

SUMMARY

- Creative and goal-driven system, big data analytics and machine learning professional with 11+ years of experience at IBM Research, JD.com and Apple. **Areas of expertise** include: natural language processing, big data analytics, deep learning, and machine learning.
- Rich industrial experience in building large scale recommendation systems of JD Finance App content feed with hundreds of millions of users collaborating across multiple organizations at JD Technology.
- Science leader for benchmarking, resource management, and performance optimization of large scale data analytic platforms with good visibility in system community. 20+ papers in top academic conferences, 30+ patents filed.

EXPERIENCES

Sr. Machine Learning Scientist, Apple, Inc 2022.9-Present

- Designed and built improved machine learning systems with knowledge graph for downstreaming NLP applications including entity linking, question answering, machine translation, etc.

Staff Scientist, JD.com 2018.7-2022.7

- Designed and built the personalized and real time matching component for the large scale JD Finance APP Content Feed recommendation system with hundreds of millions of users
- Designed and implemented an intention detection technique to improve the efficiency of the operation of the application feedback system
- Led the design and implementation of anomaly detection technique for large scale data analytic monitoring systems.
- Filed a total of 11 patents in large scale systems and NLP areas.

Research Staff Member, IBM T.J. Watson / Almaden Research Center 2014.7-2018.6

- One of the first people in the area to design and develop a benchmark for big data analytic platform Spark which was open sourced.
- Researched on Spark performance optimization techniques, publishing multiple papers on top tier system conferences and more than five patents.

Research Intern, IBM Almaden / T.J. Watson Research Center 2011-2013 Summer

- Designed and built a MapReduce online performance tuning system.
- Designed and built a resource management system for MapReduce in the cloud based on topology aware min-cost-flow algorithm.
- Publish two related papers on top tier conferences.

Research Intern, NetApp Advance Technology Group 2010 Summer

- Designed and built a cooperative storage level deduplication component for a virtualized data center to reduce the I/O consumptions between physical hosts and a storage server.

Research Assistant, Virginia Tech 2008.8-2014.5

- Researched resource management techniques on cloud computing platforms. The PhD thesis, a resource management framework for cloud computing, was written on multiple related solid projects and top tier research papers.

EDUCATION

Ph.D. in Computer Science, Virginia Tech	2008.8-2014.5
M.S. in Computer Science, University of Science and Technology of China	2005.9-2008.7
B.S. in Computer Science, East China Normal University	2001.9-2005.6

HONORS AND AWARDS

- IBM Research Ph.D. Fellowship 2012-2014
- Travel Grant: OSDI 2008, OSDI 2010, FAST 2010-2013, Google Workshop for Women Engineer 2009, ICS 13, HPDC 14
- Student Volunteer: SC 2010, SC 2013, Grace Hopper Conference 2012
- University Scholarship, ECNU, China 2002 - 2004

PUBLICATIONS

1. Hellina Hailu Nigatu, Min Li, Maartje Ter Hoeve, Saloni Potdar, and Sarah Chasins. 2025. mRAKL: Multilingual Retrieval-Augmented Knowledge Graph Construction for Low-Resourced Languages. In Findings of the Association for Computational Linguistics: ACL 2025, pages 13072–13089, Vienna, Austria.
2. Yanzhu Guo, Simone Conia, Zelin Zhou, Min Li, Saloni Potdar, and Henry Xiao. 2025. Do Large Language Models have an English Accent? Evaluating and Improving the Naturalness of Multilingual LLMs. In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 3823–3838, Vienna, Austria.
3. Simone Conia, Min Li, Roberto Navigli, and Saloni Potdar. SemEval-2025 task 2: Entity-aware machine translation. In Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025), pp. 2535-2557. 2025.
4. Zelin Zhou, Simone Conia, Daniel Lee, Min Li, Shenglei Huang, Umar Farooq Minhas, Saloni Potdar, Henry Xiao, and Yunyao Li. 2025. KG-TRICK: Unifying Textual and Relational Information Completion of Knowledge for Multilingual Knowledge Graphs. In Proceedings of the 31st International Conference on Computational Linguistics, pages 9096–9111, Abu Dhabi, UAE.
5. Simone Conia, Daniel Lee, Min Li, Umar Farooq Minhas, Saloni Potdar, and Yunyao Li. 2024. Towards Cross-Cultural Machine Translation with Retrieval-Augmented Generation from Multilingual Knowledge Graphs. In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, pages 16343–16360, Miami, Florida, USA. Association for Computational Linguistics.
6. Simone Conia, Daniel Lee, Min Li, Umar Farooq Minhas, and Yunyao Li. "Enhancing Machine Translation Experiences with Multilingual Knowledge Graphs." In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, no. 21, pp. 23781-23783. 2024.
7. Simone Conia, Min Li, Daniel Lee, Umar Minhas, Ihab Ilyas, and Yunyao Li. 2023. Increasing Coverage and Precision of Textual Information in Multilingual Knowledge Graphs. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, pages 1612–1634, Singapore. Association for Computational Linguistics.
8. Huasong Shan, Yuan Chen, Haifeng Liu, Yunpeng Zhang, Xiao Xiao, Xiaofeng He, Min Li, Wei Ding. ?-diagnosis: Unsupervised and real-time diagnosis of small-window long-tail latency in large-scale microservice platforms. In The World Wide Web Conference 2019 May 13 (pp. 3215-3222).
9. Min Li, Marina Danilevsky, Sara Noeman, and Yunyao Li. 2018. DIMSIM: An Accurate Chinese Phonetic Similarity Algorithm Based on Learned High Dimensional Encoding. In Proceedings of the 22nd Conference on Computational Natural Language Learning, pages 444–453, Brussels, Belgium.
10. Luna Xu, Seung-Hwan Lim, Min Li, Ali R. Butt, Ramakrishnan Kannan. Scaling Up Data-Parallel Analytics Platforms: Linear Algebraic Operation Cases. In Proceedings of the IEEE International Conference on Big Data (Big Data), Boston, MA, pages 10, Dec 2017. (AR: 18%).

11. Luna Xu, Min Li, Li Zhang, Ali R. Butt, Yandong Wang, and Zane Zhenhua Hu. MEMTUNE: Dynamic Memory Management for In-memory Data Analytic Platforms. In Proceedings of the IEEE International Parallel & Distributed Processing Symposium (IPDPS), Chicago, IL, pages 10, May 2016. (AR: 23%).
12. Arnab Kumar Paul, Wenjie Zhuang, Luna Xu, Min Li, Mustafa Rafique, and Ali R. Butt. CHOPPER: Optimizing Data Partitioning for In-Memory Data Analytics Frameworks. In Proceedings of the IEEE International Conference on Cluster Computing (Cluster), Taipei, Taiwan, September 2016. (AR: 24%).
13. Rui Zhang, Irene Manotas, Min Li, and Dean Hildebrand, Towards a Big Data Benchmarking and Demonstration Suite for the Online Social Networks Era with Realistic Workloads and Live Data. The Six Workshop on Big Data Benchmark, Performance Optimization and Emerging Hardware(BPOE-6).2015. Kohala Coast, Hawai'i.
14. Dakshi Agrawal, Ali Butt, Doshi Kshitij, Josep-L. Larriba-Pey, Min Li, Frederick R Reiss, Francois Raab, Berni Schiefer, Yinglong Xia, SparkBench: A Spark Performance Testing Suite, Seventh TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2015).Kohala Coast, Hawai'i.
15. Yandong Wang, Li Zhang, Jian Tan, Min Li, Yuqing Gao, Xavier Guerin, Xiaoqiao Meng, and Shicong Meng. 2015. HydraDB: a resilient RDMA-driven key-value middleware for in-memory cluster computing. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '15).
16. Jian Tan, Li Zhang, Min Li, and Yandong Wang. 2015. Multi-resource Fair Sharing for Multiclass Workflows. SIGMETRICS Perform. Eval. Rev. 42, 4 (June 2015), 31-37.
17. Min Li, Jian Tan, Yandong Wang, Li Zhang, and Valentina Salapura. SparkBench: a comprehensive benchmarking suite for in memory data analytic platform Spark. In Proceedings of the 12th ACM International Conference on Computing Frontiers (CF '15). ACM, New York, NY, USA, 2015.
18. Valentina Salapura, Kirk A. Beaty, Alan Bivens, Minkyong Kim, and Min Li. 2015. Towards building an analytics platform in the cloud. In Proceedings of the 12th ACM International Conference on Computing Frontiers (CF '15). ACM, New York, NY, USA.
19. Rui Zhang, Min Li, Hildebrand Dean, "Finding the Big Data Sweet Spot: Towards Automatically Recommending Configurations for Hadoop Clusters on Docker Containers," in Cloud Engineering (IC2E), 2015 IEEE International Conference on , vol., no., pp.365-368, 9-13 March 2015.
20. Luna Xu, Min Li, Ali R. Butt, GERBIL: MPI + YARN, To appear 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), Shengzhen, China, 2015. (AR: 25.7%).
21. Min Li, Liangzhao Zeng, Shicong Meng, Jian Tan, Li Zhang, Ali R. Butt, Nicholas Fuller, mrOnline: MapReduce Online Performance Tuning, In Proceedings of the ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), Vancouver, Canada, June, 2014. (AR: 16.1%).
22. Min Li, Dinesh Subhraveti, Ali R. Butt, Aleksandr Khasymski, and Prasenjit Sarkar. CAM: A Topology Aware Minimum Cost Flow Based Resource Manager for MapReduce Applications in the Cloud. In Proceedings of the ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), Delft, The Netherlands. June 2012. (AR: 16.1%).
23. Min Li, Shravan Gaonkar, Ali R. Butt, Deepak Kenchammana and Kaladhar Voruganti. Co-operative Storage-Level De-Duplication for I/O Reduction in Virtualized Data Centers, In Proceedings of the IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), Washington D.C.. August 2012.
24. Ramya Prabhakar, Sudharshan S. Vazhkudai, Youngjae Kim, Ali R. Butt, Min Li, and Mahmut Kandemir. Provisioning a Multi-Tiered Data Staging Area for Extreme-Scale Machines. In Proceedings of the 31st International Conference on Distributed Computing Systems (ICDCS), Minneapolis, MN, June 2011.

25. Min Li, Sudharshan S. Vazhkudai, Ali R. Butt, Fei Meng, Xiaosong Ma, Youngjae Kim, Christian Engelmann, and Galen Shipman. Functional Partitioning to Optimize End-to-End Performance on Many-core Architectures. In Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC10), New Orleans, LA. November 2010. (AR: 20%).
26. Min Li, Yan Xiong, Xin Jin, A Clustered-based Relative Localization Algorithm for Wireless Sensor Networks, Computer Engineering(in Chinese), 34(19):101-103, 2008
27. Xin Jin, Yan Xiong, Min Li, Lihua Yue, Connectible-Cell based Topology Control Algorithm for Wireless Sensor Networks, Journal of Computer Research and Development(in Chinese), 45(2):217-226, 2008

PATENTS (Partial List)

1. Methods and apparatus to dynamically rewire direct connections of data objects and memory resources to processing elements as efficient memory based accessing and communication, filed in Oct-16, Co-inventors: Ramasamy, H. Salapura, V. Bivens, J. Ruan, Y. Mahindru, R. Schenfeld, E
2. A system and method of enabling shuffle-less MapReduce cluster in disaggregated systems, filed in Jun-16, Co-inventors: Ramasamy, H. Salapura, V. Bivens, J. Ruan, Y. Mahindru, R. Schenfeld, E
3. Facial recognition for sub-healthy detection and alert generation, filed in Jun-16, Co-inventors: Salapura, V. Hu, K.
4. A system and method for utilizing accelerators to accelerate data analytic workloads in disaggregated systems, filed in Mar-16, Co-inventors: Ramasamy, H. Salapura, V. Bivens, J. Ruan, Y. Mahindru, R. Schenfeld, E
5. A method and system of dynamic RDD cache management for in memory data analytic platforms, filed in Nov-15, Co-inventors: Salapura V., Zhang L., Wang, Y.
6. A Method and System for Application SLA-based Elastic Resource Provisioning in Disaggregated Systems, filed in Sep-15, Co-inventors: Ramasamy, H. Salapura, V. Bivens, J. Ruan, Y. Mahindru, R. Schenfeld, E
7. Multiple concurrent simulated app solutions to problems, Jul-15, Co-inventors: Deng Y. , Li J, Tai T
8. A Method and System for SLA-based Agile and Component Level Resource Provisioning in Disaggregated Systems, Jul-15, Co-inventors: Ramasamy, H. Salapura, V. Bivens, J. Ruan, Y. Mahindru, R. Schenfeld, E
9. A Method and System of Stage-aware Performance Modeling for DAG based Data Analytic Platforms, May-15 , Co-inventors: Salapura, V. WANG, Y. Zhang, L.
10. Stage-aware Performance Modeling For Computer Cluster Sizing, filed in May-2015, Co-inventors: Salapura, V. WANG, Y. Zhang, L.
11. Cooking Helper With Bidirectional Calendar Interface, filed in May-2015, Co-inventors: KIM, M. Pickover, C. Salapura, V.
12. Baby Comforting System with Active Learning, filed in Apr-15, Co-inventors: KIM, M. Pickover, C. Salapura, V.
13. Comforting System With Active Learning ,filed in April 2015 , Co-inventors: KIM, M. Pickover, C. Salapura, V.
14. Unmanned Aerial Vehicle For Interacting With A Pet ,filed in April 2015, Co-inventors: BIVENS, J. KIM, M. Pickover, C. Salapura, V.
15. Vertical Tuning Of Distributed Analytics Clusters, ,filed in March 2015 , Co-inventors: ZHANG, R.
16. Automatic Trigger Of Integrated Game Actions For Exercise And Well Being, ,filed in March 2015, Co-inventors: KIM, M. Pickover, C. Salapura, V.

17. Television Program Optimization For User Exercise ,filed in Feb 2015, Co-inventors: KIM, M. Pickover, C. Salapura, V.
18. Visual Comparisons Using Personal Objects,filed in Jan 2015,Co-inventors: KIM, M. Pickover, C. Salapura, V.
19. Floor Covering Having Adjustable Hardness,filed in Dec 2014, Co-inventors: KIM, M. Pickover, C. Salapura, V.
20. Cognitive Alert Control Framework For Mobile Devices ,filed in Dec 2014 , Co-inventors: BIVENS, J. KIM, M. Pickover, C. Salapura, V.
21. Outcomes-based Application Monitoring ,filed in Oct 2014 , Co-inventors: BIVENS, J. KIM, M. Pickover, C. Salapura, V.
22. Container Migration And Provisioning ,filed in Oct 2014 , Co-inventors: Jain, R. ROUTRAY, R. R. SONG, Y. Tan, C.
23. Dynamic Tuning Of Memory In Mapreduce Systems ,filed in Jul 2013 , Co-inventors: FULLER, N. C. MENG, S. TAN, J. Zeng, L. Zhang, L.
24. Dynamic Resource Allocation In Mapreduce ,filed in Jul 2013 , Co-inventors: FULLER, N. C. MENG, S. TAN, J. Zeng, L. Zhang, L.
25. Enabling Dynamic Job Configuration In Mapreduce ,filed in July 2013 ,Co-inventors: FULLER, N. C. KIM, M. MENG, S. TAN, J. Zeng, L. Zhang, L.
26. Managing Virtual Machine Placement In A Virtualized Computing Environment ,filed in Aug 2012 ,Co-inventors: BUTT Chavda, K. ZHOU, P.
27. A System And Method For De-virtualized Mapreduce I/o In The Cloud ,published in Dec 2012, Co-inventors: BUTT JAIN, R. Sarkar, P. Subhraveti, D.
28. Efficient Execution Of Jobs In A Shared Pool Of Resources, Issued as Patent 8972983 in US in March 2015, Co-inventors: Sarkar, P. Subhraveti, D.

PROFESSIONAL EXPERIENCE

Technical Program Committee: IEEE BigData17, SC15, NAS15, FIT15, NVMSA15, PSDW15, IGCC14 WS.

Publicity Chair: MASCOTS15, APC15

Peer Reviewer: CLUSTER14, CCGrid14, FAST13, MASCOTS12, ICDCS12,ICDE2016, IPDPS2016, TPDS, Cluster Computing, TOCC, JCST.

TECHNICAL SKILLS

Programming languages: Python, Java, C/C++, Bash.

Machine learning libraries: Tensorflow, Pytorch, Scikit-learn

Big data analytic frameworks: Hadoop, HDFS, Spark.

Software tools: Latex, Eclipse, IntelliJ, MATLAB, AMPL.

Operating systems: Linux, Windows