

Jee Whan Choi

Curriculum Vitae

February 2024

📍 Computer & Information Science, University of Oregon, Eugene OR, USA.
📞 +1 (404) 729-0795
✉ jeec@uoregon.edu
🏠 jeewhanchoi.com

Research Interest

My research interest lies in high-performance computing (HPC), with emphasis on predicting the behavior of complex software-hardware interactions. I have developed and optimized numerous scientific kernels on the latest CPUs and accelerators, and modeled their performance and energy efficiency. I am currently studying multi-dimensional tensor decomposition algorithms as means of informed learning and Big Data analytics, as well as developing and optimizing large-scale implementations on the latest supercomputers.

Education

- May 2015 **Ph.D.** Electrical & Computer Engineering Georgia Institute of Technology
THESIS - *Power and Performance Modeling for High-Performance Computing Algorithms*
- May 2004 **M.S.** Electrical & Computer Engineering Georgia Institute of Technology
THESIS - *Reducing Communication Through Buffers on a SIMD Architecture*
- May 2000 **B.Sc.** Electrical & Computer Engineering Georgia Institute of Technology
Highest Honors

Employment History

- Dec. 2018 – **Assistant Professor**
Current Computer & Information Science, *University of Oregon*
- Nov. 2017 – **Research Staff Member**
Dec. 2018 IBM T. J. Watson Research Center, *IBM Research*
- May 2015 – **Postdoctoral Researcher**
Nov. 2017 IBM T. J. Watson Research Center, *IBM Research*
- Aug. 2006 – **Graduate Research/Teaching Assistant**
May 2015 Computational Science & Engineering, *Georgia Institute of Technology*
- May – **Graduate Intern**
Dec. 2010 Parallel Computing Lab, *Intel*
- Aug. 2004 – **Researcher**
Aug. 2006 Network Research Center, *Samsung Electronics*
- Aug. 2000 – **Graduate Teaching Assistant**
May 2004 Electrical and Computer Engineering, *Georgia Institute of Technology*

Publications

Google Scholar <https://scholar.google.com/citations?user=yCA985IAAAAJ&hl=en>
Citations 1650
H-index 17
i10-index 19

Work in progress

1. Chen, A, BA Erickson, JE Kozdon, and JW Choi (2024). Matrix-free SBP-SAT finite difference methods and the multigrid preconditioner on GPUs. In: *Submitted to ICS'24*.
2. Laukemann, J, AE Helal, SIG Anderson, F Checconi, P Soh, JJ Tithi, T Ranadive, F Petrini, and JW Choi (2024). Accelerating Sparse Tensor Operations Using Adaptive Linearized Representation. In: *Submitted to TPDS*.
3. McLaughlin, J and JW Choi (2024). A high-performance characterization of hybridized PDEs. In: *Submitted to ICS'24*.
4. Soh, Y, R Kannan, P Sao, and JW Choi (2024). Accelerating Constrained Sparse Tensor Factorization on Massively Parallel Architectures. In: *Submitted to ICS'24*.

Papers in refereed conference proceedings

1. Dutta, A, J Choi, and A Jannesari (2023). Power Constrained Autotuning using Graph Neural Networks. In: *The 37th IEEE International Parallel and Distributed Processing Symposium (IPDPS'23)*.
2. McLaughlin, J, J Choi, and R Durairajan (2023). xGrid: A Location-Oriented Topology Design for LEO Satellites. In: *Proceedings of the 1st ACM Workshop on LEO Networking and Communication (LEO-NET'23)*.
3. Soh, P, AE Helal, J Laukemann, F Checconi, JJ Tithi, T Ranadive, F Petrini, and JW Choi (2023). Dynamic Tensor Linearization and Time Slicing for Efficient Factorization of Infinite Data Streams. In: *The 37th IEEE International Parallel and Distributed Processing Symposium (IPDPS'23)*.
4. Nguyen, A, AE Helal, F Checconi, J Laukemann, JJ Tithi, Y Soh, T Ranadive, F Petrini, and JW Choi (2022). Efficient, out-of-Memory Sparse MTTKRP on Massively Parallel Architectures. In: *Proceedings of the 36th ACM International Conference on Supercomputing (ICS'22)*.
5. Anderson, SIG, K Teranishi, DM Dunlavy, and JW Choi (Sept. 2021). Extended Abstract: Performance-Portable Sparse Tensor Decomposition Kernels on Emerging Parallel Architectures. In: *The 25th Annual IEEE Conference on High Performance Extreme Computing (HPEC'21)*.
6. Helal, AE, J Laukemann, F Checconi, JJ Tithi, T Ranadive, F Petrini, and JW Choi (June 2021). ALTO: Adaptive Linearized Storage of Sparse Tensors. In: *The 35th ACM International Conference on Supercomputing (ICS'21)*.
7. Soh, Y, P Flick, X Liu, S Smith, F Checconi, F Petrini, and JW Choi (May 2021). High Performance Streaming Tensor Decomposition. In: *The 35th IEEE International Parallel and Distributed Processing Symposium (IPDPS'21)*.
8. Chakaravarthy, VT, JW Choi, DJ Joseph, P Murali, Y Sabharwal, S Shivmaran, and D Sreedhar (June 2018). On Optimizing Distributed Tucker Decomposition for Sparse Tensors. In: *The 32nd ACM International Conference on Supercomputing (ICS'18)*.
9. Choi, JW, X Liu, and VT Chakaravarthy (Nov. 2018). High-performance Dense Tucker Decomposition on GPU Clusters. In: *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'18)*.
10. Choi, JW, X Liu, S Smith, and T Simon (May 2018). Blocking Optimization Techniques for Sparse Tensor Computation. In: *The 32nd IEEE International Parallel and Distributed Processing Symposium (IPDPS'18)*.
11. Buono, D, F Artico, F Checconi, JW Choi, X Que, and L Schneidenbach (May 2017). Data Analytics with NVLink: An SpMV Case Study. In: *14th ACM International Conference on Computing Frontiers (CF'17)*.
12. Chakaravarthy, VT, JW Choi, DJ Joseph, X Liu, P Murali, Y Sabharwal, and D Sreedhar (May 2017). On Optimizing Distributed Tucker Decomposition for Dense Tensors. In: *The 31st IEEE International Parallel and Distributed Processing Symposium (IPDPS'17)*.
13. Li, J, JW Choi, I Perros, J Sun, and R Vuduc (May 2017). Model-Driven Sparse CP Decomposition for Higher-Order Tensors. In: *The 31st IEEE International Parallel and Distributed Processing Symposium (IPDPS'17)*.

14. Liu, X, D Buono, F Checconi, JW Choi, X Que, F Petrini, JA Gunnels, and JA Stuecheli (May 2016). An Early Performance Study of Large-Scale POWER8 SMP Systems. In: *The 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS'16)*.
15. Choi, JW, M Dukhan, X Liu, and R Vuduc (May 2014). Algorithmic Time, Energy, and Power on Candidate HPC Compute Building Blocks. In: *The 28th IEEE International Parallel and Distributed Processing Symposium (IPDPS'14)*.
16. Choi, JW, D Bedard, R Fowler, and R Vuduc (May 2013). A Roofline Model of Energy. In: *The 27th IEEE International Symposium on Parallel and Distributed Processing (IPDPS'13)*.
17. Chandramowliswaran, A, JW Choi, K Madduri, and R Vuduc (June 2012). Brief Announcement: Towards a Communication Optimal Fast Multipole Method and Its Implications at Exascale. In: *The 24th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'12)*.
18. Smelyanskiy, M, K Vaidyanathan, JW Choi, B Joó, J Chhugani, MA Clark, and P Dubey (Nov. 2011). High-performance lattice QCD for multi-core based parallel systems using a cache-friendly hybrid threaded-MPI approach. In: *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'11)*.
19. Choi, JW, A Singh, and RW Vuduc (Jan. 2010). Model-driven autotuning of sparse matrix-vector multiply on GPUs. In: *The 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'10)*.
20. Vuduc, R, A Chandramowliswaran, JW Choi, M Guney, and A Shringarpure (June 2010). On the limits of GPU acceleration. In: *The 2nd USENIX conference on Hot topics in parallelism (HotPar'10)*.
21. Choi, J, S Apewokin, B Valentine, DS Wills, and L Wills (Jan. 2008). Edge noise removal in multimodal background modeling techniques. In: *Electronic Imaging*.
22. Valentine, B, JW Choi, S Apewokin, L Wills, and S Wills (Sept. 2008). Bypassing BigBackground: An efficient hybrid background modeling algorithm for embedded video surveillance. In: *The 2nd ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC'08)*.

Papers in refereed workshop proceedings

1. Choi, JW and RW Vuduc (May 2016). Analyzing the Energy Efficiency of the Fast Multipole Method Using a DVFS-Aware Energy Model. In: *The 30th IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW'16)*.
2. Choi, JW, A Chandramowliswaran, K Madduri, and R Vuduc (Mar. 2014). A CPU: GPU Hybrid Implementation and Model-Driven Scheduling of the Fast Multipole Method. In: *The 7th Workshop on General Purpose Processing Using GPUs (GPGPU'14)*.
3. Choi, JW and RW Vuduc (May 2012). Modeling and Analysis for Performance and Power. In: *The 26th IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW'12)*.
4. Vuduc, RW, K Czechowski, A Chandramowliswaran, and JW Choi (May 2012). Courses in High-performance Computing for Scientists and Engineers. In: *The 26th IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW'12)*.

Papers in refereed journals

1. Apewokin, S, B Valentine, JW Choi, L Wills, and S Wills (Jan. 2011). Real-Time Adaptive Background Modeling for Multicore Embedded Systems. *Journal of Signal Processing Systems* **62** (1), 65–76.

Ph.D. & M.S. thesis

1. Choi, JW (May 2015). “Power and performance modeling for high-performance computing algorithms”. PhD thesis. Georgia Institute of Technology.
2. Choi, JW (May 2004). “Reducing Communication Through Buffers on a SIMD Architecture”. MA thesis. Georgia Institute of Technology.

Articles

1. Choi, JW and RW Vuduc (Mar. 2013). How Much (Execution) Time and Energy Does My Algorithm Cost? *XRDS* **19**(3), 49–51.

Books & book chapters

1. Vuduc, R and JW Choi (Nov. 2013). "A Brief History and Introduction to GPGPU". In: *Modern Accelerator Technologies for Geographic Information Science*. Ed. by X Shi, V Kindratenko, and C Yang. Boston, MA: Springer US, pp. 9–23.
2. Kim, H, R Vuduc, S Baghsorkhi, and JW Choi (Nov. 2012). *Performance Analysis and Tuning for General Purpose Graphics Processing Units (GPGPU)*.
3. Williams, SW, N Bell, JW Choi, M Garland, L Oliker, and R Vuduc (Dec. 2010). "Sparse Matrix-Vector Multiplication on Multicore and Accelerators". In: *Scientific Computing with Multicore and Accelerators*. Boca Raton, FL: CRC Press, pp. 1–27.

Posters in peer reviewed venues

1. Tsai, YM, JW Choi, X Liu, and W Wang (Nov. 2017). Accelerating the Higher Order Singular Value Decomposition Algorithm for Big Data with GPUs. In: *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'17)*.
2. Choi, JW and R Vuduc (Feb. 2014). What Can the Roofline Model of Energy Tell Us About How to Build the Next Supercomputer? In: *SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP'14)*.
3. Choi, JW and R Vuduc (Nov. 2013). Power and Performance Modeling for High Performance Computing Algorithms. In: *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'17)*.
4. Choi, JW and R Vuduc (May 2012). Modeling and Analysis for Performance and Power. In: *The 26th IEEE International Symposium on Parallel and Distributed Processing (IPDPS'12)*.
5. Choi, JW and R Vuduc (Jan. 2010). Model-driven Autotuning of Sparse Matrix-Vector Multiply on GPUs. In: *The 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'10)*.

Grants

I have acquired approximately \$1.02M in external research grants since 2020 from the National Science Foundation, Department of Energy, and industry collaborators.

Source	Status	Project Period	Title	Amount
DoE	Pending		Pre-Application: AI-Driven Code Generation for Code Verification and Performance Optimization using Large Language Models	
DoE	Pending		Pre-Application: Julia-Based End-to-End Framework for Data Reduction using Randomized Tensor Decomposition Algorithms	
DoE	Rejected		Pre-Application: Probabilistic Modeling of Large-Scale Distributed Systems using Hyperbolic Functions and Algorithm-Based Fault-Tolerant Application Design	
Intel	Funded	07/2023 – 06/2024	Linearized Sparse Tensor Operations for Large-scale Systems	\$ 157,568
NSF	Rejected	10/2023 – 09/2028	Mid-scale RI-1 (M1:IP): National Software Sustainability Infrastructure	\$ 2,999,999
NSF	Rejected	10/2022 – 09/2025	Collaborative Research: SHF: Small: Energy and Power Aware Parallelization for Multicore Processors	\$ 298,310
DoE	Funded	10/2022 – 09/2025	In Situ Data Analysis Through Physics-constrained and Goal-oriented Tensor Decomposition	\$ 300,000
Intel	Funded	04/2022 – 03/2023	Optimizing Streaming Tensor Decomposition	\$ 158,669
NSF	Funded	10/2021 – 9/2021	Extreme-scale Sparse Data Analytics	\$ 75,000
Intel	Funded	11/2020 – 10/2021	Optimizing Streaming Tensor Decomposition	\$ 157,500
DoE	Funded	08/2020 – 07/2021	Analysis of High Performance Portable Tensor Decomposition Implementations	\$ 81,500
Intel	Funded	01/2020 – 07/2020	Optimizing Streaming Tensor Decomposition	\$ 88,500
Total				\$ 1,018,737

Teaching

2024	Spring	CIS 330	<i>C/C++ and Unix</i>
		CIS 607	<i>High-Performance Tensor Decomposition and Data Analytics</i>
	Winter	CIS 330	<i>C/C++ and Unix</i>
2023	Fall	CS 531/631	<i>Introduction to Parallel Computing</i>
2022	Spring	CIS 330	<i>C/C++ and Unix</i>
	Winter	CIS 330	<i>C/C++ and Unix</i>
		CIS 607	<i>High-Performance Tensor Decomposition and Data Analytics</i>
2021	Fall	CIS 431/531	<i>Introduction to Parallel Computing</i>
	Spring	CIS 330	<i>C/C++ and Unix</i>
	Winter	CIS 330	<i>C/C++ and Unix</i>
		CIS 607	<i>High-Performance Tensor Decomposition and Data Analytics</i>
2020	Fall	CIS 631	<i>Parallel Processing</i>
	Spring	CIS 330	<i>C/C++ and Unix</i>
		CIS 607	<i>High-Performance Tensor Decomposition and Data Analytics</i>
	Winter	CIS330	<i>C/C++ and Unix</i>
2019	Fall	CIS 631	<i>Parallel Processing</i>
	Spring	CIS 330	<i>C/C++ and Unix</i>

Services

Departmental Committees

I have served on the following departmental committees.

- ▶ **Graduate Education** (Fall 2021 – Current) Responsible for the Annual Ph.D. Student Review; helped develop the GIS system for web-based Annual Ph.D. Student Review process; helped develop the Graduate Student Annual Review Handbook
- ▶ **Diversity** (Fall 2020 – Spring 2021) Responsible for the NSF BPC Plan & Data Buddies Survey
- ▶ **Programming Competitions** (Fall 2020 – Current)
- ▶ **Graduate Education** (Winter 2019 – Spring 2020) Responsible for the Annual Ph.D. Student Review
- ▶ **Computing Resources** (Winter 2019 – Spring 2020)

Student Advising & Committees

I have advised and served on the committees for the following students at the University of Oregon.

- ▶ **Joseph McLaughlin** (09/2021 – Current) *Ph.D. Advisor*
- ▶ **Yongseok Soh** (09/2019 – Current) *Ph.D. Advisor*
- ▶ **Isaac Geronimo Anderson** (09/2019 – Current) *Ph.D. Advisor*
- ▶ **Konstantin Shvedov** (11/2021 – 06/2022) Master's Thesis Advisor
- ▶ **Andy Nguyen** (06/2020 – Current) Undergraduate Research Advisor & Thesis Advisor
- ▶ **Hammad Ather** (02/2024 – Current) Dissertation Advisory Committee & Area Exam Committee
- ▶ **Sudharshan Srinivasan** (06/2021 – Current) Directed Research Project & Dissertation Advisory Committee & Area Exam Committee

- *Alexandre Chern* (02/2021 – Current) Dissertation Advisory Committee
- *Kewen Meng* (02/2020 – 06/2021) Dissertation Advisory Committee
- *Abhishek Yenpure* (04/2019 – 12/2022) Directed Research Project & Dissertation Advisory Committee & Area Exam Committee

Academic Services

Technical program committee for conferences

- | | |
|------|--|
| 2024 | <ul style="list-style-type: none"> ● The 24th IEEE/ACM international Symposium on Cluster, Cloud and Internet Computing (CCGRID'24) May 2024 ● The 38th IEEE International Parallel and Distributed Processing Symposium (IPDPS'23) May 2024 ● SIAM Conference on Parallel Processing for Scientific Computing (PP24) Mar 2024 ● The International Conference on High-Performance Computing in Asia-Pacific Region (HPC Asia 2024) Jan. 2024 |
| 2023 | <ul style="list-style-type: none"> ● 30th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'23), Program Vice-Chair, Algorithms Dec. 2023 ● The 37th IEEE International Parallel and Distributed Processing Symposium (IPDPS'23) May 2023 ● The 52nd International Conference on Parallel Processing (ICPP'23) Aug. 2023 |
| 2022 | <ul style="list-style-type: none"> ● 29th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'22) Dec. 2022 ● The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'22) Nov. 2022 ● IEEE International Conference on Cluster Computing Sep. 2022 ● The 51st International Conference on Parallel Processing (ICPP'22) Aug. 2022 ● The 36th IEEE International Parallel and Distributed Processing Symposium (IPDPS'22) May 2022 ● High Performance Computing, Machine Learning and Data Analytics (ISC'22), <i>Poster Session</i> May 2022 |
| 2021 | <ul style="list-style-type: none"> ● 28th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'21) Dec. 2021 ● The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'21) Nov. 2021 ● The 35th IEEE International Parallel and Distributed Processing Symposium (IPDPS'21) May 2021 |
| 2020 | <ul style="list-style-type: none"> ● 28th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'20) Dec. 2020 ● The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'20) Nov. 2020 |
| 2018 | <ul style="list-style-type: none"> ● The 47th International Conference on Parallel Processing (ICPP'18) Aug. 2018 ● The 11th International Conference on Contemporary Computing (IC3'18) Aug. 2018 ● The 32nd IEEE International Parallel and Distributed Processing Symposium (IPDPS'18) May 2018 |
| 2017 | <ul style="list-style-type: none"> ● The 47th International Conference on Parallel Processing (ICPP'18) Aug. 2018 ● The 10th International Conference on Contemporary Computing (IC3'17) Aug. 2017 |
| 2016 | <ul style="list-style-type: none"> ● 23rd IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'16) Dec. 2016 |
| 2014 | <ul style="list-style-type: none"> ● The 7th International Conference on Contemporary Computing (IC3'14) Aug. 2014 |

Technical program committee for workshops

- 2021** • International Workshop on Parallel Programming Models in High-Performance Cloud (ParaMo'21) Aug. 2021
- 2020** • Special Session: Auto-Tuning for Multicore and GPU (ATMG'20) Dec. 2020
 • International Workshop on Parallel Programming Models in High-Performance Cloud (ParaMo'20) Aug. 2020
- 2019** • Special Session: Auto-Tuning for Multicore and GPU (ATMG'19) Dec. 2019
 • International Workshop on Parallel Programming Models in High-Performance Cloud (ParaMo'19) Aug. 2019
 • The First International Workshop on the Intersection of High Performance Computing and Machine Learning (HPCaML'19) Feb. 2019
- 2017** • The 13th Workshop on High-Performance, Power-Aware Computing (HP-PAC'17) May 2017
- 2015** • The 3rd International Workshop on Energy Efficient Supercomputing (E2SC'15) May 2015

Organizing Committee

I have organized the following meetings.

- Minisymposium @ SIAM Conference on Computational Science and Engineering (2021, 2019)
- Minisymposium @ SIAM Conference on Parallel Processing for Scientific Computing (2020, 2018)
- Workshop on Power and Energy Aspect of Computation (2015)

Journal Review

I am a regular reviewer for the following journals.

- ACM Computing Surveys (CSUR) (2021, 2017, 2016)
- Concurrency and Computation: Practice and Experience (CCPE) (2020, 2018, 2017)
- IEEE Transaction on Computers (TC) (2015)
- IEEE Transaction on Parallel and Distributed Systems (TPDS) (2024 – 2017, 2015)
- IET Computers & Digital Techniques (2016, 2015)
- International Journal of High Performance Computing Applications (IJHPCA) (2019, 2017, 2016)
- Journal of Parallel and Distributed Computing (JPDC) (2023, 2022, 2018, 2014)
- Knowledge and Information Systems (KAIS) (2019, 2018)
- Scientific Programming (2017)
- SIAM Journal on Scientific Computing (SISC) (2015)
- Transactions on Architecture and Code Optimization (TACO) (2024, 2020 – 2018, 2016, 2015)
- Transaction on Mathematical Software (TOMS) (2022, 2021, 2017)

Miscellaneous Activities

- I have served as a reviewer for the 2023 *The Computer and Information Science and Engineering Graduate Fellowships (CSGrad4US)*.
- I have served as a reviewer for the 2021 — 2023 *SIGHPC Computational & Data Science Fellowship*.
- I have served as an external reviewer for many other conferences and workshops, most notably *IEEE Micro, ICPP, IPDPS, and SC*.

Invited Talks

- *Portable High-Performance Format for Sparse Tensor Computation*, Workshop on Sparse Tensor Computations, Chicago, IL Oct. 2023.
- *Practical, Out-of-Memory Sparse MTTKRP on Massively Parallel Architectures*, SIAM Parallel Processing, Seattle WA, Feb. 2022.
- *Tensor Decomposition for Topic Modeling in AI*, Communicating AI: Theory, Research, and Practice, Los Angeles, CA, Feb. 2020.
- *Malware Analysis Using Tensor Decomposition*, Workshop on AI and Tensor Factorization for Physical, Chemical, and Biological Systems, Santa Fe, NM, Sep. 2019.
- *Optimizing Tensor Decomposition on HPC Systems - Challenges and Approaches*, Workshop on High-Level Parallel Programming Models and Supportive Environments, Rio de Janeiro, Brazil, May 2019.
- *Blocking Optimizations for Sparse MTTKRP*, Invited Workshop on Compiler Techniques for Sparse Tensor Algebra, Cambridge, MA, Jan. 2019.

Mentoring Experience

I was a mentor to *nine* undergraduate and M.S. students before joining the University of Oregon, including many from underrepresented minority groups. Most notably, I mentored *Jiajia Li* during her summer internship at IBM Research, which led to her receiving the prestigious **IBM Fellowship**. Jiajia went on to become a tenure-track faculty at the College of William and Mary in 2021.

Professional Memberships

I am a member of the following organizations.

- Association of Computing Machinery (ACM)
- Institute of Electrical and Electronics Engineers (IEEE)
- Society for Industrial and Applied Mathematics (SIAM)