

# Feng Yan

## **Contact Information**

---

**E-mail:** fyan5@central.uh.edu

**Address:** PGH 589, Department of Computer Science  
3551 Cullen Blvd., Houston, TX 77204-3010

**Homepage:** <https://www2.cs.uh.edu/~fyan/>

## **Professional Experience**

---

**Director** July 2016 - now  
Intelligent Data and Systems Lab (IDS Lab)

**Associate Professor** September 2022 - now  
Department of Computer Science  
Department of Electrical and Computer Engineering  
University of Houston

**Associate Professor** July 2022 - August 2022  
Department of Computer Science & Engineering  
University of Nevada, Reno

**Assistant Professor** July 2016 - June 2022  
Department of Computer Science & Engineering  
University of Nevada, Reno

**Research Intern** June 2014 - September 2014 (Full Time)  
Microsoft Research October 2014 - July 2015 (Part Time)  
Supervisor: Yuxiong He

**Research Associate** July 2013 - October 2013 (Full Time)  
Systems Research Lab (SIMPL), HP Labs November 2013 - May 2014 (Part Time)  
Supervisor: Lucy Cherkasova

## **Education**

---

**PhD in Computer Science** April 2016  
Department of Computer Science, College of William and Mary  
Advisor: Evgenia Smirni

**MS in Computer Science** May 2011  
Department of Computer Science, College of William and Mary  
Advisor: Evgenia Smirni

## **Research Interests**

---

My research lies at the intersection of Data, AI, and Systems. I work in the areas of Big Data, Machine Learning/Deep Learning, Cloud Computing, High Performance Computing, Performance Evaluation, Storage Systems, Interdisciplinary Applications.

*Some recently focused topics: Large Language Models (LLM), Large-scale Machine Learning, Machine Learning as a Service (MLaaS), AutoML/NAS, Serverless Computing and Storage, Federated Learning, Workflow Optimization, Blockchain Systems, Big Data/AI-driven applications in Wildfire and Smoke Forecasting, Material Science, and Pavement Engineering.*

## **Highlights**

---

- IEEE Exemplary Editor Award
- Regents' Rising Researcher Award
- Outstanding Service Award of IEEE ACSOS
- NSF CAREER Award
- CSE Best Researcher Award
- NSF CRII Award
- PI/Site-PI of over **\$3M** external funding
- Co-PI/Senior Personnel of over **\$1.38M** external funding
- Received **3** Best Paper Awards
- NIPS 2017 Oral paper (40 out of 3240 submissions, 1.2%)
- Published **27** papers in prestigious venues chosen by CSRankings.org
- Published **12** prestigious ACM/IEEE Transaction journals
- Committee Chair of **14** PhD students and **7** MS students
- Advised/Advising **23** undergraduate students and **3** K-12 students
- Serve as Social Media Chair of the ACM SIGMETRICS community
- Lead ACM SIGMETRICS Slack Channel Initiative
- Serve as Secretary, IEEE Northern Nevada Section
- Associate Editor, IEEE Transactions on Networking
- Associate Editor, IEEE Open Journal of the Communications Society
- Served in more than **40** Technical Program Committees
- Reviewed more than **500** conference and journal paper submissions
- Chair of MS in AI Initiative at UH
- Serve as Co-founder and Co-director of CS Undergraduate Research Program at UH
- Serve as Co-founder and Co-director of the Big Data Minor program at UNR

## **Research Projects and Publications**

---

***External Grants (Total: \$4,364,018.34, PI/Site-PI: \$2,983,696.34, Co-PI: \$699,249, Senior Personnel: \$681,073):***

[G1] **UH-PI** (Lead PI: Lei Yang, UC Davis PI: Junshan Zhang), NSF Collaborative Research: NeTS: Medium: CAML: Communication-constrained Adaptive Machine Learning in Heterogeneous Edge Networks, 2025-2028, **\$1,200,000** (Feng's share: **\$360,000**)

[G2] **Sole-PI**, NSF CAREER: Automated and Efficient Machine Learning as a Service, 2021-2026, **\$517,459**

[G3] **Lead-PI** (UNR Co-PIs: Lei Yang, Heather Holmes, W&M-PI: Evgenia Smirni), NSF BIGDATA: IA: Collaborative Research: Protecting Yourself from Wildfire Smoke: Big Data Driven Adaptive Air Quality Prediction Methodologies, 2018-2022, **\$1,401,247.34** (\$1,042,676 from NSF, \$358,571.34 from Amazon, UNR's Share: \$1,043,260.34)

[G4] **Sole-PI**, NSF CRII: SHF: Optimizing Deep Learning Training through Modeling and Scheduling Support, 2018-2021, **\$174,990**

[G5] **Sole-PI**, Amazon Web Services (AWS) Cloud Credits for Research Award, 2019, **\$30,000**

[G6] **UNR PI** (PI: Nikhilesh Chawla), NSF AI Institute: Planning: Novel Neural Architectures for 4D Materials Science, 2020-2022, **\$500,000**

[G7] **Co-PI** (PI: Lei Yang), NSF REU Site: Cross-disciplinary Research Experience for Undergraduates on Big Data Analytics in Smart Cities, 2020-2023, **\$324,000**

[G8] **Co-PI** (PI: Peter Sebaaly), FAA Update Backcalculation Software (BAKFAA), 2020-2023, **\$375,249**

[G9] **Senior Personnel** (PI: Kostas Alexis), NSF RET Site: Cross-disciplinary Research Experiences on Smart Cities for Nevada Teachers: Integrating Big Data into Robotics, 2018-2021, **\$581,073**

[G10] **Senior Personnel** (PI: Tin C Nguyen), Nevada NASA Space Consortium Grant: 2020-2021, **\$50,000**

[G11] **Senior Personnel** (PI: Huang La), Nevada NASA Space Consortium Grant, 2018-2019, **\$50,000**

***Internal Grants (Total: \$93,800, PI: \$53,800, Co-PI: \$40,000):***

[G1] **PI** (Co-PIs: Fred Harris, Sergiu Dascalu, Sushil Louis, George Bebis, Dave Feil- Seifer, Jim La, Siming Liu, Dongfang Zhao), UNR CSE DF Grant: Developing GPU Infrastructure for UNR Big Data and Robotics Training and Research, 2018, **\$38,800**

[G2] **PI**, UNR GSA Reach Grant Award: MLaaS for Serving Intelligent Internet of Things, 2020, **\$2,500**

[G3] **PI**, UNR GSA Reach Grant Award: Towards Efficient Deep Learning Serving, 2017, **\$2,500**

[G4] **PI** (Co-PIs: Lei Yang, Dongfang Zhao, George Bebis), UNR CSE DF Grant: Enhancing GPU Cluster for UNR Big Data Training and Research, 2017, **\$10,000**

[G5] **Co-PI** (PI: Lei Yang), UNR CoN DF Grant: Building GPU Cluster for UNR Big Data Training and Research, 2017, **\$40,000**

**Peer-reviewed Research Publications (only those during independent career are listed here):**

*We are actively publishing at the most prestigious venues in computer system area (e.g., SOSP, SC, HPDC, VLDB, USENIX ATC, FAST, EuroSys, etc.) and machine learning area (e.g., NIPS/NeurIPS, ICLR, KDD, AAAI etc.)*

**top-cs:** top-tier Computer Science publication venue according to [CSRankings.org](https://www.csrankings.org)

**top-cr:** top-tier Computer Science publication venue according to [Conference Ranks](#) or [CORE](#)

**top-w:** workshop in top-tier publication venue

**1. ICLR 2026 (top-cs, top-cr)**

Fan Shu, Yite Wang, Ruofan Wu, Boyi Liu, Zhewei Yao, Yuxiong He, Feng Yan, DARE-bench: Evaluating Modeling and Instruction Fidelity of LLMs in Data Science, in Proceedings of The Fourteenth International Conference on Learning Representations (ICLR 2026), Rio de Janeiro, Brazil, April, 2026.

**2. Neurocomputing 2025 (top journal, impact factor 6.5)**

Jianping Zhang, Lei Yang, Seyed Mahmoud Sajjadi Mohammadabadi, Feng Yan, A survey on self-supervised learning: Recent advances and open problems, in Journal of ACM Transactions on Recommender Systems (Neurocomputing 2025).

**3. TORS 2025 (top journal)**

Tunhou Zhang, Dehua Cheng, Yuchen He, Zhengxing Chen, Xiaoliang Dai, Liang Xiong, Yudong Liu, Feng Cheng, Yufan Cao, Feng Yan, Hai Li, Yiran Chen, Wei Wen, Towards Automated Model Design on Recommender Systems, in Journal of ACM Transactions on Recommender Systems (TORS 2025).

**4. TOS 2025 (top journal)**

Xinyang Shao, Yiduo Wang, Cheng Li, Hengyu Liang, Chenhan Wang, Feng Yan, Yinlong Xu, Towards Agile and Judicious Metadata Load Balancing for Ceph File System via Matrix-based Modeling, in Journal of ACM Transactions on Storage (TOS 2025).

**5. Fire 2025 (top journal in the domain, impact factor 2.7)**

Tina Samavat, Amirhessam Yazdi, Feng Yan, Lei Yang, Early-Stage Wildfire Detection: A Weakly Supervised Transformer-Based Approach, in Journal of Fire (Fire 2025).

**6. IoT 2024 (top journal, impact factor: 8.2)**

Seyed Mahmoud Sajjadi Mohammadabadi, Syed Zawad, Feng Yan, and Lei Yang, Speed Up Federated Learning in Heterogeneous Environments: A Dynamic Tiering Approach, in IEEE Internet of Things Journal (IoT 2024).

**7. PEVA 2024 (CORE-B, h-index: 71)**

Ahsan Ali\*, Xiaolong Ma\*, Syed Zawad, Paarijaat Aditya, Istemi Ekin Akkus, Ruichuan Chen, Lei Yang, Feng Yan, Enabling scalable and adaptive machine learning training via serverless computing on public cloud, in Performance Evaluation Journal, special issue on Performance Analysis and Evaluation of Systems for Artificial Intelligence (PEVA 2024) (\*: Equal Contribution and Co-First Authors).

**8. PEVA 2024 (CORE-B, h-index: 71)**

Syed Zawad\*, Xiaolong Ma\*, Jun Yi\*, Cheng Li, Minjia Zhang, Lei Yang, Feng Yan, Yuxiong He, FedCust:

Offloading Hyperparameter Customization for Federated Learning, in Performance Evaluation Journal, special issue on Performance Analysis and Evaluation of Systems for Artificial Intelligence (PEVA 2024) (\*: Equal Contribution and Co-First Authors).

9. **ICLR 2024**

Guanhua Wang\*, Heyang Qin\*, Sam Ade Jacobs, Xiaoxia Wu, Connor Holmes, Zhewei Yao, Samyam Rajbhandari, Olatunji Ruwase, Feng Yan, Lei Yang, Yuxiong He, ZeRO++: Extremely Efficient Collective Communication for Large Model Training, in Proceedings of the Twelfth International Conference on Learning Representations (ICLR 2024), Vienna Austria, May, 2024 (\*Equal Contribution and Co-First Authors, Paper acceptance rate: 31%).

10. **ICPE 2024**

Xiaolong Ma, Feng Yan, Lei Yang, Ian Foster, Michael Papka, Zhengchun Liu and Rajkumar Kettimuthu, MalleTrain: Deep Neural Networks Training on Unfillable Supercomputer Nodes, in Proceedings of the 15th ACM/SPEC International Conference on Performance Engineering (ICPE 2024), London, UK, May, 2024 (Industry Track, Paper acceptance rate: TBD).

11. **EuroSys 2024**

Kai Ma, Cheng Li, Enzo Zhu, Ruichuan Chen, Feng Yan, Kang Chen, Noctua: Towards Practical and Automated Fine-grained Consistency Analysis, in Proceedings of the 19th European Conference on Computer Systems (EuroSys 2024), Athens, Greece, April, 2024 (Paper acceptance rate: TBD).

12. **TPDS 2023 (impact factor: 5.3)**

Hao Wu, Shiyi Wang, Youhui Bai, Cheng Li, Quan Zhou, Jun Yi, Feng Yan, Ruichuan Chen, Yinlong Xu, A Generic, High-Performance, Compression-Aware Framework for Data Parallel DNN Training, in IEEE Transactions on Parallel and Distributed Systems (TPDS 2023).

13. **JNCA 2023 (impact factor: 8.7)**

Shreshth Tuli, Fatemeh Mirhakimi, Samodha Pallegatta, Syed Zawad, Giuliano Casale, Bahman Javadi, Feng Yan, Rajkumar Buyya, Nicholas R. Jennings, AI augmented Edge and Fog computing: Trends and challenges, in Journal of Network and Computer Applications (JNCA 2023).

14. **MLSYS 2023**

Guanhua Wang\*, Heyang Qin\*, Sam Ade Jacobs, Xiaoxia Wu, Connor Holmes, Zhewei Yao, Samyam Rajbhandari, Olatunji Ruwase, Feng Yan, Lei Yang, Yuxiong He, ZeRO++: Extremely Efficient Collective Communication for Large Model Training, in Workshop on ML for Systems at NeurIPS 2023 (MLSYS 2024), Vienna Austria, May, 2024 (\*Equal Contribution and Co-First Authors, Paper acceptance rate: TBD).

15. **CLUSTER 2023**

Xinying Wang, Lipeng Wan, Scott Klasky, Dongfang Zhao, and Feng Yan SciLance: Mitigate Load Imbalance for Parallel Scientific Applications in Cloud Environments, in Proceedings of the 2023 IEEE International Conference on Cluster Computing (CLUSTER 2023), Santa Fe, New Mexico, Oct, 2023 (Paper acceptance rate: 24.6%).

16. **CCGRID 2023 (top-cr)**

Syed Zawad, Ali Anwar, Yi Zhou, Nathalie Baracaldo, and Feng Yan, HDFL: A Heterogeneity and Client Dropout-Aware Federated Learning Framework, in Proceedings of the 23rd IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID 2023), Bangalore, India, May, 2023 (Paper acceptance rate: TBD).

17. **VLDB 2023 (top-cs, top-cr)**

Jingyuan Zhang, Ao Wang, Xiaolong Ma, Benjamin Carver, Nicholas John Newman, Ali Anwar, Lukas Rupprecht, Dimitrios Skourtis, Vasily Tarasov, Feng Yan, and Yue Cheng, SION: Elastic Serverless Cloud Storage, in Proceedings of the 49th International Conference on Very Large Data Bases (VLDB 2023), Vancouver, Canada, Aug, 2023 (Paper acceptance rate: TBD).

18. **WWW 2023 (top-cs, top-cr)**

Tunhou Zhang, Dehua Cheng, Yuchen He, Zhengxing Chen, Xiaoliang Dai, Liang Xiong, Feng Yan, Hai Li, Yiran Chen, Wei Wen, NASRec: Weight Sharing Neural Architecture Search for Recommender Systems, in Proceedings of the 2023 ACM Web Conference (WWW 2023), Austin, Texas, USA, May, 2023 (Paper acceptance rate: 19.2%).

19. **ICLR 2023 (top-cs, top-cr)**

Syed Zawad, Cheng Li, Zhewei Yao, Elton Zheng, Yuxiong He, Feng Yan, DySR: Adaptive Super-Resolution via Algorithm and System Co-design, in Proceedings of The Eleventh International Conference on Learning Representations (ICLR 2023), Kigali Rwanda, May, 2023 (Acceptance rate: 31.8%).

20. **HPCA 2023 (top-cs, top-cr)**

Quan Zhou, Haiquan Wang, Xiaoyan Yu, Cheng Li, Youhui Bai, Feng Yan, Yinlong Xu, MPress: Democratizing Billion-Scale Model Training on Multi-GPU Servers via Memory-Saving Inter-Operator Parallelism, in Proceedings of The 29th IEEE International Symposium on High-Performance Computer Architecture (HPCA 2023), Montreal, Canada, Feb, 2023 (Acceptance rate:  $91/364=25\%$ ).

21. **WACV 2023 (top-cr)**

Tunhou Zhang, Mingyuan Ma, Feng Yan, Hai Li, Yiran Chen, PIDS: Joint Point Interaction-Dimension Search for 3D Point Cloud, in Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV 2023), Waikoloa, Hawaii, USA, Jan, 2023 (Paper acceptance rate: TBD).

22. **VLDB 2022 (top-cs, top-cr)**

Ahsan Ali, Riccardo Pincioli, Feng Yan, Evgenia Smirni, "Optimizing Inference Serving on Serverless Platforms", in Proceedings of the 48th International Conference on Very Large Data Bases (VLDB 2022), Sydney, Australia, September, 2022 (Acceptance rate: TBD).

23. **FL-AAAI 2022 (top-w)**

Jingoo Han, Ahmad Faraz Khan, Syed Zawad, Ali Anwar, Nathalie Baracaldo Angel, Yi Zhou, Feng Yan, and Ali R. Butt, "Tokenized Incentive for Federated Learning", in Proceedings of the International Workshop on Trustable, Verifiable and Auditable Federated Learning in Conjunction with AAAI 2022 (**FL-AAAI 2022**), Vancouver, BC, Canada, March, 2022.

24. **IPDPS 2022 (top-cr)**

Olamide Timothy Tawose, Bin Li, Lei Yang, Feng Yan, and Dongfang Zhao, “Topological Modeling and Parallelization of Multidimensional Data on Microelectrode Arrays”, in Proceedings of the 36th IEEE International Parallel & Distributed Processing Symposium (**IPDPS 2022**), Lyon, France, June, 2022 (Acceptance rate: TBD).

25. **NeurIPS 2021 (top-cs, top-cr)**

Heyang Qin, Samyam Rajbhandari, Olatunji Ruwase, Feng Yan, Lei Yang, Yuxiong He, “SimiGrad: Fine-Grained Adaptive Batching for Large Scale Training using Gradient Similarity Measurement”, in Proceedings of the Neural Information Processing Systems 2021 (**NeurIPS 2021**), Virtual, December, 2021 (Acceptance rate:  $2371/9122=26\%$ ).

26. **SOSP 2021 (top-cs, top-cr)**

Youhui Bai, Cheng Li, Quan Zhou, Jun Yi, Ping Gong, Feng Yan, Ruichuan Chen, Yinlong Xu, “Gradient Compression Supercharged High-Performance Data Parallel DNN Training”, in Proceedings of the 28th ACM Symposium on Operating Systems Principles (**SOSP 2021**), Virtual, October, 2021 (Acceptance rate:  $54/348=15.5\%$ ).

27. **SoCC 2021 (top venue in cloud)**

Chengliang Zhang, Junzhe Xia, Baichen Yang, Huancheng Puyang, Wei Wang, Ruichuan Chen, Istemi Ekin Akkus, Paarijaat Aditya, Feng Yan, “Citadel: Protecting Data Privacy and Model Confidentiality for Collaborative Learning”, in Proceedings of the ACM Symposium on Cloud Computing 2021 (**SoCC 2021**), Seattle, WA, November, 2021 (Acceptance rate:  $46/145=31.7\%$ ).

28. **SC 2021 (top-cs, top-cr)**

Abdullah Al-Mamun, Feng Yan, Dongfang Zhao, “BAASH: Lightweight, Efficient, and Reliable Blockchain-As-A-Service for HPC Systems”, in Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis (**SC 2021**), St. Louis, MO, November, 2021 (Acceptance rate:  $86/365=23.6\%$ ).

29. **SC 2021 (top-cs, top-cr)**

Yiduo Wang, Cheng Li, Xinyang Shao, Youxu Chen, Feng Yan, Yinlong Xu, “Lunule: An Agile and Judicious Metadata Load Balancer for CephFS”, in Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis (**SC 2021**), St. Louis, MO, November, 2021 (Acceptance rate:  $86/365=23.6\%$ ).

30. **SC-W 2021 (top-w)**

Xinying Wang, Lipeng Wan, Jieyang Chen, Qian Gong, Ben Whitney, Jinzhen Wang, Ana Gainaru, Qing Liu, Norbert Podhorszki, Dongfang Zhao, Feng Yan, Scott Klasky, “Unbalanced Parallel I/O: An Often-Neglected Side Effect of Lossy Scientific Data Compression”, in Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis (**DRBSD-7**), St. Louis, MO, November, 2021 (6-pages full paper).

31. **VLDB 2021 (top-cs, top-cr)**

Jiawei Wang, Cheng Li, Kai Ma, Jingze Huo, Feng Yan, Xinyu Feng, and Yinlong Xu, “AutoGR: Automated Geo-Replication with Fast System Performance and Preserved Application Semantics”, in Proceedings of the 47th International Conference on Very Large Data Bases (**VLDB 2021**), Copenhagen, Denmark, August, 2021 (Acceptance rate: TBD).

32. **AAAI 2021 (top-cs, top-cr)**

Syed Zawad, Ahsan Ali, Pin-Yu Chen, Ali Anwar, Yi Zhou, Nathalie Baracaldo, Yuan Tian, and Feng Yan, “Curse or Redemption? How Data Heterogeneity Affects the Robustness of Federated Learning”, in Proceedings of the AAAI Conference on Artificial Intelligence (**AAAI 2021**), Virtual, Feb, 2021 (**Oral Presentation, Acceptance rate: 1692/7911=21%**).

33. **AAAI 2021 (top-cs, top-cr)**

Hsin-Pai Cheng, Tunhou Zhang, Yixing Zhang, Shiyu Li, Feng Liang, Feng Yan, Meng Li, Vikas Chandra, Hai Li, Yiran Chen, “NASGEM: Neural Architecture Search via Graph Embedding Method”, in Proceedings of the AAAI Conference on Artificial Intelligence (**AAAI 2021**), Virtual, Feb, 2021 (**Oral Presentation, Acceptance rate: 1692/7911=21%**).

34. **ICDE 2021 (top-cr)**

Abdullah Al-Mamun, Feng Yan, Dongfang Zhao, “SciChain: Blockchain-enabled Lightweight and Efficient Data Provenance for Reproducible Scientific Computing”, in Proceedings of the 37th IEEE International Conference on Data Engineering (**ICDE 2021**), Chania, Crete, Greece, April, 2021 (Short Paper, Acceptance rate: TBD).

35. **PASC 2021**

PASC 2021 Feng Li, Dali Wang, Feng Yan, Fengguang Song, “X-Composer: Enabling Cross-Environments In-Situ Workflows between HPC and Cloud”, in Proceedings of the ACM and CSCS Platform for Advanced Scientific Computing (**PASC 2021**), University of Geneva, Switzerland, July, 2021 (Acceptance rate: TBD).

36. **ITNG 2021**

Adam Cassell, Andrew Munoz, Brianna Blain-Castelli, Nikkolos Irwin, Feng Yan, Sergiu Dascalu, and Frederick C Harris, “CARS: A Containerized Amazon Recommender System”, in Proceedings of the 18th International Conference on Information Technology: New Generations (**ITNG 2021**), Las Vegas, NV, USA, April, 2021 (**Best Student Paper Award**).

37. **ICSE-W 2021 (top-w)**

Chengru Yang, Zhehao Li, Chaoyi Ruan, Guanbin Xu, Cheng Li, Ruichuan Chen, Feng Yan, “PerfEstimator: A Generic and Extensible Performance Estimator for Data Parallel DNN Training”, in Proceedings of the ICSE21 Workshop on Cloud Intelligence (**ICSE Workshops 2021**), Madrid, Spain, May, 2021 (6-page long paper).

38. **INFOCOM-W 2021 (top-w)**

Md Kamran Chowdhury Shisher, Heyang Qin, Lei Yang, Feng Yan, and Yin Sun, “The Age of Correlated Features in Supervised Learning based Forecasting”, in Proceedings of the IEEE INFOCOM Age of Information Workshop (**INFOCOM Workshops 2021**), Virtual Conference, May, 2021.

39. **SC 2020 (top-cs, top-cr)**

Ahsan Ali\*, Riccardo Pincioli\*, Feng Yan, Evgenia Smirni, “BATCH: Machine Learning Inference Serving on Serverless Platforms with Adaptive Batching”, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (**SC 2020**), Atlanta, GA, USA, Nov, 2020 (\*Equal Contribution and Co-First Authors, **Acceptance rate: 85/380=22%**).

40. **SC 2020 (top-cs, top-cr)**

Amirhesam Yazdi, Xing Lin, Lei Yang, Feng Yan, “SEFEE: Lightweight Storage Error Forecasting in Large Scale Enterprise Storage Systems”, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (**SC 2020**), Atlanta, GA, USA, Nov, 2020 (**Acceptance rate: 85/380=22%**).

41. **SIGKDD 2020 (top-cs, top-cr)**

Wei Wen, Feng Yan, Yiran Chen, Hai Li, “AutoGrow: Automatic Layer Growing in Deep Convolutional Networks”, in Proceedings of the 2020 ACM SIGKDD Conference on Knowledge Discovery and Data Mining (**SIGKDD 2020**), San Diego, CA, August, 2020 (Research Track, Oral Presentation, **Acceptance rate: 216/1279=16.8%**).

42. **USENIX ATC 2020 (top-cs, top-cr)**

Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, Yang Liu, “BatchCrypt: Efficient Homomorphic Encryption for Cross-Silo Federated Learning”, in Proceedings of the 2020 USENIX Annual Technical Conference (**USENIX ATC 2020**), Boston, MA, USA, July, 2020 (**Acceptance rate: 65/348=18.6%**).

43. **HPDC 2020 (top-cs, top-cr)**

Zheng Chai\*, Ahsan Ali\*, Syed Zawad\*, Stacey Truex, Ali Anwar, Nathalie Baracaldo, Yi Zhou, Heiko Ludwig, Feng Yan, Yue Cheng, “TiFL: A Tier-based Federated Learning System”, in Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing (**HPDC 2020**), Stockholm, Sweden, June, 2020 (\*Equal Contribution and Co-First Authors, Full paper, **Acceptance rate: 16/71=22%**).

44. **FAST 2020 (top-cs, top-cr)**

Ao Wang, Jingyuan Zhang, Xiaolong Ma, Ali Anwar, Lukas Rupperecht, Dimitrios Skourtis, Vasily Tarasov, Feng Yan, Yue Cheng, “Orchestrating Ephemeral Cloud Functions to Build A Cost-Effective Object Cache”, in Proceedings of the 18th USENIX Conference on File and Storage Technologies (**FAST 2020**), Santa Clara, CA, USA, Feb, 2020 (**Acceptance rate: 23/138=16.7%**).

45. **AAAI 2020 (Oral) (top-cs, top-cr)**

Xinying Wang, Timothy Tawose, Feng Yan, Dongfang Zhao, “HDK: Toward High-Performance Deep-Learning-Based Kirchhoff Analysis”, in Proceedings of the AAAI Conference on Artificial Intelligence (**AAAI 2020**), New York, NY, USA, Feb, 2020 (**Oral Presentation, Paper acceptance rate: 1591/7737=20.6%, Oral acceptance rate: 453/7737=5.9%**).

46. **AAAI 2020 (Oral) (top-cs, top-cr)**

Tunhou Zhang, Hsin-Pai Cheng, Zhenwen Li, Feng Yan, Chengyu Huang, Hai Li, Yiran Chen, “AutoShrink: A Topology-aware NAS for Discovering Efficient Neural Architecture”, in Proceedings of the AAAI Conference

on Artificial Intelligence (AAAI 2020), New York, NY, USA, Feb, 2020 (**Oral Presentation, Paper acceptance rate: 1591/7737=20.6%**, Oral acceptance rate: 453/7737=5.9%).

47. **IPDPS 2020 (top-cr)**

Jun Yi, Chengliang Zhang, Wei Wang, Cheng Li, Feng Yan, “Not All Explorations Are Equal: Harnessing Heterogeneous Profiling Cost for Efficient MLaaS Training”, in Proceedings of the 2020 IEEE International Parallel and Distributed Processing Symposium (**IEEE IPDPS 2020**), New Orleans, LA, USA, May, 2020 (**Acceptance rate: 110/446=24.7%**).

48. **IoT 2020 (top journal, impact factor: 9.515)**

Heyang Qin, Syed Zawad, Yanqi Zhou, Lei Yang, Sanjay Padhi, Feng Yan, “Reinforcement Learning Empowered MLaaS Scheduling for Serving Intelligent Internet of Things”, in IEEE Internet of Things Journal (**IoT 2020**) (**Impact Factor: 9.515**).

49. **TBD 2020 (top journal, impact factor: 5.67)**

Xinying Wang, Cong Xu, Ke Wang, Feng Yan, Dongfang Zhao, “Memory Scaling of Cloud-based Big Data Systems: A Hybrid Approach”, in IEEE Transactions on Big Data (**TBD 2020**) (**Impact Factor: 5.67**).

50. **TCC 2020 (top journal, impact factor: 5.967)**

Chengliang Zhang, Minchen Yu, Wei Wang, Feng Yan, “Enabling Cost-Effective, SLO-Aware Machine Learning Inference Serving on Public Cloud”, in IEEE Transactions on Cloud Computing (**TCC 2020**) (**Impact Factor: 5.967**).

51. **TNSM 2020 (top journal, impact factor: 4.682)**

Riccardo Pincioli, Ahsan Ali, Feng Yan, Evgenia Smirni, “CEDULE+: Resource Management for Burstable Cloud Instances Using Predictive Analytics”, in IEEE Transactions on Network and Service Management (**TNSM 2020**), 2020. (**TNSM 2020**) (**Impact Factor: 4.682**).

52. **TMC 2020 (top journal, impact factor: 4.474)**

Lixing Yu, Ming Li, Wenqiang Jin, Yifan Guo, Qianlong Wang, Feng Yan, Pan Li, “STEP: A Spatio-Temporal Fine-Granular User Traffic Prediction System for Cellular Networks”, in IEEE Transactions on Mobile Computing (**TMC 2020**) (**Impact Factor: 4.474**).

53. **CVPR-W 2020 (top-w)**

Huanrui Yang, Minxue Tang, Wei Wen, Feng Yan, Daniel Hu, Ang Li, Hai Li, Yiran Chen, “Learning Low-rank Deep Neural Networks via Singular Vector Orthogonality Regularization and Singular Value Sparsification”, in Joint Workshop on Efficient Deep Learning in Computer Vision (**CVPR Workshops 2020**), Seattle, WA, USA, June, 2020.

54. **SC 2019 (top-cs, top-cr)**

Heyang Qin, Syed Zawad, Yanqi Zhou, Lei Yang, Dongfang Zhao, Feng Yan, “Swift Machine Learning Model Serving Scheduling: A Region Based Reinforcement Learning Approach”, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (**SC 2019**), Denver, CO, USA, November 2019 (**Acceptance rate: 78/344=22%**).

55. **USENIX ATC 2019 (top-cs, top-cr)**  
Chengliang Zhang, Minchen Yu, Wei Wang, Feng Yan, “MARk: Exploiting Cloud Services for Cost-Effective, SLO-Aware Machine Learning Inference Serving”, in Proceedings of the 2019 USENIX Annual Technical Conference (**USENIX ATC 2019**), Renton, WA, USA, July 2019 (**Acceptance rate: 71/356=20%**).
56. **EuroSys 2019 (top-cs, top-cr)**  
Connor Holmes, Daniel Mawhirter, Yuxiong He, Feng Yan, Bo Wu, “GRNN: Low-Latency and Scalable RNN Inference on GPUs”, in Proceedings of the 14th European Conference on Computer Systems (**EuroSys 2019**), Dresden, Germany, March 2019 (**Acceptance rate: 45/207=21%**).
57. **Middleware 2019 (top-cr)**  
Ahsan Ali, Riccardo Pinciroli, Feng Yan, Evgenia Smirni, “It’s not a Sprint, it’s a Marathon: Stretching Multi-resource Burstable Performance in Public Clouds”, in Proceedings of the ACM/IFIP/USENIX Middleware 2019 (**Middleware 2019**), Davis, CA, USA, December, 2019 (Industry track, **acceptance rate: 32%**).
58. **BMVC 2019 (top-cr)**  
Yanqi Zhou, Peng Wang, Sercan Arik, Haonan Yu, Syed Zawad, Feng Yan, Greg Diamos, “EPNAS: Efficient Progressive Neural Architecture Search”, in Proceedings of the 2019 30th British Machine Vision Conference (**BMVC 2019**), Cardiff, UK, September 2019 (**Acceptance rate: 231/815=28%**).
59. **ICCV-W 2019 (top-w)**  
Hsin-Pai Cheng, Tunhou Zhang, Yukun Yang, Feng Yan, Harris Teague, Hai Li, Yiran Chen, “MSNet: Structural Wired Neural Architecture Search for Internet of Things”, in Proceedings of **ICCV 2019** Neural Architects Workshop, Seoul, Korea, Oct, 2019.
60. **CLOUD 2019 (Best Paper Award) (flagship venue in cloud)**  
Xinying Wang, Abdullah Al-Mamun, Feng Yan, Dongfang Zhao, “BlockLite: Toward Accurate and Efficient Emulation of Public Blockchains in the Cloud”, in Proceedings of the 2019 International Conference on Cloud Computing (**CLOUD 2019**), San Diego, CA, USA, June 2019 (**Best Paper Award, Paper Acceptance rate: 26.7%**).
61. **CLOUD 2019 (flagship venue in cloud)**  
Hsin-Pai Cheng\*, Patrick Yu\*, Haojing Hu\*, Syed Zawad\*, Feng Yan, Shiyu Li, Hai Li, Yiran Chen, “Towards Decentralized Deep Learning with Differential Privacy”, in Proceedings of the 2019 International Conference on Cloud Computing (**CLOUD 2019**), San Diego, CA, USA, June 2019 (\*Equal Contribution and Co-First Authors, **Acceptance rate: 26.7%**).
62. **AGU 2019**  
Jingting Huang, Amir Ghasemkhani, Sandra Marcela Loria Salazar, Feng Yan, Lei Yang, Evgenia Smirni, Jens Redemann, Heather Holmes, “Using Novel Machine Learning Algorithms to Improve the Spatiotemporal Coverage of Satellite Aerosol Optical Depth”, in AGU Fall Meeting 2019 (**AGU 2019**), San Francisco, CA, USA, Dec, 2019.
63. **ITNG 2019**

Syed Zawad, Feng Yan, Rui Wu, Lee Barford, Frederick C Harris, “Randomized Benchmarking of Quantum Gates on a GPU”, in 16th International Conference on Information Technology-New Generations (**ITNG 2019**), Las Vegas, NV, USA, April, 2019.

64. **CLOUD 2018 (Best Student Paper Award) (flagship venue in cloud)**

Xinying Wang, Cong Xu, Ke Wang, Feng Yan, Dongfang Zhao, “Toward Cost-effective Memory Scaling in Clouds: Symbiosis of Virtual and Physical Memory”, in Proceedings of the 10th IEEE International Conference on Cloud Computing (**IEEE CLOUD 2018**), San Francisco, CA, USA, July 2018 (**Best Student Paper Award (1 out of 300+)**, **Paper acceptance rate: 20%**).

65. **ICDCS 2018 (top-cr)**

Chengliang Zhang, Huangshi Tian, Wei Wang, Feng Yan, Stay Fresh: “Speculative Synchronization for Fast Distributed Machine Learning”, in Proceedings of the 38th IEEE International Conference on Distributed Computing Systems (**ICDCS 2018**), Vienna, Austria, July 2018 (**Acceptance rate: 78/378=20%**).

66. **ICAC 2018 (top-cr)**

Ahsan Ali, Riccardo Pinciroli, Feng Yan, Evgenia Smirni, “CEDULE: A Scheduling Framework for Burstable Performance in Cloud Computing”, in Proceedings of the 15th IEEE International Conference on Autonomic Computing (**ICAC 2018**), Trento, Italy, September 2018 (**Acceptance rate: 15/49=30%**).

67. **TNSM 2018 (top journal, impact factor: 4.682)**

Feng Yan, Yuxiong He, Olatunji Ruwase, Evgenia Smirni, “Efficient Deep Neural Network Serving: Fast and Furious”, in IEEE Transactions on Network and Service Management (**TNSM 2018**), 2018 (**Impact Factor: 4.682**).

68. **NIPS-W 2018 (top-w)**

Hsin-Pai Cheng\*, Patrick Yu\*, Haojing Hu\*, Feng Yan, Shiyu Li, Hai Li, Yiran Chen, “LEASGD: an Efficient and Privacy-Preserving Decentralized Algorithm for Distributed Learning”, in Proceedings of **NIPS 2018** Workshop on Privacy Preserving Machine Learning (**PPML**), Montreal, Canada, Dec 2018 (\*Equal Contribution and Co-First Authors).

69. **NIPS-W 2018 (top-w)**

Hsin-Pai Cheng\*, Yuanjun Huang\*, Xuyang Guo\*, Feng Yan, Yifei Huang, Wei Wen, Hai Li, Yiran Chen, “Differentiable Fine-grained Quantization for Deep Neural Network Compression”, in Proceedings of **NIPS 2018** Workshop on Compact Deep Neural Networks with Industrial Applications (**CDNNRIA**), Montreal, Canada, Dec 2018 (\*Equal Contribution and Co-First Authors).

70. **CoDA 2018**

Panika Valecha, Huiping Cao, Qixu Gong, Mai Zheng, Feng Yan, Xing Lin, and Art Harkin, “Analysis and Prediction of Storage Error Events for High Performance Computing Systems”, at Department of Energy (DOE) Conference on Data Analysis (**CoDA 2018**), Santa Fe, NM, USA, March 2018.

71. **NIPS 2017 (Oral) (top-cs, top-cr)**

Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, Hai Li, “TernGrad: Ternary Gradients to Reduce Communication in Distributed Deep Learning”, in Proceedings of the Neural Information

Processing Systems 2017 (**NIPS 2017**), Long Beach, CA, USA, Dec 2017 (**Oral Presentation, Paper acceptance rate: 20%, Oral acceptance rate: 40/3240=1.2%**).

72. **Middleware 2017 (top-cr)**

Jeff Rasley, Yuxiong He, Feng Yan, Olatunji Ruwase, Rodrigo Fonseca, “HyperDrive: Exploring Hyperparameters with POP Scheduling”, in Proceedings of the ACM/IFIP/USENIX Middleware 2017 (**Middleware 2017**), Las Vegas, NV, USA, Dec 2017 (**Acceptance rate: 20/85=23%**).

73. **CLOUD 2017 (flagship venue in cloud)**

Feng Yan, Lihua Ren, Daniel Dubois, Giuliano Casale, Jiawei Wen, Evgenia Smirni, “How to Supercharge the Amazon T2: Observations and Suggestions”, in Proceedings of the 10th IEEE International Conference on Cloud Computing (**IEEE CLOUD 2017**), Honolulu, Hawaii, USA, June 2017 (**Acceptance rate: 18%**).

74. **TCC 2017 (top journal, impact factor: 5.967)**

Feng Yan, Ludmila Cherkasova, Zhuoyao Zhang, Evgenia Smirni, “DyScale: a MapReduce Job Scheduler for Heterogeneous Multicore Processors”, in IEEE Transaction on Cloud Computing (**TCC 2017**), 5(2): 317-330, IEEE, 2017, DOI: 10.1109/TCC.2015.2415772 (**Impact Factor: 5.967**).

75. **SOSP-W 2017 (top-w)**

Jeff Rasley, Yuxiong He, Feng Yan, Olatunji Ruwase, Rodrigo Fonseca, “HyperDrive: Flexible and Efficient Parallel Hyperparameter Exploration”, in Proceedings of Workshop on AI Systems (**AI Systems**) at Symposium on Operating Systems Principles (**SOSP 2017**), Shanghai, China, Oct 2017.

76. **SC 2016 (top-cs, top-cr)**

Feng Yan, Yuxiong He, Olatunji Ruwase, Evgenia Smirni, “SERF: Efficient Scheduling for Fast Deep Neural Network Serving via Judicious Parallelism”, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (**SC 2016**), Salt Lake City, USA, November 2016 (Acceptance rate: **82/446=18%**).

*Before 2016:*

77. **TOMPECS 2016 (top journal, impact factor: TBD)**

Feng Yan, Xenia Mountroudou, Alma Riska, Evgenia Smirni, “PREFigure: an Analytic Framework for HDD Management”, in ACM Transaction on Modeling and Performance Evaluation of Computing Systems (**TOMPECS 2016**), pp. 10:1-10:27, ACM, 2016.

78. **Cluster Computing 2016 (journal, impact factor: 1.851)**

Ji Xue, Feng Yan, Alma Riska, Evgenia Smirni, “Scheduling Data Analytics Work with Performance Guarantees: Queuing and Machine Learning Models in Synergy”, in the Cluster Computing journal (**Cluster Computing 2016**), 2016 (**Impact Factor: 1.851**).

79. **NOMS-D 2016**

Feng Yan, Evgenia Smirni, “Workload Interleaving with Performance Guarantees in Data Centers”, in Proceedings of the IEEE/IFIP Network Operations and Management Symposium (**NOMS 2016**), pp. 967-972, Istanbul, Turkey, April 2016 (Dissertation Paper).

80. **SIGKDD 2015 (Oral) (top-cs, top-cr)**

Feng Yan, Olatunji Ruwase, Yuxiong He, Trishul Chilimbi, “Performance Modeling and Scalability Optimization of Distributed Deep Learning Systems”, in Proceedings of the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (**SIGKDD 2015**), pp. 1355-1364, ACM, Sydney, Australia, August 2015, DOI: 10.1145/2783258.2783270 (**Oral Presentation, Paper acceptance rate: 19%**).

81. **CNSM 2015**

Ji Xue, Feng Yan, Robert Birke, Lydia Y. Chen, Thomas Scherer, Evgenia Smirni, “PRACTISE: Robust Prediction of Data Center Time Series”, in Proceedings of the 11th International Conference on Network and Service Management (**CNSM 2015**), pp. 1-9, IFIP Digital Library, Barcelona, Spain, November 2015, (**Acceptance rate: 18/102=17%**).

82. **CAC 2015 (top-cr)**

Ji Xue, Feng Yan, Alma Riska, Evgenia Smirni, “Proactive Management of Systems via Hybrid Analytic Techniques”, in Proceedings of 2015 IEEE International Conference on Cloud and Autonomic Computing (**ICAC 2015**), pp. 137-148, IEEE Press, Cambridge, MA, September 2015, DOI: 10.1109/ICAC.2015.31 (**Acceptance rate: 33%**).

83. **UCC-D 2015**

Thomas Scherer, Ji Xue, Feng Yan, Robert Birke, Lydia Y. Chen, Thomas Scherer, Evgenia Smirni, “PRACTISE – Demonstrating a Neural Network based Framework for Robust Prediction of Data Center Workload”, in IEEE/ACM International Conference on Utility and Cloud Computing (**UCC 2015**), Limassol, Cyprus, December 2015 (Demo Paper).

84. **CLOUD 2014 (flagship venue in cloud)**

Feng Yan, Ludmila Cherkasova, Zhuoyao Zhang, Evgenia Smirni, “Optimizing Power and Performance Trade-offs of MapReduce Job Processing with Heterogeneous Multi-Core Processors”, in Proceedings of the 7th IEEE International Conference on Cloud Computing (**IEEE CLOUD 2014**), pp. 240-247, IEEE Press, Alaska, USA, June 2014, DOI: 10.1109/CLOUD.2014.41 (**Acceptance rate: 18%**).

85. **ICAC 2014 (top-cr)**

Ji Xue, Feng Yan, Alma Riska, Evgenia Smirni, “Storage Workload Isolation via Tier Warming: How Models Can Help”, in Proceedings of the 11th USENIX International Conference on Autonomic Computing (**ICAC 2014**, in conjunction with USENIX ATC 2014), pp. 1-11, Philadelphia, PA, June 2014 (**Acceptance rate: 22%**).

86. **ICPE 2014**

Feng Yan, Shannon Hughes, Alma Riska, Evgenia Smirni, “Agile Middleware for Scheduling: Meeting Competing Performance Requirements of Diverse Tasks”, in Proceedings of the 5th ACM/SPEC International

Conference in Performance Engineering (**ICPE 2014**), pp. 185-196, ACM, Dublin, Ireland, March 2014, DOI: 10.1145/2568088.2568104 (**Acceptance rate: 30%**).

87. **NOMS 2014 (top-cr)**

Feng Yan, Ludmila Cherkasova, Zhuoyao Zhang, Evgenia Smirni, “Heterogeneous Cores For MapReduce Processing: Opportunity or Challenge?”, in Proceedings of the IEEE/IFIP Network Operations and Management Symposium (**NOMS 2014**), pp. 1-4, IEEE Press, Krakow, Poland, May 2014, DOI: 10.1109/NOMS.2014.6838339.

88. **MASCOTS 2013 (top-cr)**

Feng Yan, Shannon Hughes, Alma Riska, Evgenia Smirni, “Overcoming Limitations of Off-the-shelf Priority Schedulers in Dynamic Environments”, in Proceedings of the 21th ACM/IEEE Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (**MASCOTS 2013**), pp. 505-514, IEEE Press, San Francisco, CA, August 2013, DOI: 10.1109/MASCOTS.2013.72 (**Acceptance rate: 27%**).

89. **ICAC-P 2013**

Feng Yan, Evgenia Smirni, “Toward Automating Disk Power Savings with Performance Guarantees”, in **ICAC 2013** Ph.D. Forum (in conjunction with ATC 2013), San Jose, CA, June 2013.

90. **ICAC 2012 (top-cr)**

Feng Yan, Alma Riska, Evgenia Smirni, “Toward Fast Eventual Consistency with Performance Guarantees”, in Proceedings of the 9th ACM International Conference on Autonomic Computing (**ICAC 2012**), pp. 167-172, ACM, San Jose, CA, September 2012, DOI: 10.1145/2371536.2371566.

91. **ICPE 2012**

Feng Yan, Alma Riska, Evgenia Smirni, “Busy Bee: How to Use Traffic Information for Better Scheduling of Background Tasks”, in Proceedings of the 3rd ACM/SPEC International Conference in Performance Engineering (**ICPE 2012**), pp.145-156, ACM, Boston, USA, April 2012, DOI: 10.1145/2188286.2188308 (**Acceptance rate: 28%**).

92. **HotPower 2012 (top-w)**

Feng Yan, Xenia Mountroudou, Alma Riska, Evgenia Smirni, “Quantitative Estimation of the Performance Delay with Propagation Effects in Disk Power Savings”, in **USENIX HotPower 2012** (in conjunction with USENIX OSDI 2012), Hollywood, CA, October 2012 (**Acceptance rate: 25%**).

93. **ICDCSW 2012 (top-w)**

Feng Yan, Alma Riska, Evgenia Smirni, “Fast Eventual Consistency with Performance Guarantees for Distributed Storage”, in Proceedings of the 32nd International Conference on Distributed Computing Systems Workshops (**ICDCSW 2012**), pp. 23-28, IEEE Press, in **DCPerf 2012** (in conjunction with **ICDCS 2012**), Macau, China, June 2012, DOI: 10.1109/ICDCSW.2012.21.

94. **SIGMETRICS-W 2011 (top-w)**

Feng Yan, Xenia Mountroudou, Alma Riska, Evgenia Smirni, “Copy Rate Synchronization with Performance Guarantees for Work Consolidation in Storage Clusters”, in Proceedings of the ACM SIGMETRICS

Performance Evaluation Review (PER), Volume 39, Issue 3, pp. 82-86, in **GreenMetrics** 2011 (in conjunction with **SIGMETRICS** 2011), San Jose, CA, June 2011, DOI: 10.1145/2160803.2160866.

95. **MASCOTS 2011 (top-cr)**

Feng Yan, Xenia Mountroudou, Alma Riska, Evgenia Smirni, “Toward Automating Work Consolidation with Performance Guarantees in Storage Clusters”, in Proceedings of the 19th ACM/IEEE Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (**MASCOTS 2011**), pp.326-335, IEEE Press, Singapore, July 2011, DOI: 10.1109/MASCOTS.2011.32 (**Acceptance rate: 26%**).

*Invited Article:*

96. Feng Yan, Xenia Mountroudou, Alma Riska, Evgenia Smirni, “PREFigure: a Performance, Power, and Reliability Framework for Disk Drives”, in the IEEE Computer Society Special Technical Community on Sustainable Computing Newsletter (**STC**), Volume 2, Issue 2, April 2013.

*Preprint Research Publications:*

97. Ruofan Wu, Youngwon Lee, Fan Shu, Danmei Xu, Seung-won Hwang, Zhewei Yao, Yuxiong He, Feng Yan, ComposeRAG: A Modular and Composable RAG for Corpus-Grounded Multi-Hop Question Answering, preprint 2025

98. Youngwon Lee, Seung-won Hwang, Ruofan Wu, Feng Yan, Danmei Xu, Moutasem Akkad, Zhewei Yao, Yuxiong He, Agentic Verification for Ambiguous Query Disambiguation, preprint 2025

99. Xiaolong Ma, Xu Dong, Ashley Tarrant, Lei Yang, Rao Kotamarthi, Jiali Wang, Feng Yan, Rajkumar Kettimuthu, Diffusion-Based, Data-Assimilation-Enabled Super-Resolution of Hub-height Winds, preprint 2025

100. Chengliang Zhang, Junzhe Xia, Baichen Yang, Huancheng Puyang, Wei Wang, Ruichuan Chen, Istemi Ekin Akkus, Paarijaat Aditya, Feng Yan, “Citadel: Protecting Data Privacy and Model Confidentiality for Collaborative Learning with SGX”, preprint 2021

101. Hsin-Pai Cheng, Feng Liang, Meng Li, Bowen Cheng, Feng Yan, Hai Li, Vikas Chandra, Yiran Chen, “ScaleNAS: One-Shot Learning of Scale-Aware Representations for Visual Recognition”, preprint 2020.

102. Feng Li, Dali Wang, Feng Yan, Fengguang Song, “ElasticBroker: Combining HPC with Cloud to Provide Realtime Insights into Simulations”, preprint 2020.

103. Xinying Wang, Olamide Timothy Tawose, Feng Yan, Dongfang Zhao, “Distributed Nonblocking Commit Protocols for Many-Party Cross-Blockchain Transactions”, preprint 2020.

104. Hsin-Pai Cheng, Tunhou Zhang, Yukun Yang, Feng Yan, Shiyu Li, Harris Teague, Hai Li, Yiran Chen, “SwiftNet: Using Graph Propagation as Meta-knowledge to Search Highly Representative Neural Architectures”, preprint 2019.

105. Wei Wen, Yandan Wang, Feng Yan, Cong Xu, Yiran Chen, Hai Li, “SmoothOut: Smoothing Out Sharp Minima for Generalization in Large-Batch Deep Learning”, preprint 2018.

### **Selected Awards**

---

- IEEE Exemplary Editor Award, 2025
- Regents' Rising Researcher Award, 2022
- Outstanding Service Award of IEEE ACSOS, 2022
- NSF CAREER Award, 2021
- NSF CRII Award, 2018
- Best Student Paper Award, IEEE CLOUD 2018 (my PhD student Xinying Wang is the first author)
- Best Paper Award, CLOUD 2019 (my PhD student Xinying Wang is the first author)
- Best Student Paper Award, ITNG 2021 (the paper is developed from a course project in my big data course)
- SIGSOFT CAPS Grant Award, ACM, 2012, 2014
- Student Scholarship Award, SOSP 2013, USENIX
- Student Grant Award for OSDI 2012, USENIX
- Student Grant Award for ICDCS 2012, IEEE and IBM
- Support Grant Award for ICPE 2012, ACM
- Student Grant Award for SIGMETRICS Conference 2011, ACM

### **Professional Activities**

---

**Program Committees:** Sigmetrics (2026, 2025, 2024, 2023), ICS(2026, 2025), NeurIPS (2025, 2024, 2023, 2022), ICLR (2026), ICML (2026, 2022), AAI (2024, 2023, 2022, 2021), CVPR (2023, 2022, 2021), ECCV (2022), ICCV (2021), ICML-FL (2023, 2022, 2021), DSN (2025, 2024, 2022, 2021, 2020, 2017), MASCOTS (2023), IEEE ICDCS (2021, 2018), CLOUD (2023, 2022, 2021, 2020, 2019), IPCCC (2024, 2023, 2022, 2021, 2020, 2019, 2018), IEEE ACSOS/ICAC (2021, 2020, 2019, 2017), HPDC (2019), ICCCN (2019, 2018), ACM/SPEC ICPE (2019, 2018, 2017), HIDS (2023, 2022), ICICS (2021), PAISE (2022, 2021), IEEE Cloud Summit (2022, 2021, 2020), DSS (2020, 2018, 2017), PAISE (2022), BDCloud (2020), HPBD&IS(2020, 2019), SOSE(2021, 2020, 2019), ChinaSys (2020), CLOUD COMPUTING (2024, 2023, 2022, 2021, 2020, 2019, 2018, 2017), ADVCOMP (2024, 2023, 2022, 2021, 2020, 2019), ICMLDS (2019, 2018, 2017), BIGCOM(2019), IEEE VHPC (2018), ICIT (2018), IEEE/IFIP IEEE BigData (2017, 2016, 2014, 2013), PDCTA (2017), ALLDATA (2017, 2015), DATA ANALYTICS (2017, 2016), ACM ESEC/FSE (artifact evaluation track) (2015).

**Invited Journal and Conference Review:** IEEE TNLS, IEEE TPDS, JPDC, IEEE TDSC, ACM TOMACS, IEEE TC, IEEE TCC, IEEE TSC, ACM TOS, IEEE TBD, ACM TACO, IEEE TNSM, IEEE TII, IEEE Access, ACM ToMPECS, ELSEVIER PEVA, ELSEVIER SMPAT, SC, ACM SIGMETRICS, IFIP Performance, IEEE/ACM CCGrid, IEEE IC2E, QEST, WSC.

**Other Service and Activities:**

- Social Media Chair, ACM SIGMETRICS community, 2019 - now
- Lead ACM SIGMETRICS Slack Channel Initiative, 2025 - now
- Associate Editor, IEEE Transactions on Networking, 2026 - now
- Associate Editor, IEEE Open Journal of the Communications Society, 2025 – now
- Distributed Systems for AI/ML Track Chair, ICDCS 2025
- Corporate Sponsorship Chair, ACM SIGMETRICS, 2023
- Corporate Sponsorship Chair, ACM HotStorage, 2023
- Chair of MS in AI Initiative, UH Katy, 2025 - now
- Co-founder and Co-director of CS Undergraduate Research Program at UH, 2023 - now
- UH CS Chair Search Committee, UH, 2025 - now
- UH CS Faculty Annual Performance Review Committee, UH, 2025 - now
- Chair of UH CS Faculty Hiring Committee, UH, 2023 - 2024
- UH CS Faculty Hiring Committee, UH, 2022 - 2023
- UH CS Graduate Admission Committee, UH, 2022 - 2025
- UH CS Graduate Studies Committee, UH, 2024 - 2025
- Judge of UH Summer AI Student Showcase, 2024
- Invited Speaker and Judge of UH Data Science Showcase, 2022
- Secretary, IEEE Northern Nevada Section, 2017 - 2022
- Co-founder and Co-director of the Big Data Minor program at UNR, 2016 - 2022
- UNR High Performance Computing Committee, UNR, 2017 - 2022
- GSA Research Grant Committee, UNR, 2017 - 2022
- Faculty Representative for Google University Programs, 2016 - 2022
- USENIX Campus Representative, 2019 - 2022
- CSE Graduate Committee, UNR, 2016- 2018, 2021 - 2022
- NSF Panel Reviewer, 2020
- Panel reviewer for Italian Ministry for Education, University and Research Grant, 2018
- Publicity Chair, ACSOS, 2021
- Publicity Chair, SRDS, 2019
- Chair, CSE Colloquium Committee, UNR, 2018 - 2021
- UNR Pack Research Experience Program (PREP) Mentor, 2021
- Hosted 2 ACM Distinguished Speaker events and 1 Hats Off event, 2017, 2018
- Hosted Amazon Web Services Workshop for Students and Researchers, 2019
- Hosted 2 Google Campus events, 2017, 2018
- Invited Faculty of NetApp University Day, 2018 - 2020
- UNR Engineering Reception, 2018
- CSE Colloquium Committee, UNR, 2016- 2018
- USENIX ATC Summarizer 2014
- USENIX SOSP Summarizer 2013
- USENIX ICAC Ph.D. Forum 2013
- USENIX OSDI Summarizer 2012
- Google Doctoral Forum 2012

- ACM SIGMETRICS Volunteer 2011
- The International Conference on Quantitative Evaluation of Systems (QEST) Volunteer 2010

## **Students**

---

### ***Graduate Students Advised (Chair)***

1. Patrick Watters, PhD student, 08/2021 -
2. Mina Yazdani, PhD student, 01/2022 -
3. Xu Dong, MS/PhD student, 01/2023 -
4. Ruofan Wu, PhD student, 08/2023 -
5. Xinyang Li, PhD student, 08/2023 -
6. Anjir Ahmed Chowdhury, PhD student, 08/2023 -
7. Fan Shu, PhD student, 08/2024 -
8. *Xiaolong Ma, MS/PhD student, graduated 12/2024, first placement: Argonne National Laboratory*
9. *Heyang Qin, PhD student, graduated 12/2022, first placement: Microsoft Research*
10. *Syed Zawad, PhD student, graduated 12/2022, first placement: IBM Research*
11. *Amirhesam Yazdi, PhD student, graduated 12/2022, first placement: Teaching Associate Professor, IS Department, University of Nevada, Reno*
12. *Jun Yi, PhD student, graduated 08/2022, first placement: MathWorks*
13. *Xinying Wang, PhD student, graduated 05/2022, first placement: Oak Ridge National Laboratory*
14. *Ahsan Ali, PhD student, graduated 05/2021, first placement: Argonne National Laboratory*
15. *Nathan Wiseman, MS student, graduated 08/2020, first placement: Founder of Virtu.ai.org*
16. *Cayler Miley, MS student, graduated 08/2021, first placement: Clear Capital*
17. *Ke Xu, MS student, graduated 12/2019, first placement: Facebook*
18. *Di Li, MS student, graduated 12/2018*
19. *Gautham Yerroju, MS student, graduated 08/2018, first placement: Odo*

### ***Graduate Students Committee Member***

1. Martin Lee, PhD student (Chair: Zhiang Deng)
2. Adnane Gdihi, PhD student (Chair: Lennart Johnsson)
3. Salah Uddin Kadir, PhD student (Chair: Stephen Huang)
4. Aryo Yarahmadi, PhD student (Chair: Omprakash Gnawali)
5. Vu Minh Hoang Dang, PhD student (Chair: Rakesh M. Verma)
6. Yang Lu, PhD student (Chair: Stephen Huang)
7. Salma Sultana, PhD student (Chair: Dr. Zhu Han)
8. Mahmoud Sajjadi, PhD student (Chair: Dr. Lei Yang)
9. Amir Ghasemkhani, PhD student (Chair: Dr. Lei Yang)
10. Yunchuan Liu, PhD student (Chair: Dr. Lei Yang)
11. Abdullah Al-Mamun, PhD student (Chair: Dr. Dongfang Zhao)
12. Timothy Tawose, PhD student (Chair: Dr. Dongfang Zhao)
13. Khalid Bakhshaliyev, PhD student (Chair: Dr. Mehmet Gunes)

14. Ehsan Mosadegh, PhD student (Chair: Dr. Anne W Nolin)
15. Jingting Huang, PhD student (Chair: Dr. Heather Holmes)
16. Jalal Kiswanij, PhD student (Chair: Dr. Sergiu M. Dascalu)
17. Shireesh Kumar Poral Ashok Kumar, MS student (Chair: Dr. Shishir Shah)
18. Jennifer Csicsery-Ronay, MS student (Chair: Dr. Christoph F. Eick)
19. Athira Pillai, MS student (Chair: Dr. Shamik Sengupta)
20. Walker Spurgeon, MS student (Chair: Dr. Eelke Folmer)
21. Jiwan Bhandari, MS student (Chair: Dr. Eelke Folmer)

### ***Undergraduate Students Advised***

1. Siyu Gan (UH, 2025 Fall)
2. Meghana Kotta (UH, 2025 Fall)
3. Ben Jose Tuason (UH, 2024 Spring)
4. Alex Serna (CAHSI LREU, 2024 Spring)
5. Pamella May Nipay (UNR PREP, 2021 Fall)
6. Connor Jordan (intern student, 2021 Summer, Fall)
7. Agui Navarro (NSF REU Site, 2021 Summer)
8. Christina Sherpa (NSF REU Site, 2021 Summer)
9. Yongyi Zhou (NSF REU Site, 2021 Summer)
10. Connor Jordan (NSF REU Site, 2021 Summer)
11. Timothy Kashi (UNR CAPSTONE, 2019 Fall to 2020 Spring)
12. Allan Moreno (UNR CAPSTONE, 2019 Fall to 2020 Spring)
13. Carlos Varela (UNR CAPSTONE, 2019 Fall to 2020 Spring)
14. Johan Yamssi (UNR CAPSTONE, 2019 Fall to 2020 Spring)
15. Gabe Petersen (UNR CAPSTONE, 2019 Fall to 2020 Spring)
16. William Willia (UNR CAPSTONE, 2019 Fall to 2020 Spring)
17. Zachary Black (UNR CAPSTONE, 2019 Fall to 2020 Spring)
18. Desislava Nacheva (UNR CAPSTONE, 2018 Fall to 2019 Spring)
19. Masoud Modaressi (UNR CAPSTONE, 2018 Fall to 2019 Spring)
20. Alexander Kharag (UNR CAPSTONE, 2018 Fall to 2019 Spring)
21. Lance Cantu (UNR CAPSTONE, 2018 Fall to 2019 Spring)
22. Michael Knight (UNR CAPSTONE, 2018 Fall to 2019 Spring)
23. Emily Gentry (UNR CAPSTONE, 2018 Fall to 2019 Spring)
24. Mitanshu Chandna (UNR CAPSTONE, 2018 Fall to 2019 Spring)
25. Zixuan Wang (intern student, now in University of San Diego as graduate student)

### ***K-12 Students Advised***

1. Katherine Cheng, K-12 student (intern student from Mission San Jose High School)
2. Michelle Liu, K-12 student (intern student from Davidson Academy, now at Brown University)
3. Felix Zhan, K-12 student (intern student from Ed. W. Clark High School, now at Stanford University)

## **Teaching**

---

COSC 6376 Cloud Computing, University of Houston, Since 2023-Spring  
COSC 6339 Big Data Analytics, University of Houston, Since 2022-Fall  
COSC 2306 Data Programming, University of Houston, Since 2023-Fall  
CS 431/631 Introduction to Big Data, University of Nevada, Reno, 2021-Fall, 2020-Fall, 2019-Fall  
CS 436/636 Big Data Systems, University of Nevada, Reno, 2022-Spring, 2020-Spring, 2019-Spring  
CS 791 Performance Optimization of Computer Systems, University of Nevada, Reno, 2022-Spring, 2021-Spring, 2020-Spring  
CS 491/691 Simulation and Performance Evaluation, University of Nevada, Reno, 2018-Fall, 2017-Fall  
CS 491/691 Big Data, University of Nevada, Reno, 2018-Spring, 2017-Spring

## **References**

---

Dr. Evgenia Smirni  
Sidney P. Chockley Professor  
Department of Computer Science  
College of William and Mary  
Williamsburg, VA, 23187  
Email: [esmirni@cs.wm.edu](mailto:esmirni@cs.wm.edu)  
Phone: +1-757-221-3580

Dr. Yiran Chen  
John Cocke Distinguished Professor  
Electrical and Computer Engineering  
Duke University  
Durham, NC, 27708  
Email: [yiran.chen@duke.edu](mailto:yiran.chen@duke.edu)  
Phone: +1-919-660-1372

Dr. Hai (Helen) Li  
Clare Boothe Luce Professor  
Electrical and Computer Engineering  
Duke University  
Durham, NC, 27708  
Email: [hai.li@duke.edu](mailto:hai.li@duke.edu)  
Phone: +1-919-660-1373

Dr. Giuliano Casale  
Professor  
Department of Computing  
Imperial College London  
London, SW7 2AZ; UK  
Email: [g.casale@imperial.ac.uk](mailto:g.casale@imperial.ac.uk)  
Phone: +44 20 759 42920

Dr. Yuxiong He  
Distinguished AI Engineer  
Snowflake  
929 108th Ave NE Suite 600,  
Bellevue, WA 98004  
Email: [yuxiong.he@snowflake.com](mailto:yuxiong.he@snowflake.com)  
Phone: +1-425-421-0933

Dr. Lucy Cherkasova  
Principal Research Scientist  
ARM Research  
150 Rose Orchard Way  
San Jose, CA, 95134  
Email: [lucy.cherkasova@gmail.com](mailto:lucy.cherkasova@gmail.com)  
Phone: +1-408-420-4340