

NEIL SHAH

Director of Research, Senior Principal Scientist
Snap Inc.
Seattle, WA

Email: nshah171[at]gmail[dot]com
Last Update: April 30, 2026

EDUCATION

- **PhD:** Computer Science, Carnegie Mellon University, August 2013 - October 2017
Advisor: Prof. Christos Faloutsos. Degree received: December 20, 2017.
- **MS:** Computer Science, Carnegie Mellon University, August 2013 - May 2017
- **BS:** Computer Science (Minor in Mathematics), North Carolina State University, August 2010 - May 2013
GPA: **4.0** (class rank #1), *Summa Cum Laude with Honors*.

POSITIONS

- **Snap Inc.**, Director of Research, Senior Principal Scientist.
Lead a horizontal team of researchers and engineers on academic, applied research, and engineering initiatives in user modeling and personalization, with applications to growth, content, ads, and safety.
December 2017 - present
- **Carnegie Mellon University**, Graduate Researcher.
Computer Science Department. Worked on algorithms and applications for anomaly detection in large social graphs.
August 2013 - October 2017
- **Twitch**, Visiting Researcher.
Worked on anti-abuse technologies as a member of the Science team.
January 2016 - May 2016
- **Microsoft Research**, Research Intern.
Improved metrics and methods for measuring research impact for Microsoft Academic Search.
June 2015 - August 2015
- **Lawrence Livermore National Laboratory**, Research Intern.
Developed algorithms to automatically identify patterns and anomalies in time-evolving graphs.
June 2014 - August 2014
- **IBM Silicon Valley Lab**, Software Intern.
Worked in the IBM BigInsights group on indexing and analytics of system log data.
May 2012 - January 2013
- **North Carolina State University**, Undergraduate Researcher.
Department of Computer Science. Worked on compressing and indexing large scientific datasets.
June 2009 - April 2013

AWARDS & DISTINCTIONS

- ACM CIKM Best Paper Award, 2025
- North Carolina State University Department of Computer Science Rising Star Award, 2023
- ACM WSDM Outstanding Service Award, 2022
- ACM SIGCHI Best Research Paper Honorable Mention Award, 2019
- Symantec Graduate Research Fellowship Finalist, 2017
- ACM SIGKDD Best Research Paper Award, 2016
- National Science Foundation Graduate Research Fellowship, 2013
- North Carolina State University College of Engineering Senior Award for Scholarly Achievement, 2013

- North Carolina State University Department of Computer Science Senior Faculty Scholar, 2012
- National Science Foundation Research Experience for Undergraduates Grant, 2011
- North Carolina State University College of Engineering Dean’s Research Assistantship, 2011
- North Carolina State University Caldwell Fellowship, 2011
- Coca-Cola Scholarship, 2010
- Zinch Scholarship, 2010
- National Merit Scholarship, 2010
- CompTIA Information Technology Merit Award, 2010
- 2nd place, National Siemens Competition in Math, Science and Technology, 2009
- 1st place, Regional Siemens Competition in Math, Science and Technology, 2009

PUBLICATIONS

Conference Publications

91. X. Ding, X. Huang, C. Ju, L. Collins, Y. Liu, L. Akoglu, N. Shah, T. Zhao. *Hierarchical Token Prepending: Enhancing Information Flow in Decoder-based LLM Embeddings*. **ACL 2026**.
90. X. Huang, X. Ding, M. Ju, Y. Liu, N. Shah, T. Zhao. *Threshold Differential Attention for Sink-Free, Ultra-Sparse, and Non-Dispersive Language Modeling*. **ACL 2026**.
89. J. Zhu, M. Ju, Y. Liu, S. Vij, D. Koutra, N. Shah, T. Zhao. *Beyond Unimodal Perspectives: Generative Retrieval with Multimodal Semantics*. **SIGIR 2026**.
88. C. Ju, T. Zhao, L. Neves, L. Collins, B. Kumar, et al. *Semantic IDs for Recommender Systems at Snapchat: Use Cases, Technical Challenges, and Design Choices*. **SIGIR 2026**.
87. Q. Truong, Y. Song, D. Loveland, C. Ju, T. Zhao, N. Shah, J. Tang. *Plain Transformers are Surprisingly Powerful Link Predictors*. **ICML 2026**.
86. G. Lee, B. Kumar, C. Ju, T. Zhao, K. Shin, N. Shah, L. Collins. *Sequential Data Augmentation for Generative Recommendation*. **WSDM 2026**.
85. Z. Zheng, Y. Zhu, H. Liu, C. Ju, T. Zhao, N. Shah, J. Li. *Pretrained Language Model based Cold-Start Recommendation with Meta-Item Embeddings*. **CIKM 2025**.
84. C. Ju, L. Collins, L. Neves, B. Kumar, L. Wang, T. Zhao, N. Shah. *Generative Recommendation with Semantic IDs: A Practitioner’s Handbook*. **CIKM 2025. Best Paper Award**.
83. Q. Truong, Z. Chen, C. Ju, T. Zhao, N. Shah, J. Tang. *A Pre-Training Framework for Relational Data with Information Theoretic Principles*. **NeurIPS 2025**.
82. D. Loveland, M. Ju, T. Zhao, N. Shah, D. Koutra. *On the Role of Weight Decay in Collaborative Filtering: A Popularity Perspective*. **KDD 2025**.
81. T. Zhao, Y. Liu, M. Kolodner, K. Montemayor, E. Ghazizadeh, et al. *GiGL: Large-Scale Graph Neural Networks at Snapchat*. **KDD 2025**.
80. M. Ju, L. Neves, B. Kumar, L. Collins, T. Zhao, et al. *Revisiting Self-Attention for Cross-Domain Sequential Recommendation*. **KDD 2025**.
79. M. Ju, L. Neves, B. Kumar, L. Collins, T. Zhao, et al. *Learning Universal User Representations Leveraging Cross-domain User Intent at Snapchat*. **SIGIR 2025**.
78. R. Xue, T. Zhao, N. Shah, X. Liu. *Haste Makes Waste: A Simple Approach for Scaling Graph Neural Networks*. **ICML 2025**.
77. N. Bui, M. Yang, R. Chen, L. Neves, C. Ju, R. Ying, N. Shah, T. Zhao. *Learning Along the Arrow of Time: Hyperbolic Geometry for Backward-Compatible Representation Learning*. **ICML 2025**.
76. D. Loveland, X. Wu, T. Zhao, D. Koutra, N. Shah, M. Ju. *Understanding and Scaling Collaborative Filtering Optimization from the Perspective of Matrix Rank*. **WWW 2025**.
75. X. Wu, D. Loveland, R. Chen, Y. Liu, X. Chen, et al. *GraphHash: Graph Clustering Enables Parameter Efficiency in Recommender Systems*. **WWW 2025**.
74. J. Zhu, Y. Zhou, S. Qian, Z. He, T. Zhao, N. Shah, D. Koutra. *Mosaic of Modalities: A Comprehensive Benchmark for Multimodal Graph Learning*. **CVPR 2025**.

73. J. Liu, H. Mao, Z. Chen, T. Zhao, N. Shah, J. Tang. *Towards Neural Scaling Laws on Graphs*. **LoG 2024**.
72. M. Ju, W. Shiao, Z. Guo, Y. Ye, Y. Liu, N. Shah, T. Zhao. *Test-time Aggregation For Collaborative Filtering*. **NeurIPS 2024**.
71. P. Kung, Z. Fan, T. Zhao, Y. Liu, Z. Lai, J. Shi, Y. Wu, J. Yu, N. Shah, G. Venkataraman. *Improving Embedding-Based Retrieval in Friend Recommendation with ANN Query Expansion*. **SIGIR 2024**.
70. R. Chen, T. Zhao, A. K. Jaiswal, N. Shah, Z. Wang. *LLaGA: Large Language and Graph Assistant*. **ICML 2024**.
69. H. Mao, Z. Chen, W. Tang, J. Zhao, Y. Ma, T. Zhao, N. Shah, M. Galkin, J. Tang. *Graph Foundation Models*. **ICML 2024**.
68. A. Calabrese, L. Neves, N. Shah, M. Bos, B. Ross, M. Lapata, F. Barbieri. *Explainability and Hate Speech: Structured Explanations Make Social Media Moderators Faster*. **ACL 2024**.
67. H. Mao, J. Li, H. Shomer, B. Li, W. Fan, Y. Ma, T. Zhao, N. Shah, J. Tang. *Revisiting Link Prediction: A Data Perspective*. **ICLR 2024**.
66. Y. Wang, T. Zhao, Y. Zhao, Y. Liu, X. Cheng, N. Shah, T. Derr. *A Topological Perspective on Demystifying GNN-based Link Prediction Performance*. **ICLR 2024**.
65. T. Zhao, N. Shah, E. Ghazizadeh. *Learning from Graphs without Explicit Graph Machine Learning Models*. **ICLR 2024**.
64. J. Li, H. Shomer, H. Mao, S. Zeng, Y. Ma, N. Shah, J. Tang, D. Yin. *Evaluating Graph Neural Networks for Link Prediction: Current Pitfalls and New Benchmarking*. **NeurIPS 2023**.
63. H. Mao, Z. Chen, W. Jin, H. Han, Y. Ma, T. Zhao, N. Shah, J. Tang. *Demystifying Structural Disparity in Graph Neural Networks: Can One Size Fit All?*. **NeurIPS 2023**.
62. M. Ju, T. Zhao, W. Yu, N. Shah, F. Ye. *GraphPatcher: Mitigating Degree Bias for Graph Neural Networks via Test-time Node Patching*. **NeurIPS 2023**.
61. W. Shiao, U. S. Saini, Y. Liu, T. Zhao, N. Shah, E. Papalexakis. *CARL-G: Clustering-Accelerated Representation Learning on Graphs*. **KDD 2023**.
60. S. Bhatia, M. Wadhwa, K. Kawaguchi, N. Shah, P. S. Yu, B. Hooi. *Sketch-based Anomaly Detection in Streaming Graphs*. **KDD 2023**.
59. Z. Guo, W. Shiao, S. Zhang, Y. Liu, N. Chawla, N. Shah, T. Zhao. *Linkless Link Prediction via Relational Distillation*. **ICML 2023**.
58. J. Shi, V. Chaurasiya, Y. Liu, S. Vij, Y. Wu, S. Kanduri, N. Shah, P. Yu, N. Srivastava, L. Shi, G. Venkataraman, J. Yu. *Embedding-based Retrieval in Friend Recommendation*. **SIGIR 2023**.
57. J. Li, H. Shomer, J. Ding, Y. Wang, Y. Ma, N. Shah, J. Tang, D. Yin. *Are Message Passing Neural Networks Really Helpful for Knowledge Graph Completion?*. **ACL 2023**.
56. X. Han, T. Zhao, Y. Liu, X. Hu, N. Shah. *MLPInit: Embarrassingly Simple GNN Training Acceleration with MLP Initialization*. **ICLR 2023**.
55. M. Ju, T. Zhao, Q. Wen, W. Yu, N. Shah, Y. Ye, C. Zhang. *Multi-task Self-supervised Graph Neural Networks Enable Stronger Task Generalization*. **ICLR 2023**.
54. W. Shiao, Z. Guo, T. Zhao, V. Papalexakis, Y. Liu, N. Shah. *Link Prediction with Non-Contrastive Learning*. **ICLR 2023**.
53. W. Jin, T. Zhao, J. Ding, Y. Liu, J. Tang, N. Shah. *Empowering Graph Representation Learning with Test-Time Graph Transformation*. **ICLR 2023**.
52. Y. Wang, B. Hooi, Y. Liu, T. Zhao, Z. Guo, N. Shah. *Flashlight: Scalable Link Prediction with Effective Decoders*. **LoG 2023**.
51. Y. Wang, B. Hooi, Y. Liu, N. Shah. *Graph Explicit Neural Networks: Explicitly Encoding Graphs for Efficient and Accurate Inference*. **WSDM 2023**.
50. R. Baten, Y. Liu, H. Peters, F. Barbieri, N. Shah, L. Neves, M. Bos. *Predicting Future Location Categories of Users in a Large Social Platform*. **ICWSM 2023**.
49. S. Zhang, Y. Liu, N. Shah, Y. Sun. *Explaining Graph Neural Networks with Structure-Aware Cooperative Games*. **NeurIPS 2022**.
48. L. Zhao, L. Härtel, N. Shah, L. Akoglu. *A Practical, Progressively Expressive Graph Neural Network*. **NeurIPS 2022**.
47. Y. Wang, Y. Zhao, N. Shah, T. Derr. *Imbalanced Graph Classification via Graph-of-Graph Neural Networks*. **CIKM 2022**.

46. J. Jiang, N. Murrugarra-Llerena, M. Bos, Y. Liu, N. Shah, L. Neves, F. Barbieri. *Sunshine with a Chance of Smiles: How does Weather Impact Sentiment on Social Media?*. **ICWSM 2022**.
45. S. Zhang, Y. Liu, Y. Sun, N. Shah. *Graph-less Neural Networks: Teaching Old MLPs New Tricks via Distillation*. **ICLR 2022**.
44. W. Jin, L. Zhao, S. Zhang, Y. Liu, J. Tang, N. Shah. *Graph Condensation for Graph Neural Networks*. **ICLR 2022**.
43. L. Zhao, W. Jin, L. Akoglu, N. Shah. *From Stars to Subgraphs: Uplifting Any GNN with Local Structure Awareness*. **ICLR 2022**.
42. Y. Ma, X. Liu, N. Shah, J. Tang. *Is Homophily a Necessity for Graph Neural Networks?*. **ICLR 2022**.
41. W. Jin, X. Liu, X. Zhao, Y. Ma, N. Shah, J. Tang. *Automated Self-Supervised Learning for Graphs*. **ICLR 2022**.
40. X. Tang, Y. Liu, X. He, S. Wang, N. Shah. *Ranking Friend Stories on Social Platforms with Edge-Contextual Local Graph Convolutions*. **WSDM 2022**.
39. S. Sikdar, N. Shah, T. Weninger. *Attributed Graph Modeling with Vertex Replacement Grammars*. **WSDM 2022**.
38. H. Shin, T. Kwon, N. Shah, K. Shin. *Finding a Concise, Precise and Exhaustive Set of Near Bi-Cliques in Dynamic Graphs*. **WSDM 2022**.
37. T. Zhao, B. Ni, W. Yu, Z. Guo, N. Shah, M. Jiang. *Action Sequence Augmentation for Early Graph-based Anomaly Detection*. **CIKM 2021**.
36. Y. Ma, X. Liu, T. Zhao, Y. Liu, J. Tang, N. Shah. *A Unified View on Graph Neural Networks as Graph Signal Denoising*. **CIKM 2021**.
35. E. Gujral, L. Neves, E. Papalexakis, N. Shah. *Niche Detection in User Content Consumption Data*. **CIKM 2021**.
34. S. Shekhar, N. Shah, L. Akoglu. *FairOD: Fairness-aware Outlier Detection*. **AIES 2021**.
33. Q. Yang, W. Wang, L. Pierce, R. Vaish, X. Shi, N. Shah. *Online Communication Shifts in the Midst of the Covid-19 Pandemic: A Case Study on Snapchat*. **ICWSM 2021**.
32. F. A. Chowdhury, Y. Liu, K. Saha, N. Vincent, L. Neves, N. Shah, M. Bos. *CEAM: The Effectiveness of Cyclic and Ephemeral Attention Models of User Behavior on Social Platforms*. **ICWSM 2021**.
31. A. Sankar, Y. Liu, J. Yu, N. Shah. *Graph Neural Networks for Friend Ranking in Large-scale Social Platforms*. **WWW 2021**.
30. K. Saha, Y. Liu, N. Vincent, F. A. Chowdhury, L. Neves, N. Shah, M. Bos. *AdverTiming Matters: Examining User Ad Consumption for Effective Ad Allocations on Social Media*. **CHI 2021**.
29. T. Zhao, Y. Liu, L. Neves, O. Woodford, M. Jiang, N. Shah. *Data Augmentation for Graph Neural Networks*. **AAAI 2021**.
28. B. Joshi, F. Barbieri, N. Shah, L. Neves. *The Devil is in the Details: Evaluating Limitations of Transformer-based Methods for Granular Tasks*. **COLING 2020**.
27. P. Kaghazgaran, M. Bos, L. Neves, N. Shah. *Social Factors in Closed-Network Content Consumption*. **CIKM 2020**.
26. S. Abdali, R. Gurav, S. Menon, D. Fonseca, N. Entezari, N. Shah, E. Papalexakis. *Identifying Misinformation from Website Screenshots*. **ICWSM 2021**.
25. S. Abdali, N. Shah, E. Papalexakis. *Semi-Supervised Multi-aspect Misinformation Detection with Hierarchical Joint Decomposition*. **ECML-PKDD 2020**.
24. X. Tang, Y. Liu, N. Shah, X. Shi, P. Mitra, S. Wang. *Knowing your FATE: Friendship, Action and Temporal Explanations for User Engagement Prediction on Social Apps*. **KDD 2020**.
23. N. Shah. *FARE: Schema-Agnostic Anomaly Detection in Social Event Logs*. **DSAA 2019**.
22. H. Nilforoshan, N. Shah. *SliceNDice: Mining Suspicious Multi-attribute Entity Groups with Multi-view Graphs*. **DSAA 2019**.
21. H. Lamba, N. Shah. *Modeling Dwell Time Engagement on Visual Multimedia*. **KDD 2019**.
20. H. Habib, N. Shah, R. Vaish. *Impact of Contextual Factors on Public Snapchat Sharing*. **CHI 2019. Best Paper Honorable Mention Award**.
19. S. Jain, D. Niranjana, H. Lamba, N. Shah, P. Kumaraguru. *Characterizing and Detecting Livestreaming Chatbots*. **ASONAM 2019**.
18. G. B. Guacho, S. Abdali, N. Shah, E. Papalexakis. *Semi-Supervised Content-based Detection of Misinformation via Tensor Embeddings*. **ASONAM 2018**.
17. N. Gupta, D. Eswaran, N. Shah, L. Akoglu, C. Faloutsos. *Beyond Outlier Detection: LookOut for Pictorial Explanation*. **ECML-PKDD 2018**.

16. N. Shah, H. Lamba, A. Beutel, C. Faloutsos. *The Many Faces of Link Fraud*. **ICDM 2017**.
15. D.-C. Juan, N. Shah, M. Tang, Z. Qian, D. Marculescu, C. Faloutsos. *M3A: Model, MetaModel, and Anomaly Detection in Web Searches*. **DSAA 2017**.
14. N. Shah. *FLOCK: Combating Astroturfing on Livestreaming Platforms*. **WWW 2017**.
13. B. Hooi, H. A. Song, A. Beutel, N. Shah, K. Shin, C. Faloutsos. *FRAUDAR: Bounding Graph Fraud in the Face of Camouflage*. **KDD 2016. Best Paper Award**.
12. B. Hooi, N. Shah, A. Beutel, S. Günnemann, L. Akoglu, M. Kumar, D. Makhija, C. Faloutsos. *BIRDNEST: Bayesian Inference for Ratings-Fraud Detection*. **SDM 2016**.
11. N. Shah, D. Koutra, T. Zou, B. Gallagher, C. Faloutsos. *TimeCrunch: Interpretable Dynamic Graph Summarization*. **KDD 2015**.
10. M. Giatsoglou, D. Chatzakou, N. Shah, A. Beutel, S. Guenneman, C. Faloutsos, A. Vakali. *ND-Sync: Detecting Synchronized Fraud Activities*. **PAKDD 2015**.
9. M. Giatsoglou, D. Chatzakou, N. Shah, C. Faloutsos, A. Vakali. *Retweeting Activity on Twitter: Signs of Fraud*. **PAKDD 2015**.
8. N. Shah, A. Beutel, B. Gallagher, C. Faloutsos. *Spotting Suspicious Link Behavior with fBox: An Adversarial Perspective*. **ICDM 2014**.
7. N. Shah, E. Schendel, S. Pendse, S. Lakshminarasimhan, T. Rogers, N. Samatova. *Improving I/O Throughput with PRIMACY: Preconditioning ID-Mapper for Compressing Incompressibility*. **CLUSTER 2012**.
6. E. Schendel, Y. Jin, N. Shah, J. Chen, C.-S. Chang, S.-H. Ku, S. Ethier, S. Klasky, R. Latham, R. Ross, N. Samatova. *ISOBAR Preconditioner for Effective and High-throughput Lossless Data Compression*. **ICDE 2012**.
5. I. Arkatkar, J. Jenkins, S. Lakshminarasimhan, N. Shah, E. Schendel, S. Ethier, et al. *Analytics-driven Lossless Data Compression for Rapid In-situ Indexing, Storing and Querying*. **DEXA 2012**.
4. Y. Jin, S. Lakshminarasimhan, N. Shah, Z. Gong, C.-S. Chang, J. Chen, et al. *S-preconditioner for Multi-fold Data Reduction with Guaranteed User-controlled Accuracy*. **ICDM 2011**.
3. S. Lakshminarasimhan, N. Shah, S. Ethier, S. Klasky, R. Latham, R. Ross, N. Samatova. *Compressing the Incompressible with ISABELA: In-situ Reduction of Spatio-Temporal Data*. **EuroPar 2011**.
2. N. Shah, Y. Shpanskaya, C.-S. Chang, S.-H. Ku, A. Melechko, N. Samatova. *Automatic and Statistically Robust Spatio-temporal Detection and Tracking of Fusion Plasma Fronts*. **SciDAC 2010**.
1. P. Breimyer, G. Kora, W. Hendrix, N. Shah, N. Samatova. *pR: Automatic Parallelization of Data-parallel Statistical Computing Codes for R in Hybrid Multi-node and Multi-core Environments*. **IADIS 2009**.

Journal Publications

11. Z. Guo, T. Zhao, Y. Liu, K. Dong, W. Shiao, N. Shah, N. Chawla. *Node Duplication Improves Cold-start Link Prediction*. **TMLR 2025**.
10. Y. Dong, W. Shiao, Y. Liu, J. Li, N. Shah, T. Zhao. *Do Graph Neural Networks Improve Node Representation Learning for All?*. **DMLR 2025**.
9. F. Xia, R. Lambiotte, N. Shah, H. Tong, I. King. *Guest Editorial: Special Issue on Graph Learning*. **IEEE TNNLS 2024**.
8. D. Gomez-Zara, Y. Liu, L. Neves, N. Shah, M. Bos. *Unpacking the Exploration-Exploitation Tradeoff on Snapchat*. **Computers in Human Behavior 2023**.
7. T. Zhao, T. Jiang, N. Shah, M. Jiang. *A Synergistic Approach for Graph Anomaly Detection with Pattern Mining and Feature Learning*. **IEEE TNNLS 2021**.
6. Y. Liu, T. Safavi, N. Shah, D. Koutra. *Reducing Large Graphs to Small Supergraphs: A Unified Approach*. **SNAM 2018**.
5. B. Hooi, K. Shin, H. A. Song, A. Beutel, N. Shah, C. Faloutsos. *Graph-based Fraud Detection in the Face of Camouflage*. **TKDD 2017**.
4. N. Shah, D. Koutra, L. Jin, T. Zou, B. Gallagher, C. Faloutsos. *On Summarizing Large-Scale Dynamic Graphs*. **Data Engineering Bulletin 2017**.
3. D. Koutra, N. Shah, J. T. Vogelstein, B. Gallagher, C. Faloutsos. *DeltaCon: A Principled Massive-Graph Similarity Function with Attribution*. **TKDD 2015**.
2. J. Jenkins, I. Arkatkar, S. Lakshminarasimhan, D. Boyuka, E. Schendel, N. Shah, et al. *ALACRITY: Analytics-*

driven Lossless Data Compression for Rapid In-situ Indexing, Storing, and Querying. TLDKS 2013.

1. S. Lakshminarasimhan, N. Shah, S. Ethier, S. Klasky, R. Latham, R. Ross, N. Samatova. *ISABELA for Effective In-situ Compression of Scientific Data. Concurrency and Computation: Practice and Experience 2011.*

Workshop Publications

9. *Improving Out-of-Vocabulary Handling in Recommendation Systems. TheWebConf RelWeb 2025.*
8. M. Kolodner, M. Ju, Z. Fan, T. Zhao, E. Ghazizadeh, Y. Wu, N. Shah, Y. Liu. *Robust Training Objectives Improve Embedding-based Retrieval in Industrial Recommendation Systems. RecSys RobustRecSys 2024.*
7. M. Ju, T. Zhao, W. Yu, N. Shah, Y. Ye. *GraphPatcher: Mitigating Degree Bias for Graph Neural Networks via Test-time Augmentation. TheWebConf DCAI 2024.*
6. P. Kung, Z. Fan, T. Zhao, Y. Liu, L. Lai, J. Shi, Y. Wu, N. Shah, J. Yu. *Improving Embedding-Based Retrieval in Friend Recommendation with ANN Query Expansion. TheWebConf DCAI 2024.*
5. N. Shah. *Scale-Free, Attributed and Class-Assortative Graph Generation to Facilitate Introspection of Graph Neural Networks. KDD MLG 2020.*
4. R. Kumar, M. Kumar, N. Shah, C. Faloutsos. *Did We Get It Right? Predicting Query Performance in E-commerce Search. SIGIR eCom 2018.*
3. Y. Liu, T. Safavi, N. Shah, D. Koutra. *Reducing Million-Node Graphs to a Few Structural Patterns: A Unified Approach. KDD MLG 2016.*
2. N. Shah, A. Beutel, B. Hooi, L. Akoglu, S. Günnemann, D. Makhija, M. Kumar, C. Faloutsos. *EdgeCentric: Anomaly Detection in Edge-Attributed Networks. ICDM DMCS 2016.*
1. Y. Liu, N. Shah, D. Koutra. *An Empirical Comparison of the Summarization Power of Graph Clustering Methods. NIPS NSIS 2015.*

Surveys

3. H. Han, Y. Wang, H. Shomer, K. Guo, J. Ding, Y. Lei, M. Halappanavar, R. A. Rossi, et al. *Retrieval-Augmented Generation with Graphs (GraphRAG). arXiv 2024.*
2. T. Zhao, W. Jin, Y. Liu, Y. Wang, G. Liu, S. Günnemann, N. Shah, M. Jiang. *Graph Data Augmentation for Graph Machine Learning. IEEE Data Engineering Bulletin 2023.*
1. S. Kumar, N. Shah. *False Information on the Web and Social Media. arXiv 2018.*

Tutorials

5. Y. Liu, T. Zhao, M. Kolodner, K. Montemayor, S. Vij, N. Shah. *Training Industry-Scale Graph Neural Networks with GiGL. KDD 2025.*
4. R. Xue, H. Han, T. Zhao, N. Shah, J. Tang, X. Liu. *Large-Scale Graph Neural Networks: The Past and New Frontiers. AAAI 2024.*
3. R. Xue, H. Han, T. Zhao, N. Shah, J. Tang, X. Liu. *Large-Scale Graph Neural Networks: The Past and New Frontiers. SDM 2024.*
2. R. Xue, H. Han, T. Zhao, N. Shah, J. Tang, X. Liu. *Large-Scale Graph Neural Networks: The Past and New Frontiers. KDD 2023.*
1. T. Zhao, K. Ding, W. Jin, G. Liu, M. Jiang, N. Shah. *Augmentation Methods for Graph Learning. SDM 2023.*

Book Chapters

3. S. Abdali, G. Bastidas, N. Shah, E. Papalexakis. *Tensor Embeddings for Content-Based Misinformation Detection with Limited Supervision.* In: *Disinformation, Misinformation, and Fake News in Social Media.*
2. N. Shah. *Introduction to R.* In: *Practical Graph Mining with R.*
1. K. Padmanabhan, S. Lakshminarasimhan, Z. Gong, J. Jenkins, N. Shah, E. Schendel, I. Arkatkar, R. Ross, S. Klasky, N. Samatova. *In-situ Analysis in Support of Exploratory Scientific Discovery in Data-Intensive Science.* In: *Data-Intensive Science.*

ACADEMIC EXPERIENCE

Invited Talks

- Panelist, *AI-Augmented Academia: Research Innovations and Career Planning for Tomorrow's Workforce*, KDD Doctoral Consortium, 2025
- Panelist, *From Lists to Dialogues: Rethinking Personalization with Generative AI*, KDD Workshop on Generative AI for Personalization, 2025
- Panelist, *From Research to Product*, SIGIR Workshop on E-Commerce, 2025
- Panelist, *Is Search Dead? The Rise or Demise of Search in the Era of LLMs*, SIGIR, 2025
- Panelist, UCR exploreCSR Workshop, 2024
- Keynote Speaker, KDD Undergraduate Consortium, 2024
- Panelist, KDD Graph Learning Benchmarks (GLB) Workshop, 2023
- Invited Speaker, Samsung Research, 2023
- Keynote Speaker, RE-WORK AI Summit West, 2023
- Panel Moderator, KDD Misinformation and Misbehavior (MIS²-TrueFact) Workshop, 2022
- Panelist, KDD Deep Learning on Graphs (DLG) Workshop, 2022
- Keynote Speaker, KDD Deep Learning on Graphs (DLG) Workshop, 2022
- Panel Moderator, TigerGraph AI Summit, 2022
- Invited Speaker, WSDM, 2022
- Panelist, The Knowledge Graph Conference, 2022
- Invited Speaker, The Knowledge Graph Conference, 2022
- Invited Speaker, UC Riverside CSE Colloquium, 2022
- Keynote Speaker, WSDM Machine Learning on Graphs (MLoG) Workshop, 2022
- Panelist, UCR exploreCSR Workshop, 2022
- Invited Speaker, Pinterest Trust and Safety Summit, 2021
- Panel Moderator, KDD Outlier Detection and Discovery (ODD) Workshop, 2021
- Keynote Speaker, KDD Machine Learning in Finance (MLF) Workshop, 2021
- Keynote Speaker, SDM Doctoral Consortium, 2021
- Keynote Speaker, SDM Minisymposium on Dynamic Networks, 2020
- Keynote Speaker, ICDM Doctoral Consortium, 2019
- Keynote Speaker, WWW CyberSafety Workshop, 2018
- Keynote Speaker, KDD Outlier Detection De-constructed Workshop, 2018
- Keynote Speaker, ECML-PKDD PhD Forum, 2018

Service

Conference Organization

- Program Chair, CODS-COMAD, 2025
- Sponsorship Chair, SDM, 2025
- Hands-on Tutorial Chair, KDD, 2023, 2024
- Sponsorship Chair, ICWSM, 2023
- Cup Chair, WSDM, 2020, 2022
- Program Chair, ASONAM, 2019
- Organizer, KDD Generative AI for Recommender Systems and Personalization (GenAI4RecP) Workshop, 2025
- Organizer, KDD Federated Learning with Graph Data (FedGraph) Workshop, 2024, 2025
- Organizer, ICDM Mining and Learning on Graphs (MLoG) Workshop, 2022, 2023
- Organizer, KDD Mining and Learning with Graphs (MLG) Workshop, 2022, 2023
- Organizer, CIKM Federated Learning with Graph Data (FedGraph) Workshop, 2022
- Organizer, KDD Misinformation and Misbehavior Mining (MIS²) Workshop, 2021, 2022
- Organizer, TheWebConf CyberSafety Workshop, 2019, 2020
- Session Chair, TheWebConf (Security and Privacy, 2018; Graph Models, 2021)

- Session Chair, KDD (Graph Algorithms, 2020; Graphs and Networks, 2021; Graph Learning and Social Network, 2022)
- Session Chair, DSAA (Subgraphs, 2019)
- Session Chair, ICDM (Social track, 2016)

Peer Review

- Area Chair, NeurIPS, 2025
- Area Chair, ICLR, 2025, 2026
- Area Chair, KDD, 2022, 2025, 2026
- Area Chair, TheWebConf, 2024, 2025, 2026
- Area Chair, LoG, 2022, 2023, 2024, 2025
- Senior Program Committee, AAAI, 2023, 2024, 2025, 2026
- Senior Program Committee, PAKDD, 2023
- Senior Program Committee, SDM, 2022–2025
- Senior Program Committee, WSDM, 2022–2025
- Senior Program Committee, CIKM, 2021–2025
- Early Career Data Mining Award Committee, SDM, 2023, 2024
- Best Paper Award Committee, SDM, 2023, 2024
- Program Committee, ICLR, 2024
- Program Committee, NeurIPS, 2023, 2024
- Program Committee, ICDM, 2022
- Program Committee, ASONAM, 2022
- Program Committee, WSDM, 2019–2021
- Program Committee, KDD, 2019–2021
- Program Committee, TheWebConf, 2015, 2018, 2020–2022
- Program Committee, SDM, 2018–2021
- Program Committee, CIKM, 2017, 2020
- Reviewer, ACM TKDD, 2018–2020; Springer DAMI, 2018, 2019; ACM TSOC, 2018, 2019; ACM TKDE, 2016, 2017
- Reviewer, CSCW, 2019; CHI, 2019
- Reviewer, WISE, 2014; IPDPS, 2011
- Program Committee, numerous workshops (MLG, GLB, HeteroNAM, MIS², GreS, PhD Forums, Demo Sessions) at KDD, WSDM, WWW, ICML, ICDM, RecSys, CIKM

Mentorship

Internships and Student Advisory Roles

- Jingzhe Liu (intern at Snap Research, 2025)
- Geon Lee (intern at Snap Research, 2025)
- Kulin Shah (intern at Snap Research, 2025)
- Xingyue Huang (intern at Snap Research, 2025)
- Xueying Ding (intern at Snap Research, 2025)
- Ngoc Bui (intern at Snap Research, 2024)
- Xinyi Wu (intern at Snap Research, 2024)
- Donald Loveland (intern at Snap Research, 2024)
- Runjin Chen (intern at Snap Research, 2024)
- Jing Zhu (intern at Snap Research, 2024)
- Haitao Mao (intern at Snap Research, 2024)
- Agostina Calabrese (intern at Snap Research, 2023)
- Yushun Dong (intern at Snap Research, 2023)
- Vijay Prakash Dwivedi (intern at Snap Research, 2023)
- Mingxuan Ju (intern at Snap Research, 2023)
- Zhichun Guo (intern at Snap Research, 2022)

- Vedant Bhatia (intern at Snap, 2022)
- Yiwei Wang (intern at Snap Research, 2022)
- William Shiao (intern at Snap Research, 2022 & 2023)
- Xiaotian Han (intern at Snap Research, 2022)
- Cazamere Comrie (intern at Snap Research, 2021)
- Lingxiao Zhao (intern at Snap Research, 2021)
- Wei Jin (intern at Snap Research, 2021 & 2022)
- Shichang Zhang (intern at Snap Research, 2021)
- Yingtong Dou (intern at Snap Research, 2021)
- Yozen Liu (RA at Snap Research, 2020)
- Qi Yang (intern at Snap Research, 2020)
- Satyaki Sikdar (intern at Snap Research, 2020)
- Yao Ma (intern at Snap Research, 2020)
- Aravind Sankar (intern at Snap Research, 2020)
- Tong Zhao (intern at Snap Research, 2020)
- Nicholas Vincent (intern at Snap Research, 2020)
- Farhan Asif Chowdhury (intern at Snap Research, 2020)
- Koustuv Saha (intern at Snap Research, 2020)
- Brihi Joshi (intern at Snap Research, 2019)
- Shiyang Yan (intern at Snap Research, 2019)
- Xianfeng Tang (intern at Snap Research, 2019 & 2020)
- Parisa Kaghazgaran (intern at Snap Research, 2019)
- Himel Dev (intern at Snap Research, 2019)
- Anis Zaman (intern at Snap Research, 2019)
- Can Liu (intern at Snap Research, 2019)
- Dipankar Niranjana (BS student, IIT Delhi, 2018)
- Shreya Jain (BS student, IIT Delhi, 2018)
- Hamed Nilforoshan (intern at Snap Research, 2018)
- Hana Habib (intern at Snap Research, 2018)
- Hemank Lamba (intern at Snap Research, 2018)
- Rohan Kumar (visiting CS student at CMU, 2017)
- Qicheng Huang (EE PhD student at CMU, 2017)
- Chenlei Fang (EE PhD student at CMU, 2017)
- Tianmin Zou (CS MS student at CMU, 2017)

Thesis Supervision

- Committee Member, Haitao Mao, 2024
- Committee Member, Harry Shomer, 2024
- Committee Member, Lingxiao Zhao, 2024
- Committee Member, William Shiao, 2024
- Committee Member, Zhichun Guo, 2023
- Committee Member, Wei Jin, 2023
- Committee Member, Yingtong Dou, 2021
- Committee Member, Aravind Sankar, 2021
- Committee Member, Tong Zhao, 2020

Funding

- Contributed towards Flipkart faculty grant (with CMU: Bryan Hooi, Dhivya Eswaran, Christos Faloutsos)
- Contributed towards Wharton Customer Analytics Initiative proposal *Fraud Detection through Mining Dynamic Behavior for Group Anomalies* (with CMU: Alex Beutel, Christos Faloutsos)
- Contributed towards PNC Center for Financial Services proposal PF15003: *Fraud Detection in Financial Data* (with CMU: Alex Beutel, Christos Faloutsos)

- Contributed towards DOE-NNSA-30788.1.1990222: *Quantifying Network Changes* (with CMU: Danai Koutra, Christos Faloutsos)
- Contributed towards NSF IIS-1028746: *Collaborative Research: Understanding Climate Change — A Data Driven Approach* (with NCSU: Nagiza Samatova, Fredrick Semazzi)

Teaching

- Guest Lecture, *Improving the Scalability of Graph Neural Networks*, CMU 11-741 Machine Learning for Text and Graph-based Mining (Prof. Yiming Yang), 2024
- Seminar Talk, *Improving the Scalability of Graph Neural Networks*, Brandeis University Machine Learning Seminar, 2024
- Guest Lecture, *Improving the Scalability of Graph Neural Networks*, Georgia Tech ISYE 4803 Network Science and Analysis (Prof. Tejas Santanam), 2023
- Guest Lecture, *Improving the Scalability of Graph Neural Networks*, Georgia Tech CSE 6240 Web Search and Text Mining (Prof. Srijan Kumar), 2023
- Guest Lecture, *Improving the Scalability of Graph Neural Networks*, Michigan State CSE 482 Big Data Analysis (Prof. Jiliang Tang), 2022
- Guest Lecture, *Machine Learning on Graphs with Scarce Labels*, RPI MGMT-6560-02 Introduction to Machine Learning Applications (Prof. Lydia Manikonda), 2021
- Guest Lecture, *Mining Misbehavior on Large-Scale Social Platforms*, Vanderbilt CS-5981-06 Social Network Analysis (Prof. Tyler Derr), 2020
- Guest Lecture, *A Foray into Graph Mining*, USC CSCI-699 Introduction to Information Extraction (Prof. Xiang Ren), 2019
- Guest Lecture, *Graph Mining for Fraud Detection*, CMU 15-300 Research and Innovation in Computer Science (Prof. Todd Mowry), 2015
- Teaching Assistant, CMU 15-300 Research and Innovation in Computer Science (Prof. Todd Mowry), 2015
- Teaching Assistant, CMU 15-826 Multimedia Databases and Data Mining (Prof. Christos Faloutsos), 2014

TECHNICAL SKILLS

- *Cloud Platforms*: GCP, AWS
- *Languages/Tools*: Python, Java, C, C++, x86 Assembly
- *Web*: SQL, HTML, PHP, JavaScript, CSS, WordPress
- *Engineering Tools*: MATLAB, R
- *Typesetting*: LaTeX, Microsoft Office, LibreOffice
- *Version Control*: Git, Subversion
- *Operating Systems*: macOS, Ubuntu Linux, Microsoft Windows

REFERENCES

Available upon request.