workshop on big data, streams and heterogeneous source mining: algorithms, systems, programming models and applications*, ACM, 2012, pp. 53–60.

62 A. Farahat, **N. K. Ahmed**, and U. Dholakia, "Does a daily deal promotion signal a distressed business? an empirical investigation of small business survival," in *ACM Conference on Web Search and Data Mining (WSDM)*, 2012.

63 **N. K. Ahmed**, J. Neville, and R. R. Kompella, "Network sampling via edge-based node selection with graph induction," in *Purdue University Tech Report 11-016*, 2011.

64 **N. K. Ahmed**, F. Berchmans, J. Neville, and R. Kompella, "Time-based sampling of social network activity graphs," in *Proceedings of the eighth workshop on mining and learning with graphs*, ACM, 2010, pp. 1–9.

65 **N. K. Ahmed**, J. Neville, and R. Kompella, "Reconsidering the foundations of network sampling," in *Proceedings of the 2nd Workshop on Information in Networks*, 2010.

66 **N. K. Ahmed**, A. F. Atiya, N. El Gayar, and H. El-Shishiny, "A combined neural network/gaussian process regression time series forecasting system for the nn3 competition," in *NN3 Neural Network Competition*, 2007.

67 **N. K. Ahmed**, N. Hamdy, and S. Ahmed, "A proposed intrusion detection system for encrypted computer networks," in *Third International Conference on Informatics and Systems*, 2005, pp. 19–22.

Recent Preprints

1 L. Chen, **N. K. Ahmed**, A. Dutta, A. Bhattacharjee, S. Yu, Q. I. Mahmud, W. Abebe, H. Phan, A. Sarkar, B. Butler, *et al.*, *The landscape and challenges of HPC research and LLMs*, arXiv:2402.02018, 2024.

2 K. Cheng, **N. K. Ahmed**, T. Willke, and Y. Sun, *Structure guided prompt: Instructing large language model in multi-step reasoning by exploring graph structure of the text*, arXiv:2402.13415, 2024.

3 S. Duan, H. Ping, N. Kanakaris, X. Xiao, P. Zhang, P. Kyriakis, **N. K. Ahmed**, G. Ma, M. Capota, S. Nazarian, *et al.*, *A structure-aware framework for learning device placements on computation graphs*, arXiv:2405.14185, 2024.

4 L. Chen, A. Bhattacharjee, **N. K. Ahmed**, N. Hasabnis, G. Oren, B. Lei, and A. Jannesari, *Compcodevet: A compiler-guided validation and enhancement approach for code dataset*, arXiv:2311.06505, 2023.

5 S. Duan, N. Kanakaris, X. Xiao, H. Ping, C. Zhou, **N. K. Ahmed**, G. Ma, M. Capota, T. L. Willke, S. Nazarian, *et al.*, *Leveraging reinforcement learning and large language models for code optimization*, arXiv:2312.05657, 2023.

6 T. Kadosh, N. Hasabnis, V. A. Vo, N. Schneider, N. Krien, M. Capota, A. Wasay, **N. K. Ahmed**, T. Willke, G. Tamir, *et al.*, *Domain-specific code language models: Unraveling the potential for hpc codes and tasks*, arXiv:2312.13322, 2023.

7 Q. I. Mahmud, A. TehraniJamsaz, H. D. Phan, **N. K. Ahmed**, and A. Jannesari, *Autoparllm: Gnn-guided automatic code parallelization using large language models*, arXiv:2310.04047, 2023.

8 R. Aponte, R. A. Rossi, S. Guo, J. Hoffswell, N. Lipka, C. Xiao, G. Chan, E. Koh, and **N. K. Ahmed**, *A hypergraph neural network framework for learning hyperedge-dependent node embeddings*, arXiv:2212.14077, 2022.

9 A. Chen, R. Rossi, N. Lipka, J. Hoffswell, G. Chan, S. Guo, E. Koh, S. Kim, and **N. K. Ahmed**, *Graph learning with localized neighborhood fairness*, arXiv:2212.12040, 2022.

10 S. Aananthakrishnan, **N. K. Ahmed**, V. Cave, M. Cintra, Y. Demir, K. D. Bois, S. Eyerman, J. B. Fryman, I. Ganev, W. Heirman, *et al.*, *PIUMA: Programmable integrated unified memory architecture*, arXiv:2010.06277, 2020.

11 A. Haj-Ali, **N. K. Ahmed**, T. Willke, J. E. Gonzalez, K. Asanovic, and I. Stoica, *A view on deep reinforcement learning in system optimization*, arXiv:1908.01275, 2019.

Tutorials

1. M. Al Hasan, **N. K. Ahmed**, and J. Neville, "Methods and applications of network sampling," in *Tutorial at SIAM International Conference on Data Mining (SDM)*, 2015.
2. M. Al Hasan, **N. K. Ahmed**, and J. Neville, "Methods and applications of network sampling," in *Tutorial at IEEE International Conference on Data Mining (ICDM)*, 2013.
3. M. Al Hasan, **N. K. Ahmed**, and J. Neville, "Network sampling: Methods and applications," in *Tutorial at ACM SIGKDD International Conference on Knowl- edge Discovery and Data Mining (SIGKDD)*, 2013.

Patents

1. M. Tepper, B. Keller, M. Capota, V. Vo, **N. K. Ahmed**, and T. Willke, "System to analyze and enhance software based on graph attention networks," US11640295B2, Intel Corp, May 2023. URL: https://patents.google.com/patent/US11640295B2/, Granted.
2. M. Capota, G. MA, S. Zhou, N. Hasabnis, and **N. K. Ahmed**, "Concept for placing an execution of a computer program," US20220107793A1, Intel Corp, Apr. 2022. URL: https://patents.google.com/patent/US20220107793A1/, Published.
3. **N. K. Ahmed**, I. J. Alvarez, R. Balakrishnan, H. Mostafa, G. Raffa, and N. Himayat, "Personalized mobility as a service," US20210108939A1, Intel Corp, Apr. 2021. URL: https://patents.google.com/patent/US20210108939A1/, Published.
4. **N. K. Ahmed**, D. Sengupta, T. Anderson, and T. Willke, "Similarity search using guided reinforcement learning," US20200327118A1, Intel Corp, Oct. 2020. URL: https://patents.google.com/patent/US20200327118A1/, Published.
5. G. Ma, N. Beckage, **N. K. Ahmed**, and I. Alvarez, "Technology to apply driving norms for automated vehicle behavior prediction," US20200324794A1, Intel Corp, Oct. 2020. URL: https://patents.google.com/patent/%7BUS20200324794A1%7D/, Published.
6. A. O. Farahat and **N. K. Ahmed**, "Mining semi-structured social media," US9002852B2, Adobe Systems, Apr. 2015. URL: https://patents.google.com/patent/US9002852B2/, Granted.
7. A. O. Farahat and **N. K. Ahmed**, "Predictive tool utilizing correlations with unmeasured factors influencing observed marketing activities," US20140136280A1, Adobe Systems, May 2014. URL: https://patents.google.com/patent/%7BUS20140136280A1%7D/, Published.

Recent Keynote and Invited Talks

Feb 2023	How to Leverage Network Structure to Improve Reasoning and Prediction — Invited talk at AAAI workshop on Graphs and more Complex structures for Learning and Reasoning (GCLR).
Apr 2023	AI-based Compiler Optimization using Reinforcement Learning — Invited talk at Intel oneAPI Forum and Fellow Summit.
Sep 2022	AI for systems and Beyond — Invited talk at the MIT College of Computing, Breakthrough tech AI initiative.
Dec 2022	Graph Learning and Systems: New Opportunities and Challenges — Keynote at IEEE big data workshop on High Performance Big Graph Data Management, Analysis, and Mining.
Aug 2022	Graph Representation Learning: Opportunities and Challenges — Guest Lecture at the Technical University of Denmark (DTU), Advanced Topics in Machine Learning Course.
Apr 2022	Structured Representation Learning and Reasoning —- Invited talk at Intel Labs Staff forum.
Aug 2019	Statistical Network Models: How to Leverage Network Structure to Improve Estimation and Prediction — Keynote at the 2nd SIGKDD Workshop on on Offline and Online Evaluation of Interactive Systems.

	Role2vec: Role-based Network Embeddings — Presentation at the SIGKDD Workshop on Deep Learning for Graphs (DLG).
Feb 2019	Representation Learning in Large Attributed and Dynamic Graphs — Invited talk at SIAM Conference on Computational Science and Engineering.
Jun 2018	Sampling from Massive Graph Streams: A Unifying Framework — Invited talk at Dagstuhl Seminar on High-Performance Graph Algorithms, Germany
May 2018	The power of motif counting: Theory, Algorithms, and Applications for Large Graphs — Keynote at the IPDPS workshop on the Intersection of Graph Algorithms and Machine Learning (GraPL)
Sep 2018	High-Performance Graph Analytics — Invited talk at The Southern Data Science Conference, Atlanta, GA.

Outreach

2024	Invited Mentor at the University of Santa Cruz Capstone Program - NLP Masters students
2023	Invited Mentor at the Break Through Tech AI initiave hosted by MIT college of computing to bridge the talent gap for women and underrepresented genders
2021	Invited Mentor at the Intel Speed Mentoring for Summer Interns
	Invited Mentor at the AAAI Women Mentoring Session
	Invited Panelist at the IPDPS Career Panel
	Invited Mentor at the Intel MentorNet Program for STEM Students
2019	Invited Panelist at the Workshop on Irregular Applications: Architecture and Algorithms at Super Computing
	Invited Mentor, First Open NetSci Hackathon at NetSci
2018	Invited Mentor, Women in Machine Learning at NIPS
	Circle Leader, Intel Corp - Pay It Forward Program, mentorship and development opportunities for women
2017	AI Track Committee Member, Grace Hopper Conference of Women in Computing (GHC17)
	Program Committee Member, Workshop for Women in Machine Learning at NIPS (WiML)
2016	Invited Panelist/Mentor at Stanford Artificial Intelligence Laboratory's Outreach Summer Program for Rising 10th Grade Girls (SAILORS), Stanford University - Goal: Expose high school students in underrepresented populations to the field of Artificial Intelligence
	Invited Mentor at the Women in Machine Learning Breakfast Event, at the AAAI conference on Artificial Intelligence, Phoenix, AZ
	Invited Panelist, Machine Learning Panel, Intel Software Professionals Conference (SWPC), Intel Corp.
2014	Invited Mentor at Purdue University Graduate Mentoring Program - Goal: Assist new incoming CS graduate students with their transition into graduate school and help them in developing their social and academic networks

Honors and Awards

2023	Featured research by Montreal AI Ethics Institute — Bias and Fairness in Large Language Models.
2022	Intel Labs Recognition Award for Outstanding Contributions.

2020	Senior member, Institute of Electrical and Electronics Engineers (IEEE).
	Intel Labs Recognition Award for Outstanding Contributions.
	Intel Labs Recognition Inventor Award.
2019	Intel Labs Recognition Award for Outstanding Contributions.
2018	Nominated for ACM Practitioners Board.
2017	Finalist award in Traingle Counting — IEEE/Amazon/DARPA Graph Challenge for High-performance Graph Analytics, IEEE High Performance Extreme Computing Conference.
	Finalist award in Truss Decomposition — IEEE/Amazon/DARPA Graph Challenge for High-performance Graph Analytics, IEEE High Performance Extreme Computing Conference.
2015	Best Paper Candidate, IEEE conference on data mining (ICDM)
	Travel award to attend CRA-W women career workshop, Computing Research Association for Women.
2014	Rising Stars Award from the University of California Berkeley — Awarded to top female Ph.D. candidates and postdocs in electrical engineering and computer science.
2024	Featured research by MIT Technology Review — Groupon, Daily Deals and the Complex Question of Business Failure.
2008 - 2015	Graduate Research Assistantship, Purdue University, Computer Science.
2005–2008	Research fellowship in data mining and machine learning — United Nations Development Program (UNDP) — Awarded to top egyptian graduate students.
1997–2001	Cairo University fellowship, Awarded to outstanding undergraduate students.
2013	Intel Ph.D. fellowship nomination by Purdue University for research in algorithms for big graphs.
2012	Best Paper Runner-Up Award, SIGKDD Workshop on big data, streams and source mining (Bigmine)
	WiML Travel award for the women in machine learning workshop at Neurips
2011	Purdue Travel award to attend the Grace Hopper celebration of women in computing.
2010	SIGKDD Travel award to attend the SIGKDD conference on knowledge discovery and data mining.
2007	Ranking 5th in the Artificial Neural Network and Computational Intelligence Forecasting Competition (NN3) for timeseries prediction using ML, organized by SAS and International Institute of Forecasters.
2001	Ranking 2nd among over 500 students , Faculty of Computer Science and Information, Cairo University.

Research Awards & Funding

2022-2024	Machine Learning or Code Generation. Intel Labs, Total: $180,000. PI: Nesreen Ahmed.
2023-2025	Device Placement Optimization using AI Methods. Intel Labs, Total: $225,000. PI: Nesreen Ahmed.
2016	Approximation Methods for Massive Graph Analytics. Intel Labs, Total: $30,000. PI: Nesreen Ahmed.
2022	GNN-based Code Representation and Compiler Optimization. Intel Labs, Total: $30,000. PI: Nesreen Ahmed.
2020	Programmable Graph-based Learning for Heterogeneous Computing Systems. Intel Labs, Total: $10,000. PI: Nesreen Ahmed.

2018-2023	NSF, Collaborative Research: Framework: Software: CINES: A Scalable Cyberinfras- tructure for Sustained Innovation in Network Engineering and Science, PIs: Jure Leskovec (Stanford), Geoffrey Fox (IU), and Madhav Marathe (VT), Total: $4M.
2020-2022	NSF, EAGER: Adaptive Sampling of Massive Graph Streams, PI: Nick Duffield, TAMU, Total: $200,000

Professional Service

Conference/Workshop Organization

2024	Workshop Co-Chair and Organizer — Machine Learning on Graphs (MLoG), co-located with the ACM International Conference on Web Search and Data Mining (WSDM).
2023	Seminar Organizer — Dagstuhl Seminar on Scalable Graph Mining and Learning.
	PhD Consortium Co-Chair — ACM conference on knowledge discovery and data mining (SIGKDD).
	PC Co-Chair — Workshop on Irregular Applications: Architectures and Algorithms, co-located with SuperComputing
2022	Demo track Co-Chair — ACM International Conference on Information and Knowledge Management (CIKM).
2021-2024	Workshop Co-Chair and Organizer — Graphs, Architectures, Programming, and Learning (GrAPL), co-located with International Parallel and Distributed Processing Symposium (IPDPS).
2019	Sponsorship Chair — SIAM International Conference on Data Mining (SDM).
2016-2023	Workshop Co-Chair and Organizer — Big Graphs Workshop on High Performance Big Graph Data Management, Analysis, and Mining, co-located with IEEE Big Data.
2018	Program Committe Chair — IEEE Big Data Conference, Industry and Government Track.

Journal Editorial Board

2020-2021	Associate Editor, IEEE Transactions on Neural Networks and Learning Systems (TNNLS).
2019-2020	Associate Editor, IEEE Signal Processing Magazine.
2018-Present	Member, Frontiers in Machine Learning and Artificial Intelligence.

Senior Program Committee

2021-2024	ACM Conference on Knowledge Discovery and Data Mining (SIGKDD).
2023	National Conference on Artificial Intelligence (AAAI).

TPC and Journal Reviewing

2008-present	Neurips, ICML, ICLR, SIGKDD, WWW, AAAI, IJCAI, ICPP, HiPC, VLDB, CIKM, DSAA.
	JMLR, TKDD, TKDE, TODS, TPDS, VLDB, DMKD, KAIS, TNNLS.

Government Panels and Other Committees

2020-present	Intel proposal funding for AI research.
2024	Intel publication award committee for AI research.
2021	NSF Proposals Review/Panel.
	SIAM nomination committee for data science.

Open Source Software & Datasets

Graph Data Repository	Network Repository is the first and largest network data repository with interactive analytics and visualization, with 500+ publicly available graph datasets, and a tool enabling graph query and search. `https://networkrepository.com`
Interactive Graph Visualization	GraphVis is a tool for interactive synthetic graph generation and visualization. `https://networkrepository.com/graphvis.php`
Open Source Graph Library	Parallel Graphlet Decomposition Library (PGD) is a library for massive graph decomposition. `http://graphlets.org`
Open Source RL for vectorization	NeuroVectorizer is a framework using deep reinforcement learning (RL) to predict optimal vectorization parameters for compiler pragmas in C and C++ for loops. `https://github.com/intel/neuro-vectorizer`
LLVM Code Graph Representation	PerfoGraph is a tool that constructs a language agnostic LLVM-based graph representation of code for program synthesis using machine learning. PerfoGraph representations are numerically aware, contain composite data structure information, and properly present variables. `https://github.com/tehranixyz/perfograph`

Skills

Languages	Strong reading, writing and speaking competencies for English and Arabic.
Coding	Python, MATLAB, C/C++, OpenMP, PHP, JavaScript
Misc.	Academic research, algorithm development, teaching, conference/journal publishing.