Prasanna Balaprakash

Director of AI Programs and Distinguished R&D Scientist
Computing and Computational Sciences Directorate
Oak Ridge National Laboratory
P.O. Box 2008, 1 Bethel Valley Road, Oak Ridge, TN 37831
https://pbalapra.github.io/
linkedin.com/in/prasannaprakash
pbalapra@ornl.gov

Summary

Since March 2023, I have been leading the Artificial Intelligence Initiative at Oak Ridge National Laboratory. In this role, I oversee laboratory research, development, and application of artificial intelligence and machine learning (AI/ML) to address national challenges. My seek to deliver foundational, scalable, and practical AI/ML solutions that support Oak Ridge National Laboratory's diverse mission, offering world-class expertise in computer and computational science, neutron science, materials science, biology and health science, nuclear engineering, isotopes, manufacturing, energy, and climate science.

My research interests encompass artificial intelligence, machine learning, optimization, and high-performance computing.

I am honored to have received the U.S. Department of Energy 2018 Early Career Award and currently lead the AI sector in DOE/ASCR RAPIDS2: A SciDAC Institute for Computer Science, Data, and Artificial Intelligence. Additionally, I lead the research and development efforts of DeepHyper, a scalable automated machine learning package designed for DOE leadership-class systems.

Degrees Attained

- Ph.D., Engineering Sciences (Computer Science), Université libre de Bruxelles, Belgium, 09/15/2004—01/20/2010. Advisors: Prof. Marco Dorigo, Dr. Thomas Stuetzle, and Dr. Mauro Birattari.
- Master of Advanced Studies, Computer Science, Université libre de Bruxelles, Belgium, 09/15/2004—09/30/2005. Advisors: Prof. Marco Dorigo, Dr. Thomas Stuetzle, and Dr. Mauro Birattari.
- Master of Science, Computer Science, Otto-von-Guericke Universität Magdeburg, Germany, 09/01/2002—04/30/2004. Advisors: Prof. Graham Horton and Dr. Sanja Lazarova-Molnar.
- Bachelor of Engineering, Computer Science, Periyar University, Salem, India, 07/01/1998—04/30/2002
- High Scool, Sengunther Higher Secondary School, Erode, India, 1996–1998

Positions Held

Director of AI Programs and Distinguished R&D Scientist March 2023–Present
Computing and Computational Sciences Directorate
Oak Ridge National Laboratory, Oak Ridge, TN

Computer Scientist (RD5)/R&D Leader (L1) April 2022–February 2023
Computer Scientist (RD4)/R&D Leader (L1) August 2021–April 2022

Computer Scientist (RD4) June 2019–July 2021

Computer Scientist (RD3) June 2017–May 2019

Assistant Computer Scientist (RD2) December 2013–May 2017

Mathematics and Computer Science (MCS) Division & Leadership Computing Facility
Argonne National Laboratory, Lemont, IL

Scientist at Large, Consortium for Advanced Science and Engineering July 2018–Present
Fellow, Computation Institute August 2014–June 2018

University of Chicago, Chicago, IL

Fellow, Northwestern-Argonne Institute of Science and Engineering January 2017–Present
Northwestern University, Evanston, IL

Co-founder and Chief Technology Officer September 2012–November 2015
inSiliTech LLC, formed out of Chicago Booth's New Venture Challenge
Chicago, IL

Postdoctoral Appointee September 2010–November 2013
Mathematics and Computer Science Division
Argonne National Laboratory, Lemont, IL

Chief Technology Officer January 2009–August 2010
Mentis Consulting Sprl, Brussels, Belgium

Fellowships

- **U.S. Department of Energy 2018 Early Career Award** funded by the Advanced Scientific Computing Research program within the DOE Office of Science. $2.5M total for 5 years.
- **F.N.R.S. chargé de recherches fellowship**, from the Belgian Funds for Scientific Research from October 2010 to September 2013 (one of the most competitive postdoctoral fellowships in Belgium). Declined in order to accept postdoc position at Argonne National Laboratory.
- **F.N.R.S. aspirant fellowship**, from the Belgian Funds for Scientific Research, from October 2006 to September 2008 (one of the most competitive Ph.D. fellowships in Belgium).
- **Marie Curie fellowship**, from the European Commission through the fifth framework human resources and mobility program, from November 2004 to September 2006 (one of the most competitive Ph.D. fellowships at the European level).
- **International student scholarship**, for the best academic performance from Otto-von-Guericke Universität Magdeburg, Germany, from April 2003 to March 2004 (scholarship awarded every year to the top two international students).

Other Honors and Awards

- **Best paper award** at the 19th IEEE International Conference on eScience, December, 2023.
- Editor's pick, Surrogate modeling of advection-dominated flows using deep convolutional autoencoders, Journal of Physics of Fluids, 2021.
- Editor's pick, Non-autoregressive time-series methods for stable parametric reduced-order models, Journal of Physics of Fluids, 2020.

- **Impact Argonne Award** for Enhancement to Argonne's Reputation, 2020.
- Highly cited article, Time-series learning of latent-space dynamics for reduced-order model closure, Journal of Physica D: Nonlinear Phenomena, 2020.
- Nominated for **2019 HPCwire Readers' Choice Awards** under Best Use of HPC in Automotive.
- **Best poster finalist**, High Performance Computing, Networking, Storage and Analysis (SC 13), Denver, CO, November 2013. (211 submissions; 84 accepted; 5 finalists)
- **Best paper award**, 21st High Performance Computing Symposia (HPC), The Society for Modeling & Simulation International, San Diego, CA, April 2013.

Selected Research Coverage

- AI director featured at White House roundtable, Capitol Hill briefings, January, 2024
- Extracting the impact of climate change from scientific literature using Snorkel-enabled NLP, Highlight talk at the Future of Data-Centric AI conference, August, 2022.
- What's new in HPC research, HPCWire, July 2022. (Coverage on asynchronous distributed Bayesian optimization)
- What's new in HPC research, HPCWire, June 2022. (Coverage on taxonomy of error sources in HPC I/O machine learning models).
- Argonne to host DeepHyper training session July 15, insideHPC, June 2022.
- DeepHyper: Scalable automated machine learning for scientific applications with Prasanna Balaprakash, Applied AI podcast, May 2022.
- Argonne cuts through the noise with SambaNova system, Next Platform, July 2021.
- Can a national lab forecast traffic jams to prevent them in future?, Radio Interview, Federal News Network, December, 2020.
- Argonne scientists use reinforcement learning to train quantum algorithm, HPCWire, August, 2020.
- Neural architecture search automates the development of deep learning-based models for cancer research, News Medical, November, 2019.
- An A.I. tool for cancer researchers, Eye on A.I. Research in Fortune magazine AI newsletter, September, 2019.
- Argonne team developing MaLTESE for onboard real-time adaptive learning and controls, Green Car Congress, September, 2019.
- Lessons machine-learned, ASCR Discovery, May, 2019.

Collaboration, Management, and Leadership

Supervision: Staffs, Postdocs, Predocs, and Research Aides

1. Adarsha Balaji, Postdoctoral Researcher 2022–2023
   Area: Neuromorphic Computing, Hardware Software Co-design
2. Hongwei Jin, Postdoctoral Researcher 2022–2023
   Area: Graph Neural Networks, Numerical Optimization, Federated Learning
3. Yixuan Sun, Postdoctoral Researcher 2022–2023
   Area: Spatiotemporal Machine Learning, Automated Machine Learning
4. Romit Maulik, Assistant Computational Scientist 2021–2023
   Area: Physics-informed Machine Learning, Computational Fluid Dynamics
   Next position: Assistant Professor, the Pennsylvania State University
5. Krishnan Ragavan, Assistant Computational Mathematician 2019–2023
   Area: Continual Learning, Deep Learning Theory
6. Tanwi Mallick, Assistant Computer Science Specialist 2018–2023
   Area: Spatiotemporal Machine Learning, Natural Language Processing
7. Romain Egele, Research Aide 2018–2023
   Area: Automated Machine Learning, Scalable Deep Learning
8. Sandeep Madireddy, Assistant Computer Scientist 2016–2023
   Area: Probabilistic Machine Learning, Neuromorphic Computing
9. Sami Khairy, Postdoctoral Researcher 2021–2022
   Area: Reinforcement Learning for Science
   Current position: Research Scientist, Microsoft, Canada
10. Jaehoon Koo, Postdoctoral Researcher 2020–2022
    Area: Physics-Informed Machine Learning for Fusion Science, Autotuning
    Current position: Assistant Professor, Hanyang University, South Korea
11. Grant Getzelman, Predoctoral Researcher, with Stefan Wild and Jeff Larson 2019–2021
    Area: Learning to Optimize
    Current position: Senior Machine Learning Scientist, Biotech Startup, U.S.

Supervision: Students

12. Gideon Idumah (Case Western Reserve U.), SRP-HPC Visiting student 2022
13. Bilas Talukdar (New Jersey Institute of Tech.), SRP-HPC Visiting student 2022
14. Akhil Akella (Northern Illinois U.), ALCF summer student, Givens Fellow 2020, 2022
15. Vincent Zhong (U. of Illinois Urbana-Champaign), Givens Fellow, with Tanwi Mallick 2022
16. Joceran Gouneau (ENSEEIHT, France), Research Internship 2021–2022
17. Scott Emmons (U. of California, Berkeley), DOE CSGF Fellow 2021

18. Alec Linot (U. of Wisconsin Madison), Givens Fellow, with Romit Maulik 2021
19. Michael Mccabe (U. of Colorado Boulder), Givens Fellow 2021
20. Felix Edward Perez (U. of Texas at Dallas), EERE HPC4Mfg Internship, with Romain Egele 2021
21. Edward Friesema (U. of Nevada), EERE Robotics Internship Program, with Sami Khairy 2021
22. Thomas Randall (Clemson U.), Givens Fellow 2020, 2022
23. Yixuan Sun (Purdue U.), Givens Fellow, with Tanwi Mallick 2020, 2021
24. Shengli Jiang (U. of Wisconsin Madison), Givens Fellow 2020
25. Yixuan Sun (Purdue U.), Givens Fellow, with Tanwi Mallick 2020
26. Andres Rodríguez Rey (U. of California, Sand Diego), NSF MSGI Fellow 2020
27. Sami Khairy (Illinois Institute of Tech.), ALCF Summer Student, Givens Fellow 2019, 2020
28. Peihong Jiang (Brown U.), NSF MSGI Fellow 2019
29. Tianchen (Eric) Zhao (U. of Michigan), Givens Fellow, with Sandeep Madireddy 2018
30. Nimish Awalgaonkar (Purdue U.), Givens Fellow 2018
31. Dipnil Chakraborty (U. of Texas at Dallas), NSF MSGI Fellow 2018
32. Dipendra Jha (Northwestern U.), Givens Fellow 2018
33. Akash Roy (U. of Texas at Dallas), Research Aide 2018
34. Andy Jin (U. of Chicago), Jeff Metcalf Undergraduate Fellow 2018
35. Tian Ma (U. of Chicago), Jeff Metcalf Undergraduate Fellow 2018
36. YiMing Yu (New Jersey Institute of Tech.), NSF MSGI Fellow, with Paul Hovland 2017
37. Prateek Agarwal (Illinois Institute of Tech.), Resident Associate, with Sven Leyffer 2017
38. Amal Fethi (ENS Paris, France), Resident Associate, with Sven Leyffer 2017
39. Salvador Aguinaga (U. of Notre Dame), DOE CGSR Fellow 2017
40. Juan Li (U. of Chicago), Resident Associate, with Sven Leyffer 2016
41. Amit Roy (U. of Utah), Givens Associate, with Paul Hovland 2015
42. Arnamoy Bhattacharyya (ETH Zurich, Switzerland), Givens Associate 2014

Project Leadership and Management

1. RAPIDS2: A SciDAC Institute for Computer Science and Data. Funding agency: DOE/ASCR; Role: AI lead; Team size: 4 (direct); 20+ (coordination). 2020–2025.
2. PosEiDon: Platform for Explainable Distributed Infrastructure. Funding agency: DOE/ASCR; Role: Argonne PI; Team size: 2. 2021–2024.
3. Probabilistic Machine Learning for Rapid Large-scale and High-rate Aerostructure Manufacturing. Funding agency: DOE/EERE; Role: Lead PI; Team size: 10 (including external partners from GE Research and Edison Welding Institute). 2021–2023.
4. Big Data and AI for Smart Mobility. Funding agency: DOE/EERE-VTO; Role: MCS PI; Team size: 2. 2021–2023.
5. Ab-initio Guided Design and Materials Informatics for Accelerated Product Development of Next Generation Advanced High Strength Steels (AHSS). Funding agency: DOE HPC for Manufacturing; Role: Argonne PI; Team size: 1. 2021–2022.
6. PROTEAS-TUNE. Funding agency: DOE/ECP; Role: MCS lead for autotuning; Team size: 3. 2020–2023.
7. Scalable Data-Efficient Learning for Scientific Domains (Early Career Project). Funding agency: DOE/ASCR; Role: Lead PI; Team size: 2. 2018–2023.

Program Development

Research proposals funded (External)

1. Exploring the Power of Distributed Intelligence for Resilient Scientific Workflows, Funding agency: DOE ASCR; Role: Co-PI; Amount: $1.75M to ORNL; Duration: 09/23–08/28.
2. Center for Steel Electrification by Electrosynthesis Funding agency: DOE BES and ASCR; Role: Co-PI; Amount: $1M to ORNL; Duration: 09/23–08/27.
3. High-fidelity Digital Models for Fusion Pilot Plant Design Funding agency: DOE FES and ASCR; Role: Senior personnel; Amount: $1.2M to ORNL; Duration: 09/23–9/27.
4. Development of High-Fidelity Simulation Capabilities for ELM-free Design Optimization Funding agency: DOE FES and ASCR; Role: Co-PI ; Amount: $240K to ORNL; Duration: 09/23–08/27.
5. Fundamental nuclear physics at the exascale and beyond Funding agency: DOE Nuclear Physics and Advanced Scientific Computing Research; Role: Co-PI; Amount: $13M, $1.15M to MCS/Argonne; Duration: 10/22–9/27.
6. Improving Projections of AMOC and its Collapse Through advanced Simulations (ImPACTS) Funding agency: DOE Biological and Environmental Research and Advanced Scientific Computing Research; Role: Senior personnel; Amount: $15M, $2.78M to MCS/Argonne; Duration: 10/22–9/27.
7. Community Research on Climate and Urban Science (CROCUS) Funding agency: DOE Biological and Environmental Research; Role: Senior personnel; Amount: $25M, $600K to MCS/Argonne; Duration: 10/22–9/27.

8. Probabilistic Impact Scenarios for Extreme Weather Event Resilience Funding agency: DOE Office of Electricity and Advanced Scientific Computing Research; Role: Senior personnel; Amount: $700K, $275K to MCS/Argonne; Duration: 10/22–9/23.

9. PosEiDon: Platform for Explainable Distributed Infrastructure. Funding agency: DOE Advanced Scientific Computing Research; Role: PI; Amount: $725K to MCS/Argonne; Duration: 10/21–9/24.

10. Inertial neural surrogates for stable dynamical prediction. Funding agency: DOE Advanced Scientific Computing Research; Role: Co-PI; Amount: $1.2M to MCS/Argonne; Duration: 10/21–9/24.

11. Probabilistic Machine Learning for Rapid Large-scale and High-rate Aerostructure Manufacturing. Funding agency: DOE Office of Energy Efficiency & Renewable Energy; Role: Lead PI; Amount: $2.43M, $550K to MCS/Argonne; Duration: 10/21–9/23.

12. Developing Easy-to-Use Application-Based Software to Combat Social Media Misinformation–a Multi-Disciplinary Collaboration. Funding agency: Discovery Partner Institute Seed Grant; Role: Argonne PI; Amount: $150K total; Duration: 10/21–9/22.

13. RAPIDS2: A SciDAC Institute for Computer Science and Data. Funding agency: DOE Advanced Scientific Computing Research; Role: Co-PI (AI lead); Amount: $5.75M/yr, $1.15M/year to MCS/Argonne; Duration: 10/20–9/25.

14. Big Data and AI for Smart Mobility. Funding agency: DOE Office of Energy Efficiency & Renewable Energy; Role: Co-PI; Amount: $1.1M for ANL, $600K to MCS/Argonne; Duration: 10/20–9/22.

15. DAIN: Dynamic architectures through introspection and neuromodulation (Phase II). Funding agency: Defense Advanced Research Projects Agency; Role: Co-PI; Amount: $300K to MCS/Argonne; Duration: 10/20–9/22.

16. Scalable Deep Learning Framework for Optimal Control of Cascading Failures and Restoration in Power System. Funding agency: DOE Office of Electricity; Role: Co-PI; Amount: $380K to MCS/Argonne; Duration: 4/20–3/22.

17. Ab-initio Guided Design and Materials Informatics for Accelerated Product Development of Next Generation Advanced High Strength Steels (AHSS). Funding agency: DOE High-Performance Computing for Manufacturing; Role: Argonne PI; Amount: $300K total, $90K to MCS/Argonne; Duration: 10/20–9/21.

18. PROTEAS-TUNE. Funding agency: DOE Exascale Computing Project; Role: Co-PI; Amount: $2M for MCS/Argonne; Duration: 10/19–9/23.

19. Quantifying Energy Drivers in Chemical Separations. Funding agency: DOE Basic Energy Sciences; Role: Co-PI; Amount: $1.8M (out of $2.25M total) to Argonne, $675K to MCS; Duration: 9/19–9/22.

20. Scalable Data-Efficient Learning for Scientific Domains (Early Career Proposal). Funding agency: DOE Advanced Scientific Computing Research; Role: Lead PI; Amount: $2.5M to MCS/Argonne; Duration: 9/18–8/23.

21. Accelerating HEP Science: Inference and Machine Learning at Extreme Scales (SciDAC Application Partnership); Funding agency: DOE High Energy Physics & Advanced Scientific Computing Research; Role: Senior personnel; Amount for Argonne: $260K/year for years 1 and 2 and $450K/year for years 3 to 5; Duration: 10/17–9/22.

22. High Performance Computing and Big Data Solutions for Mobility Design and Planning Funding agency: DOE Office of Energy Efficiency & Renewable Energy; Role: Senior personnel; Amount: $600K to MCS/Argonne; Duration: 10/17–9/20; .

23. RAPIDS: A SciDAC Institute for Computer Science and Data (SciDAC Institute) Funding agency: DOE Advanced Scientific Computing Research; Role: Senior personnel; Amount for Argonne: $3.405M; Duration: 10/17–9/20.

24. CANDLE: Exascale Deep Learning and Simulation Enabled Precision Medicine for Cancer (ECP Application). Role: Named personnel; Funding agency: DOE Exascale Computing Project; Amount: $650K for Argonne; Duration: 10/16–9/20.

25. Autotuning Compiler Technology for Cross-Architecture Transformation and Code Generation (ECP Software Technology). Funding agency: DOE Exascale Computing Project; Role: Senior personnel; Amount for Argonne: $1.05M; Duration: 10/16–9/19.

26. Machine Learning Guided Error Resilience. Role: Co-PI; Funding agency: UChicago–ANL (Seed Grant); Amount: $8K to MCS/Argonne; Duration: 7/16–6/17.

27. Automated Machine Learning Pipeline for City of Chicago's Smart Data Platform. Role: Co-PI; Funding agency: Bloomberg/City of Chicago; Amount for ANL: $250K for MCS/Argonne; Duration: 4/16–10/17.

28. XTUNE: Autotuning for Exascale. Role: Named personnel; Funding agency: DOE Advanced Scientific Computing Research; Amount: $87K for MCS/Argonne. Duration: 1/15–31/16.

29. SDAV/SUPER Supplement: Improving Computational Science Throughput via Model-Based I/O Optimization. Role: Named personnel; Funding agency: DOE Advanced Scientific Computing Research; Amount: $320K to MCS/Argonne; Duration: 2/15–9/16.

30. OrFPGA: Empirical Performance Tuning for FPGA Designs in Space Applications. Role: Named personnel; Funding agency: NASA-SBIR; Amount: $200K to MCS/Argonne; Duration: 2/15–9/16.

31. RAMSES: Robust Analytical Models for Science at Extreme Scale. Role: Named personnel; Funding agency: DOE Advanced Scientific Computing Research; Amount: $3M for MCS/Argonne; Duration: 7/14–6/17.

Research proposals funded (Internal > $100K)

32. Community and Infrastructure Adaptation to Climate Change: An AI-Driven Research Tool. Funding agency: ANL/LDRD (Prime: Climate and Energy Action); Role: Collaborator; Amount: $560K; Duration: 10/21–9/23.

33. Towards Neighborhood Scale Climate Simulations using AI and Accelerated GPUs. Funding agency: ANL/LDRD (Prime: Climate and Energy Action); Role: Collaborator; Amount: $508K; Duration: 10/21–9/23.

34. AI-Emulator Assisted Data-Assimilation. Funding agency: ANL/LDRD (Prime: Future Computing); Role: Co-PI; Amount: $1M; Duration: 10/20–9/23.

35. Automated Model Inference for Cosmological Structure Formation Funding agency: ANL/LDRD (Prime: Future Computing); Role: Co-PI; Amount: $1M; Duration: 10/20–9/23.

36. Nuclear Quantum Monte Carlo with Machine Learning Techniques Funding agency: ANL/LDRD (Prime: Future Computing); Role: Co-PI; Amount for ANL: $750K; Duration: 10/20–9/23.

37. A Framework for Device and Architecture Co-Design via Hyperparameter Optimization and Backpropagation through Hardware; Funding agency: ANL LDRD Prime - Microelectronics; Role: Co-PI; Amount for ANL: $260K; Duration: 1/20–9/20.

38. Artificial Intelligence Assisted Safety Modeling and Analysis of Advanced Nuclear Reactors Funding agency: ANL/LDRD; Role: Co-PI; Amount for ANL: $505K; Duration: 10/19–9/21.

39. Towards the Creation of an Advanced Mobility Cybersecurity Testbed. Funding agency: ANL LDRD (Prime: Secure Energy and Critical Resources); Role: Co-PI; Amount: $360K; Duration: 10/18–10/19.

40. A.I C.D.I: Atomistically Informed Coherent Diffraction Imaging; Funding agency: ANL/LDRD; Role: Co-PI; Amount: $180K; Duration: 10/16–9/18.

41. SLIK-D: Scalable Machine Learning Infrastructures for Knowledge Discovery; Funding agency: ANL/LDRD; Role: Co-PI; Amount: $600K Duration: 10/16–9/18.

42. End-to-End Genome Annotation and Phenotype Prediction with Deep Learning. Funding agency: ANL/LDRD; Role: Co-PI; Amount: $300K Duration: 10/16–9/18.

Research proposals funded (Internal $<$ $100K)

43. Accelerating Quantum Machine Learning for Fusion Sciences with SambaNova. Role: PI; Funding agency: ANL/LDRD (CELS LDRD Expedition); Amount: $26K; Duration: 7/21–9/21.

44. Accelerating Inversion of Nuclear Responses with Uncertainty Quantification. Role: Co-PI; Funding agency: ANL/LDRD (CELS LDRD Expedition); Amount: $26K; Duration: 7/21–9/21.

45. CARIBU-matic: Automation of CARIBU beam tuning enabled by Machine Learning. Role: Co-PI; Funding agency: ANL/LDRD (SWIFT); Amount: $50K; Duration: 7/21–9/21.

46. A Unified Framework for Sample-Efficient Reinforcement Learning. Role: Collaborator; Funding agency: ANL/LDRD (Innovate); Amount: $25K; Duration: 6/21–9/21.

47. Accelerating Graph-convolution-based Deep Learning Framework for Large-Scale Highway Traffic Forecasting with SambaNova. Role: Co-PI; Funding agency: ANL/LDRD (CELS LDRD Expedition); Amount: $30K; Duration: 7/21–9/21.

48. Neuromorphic Computing for DOE Deep Learning Applications. Funding agency: ANL/LDRD (CELS LDRD Expedition); Role: PI; Amount for ANL: $25K; Duration: 1/19–31/19.

49. Neuromorphic Computing for DOE Deep Learning Applications. Funding agency: ANL/LDRD (CELS LDRD Expedition); Role: PI; Amount: $25K; Duration: 9/18–9/18.

Selected Research Products

Citations: 4,118; h-index: 33; i10-index: 79 (Source: Google Scholar; 02/28/2024)

Selected Papers in Surrogate Modeling of Complex Systems

1. R. Egele, P. Balaprakash, I. Guyon, V. Vishwanath, F. Xia, R. Stevens, Z. Liu. AgEBO-Tabular: Joint neural architecture and hyperparameter search with autotuned data-parallel training for tabular data. In *SC21: International Conference for High Performance Computing, Networking, Storage and Analysis*, 2021.

2. K. Raghavan and P. Balaprakash. Formalizing the generalization-forgetting trade-off in continual learning. *Advances in Neural Information Processing Systems*, 34:17284–17297, 2021.

3. R. Maulik, R. Egele, B. Lusch, and P. Balaprakash. Recurrent neural network architecture search for geophysical emulation. In *SC '20: IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis*, 2020.

Selected Papers in AI-based Optimization of High-Performance Computing Systems

4. M. Dorier, R. Egele, P. Balaprakash, J. Koo, S. Madireddy, S. Ramesh, A. D. Malony, and R. Ross. HPC storage service autotuning using variational-autoencoder-guided asynchronous Bayesian optimization. In *2022 IEEE International Conference on Cluster Computing (CLUSTER)*, 2022. (In press)

5. S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash. Learning to optimize variational quantum circuits to solve combinatorial problems. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(03):2367–2375, 2020.

6. P. Balaprakash, J. Dongarra, T. Gamblin, M. Hall, J. K. Hollingsworth, B. Norris, and R. Vuduc. Autotuning in high-performance computing applications. *Proceedings of the IEEE*, pages 1–16, 2018.

Selected Papers in AI for Scientific Applications

7. Y. Yao, H. Chan, S. Sankaranarayanan, P. Balaprakash, R. J. Harder, and M. J. Cherukara. AutophaseNN: unsupervised physics-aware deep learning of 3D nanoscale Bragg coherent diffraction imaging. *npj Computational Materials*, 8(1):1–8, 2022.

8. J. Dunn and P. Balaprakash, eds. *Data Science Applied to Sustainability Analysis*. Elsevier, 2021.

9. R. Maulik, B. Lusch, and P. Balaprakash. Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders. *Physics of Fluids*, 33(3):037106, 2021. (Editor's pick)

Complete List of Research Products

(Selected papers from the previous section are excluded)

Refereed Journal Articles and Book Chapters

10. J. Kim, S. Leyffer, and P. Balaprakash "Learning symbolic expressions: Mixed-integer formulations, cuts, and heuristics", *INFORMS Journal on Computing*, Publisher: INFORMS, 2023.

11. S. Bhola, S. Pawar, P. Balaprakash, and R. Maulik "Multi-fidelity reinforcement learning framework for shape optimization", *Journal of Computational Physics*, Vol. 482, pp. 112018, Publisher: Elsevier, 2023.

12. H. Jin, K. Raghavan, G. Papadimitriou, C. Wang, A. Mandal, M. Kiran, E. Deelman, and P. Balaprakash "Graph neural networks for detecting anomalies in scientific workflows", *The International Journal of High Performance Computing Applications*, pp. 10943420231172140, Publisher: SAGE Publications Sage UK: London, England, 2023.

13. A. Linot, J. Burby, Q. Tang, P. Balaprakash, M. Graham, and R. Maulik "Stabilized neural ordinary differential equations for long-time forecasting of dynamical systems", *Journal of Computational Physics*, Vol. 474, pp. 111838, Publisher: Elsevier, 2023.

14. S. Qin, S. Jiang, J. Li, P. Balaprakash, R. Van Lehn, and V. Zavala "Capturing molecular interactions in graph neural networks: a case study in multi-component phase equilibrium", *Digital Discovery*, Vol. 2, No. 1, pp. 138–151, Publisher: Royal Society of Chemistry, 2023.

15. R. Maulik, R. Egele, and K. Raghavan, and P. Balaprakash, "Quantifying uncertainty for deep learning based forecasting and flow-reconstruction using neural architecture search ensembles", *Physica D: Nonlinear Phenomena*, Vol. 454, pp. 133852, doi: 10.1016/j.physd.2023.133852, 2023.

16. N. A. Garland, R. Maulik, Q. Tang, X. Tang, and P. Balaprakash. Efficient data acquisition and training of collisional-radiative model artificial neural network surrogates through adaptive parameter space sampling. *Machine Learning: Science and Technology*, 2022.

17. D. Yang, P. Balaprakash, and S. Leyffer. Modeling design and control problems involving neural network surrogates. *Computational Optimization and Applications*, 2022.

18. Y. Liu, R. Hu, A. Kraus, P. Balaprakash, and A. Obabko. Data-driven modeling of coarse mesh turbulence for reactor transient analysis using convolutional recurrent neural networks. *Nuclear Engineering and Design*, 390:111716, 2022.

19. X. Wu, M. Kruse, P. Balaprakash, H. Finkel, P. Hovland, V. Taylor, and M. Hall. Autotuning polybench benchmarks with llvm clang/polly loop optimization pragmas using bayesian optimization. *Concurrency and Computation: Practice and Experience*, 34(20):e6683, 2022.

20. L. L. Lao, S. Kruger, C. Akcay, P. Balaprakash, T. A. Bechtel, E. Howell, J. Koo, J. Leddy, M. Leinhauser, Y. Q. Liu, S. Madireddy, J. McClenaghan, D. Orozco, A. Pankin, D. Schissel, S. Smith, X. Sun, and S. Williams. Application of machine learning and artificial intelligence to extend EFIT equilibrium reconstruction. *Plasma Physics and Controlled Fusion*, 64(7):074001, 2022.

21. F. Xia, J. Allen, P. Balaprakash, T. Brettin, C. Garcia-Cardona, A. Clyde, J. Cohn, J. Doroshow, X. Duan, V. Dubinkina, et al. A cross-study analysis of drug response prediction in cancer cell lines. *Briefings in bioinformatics*, 23(1):bbab356, 2022.

22. Y. Sun, T. Mallick, P. Balaprakash, and J. Macfarlane. A data-centric weak supervised learning for highway traffic incident detection. *Accident Analysis & Prevention*, 176:106779, 2022.

23. R. Maulik, V. Rao, J. Wang, G. Mengaldo, E. Constantinescu, B. Lusch, P. Balaprakash, I. Foster, and R. Kotamarthi. Efficient high-dimensional variational data assimilation with machine-learned reduced-order models. *Geoscientific Model Development*, 15(8):3433–3445, 2022.

24. S. Khairy, P. Balaprakash, L. X. Cai, and H. V. Poor. Data-driven random access optimization in multi-cell IoT networks with NOMA. *IEEE Transactions on Wireless Communications*, 2021.

25. P. Balaprakash and J. B. Dunn. Overview of data science and sustainability analysis. In J. Dunn and P. Balaprakash, editors, *Data Science Applied to Sustainability Analysis*, pages 1–14. Elsevier, 2021.

26. J. B. Dunn and P. Balaprakash. Research and development for increased application of data science in sustainability analysis. In J. Dunn and P. Balaprakash, editors, *Data Science Applied to Sustainability Analysis*, pages 283–292. Elsevier, 2021.

27. K. Raghavan, P. Balaprakash, A. Lovato, N. Rocco, and S. M. Wild. Machine-learning-based inversion of nuclear responses. *Physical Review C*, 103(3):035502, 2021.

28. R. Maulik, A. Mohan, B. Lusch, S. Madireddy, P. Balaprakash, and D. Livescu. Time-series learning of latent-space dynamics for reduced-order model closure. *Physica D: Nonlinear Phenomena*, 405:132368, 2020. (Highly cited article)

29. R. Maulik, B. Lusch, and P. Balaprakash. Non-autoregressive time-series methods for stable parametric reduced-order models. *Physics of Fluids*, 32(8), 2020. (Editor's pick)

30. S. Madireddy, J. H. Park, S. Lee, P. Balaprakash, S. Yoo, W.-K. Liao, C. D. Hauck, M. P. Laiu, and R. Archibald. In situ compression artifact removal in scientific data using deep transfer learning and experience replay. *Machine Learning: Science and Technology*, 2(2):025010, 2020.

31. R. Maulik, N. A. Garland, J. W. Burby, X.-Z. Tang, and P. Balaprakash. Neural network representability of fully ionized plasma fluid model closures. *Physics of Plasmas*, 27(7):072106, 2020.

32. T. Mallick, P. Balaprakash, E. Rask, and J. Macfarlane. Graph-partitioning-based diffusion convolutional recurrent neural network for large-scale traffic forecasting. *Transportation Research Record*, 2674(9):473–488, 2020.

33. J. Wang, P. Balaprakash, and R. Kotamarthi. Fast domain-aware neural network emulation of a planetary boundary layer parameterization in a numerical weather forecast model. *Geoscientific Model Development*, 12(10):4261–4274, 2019.

34. S. Madireddy, D.-W. Chung, T. Loeffler, S. K. Sankaranarayanan, D. N. Seidman, P. Balaprakash, and O. Heinonen. Phase segmentation in atom-probe tomography using deep learning-based edge detection. *Scientific Reports*, 9(1):1–10, 2019.

35. J. M. Wozniak, R. Jain, P. Balaprakash, J. Ozik, N. T. Collier, J. Bauer, F. Xia, T. Brettin, R. Stevens, J. Mohd-Yusof, C. G. Cardona, B. V. Essen, and M. Baughman. CANDLE/Supervisor: A workflow framework for machine learning applied to cancer research. *BMC Bioinformatics*, 19(18):491, 2018.

36. O. Subasi, S. Di, L. Bautista-Gomez, P. Balaprakash, O. Unsal, J. Labarta, A. Cristal, S. Krishnamoorthy, and F. Cappello. Exploring the capabilities of support vector machines in detecting silent data corruptions. *Sustainable Computing: Informatics and Systems*, 19:277–290, 2018.

37. A. Moawad, P. Balaprakash, A. Rousseau, and S. M. Wild. Novel large scale simulation process to support DOT's CAFE modeling system. *International Journal of Automotive Technology*, 17(6):1067–1077, 2016.

38. P. Balaprakash, M. Birattari, T. Stützle, and M. Dorigo. Estimation-based metaheuristics for the single vehicle routing problem with stochastic demands and customers. *Computational Optimization and Applications*, 61(2):463–487, 2015.

39. M. Birattari, Z. Yuan, P. Balaprakash, and T. Stützle. F-Race and Iterated F-Race: An overview. In T. Bartz-Beielstein, M. Chiarandini, L. Paquete, and M. Preuss, editors, *Experimental Methods for the Analysis of Optimization Algorithms*, pages 311–336. Springer Berlin Heidelberg, 2010.

40. P. Balaprakash, M. Birattari, T. Stützle, and M. Dorigo. Estimation-based metaheuristics for the probabilistic traveling salesman problem. *Computers & Operations Research*, 37(11):1939–1951, 2010.

41. P. Balaprakash, M. Birattari, T. Stützle, Z. Yuan, and M. Dorigo. Estimation-based ant colony optimization and local search for the probabilistic traveling salesman problem. *Swarm Intelligence*, 3(3):223–242, 2009.

42. P. Balaprakash, M. Birattari, T. Stützle, and M. Dorigo. Adaptive sample size and importance sampling in estimation-based local search for the probabilistic traveling salesman problem. *European Journal of Operational Research*, 199(1):98–110, 2009.

43. M. Birattari, P. Balaprakash, T. Stützle, and M. Dorigo. Estimation-based local search for stochastic combinatorial optimization using delta evaluations: A case study on the probabilistic traveling salesman problem. *INFORMS Journal on Computing*, 20(4):644–658, 2008.

44. P. Balaprakash, M. Birattari, and T. Stützle. Engineering stochastic local search algorithms: A case study in estimation-based local search for the probabilistic travelling salesman problem. In C. Cotta and J. van Hemert, editors, *Recent Advances in Evolutionary Computation for Combinatorial Optimization*, volume 153 of *Studies in Computational Intelligence*, pages 53–66. Springer Berlin Heidelberg, 2008.

45. M. Birattari, P. Balaprakash, and M. Dorigo. The ACO/F-Race algorithm for combinatorial optimization under uncertainty. In K. F. Doerner, M. Gendreau, P. Greistorfer, W. Gutjahr, R. F. Hartl, and M. Reimann, editors, *Metaheuristics - Progress in Complex Systems Optimization*, volume 39 of *Operations Research/Computer Science Interfaces Series*, pages 189–203. Springer US, 2007.

46. D. L. Prakash, P. Balaprakash, and D. Regener. Computational microstructure analyzing technique for quantitative characterization of shrinkage and gas pores in pressure die cast az91 magnesium alloys. *Computational Materials Science*, 32(3–4):480–488, 2005.

Publications in Refereed Proceedings of Conference and Workshop

47. R. Egele, I. Guyon, V. Vishwanath, and P. Balaprakash “Asynchronous Decentralized Bayesian Optimization for Large Scale Hyperparameter Optimization”, In the *19th IEEE International Conference on e-Science*, 2023.

48. S. Khairy and P. Balaprakash “Multi-Fidelity Reinforcement Learning with Control Variates”, In *31th European Symposium on Artificial Neural Networks (ESANN 2023)*, 2023.

49. R. Romain, I. Guyon, Y. Sun, and P. Balaprakash “Is One Epoch All You Need For Multi-Fidelity Hyperparameter Optimization?”, In *31th European Symposium on Artificial Neural Networks (ESANN 2023)*, 2023.

50. T. Randall, J. Koo, B. Videau, M. Kruse, X. Wu, P. Hovland, M. Hall, R. Ge, and P. Balaprakash "Transfer-learning-based Autotuning using Gaussian Copula", In *Proceedings of the 37th International Conference on Supercomputing*, pp. 37–49, 2023.

51. X. Wu, P. Balaprakash, M. Kruse, J. Koo, B. Videau, P. Hovland, V. Taylor, B. Geltz, S. Jana, and M. Hall "ytopt: Autotuning Scientific Applications for Energy Efficiency at Large Scales", In *2023 Cray User Group Conference*, 2023.

52. W. Godoy, P. Valero-Lara, K. Teranishi, and P. Balaprakash and J. Vetter "Evaluation of OpenAI Codex for HPC Parallel Programming Models Kernel Generation", In *Sixteenth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), 2023*, 2023.

53. Y. Orcun, H. Chan, K. Raghavan, W. Judge, M. Cherukara, P. Balaprakash, S. Sankaranarayanan. Continual Learning-Based Automated Defect Identification in Coherent Diffraction Imaging, In *Workshop on Artificial Intelligence and Machine Learning for Scientific Applications held in conjunction with the International Conference for High Performance Computing, Networking*, 2022.

54. H. Jin, K. Raghavan, G. Papadimitriou, C. Wang, A. Mandal, P. Krawczuk, L. Pottier, M. Kiran, E. Deelman, P. Balaprakash. Workflow Anomaly Detection with Graph Neural Networks In *WORKS – 17th Workshop on Workflows in Support of Large Scale Science held in conjunction with the International Conference for High Performance Computing, Networking*, 2022.

55. R. Egele, R. Maulik, K. Raghavan, P. Balaprakash, and B. Lusch. AutoDEUQ: Automated deep ensemble with uncertainty quantification. In *2022 26th International Conference on Pattern Recognition (ICPR)*, 2022.

56. K. Raghavan and P. Balaprakash. Meta continual learning via dynamic programming. In *2022 26th International Conference on Pattern Recognition (ICPR)*, 2022.

57. M. Isakov, E. Rosario, S. Madireddy, P. Balaprakash, P. Carns, R. Ross, and M. Kinsy. A taxonomy of error sources in HPC I/O machine learning models. In *SC '22: IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis*, 2022.

58. S. Khairy, P. Balaprakash, L. X. Cai, and Y. Cheng. Constrained deep reinforcement learning for energy sustainable multi-UAV based random access IoT networks with NOMA. *IEEE Journal on Selected Areas in Communications*, 39(4):1101–1115, 2021.

59. J. Koo, P. Balaprakash, M. Kruse, X. Wu, P. Hovland, and M. Hall. Customized Monte Carlo tree search for LLVM/Polly's composable loop optimization transformations. In *2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, pages 82–93. IEEE, 2021.

60. A. Yanguas-Gil, J. Koo, S. Madireddy, P. Balaprakash, J. W. Elam, and A. U. Mane. Neuromorphic architectures for edge computing under extreme environments. In *2021 IEEE Space Computing Conference (SCC)*, pages 39–45. IEEE, 2021.

61. Y. Liu, R. Hu, and P. Balaprakash. Uncertainty quantification of deep neural network-based turbulence model for reactor transient analysis. In *Verification and Validation*, volume 84782, page V001T11A001. American Society of Mechanical Engineers, 2021.

62. T. Mallick, P. Balaprakash, E. Rask, and J. Macfarlane. Transfer learning with graph neural networks for short-term highway traffic forecasting. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 10367–10374, 2020.

63. T. Mallick, M. Kiran, B. Mohammed, and P. Balaprakash. Dynamic graph neural network for traffic forecasting in wide area networks. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 1–10, 2020.

64. M. Isakov, E. Rosario, S. Madireddy, P. Balaprakash, P. Carns, R. Ross, and M. Kinsy. HPC I/O throughput bottleneck analysis with explainable local models. In *SC '20: IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis*, 2020.

65. Y. He, P. Balaprakash, and Y. Li. FIdelity: Efficient resilience analysis framework for deep learning accelerators. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 270–281, 2020.

66. S. Jiang and P. Balaprakash. Graph neural network architecture search for molecular property prediction. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 1346–1353, 2020.

67. N. A. Garland, R. Maulik, Q. Tang, X.-Z. Tang, and P. Balaprakash. Progress towards high fidelity collisional-radiative model surrogates for rapid in-situ evaluation. In *3rd Workshop on Machine Learning and the Physical Sciences*, 2020.

68. X. Wu, M. Kruse, P. Balaprakash, H. Finkel, P. Hovland, V. Taylor, and M. Hall. Autotuning polybench benchmarks with LLVM clang/polly loop optimization pragmas using bayesian optimization. In *2020 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, pages 61–70, 2020.

69. E. Del Rosario, M. Currier, M. Isakov, S. Madireddy, P. Balaprakash, P. Carns, R. B. Ross, K. Harms, S. Snyder, and M. A. Kinsy. Gauge: An interactive data-driven visualization tool for HPC application I/O performance analysis. In *2020 IEEE/ACM Fifth International Parallel Data Systems Workshop (PDSW)*, pages 15–21, 2020.

70. M. Isakov, E. Del Rosario, S. Madireddy, P. Balaprakash, P. Carns, R. B. Ross, and M. A. Kinsy. Toward generalizable models of I/O throughput. In *2020 IEEE/ACM International Workshop on Runtime and Operating Systems for Supercomputers (ROSS)*, pages 41–49, 2020.

71. S. Madireddy, A. Yanguas-Gil, and P. Balaprakash. Multilayer neuromodulated architectures for memory-constrained online continual learning. In *ICML Workshop on LifelongML*, 2020.

72. C. Kim, K. Kim, P. Balaprakash, and M. Anitescu. Graph convolutional neural networks for optimal load shedding under line contingency. In *IEEE Power & Energy Society General Meeting (PESGM)*, 2019.

73. S. Madireddy, P. Balaprakash, P. Carns, R. Latham, G. K. Lockwood, R. Ross, S. Snyder, and S. M. Wild. Adaptive learning for concept drift in application performance modeling. In *Proceedings of the 48th International Conference on Parallel Processing*, ICPP 2019, pages 79:1–79:11, 2019.

74. P. Balaprakash, R. Egele, M. Salim, S. Wild, V. Vishwanath, F. Xia, T. Brettin, and R. Stevens. Scalable reinforcement-learning-based neural architecture search for cancer deep learning research. In *SC '19: IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis*, 2019.

75. S. Lee, Q. Kang, S. Madireddy, P. Balaprakash, A. Agrawal, A. Choudhary, R. Archibald, and W.-K. Liao. Improving scalability of parallel cnn training by adjusting mini-batch size at run-time. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 830–839, 2019.

76. S. Madireddy, N. Li, N. Ramachandra, P. Balaprakash, and S. Habib. Modular deep learning analysis of galaxy-scale strong lensing images. In *NeurIPS Workshop on ML and the Physical Sciences*, 2019.

77. S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash. Reinforcement-learning-based variational quantum circuits optimization for combinatorial problems. In *NeurIPS Workshop on ML and the Physical Sciences*, 2019.

78. R. Maulik and P. Balaprakash. Site-specific graph neural network for predicting protonation energy of oxygenate molecules. In *NeurIPS Workshop on ML and the Physical Sciences*, 2019.

79. R. Maulik, V. Rao, S. Madireddy, B. Lusch, and P. Balaprakash. Using recurrent neural networks for nonlinear component computation in advection-dominated reduced-order models. In *NeurIPS Workshop on ML and the Physical Sciences*, 2019.

80. P. Jiang, H. Doan, S. Madireddy, R. S. Assary, and P. Balaprakash. Value-added chemical discovery using reinforcement learning. In *NeurIPS Workshop on ML and the Physical Sciences*, 2019.

81. S. M. Aithal and P. Balaprakash. MaLTESE: Large-scale simulation-driven machine learning for transient driving cycles. In M. Weiland, G. Juckeland, C. Trinitis, and P. Sadayappan, editors, *High Performance Computing*, pages 186–205, Cham, 2019.

82. V. Sreenivasan, R. Javali, M. Hall, P. Balaprakash, T. R. W. Scogland, and B. R. de Supinski. A framework for enabling OpenMP autotuning. In X. Fan, B. R. de Supinski, O. Sinnen, and N. Giacaman, editors, *OpenMP: Conquering the Full Hardware Spectrum*, pages 50–60, Cham, 2019. Springer International Publishing.

83. N. Wycoff, P. Balaprakash, and F. Xia. Neuromorphic acceleration for approximate Bayesian inference on neural networks via permanent dropout. In *ICONS19: International Conference on Neuromorphic Systems*, 2019.

84. S. Madireddy, A. Yanguas-Gil, and P. Balaprakash. Neuromorphic architecture optimization for task-specific dynamic learning. In *ICONS19: International Conference on Neuromorphic Systems*, 2019.

85. S. Madireddy, P. Balaprakash, P. Carns, R. Latham, R. Ross, S. Snyder, and S. M. Wild. Machine learning based parallel I/O predictive modeling: A case study on Lustre file systems. In *High Performance Computing*, pages 184–204, 2018.

86. Z. Liu, R. Kettimuthu, P. Balaprakash, and I. Foster. Building a wide-area data transfer performance predictor: An empirical study. In *MLN 2018: 1st International Conference on Machine Learning for Networking*, 2018.

87. M. Salim, T. Uram, J. Childers, P. Balaprakash, V. Vishwanath, and M. Papka. Balsam: Automated scheduling and execution of dynamic, data-intensive HPC workflows. In *Python for High-Performance and Scientific Computing Workshop (held in conjunction with SC18)*, 2018.

88. P. Balaprakash, M. Salim, T. D. Uram, V. Vishwanath, and S. M. Wild. DeepHyper: Asynchronous hyperparameter search for deep neural networks. In *2018 IEEE 25th International Conference on High Performance Computing (HiPC)*, pages 42–51, 2018.

89. P. Malakar, P. Balaprakash, V. Vishwanath, V. Morozov, and K. Kumaran. Benchmarking machine learning methods for performance modeling of scientific applications. In *PMBS 2018: Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (held in conjunction with SC18)*, 2018.

90. S. Lee, A. Agrawal, P. Balaprakash, A. Choudhary, and W. Liao. Communication-efficient parallelization strategy for deep convolutional neural network training. In *Machine Learning in HPC Workshop (held in conjunction with SC18)*, 2018.

91. S. Madireddy, P. Balaprakash, P. Carns, R. Latham, R. Ross, S. Snyder, and S. M. Wild. Modeling I/O performance variability using conditional variational autoencoders. In *2018 IEEE International Conference on Cluster Computing (CLUSTER)*, 2018.

92. Z. Liu, P. Balaprakash, R. Kettimuthu, and I. Foster. Explaining wide area data transfer performance. In *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing*, HPDC '17, pages 167–178, 2017.

93. S. Chunduri, P. Balaprakash, V. Morozov, V. Vishwanath, and K. Kumaran. Analytical performance modeling and validation of Intel's Xeon Phi architecture. In *Proceedings of the Computing Frontiers Conference*, CF'17, pages 247–250, 2017.

94. S. Madireddy, P. Balaprakash, P. Carns, R. Latham, R. Ross, S. Snyder, and S. M. Wild. Analysis and correlation of application I/O performance and system-wide I/O activity. In *International Conference on Networking, Architecture, and Storage (NAS)*, pages 1–10, 2017.

95. O. Subasi, S. Di, P. Balaprakash, O. Unsal, J. Labarta, A. Cristal, S. Krishnamoorthy, and F. Cappello. MACROD: Online adaptive machine learning framework for silent error detection. In *3rd Workshop on Fault Tolerance Systems (FTS'17)*, 2017.

96. O. Subasi, S. Di, L. Bautista-Gomez, P. Balaprakash, O. Unsal, J. Labarta, A. Cristal, and F. Cappello. Spatial support vector regression to detect silent errors in the exascale era. In *16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, pages 413–424, 2016.

97. P. Balaprakash, A. Tiwari, S. M. Wild, and P. D. Hovland. AutoMOMML: Automatic multi-objective modeling with machine learning. In M. J. Kunkel, P. Balaji, and J. Dongarra, editors, *31st International Conference on ISC High Performance (ISC-HPC)*, pages 219–239, 2016.

98. P. Balaprakash, V. Morozov, R. Kettimuthu, K. Kumaran, and I. Foster. Improving data transfer throughput with direct search optimization. In *45th International Conference on Parallel Processing (ICPP)*, pages 248–257, 2016.

99. A. Roy, P. Balaprakash, P. D. Hovland, and S. M. Wild. Exploiting performance portability in search algorithms for autotuning. In *IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 1535–1544, 2016.

100. F. Isaila, P. Balaprakash, S. M. Wild, D. Kimpe, R. Latham, R. Ross, and P. D. Hovland. Collective I/O tuning using analytical and machine learning models. In *IEEE International Conference on Cluster Computing (CLUSTER)*, pages 128–137, 2015.

101. T. Nelson, A. Rivera, P. Balaprakash, M. Hall, P. D. Hovland, E. Jessup, and B. Norris. Generating efficient tensor contractions for GPUs. In *2015 44th International Conference on Parallel Processing*, pages 969–978, 2015.

102. P. Balaprakash, L. A. B. Gomez, M. S. Bouguerra, S. M. Wild, F. Cappello, and P. D. Hovland. Analysis of the tradeoffs between energy and run time for multilevel checkpointing. In S. A. Jarvis, S. A. Wright, and S. D. Hammond, editors, *High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation – PMBS 2014*, volume 8966 of *Lecture Notes in Computer Science*, pages 249–263, 2015.

103. P. Balaprakash, Y. Alexeev, S. A. Mickelson, S. Leyffer, R. Jacob, and A. Craig. Machine-learning-based load balancing for community ice code component in CESM. In M. Daydé, O. Marques, and K. Nakajima, editors, *High Performance Computing for Computational Science – VECPAR 2014, Revised Selected Papers*, volume 8969 of *Lecture Notes in Computer Science*, pages 79–91, 2015.

104. A. Mametjanov, P. Balaprakash, C. Choudary, P. D. Hovland, S. M. Wild, and G. Sabin. Autotuning FPGA design parameters for performance and power. In *2015 IEEE 23rd Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 84–91, 2015.

105. P. Balaprakash, A. Tiwari, and S. M. Wild. Multi objective optimization of HPC kernels for performance, power, and energy. In S. A. Jarvis, S. A. Wright, and S. D. Hammond, editors, *High Performance Computing Systems. Performance Modeling, Benchmarking and Simulation – PMBS 2013*, Lecture Notes in Computer Science, pages 239–260, 2014.

106. P. Balaprakash, R. B. Gramacy, and S. M. Wild. Active-learning-based surrogate models for empirical performance tuning. In *IEEE International Conference on Cluster Computing (CLUSTER)*, pages 1–8, 2013.

107. P. Balaprakash, D. Buntinas, A. Chan, A. Guha, R. Gupta, S. H. K. Narayanan, A. A. Chien, P. Hovland, and B. Norris. Exascale workload characterization and architecture implications. In *Proceedings of the High Performance Computing Symposium*, HPC '13, pages 5:1–5:8, 2013. (Best Paper).

108. P. Balaprakash, K. Rupp, A. Mametjanov, R. B. Gramacy, P. D. Hovland, and S. M. Wild. Empirical performance modeling of GPU kernels using active learning. In *Parallel Computing: Accelerating Computational Science and Engineering (ParCo2013)*, Advances in Parallel Computing, pages 646–655, 2013.

109. P. Balaprakash, S. M. Wild, and P. D. Hovland. An experimental study of global and local search algorithms in empirical performance tuning. In *High Performance Computing for Computational Science - VECPAR 2012, 10th International Conference, Revised Selected Papers*, Lecture Notes in Computer Science, pages 261–269, 2013.

110. P. Balaprakash, S. M. Wild, and B. Norris. SPAPT: Search problems in automatic performance tuning. In *Proceedings of the International Conference on Computational Science, ICCS 2012*, volume 9, pages 1959–1968, 2012.

111. P. Balaprakash, S. M. Wild, and P. D. Hovland. Can search algorithms save large-scale automatic performance tuning? In *Proceedings of the International Conference on Computational Science, ICCS 2011*, volume 4, pages 2136–2145, 2011.

112. M. Birattari, Z. Yuan, P. Balaprakash, and T. Stützle. Automated algorithm tuning using F-Races: Recent developments. In S. Voss and M. Caserta, editors, *MIC 2009: The 8th Metaheuristics International Conference*, 2009.

113. Z. Yuan, A. Fügenschuh, H. Homfeld, P. Balaprakash, T. Stützle, and M. Schoch. Iterated greedy algorithms for a real-world cyclic train scheduling problem. In M. J. Blesa, C. Blum, C. Cotta, A. Fernández, J. Gallardo, A. Roli, and M. Sampels, editors, *Hybrid Metaheuristics*, volume 5296 of *Lecture Notes in Computer Science*, pages 102–116. 2008.

114. P. Balaprakash, M. Birattari, and T. Stützle. Improvement strategies for the F-Race algorithm: Sampling design and iterative refinement. In T. Bartz-Beielstein, M. Blesa Aguilera, C. Blum, B. Naujoks, A. Roli, G. Rudolph, and M. Sampels, editors, *Hybrid Metaheuristics*, volume 4771 of *Lecture Notes in Computer Science*, pages 108–122, 2007.

115. P. Balaprakash, M. Birattari, T. Stützle, and M. Dorigo. Incremental local search in ant colony optimization: Why it fails for the quadratic assignment problem. In M. Dorigo, L. Gambardella, M. Birattari, A. Martinoli, R. Poli, and T. Stützle, editors, *Ant Colony Optimization and Swarm Intelligence*, volume 4150 of *Lecture Notes in Computer Science*, pages 156–166. Springer Berlin Heidelberg, 2006.

Other Publications

Peer-reviewed extended abstracts

116. S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash. Reinforcement learning for quantum approximate optimization. In *High Performance Computing, Networking, Storage and Analysis (SCC)*, 2019. Poster Extended Abstract.

117. P. Balaprakash, V. Morozov, and R. Kettimuthu. Improving throughput by dynamically adapting concurrency of data transfer. In *High Performance Computing, Networking, Storage and Analysis (SCC)*, 2015. Poster Extended Abstract.

118. Y. Alexeev and P. Balaprakash. Heuristic dynamic load balancing algorithm applied to the fragment molecular orbital method. In *High Performance Computing, Networking, Storage and Analysis (SCC)*, 2015. Poster Extended Abstract.

119. L. A. Gomez, P. Balaprakash, M.-S. Bouguerra, S. M. Wild, F. Cappello, and P. D. Hovland. Energy-performance tradeoffs in multilevel checkpoint strategies. In *IEEE International Conference on Cluster Computing (CLUSTER)*, pages 278–279, 2014. Poster Extended Abstract.

120. Y. Zhang, P. Balaprakash, J. Meng, V. Morozov, S. Parker, and K. Kumaran. Raexplore: Enabling rapid, automated architecture exploration for full applications. In *High Performance Computing, Networking, Storage and Analysis (SCC)*, 2014. Poster Extended Abstract.

121. P. Balaprakash, D. Buntinas, A. Chan, A. Guha, R. Gupta, S. Narayanan, A. Chien, P. Hovland, and B. Norris. Exascale workload characterization and architecture implications. In *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 120–121, 2013. Poster Extended Abstract.

122. P. Balaprakash, A. Tiwari, and S. M. Wild. Framework for optimizing power, energy, and performance. In *High Performance Computing, Networking, Storage and Analysis (SCC)*, 2013. Poster Extended Abstract. Finalist for the best poster.

123. P. Balaprakash, D. Buntinas, A. Chan, A. Guha, R. Gupta, S. Narayanan, A. Chien, P. Hovland, and B. Norris. An exascale workload study. In *High Performance Computing, Networking, Storage and Analysis (SCC)*, pages 1463–1464, 2012. Poster Extended Abstract.

124. P. Balaprakash, M. Birattari, T. Stutzle, and M. Dorigo. Effective estimation-based stochastic local search algorithms for stochastic routing problems. In *Proceedings of ORBEL 24, 24th Annual Conference of the Belgian Operations Research Society*, pages 136–137, 2010. Extended Abstract.

125. P. Balaprakash, M. Birattari, T. Stützle, and M. Dorigo. Applications of estimation-based SLS algorithms to the stochastic routing problems. In P. Hansen and S. Voss, editors, *Metaheuristics 2008, Second International Workshop on Model Based Metaheuristics*, 2008. Extended Abstract.

126. P. Balaprakash, M. Birattari, T. Stutzle, and M. Dorigo. Estimation-based stochastic local search algorithms for the stochastic routing problems. In E.-G. Talbi and K. Mellouli, editors, *International Conference on Metaheuristics and Nature Inspired Computing, META'08*, 2008. Extended Abstract.

White papers

127. P. Balaprakash, R. Egele, E. Deelman, M. Kiran, and A. Mandal. Graph neural network for anomalies detection in scientific workflows. *ModSim 2022 Workshop on Modeling & Simulation of Systems and Applications*, 2022.

128. P. Balaprakash, R. Egele, E. Deelman, M. Kiran, and A. Mandal. MLops: The science of AI/ML software development. *ASCR Workshop on the Science of Scientific-Software Development and Use*, 2021.

129. S. Madireddy and P. Balaprakash. An information-theoretic view of learnable privacy-utility trade-off for scientific data. *ASCR Workshop on Cybersecurity and Privacy for Scientific Computing Ecosystems*, 2021.

130. E. Deelman, P. Balaprakash, M. Kiran, and A. Mandal. Adversarial modeling, simulation, and learning for trustworthy scientific computing ecosystems. *ASCR Workshop on Cybersecurity and Privacy for Scientific Computing Ecosystems*, 2021.

131. P. Balaprakash, K. Raghavan, and R. Yousefzadeh. Federated neural architecture search for privacy-preserving AI/ML. *ASCR Workshop on Cybersecurity and Privacy for Scientific Computing Ecosystems*, 2021.

132. T. Mallick, P. Balaprakash, D. Verner, and L. Levy. Analyzing the impact of Covid-19 on critical infrastructure using natural language processing. *AI@DOE Roundtable*, 2021.

133. P. Balaprakash, M. Emani, V. Vishwanath, R. Ross, and S. Wild. An AI system for AI codesign. *ASCR Workshop on Reimagining Codesign*, 2021.

134. S. Madireddy, A. Yanguas-Gil, and P. Balaprakash. Application-driven codesign for online and continual learning using neuromorphic architectures. *ASCR Workshop on Reimagining Codesign*, 2021.

135. A. Yanguas-Gil, S. Madireddy, and P. Balaprakash. Codesign approaches to enable and optimize computing at the extreme edge. *ASCR Workshop on Reimagining Codesign*, 2021.

136. P. Balaprakash, S. Collis, Y. Kim, P. Beckmann, M. Cadeddu, M. Gonzalez-Meler, R. Sullivan, S. Madireddy, and R. Kotamarthi. AI-enabled modex and edge-computing over 5G for improving the predictability of water cycle extremes. *Artificial Intelligence for Earth System Predictability Workshop (AI4ESP)*, 2021.

137. Y. Feng, R. Maulik, J. Wang, P. Balaprakash, W. Huang, V. Rao, P. Xue, W. Pringle, J. Bessac, and R. Sullivan. Characterization of extremes and compound impacts: Applications of machine learning and interpretable neural networks. *Artificial Intelligence for Earth System Predictability Workshop (AI4ESP)*, 2021.

138. J. Wang, R. Kotamarthi, V. Ghate, B. Lusch, P. Balaprakash, J. M. Wozniak, X. Yuan, W. Pringle, P. Xue, J. Bessac, and W. Chang. A hybrid climate modeling system using AI-assisted process emulators. *Artificial Intelligence for Earth System Predictability Workshop (AI4ESP)*, 2021.

139. S. Khairy and P. Balaprakash. Challenges and opportunities for AI-enabled science applications over 5G. In *5G Enabled Energy Innovation Workshop (5GEEIW)*, March 2020.

140. S. Khairy and P. Balaprakash. Edge intelligence meets cloud intelligence over 5G: Unmanned aerial vehicle swarm for extreme environments. In *5G Enabled Energy Innovation Workshop (5GEEIW)*, March 2020.

141. P. Balaprakash, J. Larson, V. Vishwanath, and S. Wild. Derivative-free mixed-integer optimization for automated predictive modeling using machine learning. In *SciML 2018: DOE ASCR Scientific Machine Learning Workshop*, 2018.

142. T. Munson and P. Balaprakash Dynamic adversarial games in complex systems and machine learning. 2017 DOE ASCR Applied Mathematics Meeting White Paper, 2017.

143. P. Balaprakash, V. Morozov, S. M. Wild, V. Vishwanath, P. D. Hovland, K. Kumaran, and B. Allcock. Machine learning for self-adaptive leadership-class machines. DOE ASCR Machine Learning Workshop, 2014.

Sponsor Reports

144. J. Carter, J. Feddema, D. Kothe, R. Neely, J. Pruet, R. Stevens, P. Balaprakash, P. Beckman, I. Foster, K. Iskra, A. Ramanathan, V. Taylor, R. Thakur, D. Agarwal, S. Crivelli, B. de Jong, D. Rouson, M. Sohn, M. Wetter, S. Wild, T. Bremer, M. Goldman, A. Kupresanin, L. Peterson, B. Spears, D. Stevens, B. Van Essen, R. Bent, M. Grosskopf, E. Lawrence, G. Shipman, K. Rose, R. Grout, N. Kouakpaizan, F. Omitaomu, S. Peles, P. Ramuhalli, A. Shankar, D. Womble, G. Zhang, T. Catanach, R. Oldfield, S. Rajamanickam, J. Ray, M. A. Leung, "Advanced Research Directions on AI for Science, Energy, and Security: Report on Summer 2022 Workshops", https://www.osti.gov/biblio/1986455, United States, DOI: 10.2172/1986455, 2023.

145. N. L. Hickmon, C. Varadharajan, F. M. Hoffman, S. Collis, and H. M. Wainwright. Artificial intelligence for earth system predictability (AI4ESP), DOE workshop report, 2022. (Co-authored two chapters: Data Acquisition to Distribution and Hybrid Modeling)

146. D. Verner, L. Levy, T. Mallick, P. Balaprakash, J. Hutchinson, and J. Bergerson. Community and infrastructure assessment for Covid. Technical report for Department of Homeland Security (DHS), Argonne National Lab, Lemont, IL, 2021. (Available upon request)

147. R. Stevens, V. Taylor, J. Nichols, B. Maccabe, K. Yelick, and D. Brown. AI for Science, DOE Workshop report, 2020. (Chapter co-lead, AI for Computer Science).

148. P. Beckman, C. Catlett, M. Ahmed, M. Alawad, L. Bai, P. Balaprakash, et al.. 5G Enabled Energy Innovation, DOE Workshop Report, 2020.

149. I. Foster, T. Lehman, N. Rao, B. Lyles, P. Balaprakash, K. Perumalla, S. Prowell, and R. Vatsavi. Towards new generation intelligent networking infrastructure for distributed science environments. DOE Workshop Report, 2017.

150. M. Berry, T. E. Potok, P. Balaprakash, H. Hoffmann, R. Vatsavai, and Prabhat. Machine Learning and Understanding for Intelligent Extreme Scale Scientific Computing and Discovery. DOE Workshop Report, 2015.

151. A. Moawad, S. Halbach, S. Pagerit, A. Rousseau, P. Balaprakash, and S. Wild. Novel process to use vehicle simulations directly as inputs to DOTs CAFE modeling system. Technical Report ANL/ESD-13/13, Report to Department of Transportation, 2014.

Theses

152. P. Balaprakash. *Estimation-based metaheuristics for stochastic combinatorial optimization: Case studies in stochastic routing problems*. Ph.D. thesis, Université libre de Bruxelles, Brussels, Belgium, 2010.

153. P. Balaprakash. Ant colony optimization under uncertainty. Master's thesis, Université libre de Bruxelles, Brussels, Belgium, 2005.

154. P. Balaprakash. Pre-processing of stochastic Petri nets and an improved storage strategy for proxel based simulation. Master's thesis, Otto-von-Guericke-Universität Magdeburg, Germany, 2004.

Technical Reports

155. Y. Liu, R. Hu, D. Dai, P. Balaprakash, and A. Obabko. Machine learning assisted safety modeling and analysis of advanced reactors. Technical report, Argonne National Laboratory, Lemont, IL, 2021.

156. M. Liu, G. Dong, K. G. Felker, M. Otten, P. Balaprakash, W. Tang, and Y. Alexeev. Exploration of quantum machine learning and AI accelerators for fusion science. Technical report, Argonne National Laboratory, Lemont, IL, 2021.

157. A. Mametjanov, P. Balaprakash, C. Choudary, P. D. Hovland, S. M. Wild, G. Sabin, and G. Wolfe. Improving FPGA design parameter exploration: Timing, power, and area. 2017.

158. Y. Zhang, P. Balaprakash, J. Meng, V. Morozov, S. Parker, and K. Kumaran. Raexplore: Enabling rapid, automated architecture exploration for full applications. Technical Report ANL/ALCF/TM-14/2, Argonne National Laboratory, 2014.

159. N. Wycoff, P. Balaprakash, and F. Xia. Towards on-chip Nayesian neuromorphic learning. *arXiv preprint arXiv:2005.04165*, 2020.

Invited Seminars, Colloquia, and Keynotes

1. P. Balaprakash, "AI for Modeling, Optimizing, and Controlling Complex Systems", INL AI/ML Symposium 11.0 – AI/ML in Instrumentation, Control, and Automation, April 2023.

2. P. Balaprakash, "ORNL AI Initiative: Advancing Science, Engineering, National Security Through AI: A Promising Pathway to Breakthroughs", TVC National Summit, Huntsville, AL, May 2023.

3. P. Balaprakash, "Scalable Automated Design and Development of Deep Neural Networks for Scientific and Engineering Applications", Fifth IPDPS Workshop on Scalable Deep Learning over Parallel and Distributed Infrastructure (ScaDL), May 2023.
4. P. Balaprakash, "Scalable Automated Design and Development of Deep Neural Networks for Scientific and Engineering Applications", MMICCs Center Seminar, PNNL, May 2023.
5. P. Balaprakash, "Optimizing HPC Systems for Scientific Applications: Machine Learning Approaches to Performance Tuning and Anomaly Detection", iWAPT2023: The Eighteenth International Workshop on Automatic Performance Tuning, May 2023.
6. P. Balaprakash, "Democratizing Deep Learning Development with DeepHyper", Sustainable Horizons Institute CRLC Seminar, May 2023.
7. P. Balaprakash, "Securing the Future of AI: Understanding and Mitigating Threats to AI Data, Models, and Processes", NPIC&HMIT 2023 Panel Artificial Intelligence and Machine Learning Panel: Current and Future Challenges, July 2023.
8. P. Balaprakash, "Advancing Safe and Trustworthy AI for Science: Understanding and Mitigating Threats to AI Data, Models, and Processes", Monterey Data Conference, August 2023.
9. P. Balaprakash, "Advancing Science through AI: A Promising Pathway to Breakthroughs", Oppenheimer Science and Energy Leadership Program Fellows Visit, August 2023.
10. Extracting the impact of climate change from scientific literature using Snorkel-enabled NLP. The Future of Data-Centric AI, August 2022. (Invited talk)
11. Democratizing deep learning development with DeepHyper. Princeton-DataX Workshop Series, May 2022. (Invited talk)
12. Automated machine learning with DeepHyper. Center for Mathematics and Artificial Intelligence (CMAI) Colloquium, George Mason University, February 2022. (Invited talk)
13. Scientific domain-informed machine learning. Florida Georgia Louis Stokes Alliance for Minority Participation (FGLSAMP), Florida A&M University, November 2021. (Invited talk)
14. Democratizing Deep Learning Development with DeepHyper. Brookhaven National Laboratory, October 2021. (Invited talk)
15. Artificial intelligence and machine learning. AI for Sustainability Workshop, Northwestern University, April 2020. (Invited talk)
16. Enabling ML approaches to HPC systems operations. Workshop on Monitoring and Analsis for High Performance Computing Sytems Plus Applications (HPCMASPA), September 2020. (Keynote and invited panelist)
17. The future of HPC systems in the presence of AI. Smoky Mountains Computational Sciences and Engineering Conference (SMC), August 2020. (Keynote and invited panelist)
18. Artificial intelligence for science. Transportation Research Board Executive Committee A.I. Policy Session, January 2019. (Invited talk)
19. Artificial intelligence to accelerate discovery and development. Argonne Outloud Public Lecture, Argonne, May 2019. (Invited talk)

20. Machine-learning-based performance modeling and tuning for high-performance computing. HPCaML 2019: The First International Workshop on the Intersection of High Performance Computing and Machine Learning, Washington, DC, February 2019. (Keynote)
21. Machine-learning-based-search for automatic performance tuning. OMASE 2019: Optimization, Modeling, Analysis and Space Exploration Workshop, Washington, DC, February 2019. (Keynote)
22. Machine learning in high performance computing. 7th Greater Chicago Area Systems Research Workshop, University of Chicago, April 2018. (Invited talk)
23. Automatic multi-objective modeling with machine learning. Workshop on Optimization and Machine Learning (ACNTW 17), Northwestern University, May 2017. (Invited talk)
24. Artificial intelligence for transportation and mobility. Tech Hub, SAE World Congress 2017, Detroit, MI, May 2017. (Invited talk)

Talks and Posters Presented at Conferences and Workshops

25. Taming high-performance computing platform heterogeneity with machine learning. Sky Computing – Toward Efficient Computing on the Cloud Special, September 2022.
26. Graph neural network for anomalies detection in scientific workflows. ModSim 2002, August 2022.
27. DeepHyper/AutoDEUQ: Automated deep ensemble with uncertainty quantification. Robust and efficient probabilistic deep learning for scientific data and beyond. SIAM Conference on Uncertainty Quantification (UQ22), April 2022.
28. DeepHyper: Scalable neural architecture and hyperparameter search for deep neural networks. DOE Booth Demo, November 2021.
29. AI for HPC: Everything that can be automated should be automated. JLESC workshop, AI for HPC panel, October 2021.
30. Scientific domain-informed machine learning. San Diego State University and Sustainable Horizon Institute, March 2021.
31. Graph-partitioning-based diffusion convolution recurrent neural network for large-scale traffic forecasting. DOE SMART All-Hands meeting (Invited talk), August 2021.
32. Large-scale traffic forecasting via spatial-temporal learning. DOE Workshop on Computational Mission Needs for Clean Energy and Advanced Manufacturing, June 2021.
33. Accelerating deep learning for science with SambaNova. The 2nd International Workshop on Machine Learning Hardware (Plenery talk), Co-located with ISC 2021, June 2021.
34. Neuromodulated neural architectures with local error signals for memory-constrained online continual learning. APS/CNM Users Meeting, May 2021.
35. Ytopt/SuRF: Machine-learning-based search for autotuning. 2021 ECP Annual Meeting, April 2021.

36. In situ compression artifact removal in scientific data using deep transfer learning and experience replay. Impacts of Applied Mathematics and Computer Science on DOE Computational Science, SIAM Conference on Computational Science and Engineering, March 2021.
37. Machine-learning-based automatic performance tuning. Workshop on Program Synthesis for Scientific Computing, August 2020.
38. Automated machine learning for molecular chemistry. Artificial Intelligence for Water Workshop, Argonne, September 2020. (Invited talk)
39. Scientific domain-informed machine learning. Argonne Physics Division Colloquium, Argonne, May 2019.
40. Scalable reinforcement-learning-based neural architecture search for cancer research. SC '19: IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis, November 2019.
41. DeepHyper: Scalable asynchronous hyperparameter search for deep learning. Data Enabled Modeling and Discovery in Science and Engineering Minisymposium, SIAM Conference on Computational Science and Engineering, February 2019.
42. Scientific domain-informed machine learning. Indian Institute of Science, Computational and Data Science Colloquium, India, December 2018.
43. DeepHyper: Asynchronous hyperparameter search for deep neural networks. 25th IEEE International Conference on High Performance Computing, Data, and Analytics, Bengaluru, India, December 2018.
44. Scientific domain-informed machine learning. Advanced Photon Source Colloquium, Argonne, December 2018.
45. Benchmarking machine learning methods for performance modeling of scientific applications. PMBS 2018: Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems, November 2018.
46. Spatial-temporal deep learning and hyperparameter search for traffic prediction. Workshop on Large-Scale Computing for Transportation Studies, Maui, HI, November 2018.
47. Reproducibility, portability and interpretability of deep learning. Workshop on Deep Learning for Multimessenger Astrophysics: Real-time Discovery at Scale, October 2018.
48. Scientific domain-informed machine learning. Department Seminar Series, Mechanical and Industrial Engineering, University of Illinois at Chicago, October 2018.
49. Scientific domain-informed machine learning. ALCF Simulation, Data, and Learning Workshop, Argonne, IL, October 2018.
50. Machine learning hardware. SciDAC Fusion Machine-Learning Workshop 2018, Princeton, NJ, June 2018.
51. Need for data locality in machine/deep learning. Fourth Workshop on Programming Abstractions for Data Locality (PADAL'17), Chicago, IL, August 2017.
52. Generative adversarial networks. US ATLAS Workshop, Argonne, July 2017.

53. Generative adversarial networks. The 3rd International Workshop on Data Science in High Energy Physics (DS@HEP 2017), Fermi National Accelerator Laboratory, May 2017.

54. Analytical performance modeling and validation of intel's xeon phi architecture. Computing Frontiers 2017, University of Siena, Siena, Italy, May 2017.

55. Improving data transfer throughput with direct search optimization. The 45th International Conference on Parallel Processing, ICPP 2016, Philadelphia, PA, August 2016.

56. AutoMOMML: Automatic Multi-Objective Modeling with Machine Learning. International Supercomputing Conference, ISC 2016, Frankfurt, Germany, June 2016.

57. Exploiting performance portability in search algorithms for autotuning. 11^{th} International Workshop on Automatic Performance Tuning, iWAPT 2016, Chicago , IL, May 2016.

58. Improving throughput by dynamically adapting concurrency of data transfer. High Performance Computing, Networking, Storage and Analysis (SCC), 2015 SC, Austin, TX, November 2015.

59. Heuristic dynamic load balancing algorithm applied to the fragment molecular orbital method. High Performance Computing, Networking, Storage and Analysis (SCC), 2015 SC, Austin, TX, November 2015.

60. Self-aware runtime and operating systems. 2015 ASCR Machine Learning Workshop, Rockville, MD, January 2015.

61. Automatic performance modeling and tuning. Department of Mathematics, Statistics and Computer Science Colloquium, Marquette University, October 2014.

62. Machine-learning-based load balancing for community ice code component in CESM. 11th International Meeting on High-Performance Computing for Computational Science (VECPAR 2014), Eugene, OR, July 2014.

63. Multi objective optimization of HPC kernels for performance, power, and energy. 4th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS13), Denver, CO, November 2013.

64. Framework for optimizing power, energy, and performance. High Performance Computing, Networking, Storage and Analysis (SCC), 2013 SC, Denver, CO, November 2013.

65. Active-learning-based surrogate models for empirical performance tuning. 2013 IEEE International Conference on Cluster Computing (CLUSTER), Indianapolis, IN, September 2013.

66. Search algorithms in empirical performance tuning and machine learning for computationally expensive simulations. PMaC/SDSC Lab Seminar, San Diego Supercomputig Center, July 2013.

67. Exascale workload characterization and architecture implications. 2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), Austin, TX, April 2013.

68. SPAPT: Search Problems in Automatic Performance Tuning. Workshop on Tools for Program Development and Analysis in Computational Science, Omaha, NE, June 2012.

69. Efficient optimization algorithms for empirical performance tuning. SIAM Conference on Parallel Processing (SIAM PP 2012), Savannah, GA, February 2012.

70. Can search algorithms save large-scale automatic performance tuning? Workshop on Automatic Performance Tuning, Singapore, June 2011.

71. Comparison of search strategies in empirical performance tuning of linear algebra kernels. Mini symposium on Autotuning Linear Algebra Computations (SIAM CSE), Reno, NV, March 2011.

72. An experimental study of estimation-based metaheuristics for the probabilistic traveling salesman problem. LION 2007 II: Learning and Intelligent Optimization, Trento, Italy, December 2007.

73. Improvement strategies for the F-Race algorithm: Sampling design and iterative refinement. 4th International Workshop on Hybrid Metaheuristics, Dortmund, Germany, October 2007.

74. ACO/F-Race: Ant colony optimization and racing techniques for combinatorial optimization under uncertainty. MIC 2005: The 6th Metaheuristics International Conference, Vienna, Austria, August 2005.

Other Presentations

75. Democratizing deep learning development with deephyper. AppliedAI Meetup, June 2022.

76. Overview of Argonne AI R&D. French Delegation visit at Argonne, April 2022.

77. DeepHyper: Automated deep learning at scale. Argonne/MCS-ARAMCO Seminar, March 2022.

78. Deephyper: Automated machine learning at scale. Argonne-IIT Collaboration Seminar, March 2022.

79. Deephyper for automated machine learning in VASP. Equation of State and Materials Theory group, March 2022.

80. DeepHyper: Scalable automated machine learning package. Argonne CPS Division seminar series, August 2020.

81. Neuromorphic acceleration for uncertainty quantification in deep neural networks. MCS All Hands Meeting, Argonne, IL, January 2020.

82. Scalable reinforcement-learning-based neural architecture search for scientific and engineering applications. LANS Seminar, Argonne, IL, February 2019.

83. Scientific domain-informed machine learning. Consortium for Computational Physics and Chemistry Meeting, Argonne, November 2018.

84. Tools for scientific machine learning. Army Research Lab Workshop, May 2018.

    bibitemBal2016e Machine learning for high performance computing. Midwest Big Data Opportunities and Challenges, Chicago, IL, September 2016.

85. Multi objective optimization of HPC kernels for performance, power, and energy. 6th Joint Laboratory for Extreme Scale Computing (JLESC) Workshop, November, Japan, June 2016.

86. AutoMOMML: Automatic Multi-objective Modeling with Machine Learning. 5th Joint Laboratory for Extreme Scale Computing (JLESC) Workshop, Lyon, France, June 2016.

87. Predictive modeling for large scale vehicle simulations. Urban Data Analytics/City of Chicago SmartData Platform Workshop, Chicago, IL, May 2015.

88. Active-learning-based surrogate models for empirical performance tuning. Computation Institute Seminar, University of Chicago, Chicago, IL, February 2014.
89. Active-learning-based surrogate models for empirical performance tuning. 10th workshop of the INRIA-Illinois-ANL Joint Laboratory, Urbana-Champaign, IL, November 2013.
90. Search algorithms in empirical performance tuning and machine learning for computationally expensive simulations. ANL Mathematics and Computer Science Division Seminar, Lemont, IL, December 2012.

Tutorials Presented

91. Organizer and lead, ALCF DeepHyper Automated Machine Learning Workshop. ALCF, Argonne, July 2022. (Over 240+ registrants and 90+ online attendees)
92. Organizer and lead, Transfer learning and online tuning with PROTEAS-TUNE/ytopt. ECP AHM Tutorial, May 2022.
93. DeepHyper: A hyperparameter search package for deep neural networks. ALCF Simulation, Data, and Learning Workshop, Argonne, December 2020.
94. DeepHyper and hyperparameter optimization. ATPESC 2020: Argonne Training Program on Extreme-Scale Computing, Argonne, August 2020.
95. Reinforcement learning and applications. SciDAC TDS Meeting, March 2020.
96. AI tools. PSE AI in Science and Engineering Workshop, Argonne, September 2019.
97. Deep learning basics. ATPESC 2019: Argonne Training Program on Extreme-Scale Computing, St. Charles, August 2019.
98. Deep learning basics. Artificial Intelligence for Science Workshop, Argonne, August 2019.
99. Recurrent networks for time series data. Argonne Artificial Intelligence for Science Workshop, Argonne, August 2019.
100. Hyperparameter optimization and DeepHyper. ATPESC 2019: Argonne Training Program on Extreme-Scale Computing, St. Charles, August 2019.
101. Recurrent networks for time series data. Artificial Intelligence for Science Workshop for Summer Students, Argonne, July 2019.
102. Deep learning basics. Artificial Intelligence for Science Workshop for Summer Students, Argonne, July 2019.
103. Data parallel deep learning. CANDLE-EXALEARN ECP Workshop, Argonne, October 2018.
104. Data for machine learning. Argonne Geospatial Workshop, Argonne, August 2018.
105. Machine learning overview and applications. Argonne Environmental Science Division Retreat, Argonne, May 2018.
106. Deep learning basics. Argonne Deep Learning Workshop, Argonne, January 2018.
107. Recurrent neural networks. Argonne Deep Learning Workshop, Argonne, January 2018.

108. Automatic multi-objective modeling with machine learning. Argonne Training Program on Extreme-Scale Computing (ATPESC), St. Charles, IL, August 2017.
109. Overview of machine learning methods. Argonne Training Program on Extreme-Scale Computing (ATPESC), St. Charles, IL, August 2017.
110. Introduction to unsupervised and supervised learning in Python: Hands-on tutorial. Argonne Training Program on Extreme-Scale Computing (ATPESC), St. Charles, IL, August 2017.
111. Overview of machine learning methods. Argonne Machine Learning Workshop, Argonne, July 2017.
112. Automatic multi-objective modeling with machine learning. Argonne Machine Learning Workshop, Argonne, July 2017.
113. Introduction to unsupervised and supervised learning in python: Hands-on tutorial. Argonne Machine Learning Workshop, Argonne, July 2017.
114. Automatic multi-objective modeling with machine learning. Argonne Machine Learning Workshop, Argonne, April 2017.
115. Overview of machine learning methods. Argonne Machine Learning Workshop, Argonne, April 2017.
116. Introduction to unsupervised and supervised learning in python: Hands-on tutorial. Argonne Machine Learning Workshop, Argonne, April 2017.

Software Developed

1. P. Balaprakash, R. Egele, M. Salim, V. Vishwanath, and S. M. Wild. DeepHyper: Scalable Asynchronous Neural Architecture and Hyperparameter Search for Deep Neural Networks, 2018–2022. A software package that uses learning, optimization, and parallel computing to automate the design and development of machine learning (ML) models for scientific and engineering applications. https://github.com/deephyper/deephyper
2. P. Balaprakash, R. Egele, P. Hovland, X. Wu, J. Koo, and B. Videau. ytopt: Machine-learning-based-search for autotuning, 2018–2023. A software package that uses learning algorithms for sampling a small number of input parameter configurations, evaluating them, and progressively fitting a surrogate model over the input-output space until exhausting the user-defined time or the maximum number of evaluations. https://github.com/ytopt-team/ytopt
3. P. Balaprakash, A. Tiwari, S. M. Wild, L. Carrington, and P. D. Hovland. AutoMOMML: Automatic Multi-objective Modeling with Machine Learning, 2016. An end-to-end, machine-learning-based framework to build predictive models for objectives such as performance and power. It adopts statistical approaches to reduce the modeling complexity and automatically identifies and configures the most suitable learning algorithm to model the required objectives based on hardware and application signatures.
https://xgitlab.cels.anl.gov/pbalapra/automomml
4. P. Balaprakash, S. M. Wild, and B. Norris. SPAPT: Search Problems in Automatic Performance Tuning, 2011. A set of extensible and portable search problems in automatic performance tuning whose goal is to aid in the development and improvement of search strategies and performance-improving transformations.
https://github.com/brnorris03/Orio/tree/master/testsuite/SPAPT

5. M. L. Ibanez, J. D. Lacoste, T. Stützle, M. Birattari, E. Yuan, and P. Balaprakash. The irace Package: Iterated Race for Automatic Algorithm Configuration, 2010. A package implementing the iterated racing procedure to automatically configure optimization algorithms by finding the most appropriate settings given a set of instances of an optimization problem. http://cran.r-project.org/web/packages/irace

6. P. Balaprakash, M. Birattari, and T. Stützle. ELS-PTSP: Estimation-based Local Search for the Probabilistic Traveling Salesman Problem, 2009. A software package providing a high-performance implementation of the estimation-based iterative improvement algorithm to tackle the probabilistic traveling salesman problem. A key novelty estimated by partial evaluation, adaptive, and importance sampling. https://github.com/pbalapra/els-ptsp

Services

Services to Institution

1. Member, Argonne AI Editorial Board, 2020–Present
2. Argonne session lead, AI for Science and Security Workshop series, 2022
3. Member, ANL Climate and Energy Action LDRD review committee, 2021, 2022 (Reviewed 16 proposals in total)
4. Co-lead, Argonne's AI for Science communications strategy, 2021
5. Postdoc mentor: Jordon Fox (NP, 2022–Present), Bryce Fore (NP, 2022–Present), Robert Underwood (MCS, 2022–Present), Zhen Xie (ALCF, 2022–Present), Hieu A. Doan (MSD, 2019–2021), Peco Myint (XSD, 2020–Present), Naveen Dandu (MSD, 2019–2022), Alp Dener (MCS, 2018–2021), Valentin Reis (MCS, 2019–2020), Andreas Victor Copan (CSE, 2019–2020)
6. Mentor, Early Career Research Proposals, 2019, 2020
7. Member, Red team review for research proposals, 2018–Present
8. Member, Wilkinson Fellowship Committee, 2016, 2018, 2020
9. Member, DOE Argonne Data Science Program review committee, 2017, 2018, 2019, 2020, 2021. (Reviewed approximately three proposals every year)
10. Member, Argonne's AI-X Task Force for Materials and Chemistry, 2018
11. Member, working group for ML strategy in Argonne Materials Science Division, 2018
12. Member, library committee for MCS, ANL, 2013–2015
13. Vice President, Postdoctoral Society of ANL, January 2012–December 2012

Services to Sponsors

14. Reviewer, Express: 2022 Exploratory Research for Extreme-Scale Science, Explainable AI, 2022
15. Session lead, AI for Science and Security Workshop series, 2022
16. Reviewer, DOE ASCR Early Career Research Proposals (ECRP) Review Panel, 2020, 2021, 2022

17. Member, DOE Computational Science Graduate Fellowship (CSGF) application screening committee, 2020, 2021, 2022
18. AI/ML expert reviewer for DOE NP: SBIR/STTR Phase I, 2021, 2022
19. AI/ML expert reviewer for DOE BES: Data Science to Advance Chemical and Materials Sciences, 2021; Building EPSCoR-State/National Laboratory Partnerships, 2020, 2022; SBIR/STTR Phase I, 2017, 2019, 2020, 2022
20. DOE INCITE computational readiness review committee, 2016, 2017, 2018, 2019, 2020, 2021 (Reviewed two proposals every year)
21. AI/ML expert reviewer for DOE ASCR: Continuation of Solicitation for the Office of Science Financial Assistance Program, 2019, 2020; SBIR/STTR Phase I, 2021
22. Session lead, Artificial Intelligence for Earth System Predictability Workshop, 2021
23. DOE AI Town Hall: Argonne co-lead for software environment and research and writing co-lead for AI for computer science, 2019
24. Co-lead for applied machine learning and member of the workshop organizing committee, DOE ASCR Smart Network Systems Workshop, Rockville, MD, December 2016
25. Co-lead for self-aware runtime and operating systems, and member of the workshop organizing committee, DOE ASCR Machine Learning Workshop, Rockville, MD, January 2015

Services to Profession

26. Editorial Board, Frontiers in Systems Neuroscience, 2021–Present
27. Program Committee, 18th IEEE International Conference on eScience - eScience 2022
28. Program Committee, SIAM Conference on Parallel Processing for Scientific Computing (PP22)
29. Program Committee, NeurIPS 2021 Workshop on Machine Learning and the Physical Sciences
30. Ph.D. Thesis Committee, Albert Njoroge Kahira, Polytechnic University of Catalonia, 2021
31. Program Committee, ICPP: International Conference on Parallel Processing, 2021, 2022
32. Program Committee, IEEE/ACM International Conference on High Performance Computing, Networking, Storage, and Analysis (SC), 2020, 2021, 2022
33. Program Committee, ICONS: International Conference on Neuromorphic Systems, 2020, 2021, 2022
34. Program Committee, PMBS: International Workshop in Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems, 2015, 2016, 2017, 2018, 2019, 2020, 2021, 2022
35. Program Committee, tutorials, SC '20: IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis, 2020
36. Program Committee, Applications track, ICPP '20: 49th International Conference on Parallel Processing, 2020

37. Program Committee, AI4S20: Workshop on Artificial Intelligence and Machine Learning for Scientific Applications, 2020
38. Program Committee, ANTS 2016: Tenth International Conference on Swarm Intelligence, 2016
39. Program Committee, IPDRM 2016: First Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware, 2016
40. Referee services for Nature Reviews Physics, Information Sciences, Swarm Intelligence, IEEE Transactions on Cybernetics, Algorithms, Statistica Sinica, International Transactions in Operational Research, International Journal of High Performance Computing Applications, European Journal of Operational Research, Journal of Parallel and Distributed Computing, Journal of Supercomputing, IEEE Transactions on Neural Networks and Learning Systems, IEEE Transactions on Evolutionary Computation, Flexible Services and Manufacturing, INFORMS Journal of Computing, European Journal of Operations Research, Journal of Computers and Operations Research, and Journal of Heuristics.

Selected Diversity, Equity, and Inclusion Activities

1. Mentoring two high school students from the Illinois Mathematics and Science Academy under the Student Inquiry and Research program, September 2022–Present
2. Supervised more than 30 summer graduate students from 20 different universities. May 2014–Present
3. Supervised two students under Sustainable Research Pathways for High-Performance Computing (SRP-HPC) Program, 2022–Present
4. Mentored 8 postdocs under Argonne mentoring program, 2019–Present
5. Invited several women scientists for Argonne LANS seminar series, 2021–Present
6. Was speaker at Sustainable Horizons Institute (SHI) seminar series, 2020–Present
7. Sponsored Sami Khairy, a postdoc from underrepresented group, to Argonne's Walter Massey Fellowship, 2021
8. Was speaker at the Florida Georgia Louis Stokes Alliance for Minority Participation (FGLSAMP), Florida A&M University, November 2021
9. Was speaker at the San Diego State University (Hispanic-Serving Institution), March 2021
10. Mentored two high school girl students for the Afro-Academic, Cultural, Technological and Scientific Olympics (ACT-SO) High School Research Program, September 2016–September 2018. Both of them won gold medals both at local and national level)
11. Mentored ANL postdocs on preparing the elevator pitch for LDRD proposals, PSA LDRD Workshop and Networking Event, April 2015
12. Mentored ANL postdocs on transitioning from postdoc to Argonne staff and headed the topic discussion, Job Hunting: A Peer's Advice Is a Good Weapon, April 2014
13. Organized Argonne's LANS seminar series for applied math research group (approx. 50 members), MCS, Argonne. January 2012–June 2013. Three to four seminars per month; responsible for the speakers, advertisements, etc.