

Rio Yokota

CONTACT INFORMATION

Supercomputing Research Center
Institute of Integrated Research
Institute of Science Tokyo
2-12-1 Ishikawadai Bldg. 9 I7-2 O-okayama Meguro-ku
Tokyo 152-8550, Japan

Phone: +81 (0) 3 5734-2121

E-mail: rioyokota@rio.scrs.iir.isct.ac.jp

RESEARCH INTERESTS

High performance computing; Large Scale Deep Learning; Numerical Analysis

EDUCATION

Keio University, Yokohama, Japan

Ph.D., Mechanical Engineering, July, 2009

- Dissertation Topic: “Validation of Vortex Methods as a Direct Numerical Simulation of Turbulence”
- Advisor: Shinnosuke Obi

HONORS AND AWARDS

ACM Gordon Bell prize, price/performance (2009)

ACADEMIC EXPERIENCE

- 4/2025 -** Team Principal, RIKEN
- 10/2024 -** Professor, Institute of Science Tokyo
- 4/2024 -** Scientific Director, National Institute of Informatics
- 1/2023 - 9/2024** Professor, Tokyo Institute of Technology
- 4/2015 - 12/2022** Associate Professor, Tokyo Institute of Technology
- 9/2011 - 3/2015** Research Scientist, KAUST
- 9/2010 - 8/2011** Post-doctoral Researcher, Boston University
- 2/2009 - 8/2010** Post-doctoral Researcher, University of Bristol

BOOK CHAPTERS

1. R. Yokota, M. AbdulJabbar, *N*-body methods on Xeon Phi coprocessors, *High Performance Parallelism Pearls* (2014)
2. R. Yokota and L. A. Barba, Treecode and fast multipole method for N-body simulation with CUDA, *GPU Computing Gems* (2011)

PEER REVIEWED JOURNALS

1. Satoshi Yui, Hiromichi Kobayashi, Makoto Tsubota, Rio Yokota, Quantum turbulence coupled with externally driven normal-fluid turbulence in counterflow of superfluid ^4He , *Journal of the Physical Society of Japan*, 94, 043601 (2025).
2. Takuya Akiba, Makoto Shing, Yujin Tang, Qi Sun, David Ha, Evolutionary Optimization of Model Merging Recipes, *Nature Machine Intelligence* (2025).
3. Kenta Niwa, Hiro Ishii, Hiroshi Sawada, Akinori Fujino, Noboru Harada, Rio Yokota, ,Natural Gradient Primal-Dual Method for Decentralized Learning, *Transactions on Signal and Information Processing over Networks*, (2024).

4. Qianxiang Ma, Rio Yokota, An Inherently Parallel $\mathcal{H}^2$ -ULV Factorization for Solving Dense Linear Systems on GPUs, *International Journal of High Performance Computing Applications*, (2024).
5. Hiroyuki Ootomo, Katsuhisa Ozaki, Rio Yokota, DGEMM on Integer Matrix Multiplication Unit, *International Journal of High Performance Computing Applications*, (2024).
6. Erik Daxberger, Siddharth Swaroop, Kazuki Osawa, Rio Yokota, Richard E. Turner, Jose Miguel Hernandez-Lobato, Mohammad Emtiyaz Khan, Improving Continual Learning by Accurate Gradient Reconstructions of the Past, *Transactions on Machine Learning Research*, (2023).
7. Hiroki Naganuma, Kartik Ahuja, Shiro Takagi, Tetsuya Motokawa, Rio Yokota, Kohta Ishikawa, Ikuro Sato, Ioannis Mitliagkas, Empirical Study on Optimizer Selection for Out-of-Distribution Generalization, *Transactions on Machine Learning Research*, (2023).
8. Sameer Deshmukh, Rio Yokota, George Bosilca, Cache Optimization and Performance Modeling of Batched, Small, and Rectangular Matrix Multiplication on Intel, AMD, and Fujitsu Processors, *ACM Transactions on Mathematical Software*, (2023).
9. Muhammad Ridwan Apriansyah, Rio Yokota, Parallel QR Factorization of Block Low-Rank Matrices, *ACM Transactions on Mathematical Software*, (2022). <https://doi.org/10.1145/3538647>
10. Hiroyuki Ootomo, Rio Yokota, Recovering Single Precision Accuracy from Tensor Cores While Surpassing the FP32 Theoretical Peak Performance, *The International Journal of High Performance Computing Application*, (2022).
11. Tingyu Wang, Rio Yokota, Lorena A. Barba, ExaFMM: a high-performance fast multipole method library with C++ and Python interfaces, *The Journal of Open Source Software*, Vol. 6, No. 61, pp. 3145 (2021).
12. Kazuki Osawa, Yohei Tsuji, Yuichiro Ueno, Akira Naruse, Chuan-Sheng Foo, Rio Yokota, Scalable and Practical Natural Gradient for Large-Scale Deep Learning, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020); PP:10.1109/TPAMI.2020.3004354.
13. Davoud S. Shamshirgar, Rio Yokota, Anna-Karin Tornberg, and Berk Hess, Regularizing the Fast Multipole Method for use in Molecular Simulation, *The Journal of Chemical Physics*, 151, 234113 (2019)
14. Akihiro Ida, Hiroshi Nakashima, Tasuku Hiraishi, Ichitaro Yamazaki, Rio Yokota, Takeshi Iwashita, QR Factorization of Block Low-rank Matrices with Weak Admissibility Condition, *Journal of Information Processing*, Vol. 12, No. 4 (2019)
15. Ichitaro Yamazaki, Akihiro Ida, Rio Yokota, Jack Dongarra, Distributed Memory Lattice H-matrix Factorization, *The International Journal of High Performance Computing Applications* (2019).
16. Mustafa Abduljabbar, Mohammed Al Farhan, Noha Al-Harthi, Rui Chen, Rio Yokota, Hakan Bagci, David Keyes, Extreme Scale FMM-Accelerated Boundary Integral Equation Solver for Wave Scattering, *SIAM Journal on Scientific Computing*, Vol. 41, No. 3, pp. C245–C268 (2019).
17. Naoya Maruyama, Takayuki Aoki, Kenjiro Taura, Rio Yokota, Mohamed Wahib, Motohiko Matsuda, Keisuke Fukuda, Takashi Shimokawabe, Naoyuki Onodera, Michel Müller, Shintaro Iwasaki, Highly Productive, High-Performance Application Frameworks for Post-Petascale Computing, *Advanced Software Technologies for Post-Peta Scale Computing*, pp. 77–98 (2018)
18. Rio Yokota, Huda Ibeid, David Keyes, Fast Multipole Method as a Matrix-Free Hierarchical Low-Rank Approximation *Lecture Notes in Computational Science and Engineering*, Springer, Vol. 117, pp. 227-244 (2017)
19. Huda Ibeid, Rio Yokota, Jennifer Pestana, David Keyes, Fast Multipole Preconditioners for Sparse Matrices Arising from Elliptic Equations, *Computing and Visualization in Science* (2017). <https://doi.org/10.1007/s00791-017-0287-5>

20. Rio Yokota, Communication Optimization of Distributed Memory FMM for Large Scale Boundary Element Methods, *Simulation*, Vol. 35, No. 3, pp. 23–29 (2016)
21. Rio Yokota, Tradeoff between FMM and $\mathcal{H}^2(HSS)$ -matrices, *Journal of the Japan Society for Computational Engineering and Science*, Vol. 21, No. 4, pp. 3498–3501 (2016)
22. Julio E. Castrillon-Candas, Marc G. Genton, Rio Yokota, Multi-level Restricted Maximum Likelihood Covariance Estimation and Kriging for Large Non-gridded Spatial Datasets, *Spatial Statistics*, Vol. 18, pp. 105–124 (2016)
23. Huda Ibeid, Rio Yokota, David Keyes, A Performance Model for the Communication in Fast Multipole Methods on HPC Platforms, *International Journal of High Performance Computing Applications*, Vol. 30, No. 4, pp. 423–437 (2016)
24. Abdelhalim Amer, Satoshi Matsuoka, Miquel Pericàs, Naoya Maruyama, Kenjiro Taura, Rio Yokota, Pavan Balaji, Scaling FMM with Data-Driven OpenMP Tasks on Multicore Architectures, *Lecture Notes in Computer Science*, Vol. 9903 LNCS, pp. 156–170 (2016).
25. Rio Yokota, George Turkiyyah, David Keyes, Communication Complexity of the Fast Multipole Method and its Algebraic Variants, *Supercomputing Frontiers and Innovations*, Vol. 1, No. 1, pp. 63–84 (2014)
26. Yousuke Ohno, Rio Yokota, Hiroshi Koyama, Gentaro Morimoto, Aki Hasegawa, Gen Masumoto, Noriaki Okimoto, Yoshinori Hirano, Huda Ibeid, Tetsu Narumi, Makoto Taiji, Petascale Molecular Dynamics Simulation Using the Fast Multipole Method on K Computer, *Computer Physics Communications*, Vol. 185, No. 10, pp. 2575–2585 (2014)
27. Hatem Ltaief, Rio Yokota, Data-Driven Execution of Fast Multipole Methods, *Concurrency and Computation: Practice and Experience*, Vol. 26, No. 11, pp. 1935–1946 (2014)
28. Rio Yokota, An FMM Based on Dual Tree Traversal for Many-Core Architectures, *Journal of Algorithms and Computational Technology*, Vol. 7, No. 3, pp. 301–324 (2013)
29. Rio Yokota, Lorena A. Barba, Tetsu Narumi, Kenji Yasuoka, Petascale Turbulence Simulation Using a Highly Parallel Fast Multipole Method, *Computer Physics Communications*, Vol. 184, No. 3, pp. 445–455 (2013)
30. Rio Yokota, Lorena A. Barba, FMM-based Vortex Method for Simulation of Isotropic Turbulence on GPUs, Compared with a Spectral Method, *Computers and Fluids*, Vol. 80, pp. 17–27 (2013)
31. Rio Yokota, Lorena A. Barba, A Tuned and Scalable Fast Multipole Method as a Preeminent Algorithm for Exascale Systems, *International Journal of High Performance Computing Applications*, Vol. 26, No. 4, pp. 337–346 (2012)
32. Rio Yokota, FMM Tree Construction on GPUs, *Ensemble*, Vol. 14, No. 2, pp. 85–89 (2012)
33. Rio Yokota, Lorena A. Barba, Hierarchical N-body Simulations with Auto-tuning for Heterogeneous Systems, *Computing in Science and Engineering*, Vol. 14, No. 3, pp. 30–39 (2012)
34. Rio Yokota, Jaydeep P. Bardhan, Matthew G. Knepley, Lorena A. Barba, Tsuyoshi Hamada, Biomolecular Electrostatics Using a Fast Multipole BEM on up to 512 GPUs and a Billion Unknowns, *Computer Physics Communications*, Vol. 182, No. 6, pp. 1272–1283 (2011)
35. Rio Yokota and Lorena A. Barba, Comparing the Treecode with FMM on GPUs for Vortex Particle Simulations of a Leapfrogging Vortex Ring, *Computers and Fluids*, Vol. 45, No. 1, pp. 155–161 (2011)
36. Rio Yokota and Shinnosuke Obi, Vortex Methods for the Simulation of Turbulent Flows, *Journal of Fluid Science and Technology*, Vol. 6, No. 1, pp.14–29 (2011)
37. Rio Yokota and Tsuyoshi Hamada, N-body Simulation and FMM on the Large-scale GPU Cluster at Nagasaki University, *Journal of the Japan Society for Computational Engineering and Science*, Vol. 15, No. 4 (2010)

38. Rio Yokota and Shinnosuke Obi, Comparing Vortex Methods and Finite Difference Methods in a Homogeneous Turbulent Shear Flow, *International Journal for Numerical Methods in Fluids*, Vol. 63, No. 7, pp. 828–846. (2010)
39. Rio Yokota, Lorena A. Barba, and Matthew G. Knepley, PetRBF—A Parallel $O(N)$ Algorithm for Radial Basis Function Interpolation with Gaussians, *Computer Methods in Applied Mechanics and Engineering*, Vol. 199, No. 25–28, pp. 1793–1804. (2010)
40. Rio Yokota, Tetsu Narumi, Ryuji Sakamaki, Shun Kameoka, Shinnosuke Obi, and Kenji Yasuoka, Fast Multipole Methods on a Cluster of GPUs for the Meshless Simulation of Turbulence, *Computer Physics Communications*, Vol. 180, No. 11, pp. 2066–2078. (2009)
41. Tarun K. Sheel, Rio Yokota, Kenji Yasuoka, and Shinnosuke Obi, The Study of Colliding Vortex Rings Using a Special-Purpose Computer and FMM, *Transactions of the Japan Society for Computational Engineering and Science*, 20080003 (2008).
42. Rio Yokota, Tarun K. Sheel, and Shinnosuke Obi, Calculation of Isotropic Turbulence Using a Pure Lagrangian Vortex Method, *Journal of Computational Physics*, Vol. 226, pp.1589–1606. (2007)

CONFERENCE
PROCEEDINGS
(PEER-REVIEWED)

1. Daisuke Nohara, Taishi Nakamura, Rio Yokota, Monotonic and Non-Monotonic Length-Performance Relationships in Reasoning RL, ICML 2026 Workshop AdaptFM, Jul. 2026.
2. Satoki Ishikawa, Sameer Satish Deshmukh, Sakina Fatima, Takumi Honda, Rio Yokota, Loss and Optimizer as Two Essential Mechanisms Behind Knowledge Distillation, ICML Workshop HiLD, Jul. 2026.
3. Masaki Kawamura, Nakamasa Inoue, Rintaro Yanagi, Hirokatsu Kataoka, Rio Yokota, PowerCLIP: Powerset Alignment for Fine-Grained Contrastive Pre-Training, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Jun. 2026.
4. Akiko Aizawa, Yuki Arase, Fei Cheng, Jiahao Huang, Zhiyi Huang, Junfeng Jiang, Teruhito Kanazawa, Daisuke Kawahara, Kazuma Kobayashi, Takashi Kodama, Sadao Kurohashi, Yusuke Oda, Yuma Tsuta, Zhen Wan, Zhishen Yang, Rio Yokota, Building Effective Japanese Medical LLMs with an Open Recipe for Domain Adaptation through Continued Pre-training, International Conference on Language Resources and Evaluation (LREC), May 2026.
5. Taishi Nakamura, Satoki Ishikawa, Masaki Kawamura, Takumi Okamoto, Daisuke Nohara, Jun Suzuki, Rio Yokota, Optimal Sparsity of Mixture-of-Experts Language Models for Reasoning Tasks, International Conference on Learning Representations (ICLR), Apr. 2026.
6. Kazuki Fujii, Yukito Tajima, Sakae Mizuki, Masaki Kawamura, Hinari Shimada, Taihei Shiotani, Koshiro Saito, Masanari Oi, Taishi Nakamura, Takumi Okamoto, Shigeki Ishida, Kakeru Hattori, Youmi Ma, Hiroya Takamura, Rio Yokota, Jun Sakuma, Naoaki Okazaki, Rewriting Pre-Training Data Boosts LLM Performance in Math and Code, International Conference on Learning Representations (ICLR), Apr. 2026.
7. Kazuki Fujii, Rio Yokota, Swallow LLM: Continual Pre-Training and RL for Sovereign AI, GPU Technology Conference (GTC), Mar. 2026.
8. Kazuki Fujii, Yukito Tajima, Masaki Kawamura, Sakae Mizuki, Taishi Nakamura, Hinari Shimada, Masanari Ohi, Taihei Shiotani, Koshiro Saito, Takumi Okamoto, Shigeki Ishida, Youmi Ma, Hiroya Takamura, Rio Yokota, Naoaki Okazaki, Qwen3-Swallow: Continual Pre-Training of Japanese–English Thinking LLMs with Megatron-LM, GPU Technology Conference, Mar. 2026.
9. Hiro Ishii, Kenta Niwa, Hiroshi Sawada, Akinori Fujino, Noboru Harada, Rio Yokota, FedPM: Federated Learning Using Preconditioned Mixing of Local Parameters, The 40th Annual AAAI Conference on Artificial Intelligence, Jan. 2026.

10. Akihiro Ida, Kazuya Goto, Rio Yokota, Tasuku Hiraishi, Toshihiro Hanawa, Takeshi Iwashita, Masatoshi Kawai, Satoshi Ohshima, Tetsuya Hoshino, Tensor-Core-Optimized Strategies for BLR $\times$ Tall-Skinny Matrix Multiplication in BEM, SCA/HPCAsia, Jan. 2026.
11. Yukito Tajima, Nakamasa Inoue, Yusuke Sekikawa, Ikuro Sato, Rio Yokota, Masked Gated Linear Unit, Annual Conference on Neural Information Processing Systems (NeurIPS), Dec. 2025.
12. Avrajit Ghosh, Bai Cong, Rio Yokota, Saiprasad Ravishankar, Rongrong Wang, Molei Tao, Mohammad Emtiyaz Khan, Thomas Möllenhoff, Variational Learning Finds Flatter Solutions at the Edge of Stability, Annual Conference on Neural Information Processing Systems (NeurIPS), Dec. 2025.
13. Satoshi Ohshima, Akihiro Ida, Masatoshi Kawai, Takeshi Fukaya, Rio Yokota, A Study on the Performance and Usability of Managed Memory and Unified Memory for Accelerating Numerical Calculation Program, 18th IEEE International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoc), Dec. 2025.
14. Kazuma Kobayashi, Zhen Wan, Fei Cheng, Yuma Tsuta, Xin Zhao, Junfeng Jiang, Jiahao Huang, Zhiyi Huang, Yusuke Oda, Rio Yokota, Yuki Arase, Daisuke Kawahara, Akiko Aizawa, Sadao Kurohashi, Leveraging High-Resource English Corpora for Cross-lingual Domain Adaptation in Low-Resource Japanese Medicine via Continued Pre-training, EMNLP Findings, Nov. 2025.
15. Youmi Ma, Sakae Mizuki, Kazuki Fujii, Taishi Nakamura, Masanari Ohi, Hinari Shimada, Taihei Shiotani, Koshiro Saito, Koki Maeda, Kakeru Hattori, Takumi Okamoto, Shigeki Ishida, Rio Yokota, Hiroya Takamura, Naoaki Okazaki, Building Instruction-Tuning Datasets from Human-Written Instructions with Open-Weight Large Language Models, Conference on Language Modeling (COLM), Oct. 2025.
16. Koshiro Saito, Sakae Mizuki, Masanari Ohi, Taishi Nakamura, Taihei Shiotani, Koki Maeda, Youmi Ma, Kakeru Hattori, Kazuki Fujii, Takumi Okamoto, Shigeki Ishida, Hiroya Takamura, Rio Yokota, Naoaki Okazaki, Why We Build Local Large Language Models: An Observational Analysis from 35 Japanese and Multilingual LLMs, COLM 2025 Workshop on Multilingual and Equitable Language Technologies (MELT), Oct. 2025.
17. Satoki Ishikawa, Tal Ben-Nun, Brian Van Essen, Rio Yokota, Nikoli Dryden, Lion Cub: Minimizing Communication Overhead in Distributed Lion, ICML 2025 Workshop TTODLer-FM, July 2025.
18. Taishi Nakamura, Satoki Ishikawa, Masaki Kawamura, Takumi Okamoto, Daisuke Nohara, Jun Suzuki, Rio Yokota, Optimal Sparsity of Mixture-of-Experts Language Models for Reasoning Tasks, ICML 2025 2nd AI for Math Workshop, July 2025.
19. Chen Zhuang, Lingqi Zhang, Du Wu, Peng Chen, Jiajun Huang, Xin Liu, Rio Yokota, Nikoli Dryden, Toshio Endo, Satoshi Matsuoka, Mohamed Wahib, Scaling Large-scale GNN Training to Thousands of Processors on CPU-based Supercomputers, 39th ACM International Conference on Supercomputing (ICS), June 2025.
20. Thomas Spendlhofer, Qianxiang Ma, Yasuhiro Matsumoto, Rio Yokota, On the Interplay Between Precision, Rank, Admissibility, and Iterative Refinement for Hierarchical Low-Rank Matrix Solvers, ISC High Performance, June 2025.
21. Qi Sun, Edoardo Cetin, Yujin Tang, Transformer² Self-adaptive LLMs, The 13th International Conference on Learning Representations (ICLR), Apr. 2025.
22. Edoardo Cetin, Qi Sun, Tianyu Zhao, Yujin Tang, An Evolved Universal Transformer Memory, The 13th International Conference on Learning Representations (ICLR), Apr. 2025.
23. Satoki Ishikawa, Ryo Karakida, Rio Yokota, Local Loss Optimization in the Infinite Width: Stable Parameterization of Predictive Coding Networks and Target Propagation, The 13th International Conference on Learning Representations (ICLR), Apr. 2025.

24. Satoki Ishikawa, Makoto Yamada, Han Bao, Yuki Takezawa. PhiNets: Brain-inspired Non-contrastive Learning Based on Temporal Prediction Hypothesis, The 13th International Conference on Learning Representations (ICLR), Apr. 2025.
25. Taishi Nakamura, Takuya Akiba, Kazuki Fujii, Yusuke Oda, Rio Yokota, Jun Suzuki, Drop Upcycling: Training Sparse Mixture of Experts with Partial Re-initialization, The 13th International Conference on Learning Representations (ICLR), Apr. 2025.
26. So Kuroki, Taishi Nakamura, Takuya Akiba, Yujin Tang. Agent Skill Acquisition for Large Language Models via CycleQD, The 13th International Conference on Learning Representations (ICLR), Apr. 2025.
27. Qi Sun, Marc Pickett, Aakash Kumar Nain, Llion Jones, Transformer Layers as Painters, The 39th Annual AAAI Conference on Artificial Intelligence (AAAI), Feb. 2025.
28. Chen Zhuang, Peng Chen, Xin Liu, Nikoli Dryden, Rio Yokota, Toshio Endo, Satoshi Matsuo, Mohamed Wahib, SuperGCN: General and Scalable Framework for GCN Training on CPU-powered Supercomputers, PPOPP, poster, Mar. 2025.
29. Shun Iwase, Shuya Takahashi, Nakamasa Inoue, Rio Yokota, Ryo Nakamura, Hirokatsu Kataoka, Eisaku Maeda, On the Relationship Between Double Descent of CNNs and Shape/Texture Bias Under Learning Process, 27th International Conference on Pattern Recognition (ICPR), Dec. 2024.
30. Bai Cong, Nico Daheim, Yuesong Shen, Daniel Cremers, Rio Yokota, Mohammad Emtiyaz Khan, Thomas Möllenhoff, Variational Low-Rank Adaptation Using IVON, NeurIPS 2024 Workshop on Fine-Tuning in Modern Machine Learning (FITML), Dec. 2024.
31. Kazuki Fujii, Taishi Nakamura, Rio Yokota, llm-recipes, A Framework for Seamless Integration and Efficient Continual Pre-Training of Large Language Models, SC'24 TPC Workshop, Nov. 2024.
32. Taishi Nakamura, Mayank Mishra, Simone Tedeschi, Yekun Chai, Jason T Stillerman, Felix Friedrich, Prateek Yadav, Tanmay Laud, Vu Minh Chien, Terry Yue Zhuo, Diganta Misra, Ben Bogin, Xuan-Son Vu, Marzena Karpinska, Arnav Varma Dantuluri, Wojciech Kusa, Tommaso Furlanello, Rio Yokota, Niklas Muennighoff, Suhas Pai, Tosin Adewumi, Veronika Laipala, Xiaozhe Yao, Adalberto Junior, Alpay Ariyak, Aleksandr Drozd, Jordan Clive, Kshitij Gupta, Liangyu Chen, Qi Sun, Ken Tsui, Noah Persaud, Nour Fahmy, Tianlong Chen, Mohit Bansal, Nicolo Monti, Tai Dang, Ziyang Luo, Tien-Tung Bui, Roberto Navigli, Virendra Mehta, Matthew Blumberg, Victor May, Huu Nguyen, Sampo Pyysalo, Aurora-M: The First Open Source Multilingual Language Model Red-teamed according to the U.S. Executive Order, The 31st International Conference on Computational Linguistics (COLING), Industry Track, Jan. 2025.
33. Naoaki Okazaki, Kakeru Hattori, Shota Hirai, Hiroki Iida, Masanari Ohi, Kazuki Fujii, Taishi Nakamura, Mengsay Loem, Rio Yokota, Sakae Mizuki, Building a Large Japanese Web Corpus for Large Language Models, Conference on Language Modeling COLM, Oct. 2024.
34. Kazuki Fujii, Taishi Nakamura, Mengsay Loem, Hiroki Iida, Masanari Ohi, Kakeru Hattori, Hirai Shota, Sakae Mizuki, Rio Yokota, Naoaki Okazaki, Continual Pre-Training for Cross-Lingual LLM Adaptation: Enhancing Japanese Language Capabilities, Conference on Language Modeling COLM, Oct. 2024.
35. Go Otani, Ryu Tadokoro, Ryosuke Yamada, Yuki Asano, Iro Laina, Christian Rupprecht, Nakamasa Inoue, Rio Yokota, Hirokatsu Kataoka, Yoshimitsu Aoki, Rethinking Image Super-Resolution from Training Data Perspectives, European Conference on Computer Vision (ECCV), Sep. 2024.
36. Ryo Nakamura, Ryu Tadokoro, Ryosuke Yamada, Yuki Asano, Iro Laina, Christian Rupprecht, Nakamasa Inoue, Rio Yokota, Hirokatsu Kataoka, Scaling Backwards: Minimal Synthetic Pre-training?, European Conference on Computer Vision (ECCV), Sep. 2024.

37. Ryosuke Yamada, Kensho Hara, Hirokatsu Kataoka, Koshi Makihara, Nakamasa Inoue, Rio Yokota, Yutaka Satoh, Formula-Driven Visual-Geometric Pre-training, European Conference on Computer Vision (ECCV), Sep. 2024.
38. Ryo Hayamizu, Shota Nakamura, Sora Takashima, Hirokatsu Kataoka, Ikuro Sato, Nakamasa Inoue, Rio Yokota, SIFTer: Self-improving Synthetic Datasets for Pre-training Classification Models, CVPR SynData Workshop, Jun. 2024.
39. Yuesong Shen, Nico Daheim, Gian Maria Marconi, Peter Nickl, Bai Cong, Bazan Clement Emile Marcel Raoul, Rio Yokota, Iryna Gurevych, Daniel Cremers, Mohammad Emtiyaz Khan, Thomas Möllenhoff, Variational Learning is Effective for Large Deep Networks, The 41st International Conference on Machine Learning (ICML), Jul. 2024.
40. Hiromichi Kobayashi, Satoshi Yui, Makoto Tsubota, Tomokazu Saito, Rio Yokota, Interaction between quantum turbulence and normal-fluid turbulence in superfluid helium, Thirteenth International Symposium on Turbulence and Shear Flow Phenomena (TSFP13), Jun. 2024.
41. Satoki Ishikawa, Ryo Karakida, On the Parameterization of Second-Order Optimization Effective towards the Infinite Width, The 12th International Conference on Learning Representations (ICLR), May. 2024.
42. Satoki Ishikawa, Rio Yokota, When Does Second-Order Optimization Speed Up Training? The 12th International Conference on Learning Representations (ICLR), Tiny paper, May. 2024.
43. Ryo Nakamura, Sora Takashima, Edgar Josafat Martinez-Noriega, Nakamasa Inoue, Hirokatsu Kataoka, Rio Yokota, Pre-training Vision Transformers with Very Limited Synthesized Images, IEEE/CVF International Conference on Computer Vision (ICCV), Oct. 2023.
44. Risa Shinoda, Ryo Hayamizu, Kodai Nakashima, Nakamasa Inoue, Rio Yokota, Hirokatsu Kataoka, SegRCDB: Semantic Segmentation via Formula-Driven Supervised Learning, IEEE/CVF International Conference on Computer Vision (ICCV), Oct. 2023.
45. Muhammad Ridwan Apriansyah, Rio Yokota, Computing Eigenvalue of Symmetric $\mathcal{H}^2$ -Matrices in Linear Time with Slicing the Spectrum, International Conference on Parallel Processing (ICPP), Aug. 2023.
46. Sameer Deshmukh, Rio Yokota, George Bosilca, O(N) Distributed Direct Factorization of Structured Dense Matrices Using Runtime Systems, International Conference on Parallel Processing (ICPP), Aug. 2023.
47. Hiroyuki Ootomo, Rio Yokota, Mixed-Precision Random Projection for RandNLA on Tensor Cores, Platform for Advanced Scientific Computing (PASC), Jun. 2023.
48. Sora Takashima, Ryoh Hayamizu, Nakamasa Inoue, Hirokatsu Kataoka, Rio Yokota, Visual Atoms: Pre-training Vision Transformers with Sinusoidal Waves, IEEE/CVF Conference on Computer Vision and Pattern Recognition, Jun. 2023.
49. Hiroyuki Ootomo, Hidetaka Manabe, Kenji Harada, Rio Yokota, Quantum Circuit Simulation by SGEMM Emulation on Tensor Cores and Automatic Precision Selection, ISC High Performance, May 2023.
50. Hiroyuki Ootomo, Rio Yokota, Reducing Shared Memory Footprint to Leverage High Throughput on Tensor Cores and its Flexible API Extension Library, HPC Asia, Feb. 2023.
51. Satoshi Ohshima, Akihiro Ida, Rio Yokota and Ichitaro Yamazaki, QR Factorization of Block Low-Rank Matrices on Multi-Instance GPU, The 23rd International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT'22), Dec. 2022.
52. Hiroki Naganuma, Kartik Ahuja, Ioannis Mitliagkas, Shiro Takagi, Tetsuya Motokawa, Rio Yokota, Kohta Ishikawa, Ikuro Sato, Empirical Study on Optimizer Selection for Out-of-Distribution Generalization, NeurIPS Workshop Distshift, Dec. 2022.
53. Kazuki Osawa, Satoki Ishikawa, Rio Yokota, Shigang Li, Torsten Hoefer, ASDL: A Unified Interface for Gradient Preconditioning in PyTorch, NeurIPS Workshop Order up! The Benefits of Higher-Order Optimization in Machine Learning, Dec. 2022.

54. Aoyu Li, Ikuro Sato, Kohta Ishikawa, Rei Kawakami, Rio Yokota, Informative Sample-Aware Proxy for Deep Metric Learning, ACM Multimedia Asia, Dec. 2022. (Best paper)
55. Qianxiang Ma, Sameer Deshmukh, Rio Yokota, Scalable Linear Time Dense Direct Solver for 3-D Problems Without Trailing Sub-Matrix Dependencies, The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC22), Nov. 2022.
56. Hirokatsu Kataoka, Ryo Hayamizu, Ryosuke Yamada, Kodai Nakashima, Sora Takashima, Xinyu Zhang, Edgar Josafat Martinez-Noriega, Nakamasa Inoue, Rio Yokota, Replacing Labeled Real-image Datasets with Auto-generated Contours, IEEE/CVF Conference on Computer Vision and Pattern Recognition, Jun. 2022.
57. Hana Hoshino, Kei Ota, Asako Kanezaki, Rio Yokota, OPIRL: Sample Efficient Off-Policy Inverse Reinforcement Learning via Distribution Matching, IEEE International Conference on Robotics and Automation, May 2022.
58. Shun Iwase, Xingyu Liu, Rawal Khirodkar, Rio Yokota, Kris M. Kitani, RePOSE: Real-Time Iterative Rendering and Refinement for 6D Object Pose Estimation, International Conference on Computer Vision, Oct. 2021.
59. Hikaru Nakata, Nakamasa Inoue, Rio Yokota, Self-supervised Continual Pretraining for Class Incremental Image Classification, Proc. CVPR CLVISION Workshop (Findings), Jun. 2021.
60. Ryo Karakida, Kazuki Osawa, Understanding Approximate Fisher Information for Fast Convergence of Natural Gradient Descent in Wide Neural Networks, Advances in Neural Information Processing Systems (NeurIPS), oral presentation, Dec. 2020.
61. Yuichiro Ueno, Kazuki Osawa, Yohei Tsuji, Akira Naruse, Rio Yokota. Rich Information is Affordable: A Systematic Performance Analysis of Second-order Optimization Using K-FAC. Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Aug. 2020.
62. Mikiya Shibuya, Shinya Sumikura, and Ken Sakurada. "Privacy Preserving Visual SLAM." Proceedings of the European conference on computer vision (ECCV). Oct. 2020.
63. Sameer Deshmukh and Rio Yokota. "Distributed Memory Task-Based Block Low Rank Direct Solver". ISC High Performance 2020, Germany (Research Poster), Jun. 2020.
64. Hiroyuki Ootomo, Rio Yokota, Randomized SVD on TensorCores, ISC High Performance 2020, Germany (Research Poster), Jun. 2020.
65. Rise Ooi, Takeshi Iwashita, Takeshi Fukaya, Akihiro Ida, Rio Yokota, Effect of Mixed Precision Computing on H-Matrix Vector Multiplication in BEM Analysis, HPC Asia 2020, Jan. 2020.
66. Sameer Deshmukh, Rio Yokota, Distributed Memory Task-Based Block Low Rank Direct Solver, HPC Asia 2020 (poster), Jan. 2020.
67. Muhammad Ridwan Apriansyah, Rio Yokota, QR Decomposition of Block Low-Rank Matrices, HPC Asia 2020 (poster), Jan. 2020.
68. Kazuki Osawa, Siddarth Swaroop, Anirudh Jain, Runa Eschenhagen, Richard E. Turner, Rio Yokota, Mohammad Emtiyaz Khan. Practical Deep Learning with Bayesian Principles, The 33rd Conference on Neural Information Processing Systems, Dec. 2019.
69. Qianxiang Ma, Rio Yokota, Runtime System for GPU-based Hierarchical LU factorization, The International Conference for High Performance Computing, Networking, Storage, and Analysis (poster), Nov. 2019.
70. Hiroyuki Ootomo, Rio Yokota, TSQR on TensorCores, The International Conference for High Performance Computing, Networking, Storage, and Analysis (poster), Nov. 2019.
71. Hiroki Naganuma, Rio Yokota, On Empirical Analysis of Layer-wised Learning Rate Schedule, ACML 2019 Workshop on Statistics and Machine Learning Researchers (poster), Nov. 2019.

72. Satoshi Ohshima, Ichitaro Yamazaki, Akihiro Ida, Rio Yokota. Optimization of Numerous Small Dense-Matrix-Vector Multiplications in $\mathcal{H}$ -matrix Arithmetic on GPU, Auto-Tuning for Multicore and GPU (ATMG) In conjunction with the IEEE MCSoc-19, Oct. 2019.
73. Yohei Tsuji, Kazuki Osawa, Yuichiro Ueno, Akira Naruse, Rio Yokota, Satoshi Matsuoka. Performance Optimizations and Analysis of Distributed Deep Learning with Approximated Second-Order Optimization Method, International Conference on Parallel Processing: The 1st Workshop on Parallel and Distributed Machine Learning, Proceedings of the 48th International Conference on Parallel Processing: Workshops, No. 21, Aug. 2019.
74. Kazuki Osawa, Yohei Tsuji, Yuichiro Ueno, Akira Naruse, Rio Yokota, Satoshi Matsuoka. Large-Scale Distributed Second-Order Optimization Using Kronecker-Factored Approximate Curvature for Deep Convolutional Neural Networks, IEEE/CVF Conference on Computer Vision and Pattern Recognition, Jun. 2019.
75. Yuichiro Ueno, Rio Yokota. Exhaustive Study of Hierarchical AllReduce Patterns for Large Messages Between GPUs, 19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID), May. 2019.
76. Hiroki Naganuma, Rio Yokota. A Performance Improvement Approach for Second-Order Optimization in Large Mini-batch Training, 2nd High Performance Machine Learning Workshop CCGrid2019 (HPML2019), May. 2019.
77. Ichitaro Yamazaki, Ahmad Abdelfattah, Akihiro Ida, Satoshi Ohshima, Stanimire Tomov, Rio Yokota, Jack Dongarra. Analyzing Performance of BiCGStab with Hierarchical Matrix on GPU clusters, 32nd IEEE International Parallel and Distributed Processing Symposium, May. 2018.
78. Satoshi Ohshima, Ichitaro Yamazaki, Akihiro Ida, Rio Yokota. Optimization of Hierarchical Matrix Computation on GPU, SC Asia, Mar. 2018.
79. Hiroki Naganuma, Rio Yokota. Accelerating Convolutional Neural Networks Using Low Precision Arithmetic, HPC Asia, Jan. 2018.
80. Kazuki Osawa, Rio Yokota. Evaluating the Compression Efficiency of the Filters in Convolutional Neural Networks, The 26th International Conference on Artificial Neural Networks, Sep. 2017.
81. Mustafa AbdulJabbar, Mohammed Al Farhan, Rio Yokota, David Keyes. Performance Evaluation of Computation and Communication Kernels of the Fast Multipole Method on Intel Manycore Architecture, 3rd International European Conference on Parallel and Distributed Computing, Aug. 2017.
82. Kazuki Osawa, Akira Sekiya, Hiroki Naganuma, Rio Yokota. Accelerating Matrix Multiplication in Deep Learning by Using Low-Rank Approximation, The 2017 International Conference on High Performance Computing and Simulation, Jul. 2017.
83. Mustafa AbdulJabbar, George Markomanolis, Huda Ibeid, Rio Yokota, David Keyes. Communication Reducing Algorithms for Distributed Hierarchical N-Body Problems with Boundary Distributions, 32nd International Conference, ISC High Performance, Lecture Notes in Computer Science, Vol. 10266, pp. 79–96, Jun. 2017.
84. Keisuke Fukuda, Motohiko Matsuda, Naoya Maruyama, Rio Yokota, Kenjiro Taura, Satoshi Matsuoka. Tapas: An Implicitly Parallel Programming Framework For Hierarchical N-body Algorithms, The 22nd IEEE International Conference on Parallel And Distributed Systems, Page 1100-1109, Dec. 2016.
85. Rio Yokota. Fast Multipole Method as a Matrix-free Hierarchical Low-rank Approximation, International Workshop on Eigenvalue Problems, Sep. 2016.
86. Rio Yokota, Huda Ibeid, David Keyes. Preconditioning Sparse Matrices Using a Highly Scalable Fast Multipole Method, 3rd International Workshops on Advances in Computational Mechanics, Oct. 2015.

87. Huda Ibeid, Rio Yokota, Jennifer Pestana, David Keyes. Fast Multipole Preconditioners for Sparse Linear Solvers, 11th World Congress on Computational Mechanics, Jul. 2014.
88. Hatem Ltaief, Rio Yokota. High Performance Numerical Algorithms for Seismic and Reservoir Simulations, GPU Technology Conference, Mar. 2014.
89. Rio Yokota. Fast N-body Methods as a Compute-Bound Preconditioner for Sparse Solvers on GPUs, GPU Technology Conference, Mar. 2014.
90. Abdelhalim Amer, Naoya Maruyama, Miquel Pericas, Kenjiro Taura, Rio Yokota, Satoshi Matsuoka. Fork-Join and Data-Driven Execution Models on Multi-core Architectures: Case Study of the FMM, International Supercomputing Conference, Lecture notes in computer science, LNCS, Vol. 7905, pp. 255-266, Jun. 2013.
91. Jennifer Pestana, Rio Yokota, Huda Ibeid, David Keyes. Fast Multipole Method Preconditioning, International Conference On Preconditioning Techniques For Scientific And Industrial Applications, Jun. 2013.
92. Abdul Abdelfatteh, Hatem Ltaief, Rio Yokota. Investigating New Numerical Techniques for Reservoir Simulations on GPUs, GPU Technology Conference, Mar. 2013.
93. Kenjiro Taura, Jun Nakashima, Rio Yokota, Naoya Maruyama. A Task Parallelism Meets Fast Multipole Methods, Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (Scala), Nov. 2012.
94. Rio Yokota. Petascale Fast Multipole Methods on GPUs, GPU Technology Conference Japan, Jul. 2012.
95. Hatem Ltaief, Rio Yokota. Data-Driven Fast Multipole Method on Distributed Memory Systems with Hardware Accelerators, 21st International Conference on Domain Decomposition Methods, Jun. 2012.
96. Enas Yunis, Rio Yokota, Aron Ahmadi. Scalable Force Directed Graph Layout Algorithms Using Fast Multipole Methods, The 11th International Symposium on Parallel and Distributed Computing, Jun. 2012.
97. Rio Yokota, Lorena Barba. Recent Trends in Hierarchical N-body Methods on GPUs, GPU Technology Conference, May. 2012.
98. Hoang Vu Nguyen, Rio Yokota, Georgiy Stenchikov. A Parallel Numerical Simulation of Dust Particles Using Direct Numerical Simulation, European Geosciences Union General Assembly, Apr. 2012.
99. Tetsu Narumi, Rio Yokota, Lorena Barba, Kenji Yasuoka. Petascale Turbulence Simulation Using FMM, HOKKE-19, Nov. 2011.
100. Rio Yokota, Lorena Barba. Parameter Tuning of a Hybrid Treecode-FMM on GPUs, The First International Workshop on Characterizing Applications for Heterogeneous Exascale Systems, Jun. 2011.
101. Rio Yokota, Lorena Barba. Fast Multipole Method vs. Spectral Methods for the Simulation of Isotropic Turbulence on GPUs, 23rd International Conference on Parallel Computational Fluid Dynamics, May. 2011.
102. Rio Yokota, Jaydeep Bardhan, Matthew Knepley, Lorena Barba. (Really) Fast Macromolecular Electrostatics – Fast Algorithms, Open Software and Accelerated Computing, ACS Division of Physical Chemistry 240th National Meeting, Aug. 2010.
103. Rio Yokota, Lorena Barba. Performance of the Fast Multipole Method on GPUs Using Various Kernels, 9th World Congress on Computational Mechanics, Jul. 2010.
104. Rio Yokota, Lorena Barba. Comparing the Treecode with FMM on GPUs for Vortex Particle Simulations of a Leapfrogging Vortex Ring, 22nd International Conference on Parallel Computational Fluid Dynamics, May. 2010.

105. Rio Yokota, Shinnosuke Obi. Lagrangian Simulation of Turbulence Using Vortex Methods, 2nd International Workshops on Advances in Computational Mechanics, Mar. 2010.
106. Tsuyoshi Hamada, Rio Yokota, Keigo Nitadori, Tetsu Narumi, Kenji Yasuoka, Makoto Taiji, Kyoshi Oguri. 42 TFlops Hierarchical N-Body Simulation on GPUs with Applications in Both Astrophysics and Turbulence, Supercomputing, Nov. 2009.
107. Rio Yokota, Koji Fukagata, Shinnosuke Obi. Lagrangian Vortex Methods in Turbulent Channel Flows, 12th EUROMECH European Turbulence Conference, Sep. 2009.
108. Rio Yokota, Tetsu Narumi, Ryuji Sakamaki, Kenji Yasuoka, Shinnosuke Obi. Fast Multipole Methods on GPUs for the Meshfree Simulation of Turbulence, 10th US National Congress on Computational Mechanics, Jul. 2009.
109. Rio Yokota, Tetsu Narumi, Ryuji Sakamaki, Shun Kameoka, Kenji Yasuoka, Shinnosuke Obi. DNS of Homogeneous Turbulence Using Vortex Methods Accelerated by the FMM on a Cluster of GPUs, 21st International Conference on Parallel Computational Fluid Dynamics, May. 2009.
110. Rio Yokota, Tetsu Narumi, Ryuji Sakamaki, Shun Kameoka, Kenji Yasuoka, Shinnosuke Obi. Meshfree Simulation of Turbulence Using the Fast Multipole Methods on GPUs, 22nd Symposium on Computational Fluid Dynamics, Dec. 2008.
111. Rio Yokota, Shinnosuke Obi. Direct Numerical Simulation of Homogeneous Shear Flow Using Vortex Methods, 4th International Conference on Vortex Flows and Vortex Models, Apr. 2008.
112. Rio Yokota, Shinnosuke Obi. Mesh-Free Simulation of the Homogeneous Shear Flow Using Vortex Methods, 23rd IIS Turbulence and Shear Flow Dynamics Symposium, Mar. 2008.
113. Rio Yokota, Shinnosuke Obi. Pure Lagrangian Vortex Methods for the Simulation of Decaying Isotropic Turbulence, 5th International Symposium on Turbulence and Shear Flow Phenomena, Aug. 2007.
114. Rio Yokota, Shinnosuke Obi. Vortex Flow Simulation Between Multiple Bridge Decks, Whither Turbulence Prediction and Control, Mar. 2006.
115. Rio Yokota, Shinnosuke Obi. Vortex Flow Simulation of Multiple Bluff Bodies, 3rd International Conference on Vortex Flows and Vortex Models, Nov. 2005.

CONFERENCE
PRESENTATIONS
(NOT
PEER-REVIEWED)

1. Bai Cong, Nico Daheim, Yuesong Shen, Rio Yokota, Mohammad Emtiyaz Khan, Thomas Möllenhoff, Improving LoRA with Variational Learning, arXiv preprint (CoRR 2025), Jun. 2025.
2. Hiromichi Kobayashi, Satoshi Yui, Makoto Tsubota, Tomokazu Saito, Rio Yokota, Two-way coupled simulation of quantum turbulence and normal-fluid turbulence in superfluid helium-4, APS DFD 76th Annual Meeting of the Division of Fluid Dynamics, Nov. 2023.
3. Hiromichi Kobayashi, Satoshi Yui, Makoto Tsubota, Rio Yokota, Two-phase flow of quantum turbulence and normal-fluid turbulence in superfluid helium-4, ERCOFTAC Symposium on Engineering Turbulence Modeling and Measurements (ETMM), Sep. 2023.
4. Hiromichi Kobayashi, Satoshi Yui, Makoto Tsubota, Tomokazu Saito, Rio Yokota, Vortex-filament bundle induced by normal-fluid turbulence in turbulent superfluid helium-4, International Symposium on Quantum Fluids and Solids (QFS), Aug. 2023.
5. Satoshi Yui, Hiromichi Kobayashi, Makoto Tsubota, Tomokazu Saito, Rio Yokota, Vortex-Filament Bundle Induced by Normal-Fluid Turbulence in Turbulent Superfluid Helium-4, International Symposium on Quantum Fluids and Solids (QFS), Aug. 2023.
6. Qianxiang Ma, Rio Yokota, O(N) Factorization of Dense Matrices on GPUs Without Trailing Submatrix Dependencies, SIAM Conference on Computational Science and Engineering (CSE), Feb. 2023.
7. Muhammad Ridwan Apriansyah, Rio Yokota, Parallel QR Factorization of Block Low-Rank Matrices, SIAM Conference on Computational Science and Engineering (CSE), Feb. 2023.

8. Rio Yokota, Matrices in Deep Neural Networks and How to Compute Them in Parallel, IEEE CLUSTER 2022 Keynote, Sep. 2022.
9. Sameer Deshmukh, Acceleration of $O(N)$ Solvers for Large Dense Matrices, Conference on Advance Topics and Auto Tuning in High-Performance Scientific Computing (ATAT2022), Mar. 2022.
10. Muhammad Ridwan Apriansyah, Parallel QR Factorization of Block Low-rank Matrices, Conference on Advance Topics and Auto Tuning in High-Performance Scientific Computing (ATAT2022), Mar. 2022.
11. Thomas Spendlhofer, Iterative Refinement with Hierarchical Low-rank Preconditioners Using Mixed Precision, Conference on Advance Topics and Auto Tuning in High-Performance Scientific Computing (ATAT2022), Mar. 2022.
12. Rio Yokota, Approximations of Natural Gradient Descent in Distributed Training, INFORMS Annual Meeting Session: Beyond first order methods in machine learning systems I, Oct. 2021.
13. Rio Yokota, Overview of structured low-rank approximation methods, IUTAM Symposium on Computational Methods for Large-Scale and Complex Wave Problems, Jun. 2021.
14. Rio Yokota, Overview of Distributed Memory Parallelism in Deep Learning, DD26 MS01, Learning, Algorithms, Domain Decomposition Methods, and Applications, Dec. 2020.
15. Rio Yokota, Distributed Deep Learning with Second Order Information, SPCL_Bcast (COMM_WORLD), Oct. 2020.
16. Rio Yokota, Degree of Approximation and Overhead of Computing Curvature, Information, and Noise Matrices, ICML Workshop “Beyond first order methods in machine learning systems”, July, 2020.
17. Rio Yokota, Recent Trends in Hierarchical Low-Rank Approximation Methods, Tokyo Institute of Technology and Stony Brook University Joint Science and Technology Meeting, May. 2019.
18. Rio Yokota, Yohei Tsuji, Kazuki Osawa, Second Order Optimization for Distributed Data-parallel Deep Learning on 4000 GPUs, I2R-TokyoTech Co-workshop on DL 2.0, Mar. 2019.
19. Rio Yokota: Kronecker Factorization for Second Order Optimization in Deep Learning, SIAM CSE, Feb. 2019.
20. Rio Yokota. Optimization Methods for Large Scale Distributed Deep Learning, IPAM Workshop I: Big Data Meets Large-Scale Computing, Sep. 2018.
21. Rio Yokota. Early Application Results on TSUBAME 3, Smoky Mountains Computational Sciences and Engineering Conference, Aug. 2018.
22. Rio Yokota. Scaling Deep Learning to Thousands of GPUs, HPC 2018, Jul. 2018.
23. Rio Yokota. Energy Conserving Fast Multipole Methods for the Calculation of Long-range Interactions, Mathematics in Action: Modeling and analysis in molecular biology and electrophysiology, Jun. 2018.
24. Rio Yokota. Can we use Hierarchical Low-Rank Approximation for Deep Learning?, HPC Saudi 2018, Mar. 2018.
25. Rio Yokota. Hierarchical Low-Rank Approximations at Extreme Scale, 32nd International Conference, ISC High Performance, Jun. 2017.
26. Rio Yokota. Compute-Memory Tradeoff in Hierarchical Low-Rank Approximation Methods, SIAM Conference on Computational Science and Engineering, Feb. 2017.
27. Rio Yokota. Energy Conservation of Fast Multipole Methods in Classical Molecular Dynamics Simulations, 7th AICS International Symposium, Feb. 2017.
28. Rio Yokota. Improving Data Locality of Fast Multipole Methods, Third Workshop on Programming Abstractions for Data Locality, Kobe, Oct. 2016.

29. Huda Ibeid, Rio Yokota, David Keyes. A Matrix-Free Preconditioner for Elliptic Solvers Based on the Fast Multipole Method, SIAM Conference on Parallel Processing for Scientific Computing, Apr. 2016.
30. Rio Yokota, A Common API for Fast Multipole Methods, Accelerate Data Analytics and Computing Workshop, Jan. 2016.
31. Rio Yokota, Francois-Henri Rouet, Xiaoye Sherry Li. Comparison of FMM and HSS at Large Scale, SIAM Conference on Applied Linear Algebra, Oct. 2015.
32. Rio Yokota. Various Implementations of FMM and Their Performance on Future Architectures, Multi-resolution Interactions Workshop, Aug. 2015.
33. Rio Yokota. ExaFMM – a Testbed for Comparing Various Implementations of the FMM, SIAM Conference on Computational Science and Engineering, Mar. 2015.
34. Huda Ibeid, Jennifer Pestana, Rio Yokota, David Keyes. Fast Multipole Method as Preconditioner, SIAM Conference on Computational Science and Engineering, Mar. 2015.
35. Rio Yokota, David Keyes. Communication Complexity of the Fast Multipole Method and its Algebraic Variants, CBMS-NSF Conference: Fast Direct Solvers for Elliptic PDEs, Jun. 2014.
36. Rio Yokota. Advances in Fast Multipole Methods for Scalable Electrostatics Calculations, Workshop: Electrostatics methods in Molecular Simulation, May. 2013.
37. Huda Ibeid, Rio Yokota, David Keyes. Fast Multipole Method as a Preconditioner, SIAM Conference on Computational Science and Engineering, Feb. 2013.
38. Rio Yokota. Petascale Fast Multipole Methods on GPUs, The 11th International Symposium on Parallel and Distributed Computing, Jun. 2012.
39. Rio Yokota, Tetsu Narumi, Lorena Barba, Kenji Yasuoka. Scaling Fast Multipole Methods up to 4000 GPUs, ATIP/A*CRC Workshop on Accelerator Technologies for High Performance Computing, May. 2012.
40. Rio Yokota. Running Fast Multipole Method on the Full Node of TSUBAME and K computer, Scalable Hierarchical Algorithms for Extreme Computing, Apr. 2012.
41. Rio Yokota. Fast N-body Methods on Many-core and Heterogenous Systems, International Workshop on Computational Science and Numerical Analysis, Mar. 2012.
42. Rio Yokota. Petaflops Scale Turbulence Simulation on TSUBAME 2.0, GPU@BU Workshop, Nov. 2011.
43. Rio Yokota, Lorena Barba. Large Scale Multi-GPU FMM for Bioelectrostatics, SIAM Conference on Computational Science and Engineering, Feb. 2011.
44. Rio Yokota. 12 Steps to a Fast Multipole Method on GPUs, Pan-American Advanced Studies Institute, Jan. 2011.
45. Rio Yokota, Lorena Barba. RBF Interpolation using Gaussians with Domain Decomposition on GPUs, SIAM Annual Meeting, Jul. 2010.
46. Rio Yokota. Range of Applications for the Fast Multipole Method on GPUs, Accelerated Computing, Jan. 2010.

JAPANESE
DOMESTIC
CONFERENCE
(PEER-REVIEWED)

1. 服部翔, 水木栄, 藤井一喜, 中村泰士, 塩谷泰平, 植木快, 新妻巧朗, 川畑輝, 田森秀明, Youmi Ma, 前田航希, 大井聖也, 齋藤幸史郎, 岡本拓己, 石田茂樹, 横田理央, 高村大也, 岡崎直観, 新聞記事からつくる時事と社会に強い日本語 LLM, 言語処理学会第 31 回年次大会, Mar. 2025.
2. Youmi, 水木栄, 藤井一喜, 中村泰士, 大井聖也, 島田比奈理, 塩谷泰平, 齋藤幸史郎, 前田航希, 服部翔, 岡本拓己, 石田茂樹, 横田理央, 高村大也, 岡崎直観, 模倣学習による大規模言語モデルの指示チューニング, 言語処理学会第 31 回年次大会, Mar. 2025.

3. 大川快、櫻田健、横田理央、単眼カメラを用いたリアルタイムな3次元マップの変化検出を目的とした密なバンドル調整、第27回画像の認識・理解シンポジウム (MIRU), Aug. 2024.
4. 大谷豪、田所龍、山田亮佑、Yuki M. asano, Iro Laina, Chistian Repprecht、井上中順、横田理央、片岡裕雄、青木義満、画像超解像における学習データ構築の再考、MIRU, Aug. 2024.
5. 田所龍、中村凌、山田亮佑、Yuki M. Asano, Iro Laina, Chistian Repprecht、井上中順、横田理央、片岡裕雄、Scaling Backwards: Minimal Synthetic Pre-training? MIRU, Aug. 2024.
6. 高橋秀弥、井上中順、横田理央、片岡裕雄、前田英作、学習過程における形状・テクスチャ偏重度の推移と事前学習データセットとの関係について、第26回画像の認識・理解シンポジウム (MIRU), ポスター, Jul. 2023.
7. 大川快、櫻田健、横田理央、単眼カメラを用いたリアルタイムな3次元マップの変化検出を目的とした密なバンドル調整、第26回画像の認識・理解シンポジウム (MIRU), ポスター, Jul. 2023.
8. 山田亮佑、原健翔、片岡裕雄、牧原昂志、井上中順、横田理央、佐藤雄隆、Formula-Supervised Visual-Geometric Pre-training, 第26回画像の認識・理解シンポジウム (MIRU), ロングオーラル, Jul. 2023.
9. 篠田理沙、速水亮、中嶋航大、井上中順、横田理央、片岡裕雄、数式ドリブン教師あり学習によるセマンティックセグメンテーション、第26回画像の認識・理解シンポジウム (MIRU), ロングオーラル, Jul. 2023.
10. 浅倉拓也、井上中順、横田理央、篠田浩一、受容野の自動最適化によるモードに適応的なTransformerの開発、人工知能学会全国大会, Jun. 2023.
11. 中村祥大、横田理央、ラージバッチ学習における汎化性能の低下を抑制する正則化手法、人工知能学会全国大会, Jun. 2023.
12. 高橋秀弥、井上中順、横田理央、片岡裕雄、前田英作、画像識別における形状・テクスチャ偏重度と二重降下現象の関係について、PRMU, Mar. 2023.
13. 杉山佳史、片岡裕雄、横田理央、井上中順、敵対的距離学習モジュールを用いた特徴変動に頑健な画像認識のための対照学習、パターン認識・メディア理解研究会 (PRMU), Dec. 2022.
14. 田所龍、片岡裕雄、川上玲、横田理央、井上中順、蒸留画像による事前学習効果についての検討、ビジョン技術の実利用ワークショップ (ViEW), Dec. 2022.
15. 高橋那弥、八嶋晋吾、石川康太、佐藤育郎、走行動画の大規模自己教師あり学習の検討と計画、第25回画像の認識・理解シンポジウム (MIRU), Jul. 2022.
16. 大友広幸、横田理央、Tensor コアを用いた TSQR, 日本応用数理学会年会, Sept. 2019.
17. 岩瀬駿、櫻田健、変化の属性を考慮した物体レベルのシーン変化検出、第22回画像の認識・理解シンポジウム (MIRU2019), Oct. 2019.
18. 長沼大樹、横田理央、ラージバッチ学習のための自然勾配学習法における Smoothing の有効性、The 3rd Cross-disciplinary Workshop on Computing Systems, Infrastructures, and Programming (xSIG), May. 2019.
19. 長沼大樹、岩瀬駿、郭林昇、中田光、横田理央、自然勾配近似法を用いた大規模並列深層学習におけるハイパーパラメータ最適化、第17回情報科学技術フォーラム 2018, Sep. 2018.
20. 長沼大樹、横田理央、畳み込みニューラルネットワークにおける低精度演算を用いた高速化の検証、GTC Japan, Dec. 2017.
21. 長沼大樹、関谷翠、大沢和樹、大友広幸、桑村裕二、横田理央、深層学習における低精度演算を用いた高速化及びアクセラレーターの性能評価、パターン認識・メディア理解研究会, Oct. 2017.
22. 大沢和樹、関谷翠、長沼大樹、横田理央、低ランクテンソル分解を用いた畳み込みニューラルネットワークの高速化、パターン認識・メディア理解研究会, Oct. 2017.
23. 長沼大樹、大沢和樹、関谷翠、横田理央、深層学習における半精度演算を用いた圧縮モデルの高速化、日本応用数理学会年会, Sep. 2017.

24. 大島聡史, 山崎市太郎, 伊田明弘, 横田理央. GPU クラスタ上における階層型行列計算の最適化, Summer United Workshops on Parallel, Distributed and Cooperative Processing, Jul. 2017.
25. 大沢和樹, 関谷翠, 長沼大樹, 横田理央. 畳み込みニューラルネットワークの低ランク近似を用いた高速化, 第 22 回計算工学講演会, 計算工学講演会論文集 Vol.22, May. 2017.
26. 横田理央, 小尾晋之介. 平行平板間乱流における渦法の検証, 日本流体力学会年会, Sep. 2009.
27. 横田理央, 小尾晋之介. 渦法を用いた平行平板間乱流の解析, 流体力学会年会, Sep. 2008.
28. 佐藤彰, 横田理央, 小尾晋之介. 三次元渦法による翼端渦の数値解析, 第 21 回数値流体力学シンポジウム, Dec. 2007.
29. 横田理央, 小尾晋之介. 渦法による一様せん断流の解析, 流体力学会年会, Aug. 2007.
30. 横田理央, 小尾晋之介. 渦法によるメッシュフリー乱流解析, 日本機械学会東海支部第 56 期総会・講演会, Mar. 2007.
31. 横田理央, 小尾晋之介. 3 次元渦法・境界要素法による流体-固体連成解析, 日本機械学会流体工学部門講演会, Oct. 2006.
32. 横田理央, 小尾晋之介. 渦法を用いた物体後流の 3 次元解析, 日本機械学会年次大会, Sep. 2006.
33. 横田理央, 小尾晋之介. 複数の鈍い形状物体周りの渦流れシミュレーション, 第 19 回数値流体シンポジウム, Dec. 2005.

JAPANESE
DOMESTIC
CONFERENCE (NOT
PEER-REVIEWED)

1. 安田俊吾, 横田理央, 視覚・言語モデルにおける日本語 OCR 性能の向上, 情報処理学会全国大会, Mar. 2026.
2. 野原大輔, 中村泰士, 横田理央, 大規模言語モデルの強化学習による効率的な論理推論, 情報処理学会全国大会, Mar. 2026.
3. 岡本拓己, 横田理央, One-Shot NAS による BERT のモデル圧縮, 人工知能学会全国大会 (第 38 回), May 2024.
4. 中村泰士, 横田理央, 継続学習を用いた効率の良いマルチリンガル・マルチエキスパートモデルの開発, 情報処理学会全国大会, Mar. 2024.
5. 藤井一喜, 横田理央, 大規模言語モデルの分散並列学習, 情報処理学会全国大会, Mar. 2024.
6. 岡本拓己, 横田理央, 大規模言語モデルの構造探索, 情報処理学会全国大会, Mar. 2024.
7. 岡崎直観, 服部翔, 平井翔太, 飯田大貴, 大井聖也, 藤井一喜, 中村泰士, Mengsay Loem, 横田理央, 水木栄, Swallow コーパス: 日本語大規模ウェブコーパス, 言語処理学会第 30 回年次大会, Mar. 2024.
8. 水木栄, 飯田大貴, 藤井一喜, 中村泰士, Mengsay Loem, 大井聖也, 服部翔, 平井翔太, 横田理央, 岡崎直観, 大規模言語モデルの日本語能力の効率的な強化: 継続事前学習における語彙拡張と対訳コーパスの活用, 言語処理学会第 30 回年次大会, Mar. 2024.
9. 藤井一喜, 中村泰士, Mengsay Loem, 飯田大貴, 大井聖也, 服部翔, 平井翔太, 水木栄, 横田理央, 岡崎直観, 継続事前学習による日本語に強い大規模言語モデルの構築, 言語処理学会第 30 回年次大会, Mar. 2024.
10. 大川快, 櫻田健, 横田理央, Visual SLAM を目的とした深度の一貫性を考慮した密なバンドル調整, CVIM 2024 年 1 月研究会, Jan. 2024.
11. 大谷豪, 田所龍, 片岡裕雄, 井上中順, 横田理央, 青木義満, 超解像事前学習における核心的要素の解明, ViEW2023, Dec. 2023.
12. 湯井悟志, 小林宏充, 坪田誠, 齋藤智和, 横田理央, 超流動ヘリウム 4 の量子乱流: 常流体乱流による渦糸バンドルの形成, Quantum turbulence of superfluid helium-4: formation of vortex-filament bundle by normal-fluid turbulence, 日本物理学会第 78 回年次大会, Sep. 2023.

13. 中村秋海, 横田理央, GPU と A64FX における Transformer の性能比較, 第 188 回 HPC 研究発表会, Mar 2023.
14. 齋藤智和, 横田理央, 量子渦計算の高速多重極展開法を用いた高速化, 情報処理学会全国大会, Mar 2023.
15. 石川智貴, 横田理央, 深層学習における勾配の前処理法に関する検討, 情報処理学会全国大会, Mar 2023.
16. 近江俊樹, 中村凌, 片岡裕雄, 井上中順, 横田理央, ニュートンフラクタル画像による事前学習効果, 情報処理学会全国大会, Mar 2023.
17. 大友広幸, 真鍋秀隆, 原田健自, 横田理央, Tensor コアによる単精度行列積エミュレーションの自動精度選択を用いた量子回路シミュレーション, 第 188 回ハイパフォーマンスコンピューティング研究発表会, Mar. 2023.
18. 大友広幸, 坂本亮, PEZY-SC3s プロセッサを用いた Full-state 量子回路シミュレーション, 第 187 回ハイパフォーマンスコンピューティング研究発表会, Dec. 2022.
19. 横田理央, $O(N)$ で並列性の高い密行列の LU 分解, RIMS 研究集会「数値解析が拓く次世代情報社会〜エッジから富岳まで〜」, Oct. 2022.
20. 伊田明弘, 荻田武史, 横田理央, 対称ブロック低ランク行列の精度保証付き固有値問題解法, 日本応用数理学会 2022 年度年会, Sep. 2022
21. 横田理央, 人工画像を用いた Vision Transformer の大規模事前学習, 社会的課題解決型データサイエンス・AI 研究推進体シンポジウム, Sep. 2022.
22. 大友広幸, 横田理央, Tensor コアを用いた単精度行列積エミュレーションのアプリケーションでの評価, 第 185 回ハイパフォーマンスコンピューティング研究発表会 (SWoPP2022), Jul. 2022.
23. 横田理央, 人工画像を用いた Vision Transformer の大規模事前学習, DENSO IT LAB x TOKYO TECH Discussion Night in MIRU, Jul. 2022.
24. 張新宇, 高島空良, 横田理央, ViT のファインチューニング時における NAS のモデル縮小効果, 第 36 回人工知能学会全国大会, Jun. 2022
25. 石井央, 横田理央, 深層学習における 2 次最適化の汎化性能の検証, 第 84 回情報処理学会全国大会, Mar. 2022.
26. 中村秋海, 横田理央, Vision Transformer におけるバッチサイズの汎化性能への影響, 第 84 回情報処理学会全国大会, Mar. 2022. (学生奨励賞)
27. 横田理央, 大規模並列深層学習の基礎, Nvidia HPC Weeks, (基調講演), Oct. 2021.
28. 横田理央, 階層的な低ランク近似法に関するレビュー, 第 40 回計算数理工学フォーラム, Sep. 2021.
29. 大友広幸, 横田理央, TensorCore を用いた精度補正単精度行列積, 第 180 回ハイパフォーマンスコンピューティング研究発表会 (SWoPP2021), Jul. 2021.
30. 横田理央, スパコンを用いた大規模並列分散深層学習, IBISML, 情報論的学習理論とコンピューティング基礎, Mar. 2021.
31. 横田理央, 深層学習におけるヘッセ行列, フィッシャー行列, 共分散行列の高速近似解法, AT マイクロワークショップ, Oct. 2020.
32. 中田光, 横田理央. 画像分類のための継続的な事前学習における教師なし表現学習の堅牢性に関する検証. 人工知能学会全国大会 (第 34 回), Jun. 2020.
33. 大友広幸, 横田理央 Tensor コアの API の構造解析を用いた拡張ライブラリの開発, 第 173 回ハイパフォーマンスコンピューティング研究会 (COVID-19 のため情報処理学会電子図書館にて公開のみ) Mar. 2020.
34. 所畑貴大, 長沼大樹, 横田理央, 確率的重み付け平均法のラージバッチ学習における有用性の検証, 情報処理学会全国大会, Mar. 2020.

35. 横田理央, 二次最適化を用いた巨大な言語モデルの学習および FRNN を用いたプラズマ挙動予測, ABCI グランドチャレンジ 2019 成果発表会, Feb. 2020.
36. 横田理央, 二次最適化を用いた ImageNet の大規模分散深層学習, 精密工学会 画像応用技術専門委員会, 第 5 回定例研究会「人工知能・データサイエンス」, Jan. 2020.
37. 横田理央, 近似行列分解と分散深層学習, RIMS 研究会諸科学分野を結ぶ基礎学問としての数値解析学, Nov. 2019.
38. 八島慶汰, 石川康太, 佐藤育郎, 野村哲弘, 横田理央, 松岡聡. 早期終了タイミングを予測する: 深層学習における確率勾配の分布の変化点検出, 第 22 回情報論的学習理論ワークショップ (IBIS 2019), Nov. 2019.
39. 横田理央, 深層学習に現れる密行列も構造, 自動チューニング研究会マイクロワークショップ, Oct. 2019.
40. 横田理央, 深層学習の高速化と大規模並列化, 情報処理学会連続セミナー第 3 回: AI と歩む未来 (2): 画像・映像処理の最前線, Sep. 2019
41. Peter Spalthoff, 横田理央. Flexible and Simplistic Hierarchical Matrix-Based Fast Direct Solver, 第 170 回ハイパフォーマンスコンピューティング研究発表会, Jul. 2019.
42. 大友広幸, 横田理央. Tensor コアを用いた TSQR の GPU 実装, 第 170 回ハイパフォーマンスコンピューティング研究発表会, Jul. 2019.
43. 横田理央, ImageNet ベンチマークの大規模並列深層学習, 第 14 回名工大・核融合研共同セミナー: 流体シミュレーションとディープラーニング, Jul. 2019.
44. 横田理央, 階層的 low-rank 近似法の最新研究動向と応用例, 第 22 回 AT 研究会オープンアカデミックセッション, May. 2019.
45. 横田理央, 大沢和樹, 辻陽平, 上野裕一郎, 成瀬彰. 大規模並列深層学習における 2 次の最適化手法の効果, 電子情報通信学会総合大会, Mar. 2019.
46. 長沼大樹, 横田理央. 大規模並列深層学習のための目的関数の平滑化, 第 81 回情報処理学会全国大会, Mar. 2019.
47. 大友広幸, 横田理央. Tensor コアを用いた Batched QR 分解, 第 81 回情報処理学会全国大会, Mar. 2019.
48. 中田光, 大沢和樹, 横田理央. 自然勾配法に基づく変分深層学習, 第 81 回情報処理学会全国大会, Mar. 2019.
49. 長沼大樹, 横田理央. ノイズ注入による平均化を用いたラージバッチ学習の汎化性能改善手法の検討, 電子情報通信学会総合大会, Mar. 2019.
50. 大沢和樹, 横田理央, Chuan-Sheng Foo, Vijay Chandrasekhar. Fisher 情報行列の解析に基づく大規模深層学習のための二次最適化手法, 第 81 回情報処理学会全国大会, Mar. 2019.
51. 辻陽平, 大沢和樹, 横田理央, 松岡聡, K-FAC と分散処理による深層学習の高速化, 第 165 回ハイパフォーマンスコンピューティング研究発表会 (SWoPP2018), Jul. 2018.
52. 桑村祐二, 大沢和樹, 横田理央. 自然勾配法の近似手法における学習パラメータの調整, 情報処理学会全国大会, Mar. 2018.
53. 大友広幸, 大沢和樹, 横田理央. フィッシャー情報行列のクロネッカー因子分解を用いた深層学習, 情報処理学会全国大会, Mar. 2018.
54. 大友広幸, 大沢和樹, 横田理央. フィッシャー情報行列のクロネッカー因子分解を用いた深層ニューラルネットワークの分散学習, 第 163 回ハイパフォーマンスコンピューティング研究発表会, Mar. 2018.
55. 本山義史, 遠藤敏夫, 松岡聡, 横田理央, 福田圭祐, 佐藤育郎. 低 rank 近似行列による CNN における畳み込み演算の最適化, 第 158 回ハイパフォーマンスコンピューティング研究発表会, 2017-HPC-158 No.25, Mar. 2017.

56. 関谷翠, 大沢和樹, 長沼大樹, 横田理央. 低ランク近似を用いた深層学習の行列積の高速化, 第 158 回ハイパフォーマンスコンピューティング研究発表会, 2017-HPC-158 No.25, Mar. 2017.
57. 横田理央. Fast Multipole Method を用いた多種アーキテクチャ向けスーパーコンピュータ用ライブラリの開発と分子・流体シミュレーションでの評価, 学際大規模情報基盤共同利用・共同研究拠点第 8 回シンポジウム, Jul. 2016.
58. 横田理央. FMM の性能の可搬性, 第 2 1 回計算工学講演会, May. 2016.
59. 横田理央. FMM の自動チューニング可能なパラメータについて, 第 7 回自動チューニング研究会, Dec. 2015.
60. Rio Yokota, Tetsu Narumi, Kenji Yasuoka, Toshikazu Ebisuzaki, Shinnosuke Obi. MDGRAPE-3 を用いた渦法による乱流の直接数値シミュレーション, 次世代スーパーコンピューティング・シンポジウム, Oct. 2007.
61. 横田理央, 小尾晋之介. 渦法による一様等方性乱流の解析, 第 20 回数値流体力学シンポジウム, Dec. 2006.

CONFERENCE
PROGRAM
COMMITTEE
MEMBER

1. International Conference on Parallel Architectures and Compilation Techniques (PACT 2025), publicity chair
2. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2025), proceedings vice-chair, program committee, AI4S committee
3. Conference on Neural Information Processing Systems (NeurIPS 2025), reviewer
4. IEEE International Conference on Cluster Computing (IEEE CLUSTER 2025), track co-chair
5. International Conference on Machine Learning (ICML 2025), reviewer
6. ISC High Performance (ISC 2025), program committee
7. Platform for Advanced Scientific Computing (PASC 2025), domain co-chair
8. 39th IEEE International Parallel and Distributed Processing Symposium, (IPDPS 2025), program committee, best OSS judge
9. The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR 2024), reviewer
10. 38th Conference on Neural Information Processing Systems (NeurIPS 2024), reviewer
11. The 12th International Conference on Learning Representations (ICLR 2024), reviewer
12. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2024), poster, ML track, AI4S/TPC WS, LLMs for HPC BoF
13. International Conference on Preconditioning Techniques for Scientific and Industrial Applications, (Precond 2024), program committee
14. 38th IEEE International Parallel and Distributed Processing Symposium, (IPDPS 2024), track co-chair
15. ISC High Performance (ISC 2024), posters chair
16. The 30th International Conference on Parallel Processing (Euro-Par 2024), program committee
17. ACM International Symposium on High-Performance, Parallel and Distributed Computing (HPDC 2024), program committee
18. SC Asia (SCA 2024), program committee
19. Platform for Advanced Scientific Computing (PASC 2024), program committee
20. SIAM Conference on Linear Algebra (LA24), scientific committee
21. SIAM Conference on Computational Science and Engineering (CSE23), poster judge

22. 10th International Congress on Industrial and Applied Mathematics (ICIAM 2023), program committee
23. IEEE International Conference on Cluster Computing (IEEE CLUSTER 2023), publicity chair
24. Principles and Practice of Parallel Programming (PPoPP 2023), publicity chair
25. 37th IEEE International Parallel and Distributed Processing Symposium, (IPDPS 2023), program committee
26. ISC High Performance (ISC 2023), posters deputy chair
27. Platform for Advanced Scientific Computing (PASC 2023), mini symposia and posters committee, program committee
28. The 23rd IEEE/ACM international Symposium on Cluster, Cloud and Internet Computing (CCGrid 2023), program committee
29. The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR 2023), reviewer
30. The 32nd International Conference on Parallel Architectures and Compilation Techniques (PACT 2023), publicity chair Asia
31. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2023), posters, workshops, ML tech paper, panelist, best paper
32. 37th Conference on Neural Information Processing Systems (NeurIPS 2023), reviewer
33. The Eleventh International Conference on Learning Representations (ICLR 2023), reviewer
34. European Conference on Computer Vision (ECCV 2022), reviewer
35. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2022), program committee, workshop chair
36. 49th International Conference on Parallel Processing (ICPP 2022), program committee
37. IEEE Cluster (IEEE Cluster 2022), program committee
38. ISC High Performance (ISC 2022), poster committee
39. The 36th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2022), program committee
40. International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia 2022), Applications track co-chair
41. IEEE International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), program committee
42. 48th International Conference on Parallel Processing (ICPP 2021), program committee
43. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2021), program committee
44. IEEE Cluster (IEEE Cluster 2021), program committee
45. ISC High Performance (ISC 2021), program committee, best paper committee, steering committee
46. The 35th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2021), program committee
47. International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia 2021), organizing committee
48. SIAM Computational Science and Engineering (SIAM CSE 2021), organizing committee
49. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2020), ML track Chair
50. The 26th International Conference on Parallel Processing (Euro-Par 2020), program committee

51. SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP 2020), program committee
52. International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia 2020), program committee
53. The 3rd cross-disciplinary Workshop on Computing Systems, Infrastructures, and Programming (xSIG 2019), program committee
54. IEEE 13th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoc 2019), ATMG Track Chair
55. 48th International Conference on Parallel Processing (ICPP 2019), Applications track chair, Workshop co-organizer
56. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2019), program committee
57. ISC High Performance (ISC 2019), program committee
58. IEEE Cluster (IEEE Cluster 2019), program committee
59. Platform for Advanced Scientific Computing Conference (PASC 2019), program committee
60. SIAM Conference on Computational Science and Engineering (SIAM CSE19), organizing committee
61. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2018), proceedings chair
62. ISC High Performance (ISC 2018), proceedings chair
63. IEEE Cluster (IEEE Cluster 2018), program committee
64. The 32nd IEEE International Parallel & Distributed Processing Symposium (IPDPS 2018), program committee
65. 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2018), program committee
66. SC Asia (SCA 2018), program committee
67. SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP 2018), organizing committee
68. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2017), program committee
69. Platform for Advanced Scientific Computing Conference (PASC 2017), program committee
70. The 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS 2017), program committee
71. International Conference on Supercomputing (ICS 2017), program committee
72. IEEE Cluster (IEEE Cluster 2017), program committee
73. The 23rd International Conference on Parallel Processing (Euro-Par 2017), program committee
74. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2016), program committee
75. Platform for Advanced Scientific Computing Conference (PASC 2016), program committee
76. ISC High Performance (ISC 2016), program committee
77. The 30th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2016), program committee
78. The 23rd annual IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC 2016), program committee

79. 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2016), program committee
80. International Conference on Supercomputing (ICS 2016), program committee
81. IEEE Cluster (IEEE Cluster 2016), program committee
82. 28th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD 2016), program committee
83. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2015), program committee
84. The 22nd annual IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC 2015), program committee
85. 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2015), program committee
86. IEEE Cluster (IEEE Cluster 2015), program committee
87. The 20th International Conference on Parallel Processing (Euro-Par 2014), program committee
88. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2014), program committee
89. The International Meeting on High-Performance Computing for Computational Science (VECPAR 2014), program committee
90. The annual IEEE International Conference on High Performance Computing (HiPC 2014), program committee
91. 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2014), program committee
92. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2013), program committee
93. The 28th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2013), program committee
94. 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2013), program committee
95. The annual IEEE International Conference on High Performance Computing (HiPC 2013), program committee
96. 15th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC 2013), program committee
97. The 12th International Symposium on Parallel and Distributed Computing (ISPDC 2012), program committee
98. The Third International Workshop on Frontier of GPU Computing (FCG 2012), program committee

RESEARCH GRANTS

1. 2026–2030. Realization of a next-generation high-performance linear solver to pioneer the future of computational science and engineering (Grant-in-Aid for Scientific Research (A), JSPS)
2. 2026–2029. High-performance urban model computing (Grant-in-Aid for Scientific Research (B), JSPS)
3. 2025–2030. Creation of GPU educational materials and implementation of educational programs, and promotion of advanced utilization of GPU computing resources centered on generative AI technologies (MEXT Next-Generation HPC and AI Development Support Hub Program, MEXT)

4. 2025–2028. Modernization of scientific and technical software to unlock the potential of next-generation computers (Grant-in-Aid for Scientific Research (A), JSPS)
5. 2024–2030. Red-teaming platform for large language model misalignment (JST K Program, JST)
6. 2024–2029. Research on distributed parallel pre-training of large language models (JST Program for Forming Research and Development Centers for Next-Generation AI Technologies, JST)
7. 2024–2028. Building AI foundation models under limited resources through continual pre-training using synthetic data and small amounts of real data (Fund for the Promotion of Joint International Research (Fostering Joint International Research), JSPS)
8. 2024–2027. Expanding the applicability of low-rank structured matrix methods and leveraging diverse computing architectures (Grant-in-Aid for Scientific Research (B), JSPS)
9. 2024–2027. Research toward establishing an outcome-extraction methodology for evaluating the efficacy and safety of pharmaceuticals using Large Language Models (LLM) (Health, Labour and Welfare Sciences Research Grant, Applied Research, MHLW)
10. 2023–2026. Compositional pattern recognition and understanding using deep generative models (Grant-in-Aid for Scientific Research (A), JSPS)
11. 2022–2026. Functional improvement of mixing promotion in fluid equipment by elucidating the universal statistical law of two-phase turbulence (Grant-in-Aid for Scientific Research (B), JSPS)
12. 2022–2025. Fast and accurate eigenvalue calculations by hierarchical low-rank approximation and its application to large-scale electronic structure calculations (Grant-in-Aid for Scientific Research (B), JSPS)
13. 2021–2027. Adaptive, robust, and continual artificial intelligence systems based on the Bayes duality principle (CREST, JST)
14. 2021–2024. Construction of numerical linear algebra based on lattice H-matrices and its high-performance implementation on modern architectures (Grant-in-Aid for Scientific Research (B), JSPS)
15. 2020–2023. Application of unconventional linear algebra techniques to continual learning in supergiant neural networks (Grant-in-Aid for Challenging Research (Pioneering), JSPS)
16. 2020–2023. Life-Long Deep Learning using Bayesian Principles (Grant-in-Aid for Scientific Research (B), JSPS)
17. 2018–2021. Linear solvers for machine learning hardware (Grant-in-Aid for Scientific Research (B), JSPS)
18. 2017–2020. Enhancement of an H-matrix library and optimization for next-generation supercomputers (Grant-in-Aid for Scientific Research (B), JSPS)
19. 2016–2022. Acceleration and economization of deep learning algorithms for image processing in social infrastructure (CREST, JST)
20. 2016–2019. Hierarchical particle-method framework for exascale computers balancing performance and productivity (Grant-in-Aid for Scientific Research (B), JSPS)
21. 2016–2018. Large-scale iterative solvers combining the FMM and H-matrices (Grant-in-Aid for Young Scientists (A), JSPS)
22. 2015–2016. Algebraic extension of the FMM as a preconditioner for exascale large-scale linear systems (Grant-in-Aid for Research Activity Start-up, JSPS)