Prasanna Balaprakash

Director of AI Programs | Section Head, Data and AI Systems | 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 May 2025, I am serving as the Section Head of the Data and AI Systems, Computer Science and Mathematics Division, Computing and Computational Sciences Directorate at Oak Ridge National Laboratory (ORNL). My responsibilities include providing strategic leadership for a section of 45 researchers and engineers across five specialized groups (Learning Systems, Visualization, Workflow Systems, Data Engineering, and Performance Engineering); leading advanced R&D programs while collaborating with division leadership to define research vision and priorities; and directing recruitment initiatives, performance evaluations, and career development to maintain a diverse, high-caliber scientific workforce.

Since March 2023, I have served as the Director of Artificial Intelligence (AI) Programs at ORNL. I co-direct ORNL's AI Initiative—a multi-million-dollar, lab-directed research and development portfolio. The initiative focuses on the development of secure, trustworthy, and efficient AI solutions to tackle problems of national importance. I set the strategic vision for the AI for science portfolio, develop plans, and execute them. I seek to deliver foundational, scalable, and applied AI capabilities that support the broad mission of ORNL, providing world-class solutions in areas such as computer and computational science, materials science, energy, and national security.

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

Professional Experience

Section Head, Data and AI Systems Research May 2025–Present
Computer Science and Mathematics Division
Computing and Computational Sciences Directorate
Oak Ridge National Laboratory, Oak Ridge, TN

- Provide strategic and operational leadership for the Data and AI Systems Section, comprising 45 researchers across five groups specializing in AI/ML, visualization, data workflows, engineering, and performance systems. The section includes staff at all career stages, from early-career scientists to senior technical leaders.
- Oversee and guide five research groups:
  - **Learning Systems**: Scalable artificial intelligence, machine learning, and data analytics.
  - **Visualization**: Methods, tools, and technologies for visual data analysis.
  - **Workflow Systems**: Large-scale data management, scientific workflows, and compression.
  - **Data Engineering**: Data wrangling, architecture, and systems for scientific innovation.
  - **Performance Engineering**: Engineering secure and high-performance computing systems.
- Lead and contribute to advanced R&D programs within the section's technical domains.
- Collaborate with Division Director and Group Leaders to define vision, strategy, and priorities.
- Develop and track metrics to evaluate group performance and scientific impact.
- Shape divisional research directions and investments through team building and resource alignment.
- Conduct performance evaluations and provide feedback and support for staff development.
- Implement succession planning for critical leadership roles within the section.
- Lead recruitment efforts to attract a diverse and high-caliber scientific workforce.
- Drive career advancement, manage promotions, and handle personnel concerns as needed.
- Foster a culture of research excellence, integrity, and community engagement.

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

- Co-leading an 8-million-dollar, lab-directed research and development (LDRD) portfolio—the lab's largest cross-cut LDRD investment.
- Co-directing and managing 50+ researchers across five directorates, overseeing 10+ advanced AI research projects.
- Restructured the AI initiative to emphasize secure, trustworthy, and efficient AI for scientific discovery, experimental facilities, and national security.
- Set and execute the strategic vision and research thrusts, including the development and launch of LDRD funding calls.
- Organized town halls, proposal reviews, and presentations, ensuring rigorous monthly milestones and comprehensive mid-term/year-end evaluations.

- Collaborated with the LDRD office, deputy lab director, and associate lab directors to shape and advance the research portfolio.
- Established and maintained key partnerships with academia and industry, positioning ORNL as a leader in frontier AI.

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

Major Accomplishments as ORNL's Director of AI Programs from March 2023–Present

- **ORNL AI Initiative re-prioritization:** Directed the strategic restructuring of ORNL's AI Initiative to prioritize the development of secure, trustworthy, and energy-efficient AI. This restructuring focused on three core areas—scientific discovery and complex systems, experimental facilities, and national security—providing new research directions for the initiative. Actively collaborated with the LDRD office, deputy lab director, and associate lab directors to develop a new research portfolio on secure, trustworthy, and energy-efficient AI. Currently overseeing more than 10 advanced AI projects involving over 50 researchers across the lab.
  - Accelerating scientific discovery and strengthening national security with AI
  - Very transformative: Why researchers at Oak Ridge National Laboratory are studying AI, TV Interview, 10News Today, Feburary, 2025
  - Oak Ridge National Lab on the cutting edge of AI use and research, TV Interview, WVLT News, January, 2025

- **ORNL's success in the FY24 Department of Energy (DOE) Advanced Scientific Computing Research (ASCR) AI for science portfolio:** Under the DOE ASCR AI for Science portfolio, ORNL demonstrated unparalleled leadership by successfully securing six funded proposals—three as the lead institution and three as co-lead—more than any other organization. This unique achievement includes cross-thrust funding across all five domains of the funding opportunity announcement, underscoring ORNL's comprehensive expertise. Remarkably, five of the six funded proposals were supported by ORNL's AI Initiative. Overall, the AI initiative helped secure over $40 million in external funding, reinforcing ORNL's position at the forefront of AI-driven scientific innovation.
  - ORNL projects included in $67 million from DOE for AI in science research
- **ORNL's AI academy:** Launched the ORNL's AI academy that engages over 100 researchers at ORNL through a suite of educational programs, including the AI Tutorial Series for Science, AI for Science Bootcamps, AI Workshops, AI Expo, AI Seminar Series, and AI Clinics. These initiatives are designed to promote the adoption of secure, trustworthy, and energy-efficient AI methods, fostering a collaborative environment that enables researchers to stay at the forefront of AI innovation while addressing challenges in science, engineering, and national security.
  - ORNL's AI Initiative's workforce development
- **Industry partnerships:** Started industry partnerships with leaders such as AMD, NVIDIA, Microsoft, OpenAI, Google, and startups such as SambaNova, and Cerebras, driving significant advancements in AI model development at ORNL. These partnerships have advanced access to advanced hardware and expertise, enabling ORNL to accelerate the creation of secure, trustworthy, and energy-efficient AI models. By integrating new technologies from these industry leaders, ORNL continues to push the boundaries of innovation, solidifying its role as a pioneer in applying AI to tackle complex scientific and national security challenges.
  - Oak Ridge National Laboratory selects SambaNova to expand its research in secure and energy-efficient AI
  - Microsoft DeepSpeed collaboration to accelerate scientific discovery at ORNL
- **1,000 Scientists AI Jam:** Co-organized the 1,000 Scientists AI Jam Session (Feb. 28, 2025), uniting 1,400 researchers from nine national laboratories to evaluate OpenAI's and Anthropic's latest AI models for scientific applications. Hosted Energy Secretary Chris Wright, OpenAI's president Greg Brockman, Senator Bill Hagerty, and Representative Chuck Fleischmann, highlighting Oak Ridge National Laboratory's leadership in accelerating scientific discovery through AI and industry partnerships.
  - National labs, industry host 'jam session' to test AI models
- **White House roundtable and Capitol Hill briefings**: Invited participant at a White House roundtable on Academic AI Research and Development, co-organized by OSTP and NSF in 2023, discussing the strategic directions for AI development in academia. Delivered briefings to U.S. Senate and House staffers under the Task Force on American Innovation, illustrating the impact of federal investments in critical science and technology sectors on advancing AI, thereby contributing to national policy and AI strategy. In a panel for Senate staffers, hosted by the United States Coalition for the Advancement of Supercomputing (USCAS), NVIDIA and Hewlett Packard

Enterprise, provided insight into how supercomputing, AI and meteorology can work together to provide more accurate weather predictions, especially during instances of extreme weather.

  - AI director featured at White House roundtable, Capitol Hill briefings
  - Balaprakash discusses AI's role in advancing climate and weather research at a Senate panel

- **Advising the state of Tennessee on AI opportunities and risk mitigation:** Within Tennessee state AI Advisory Council, I am leading a subcommittee dedicated to evaluating the risks posed by AI models—particularly those emerging from China—and advising the state of Tennessee on strategic AI opportunities and threats. Recognizing the urgent security implications, I recommended an immediate ban on the DeepSeek app on all state-owned machines within a week of its release. Additionally, I played a key role in developing comprehensive guidelines to safeguard state operations from vulnerabilities such as jailbreaking, ensuring protection against potential AI-driven threats. This proactive approach has strengthened Tennessee's security posture while paving the way for responsible AI adoption.

  - Balaprakash chosen for Tennessee's new AI advisory council

- **First trillion-parameter AI model training capability in non-industry setting:** Directed the development of the first trillion-parameter AI model training capability on AMD-based Frontier supercomputer system. This work has garnered significant attention across the DOE and was highlighted in a White House briefing on large-scale AI model training capabilties at DOE.

  - Frontier trained a ChatGPT-sized large language model with only 3,000 of its 37,888 Radeon GPUs — the world's fastest supercomputer blasts through one trillion parameter model with only 8 percent of its MI250X GPUs

- **ORBIT: Oak Ridge Base Foundation Model for Earth System Predictability:** Directed the development of ORBIT, an advanced vision-transformer deep learning model that scales up to 113 billion parameters. As the largest model of its kind, ORBIT surpasses the previous climate AI foundation model size by a thousandfold. Performance scaling tests conducted on the Frontier supercomputer have demonstrated that ORBIT achieves 684 petaFLOPS to 1.6 exaFLOPS sustained throughput, across 49,152 AMD GPUs. This work earned the team a finalist nomination for the Association for Computing Machinery Gordon Bell Prize for Climate Modeling. The award honors innovations in applying high-performance computing to climate modeling applications. The team also recevied the HPCWire 2024 Editors' Choice Awards – Top Supercomputing Achievement.

  - Fine-tuning forecasts: ORBIT brings long-range weather prediction within reach
  - Oak Ridge National Laboratory receives honors in 2024 HPCwire Editors' Choice award

- **Graph foundation models for atomistic materials modeling:** Directed the development of first large scale graph foundation models (GFMs) for atomistic materials modeling on 5.2 TB data materials structure data. This work developed a series of optimizations that have allowed scaling up the GFMs training to tens of thousands of GPUs on datasets consisting of hundreds of millions of graphs. Using over 154 million atomistic structures for training, we illustrated the performance of our approach on two state-of-the-art DOE supercomputers, the Perlmutter petascale system at the National Energy Research Scientific Computing Center and the Frontier exascale system at Oak Ridge Leadership Computing Facility. We achieved near-linear strong scaling performance using more than 2,000 GPUs on Perlmutter and 16,000 GPUs on Frontier.

– Ensemble of pre-trained GFMs for atomistic materials modeling in Hugging Face

Supervision: Staffs, Postdocs, PhD students, Predocs, and Research Aides

1. Romain Egele, Computer Scientist 2018–Present
   Area: Automated Machine Learning, Scalable Deep Learning
2. Prachi Jadhav, University of Tennessee, Knoxville 2024–Present
3. Linda Uncerboeck, University of Tennessee, Knoxville 2024–Present
4. Alexander Kiefer, University of Tennessee, Knoxville 2024–Present
5. Adarsha Balaji, Postdoctoral Researcher 2022–2023
   Area: Neuromorphic Computing, Hardware Software Co-design
6. Hongwei Jin, Postdoctoral Researcher 2022–2023
   Area: Graph Neural Networks, Numerical Optimization, Federated Learning
7. Yixuan Sun, Postdoctoral Researcher 2022–2023
   Area: Spatiotemporal Machine Learning, Automated Machine Learning
8. Romit Maulik, Assistant Computational Scientist 2021–2023
   Area: Physics-informed Machine Learning, Computational Fluid Dynamics
   Next position: Assistant Professor, the Pennsylvania State University
9. Krishnan Ragavan, Assistant Computational Mathematician 2019–2023
   Area: Continual Learning, Deep Learning Theory
10. Tanwi Mallick, Assistant Computer Science Specialist 2018–2023
    Area: Spatiotemporal Machine Learning, Natural Language Processing
11. Sandeep Madireddy, Assistant Computer Scientist 2016–2023
    Area: Probabilistic Machine Learning, Neuromorphic Computing
12. Sami Khairy, Postdoctoral Researcher 2021–2022
    Area: Reinforcement Learning for Science
    Current position: Research Scientist, Microsoft, Canada
13. Jaehoon Koo, Postdoctoral Researcher 2020–2022
    Area: Physics-Informed Machine Learning for Fusion Science, Autotuning
    Next position: Assistant Professor, Hanyang University, South Korea
14. Grant Getzelman, Predoctoral Researcher, with Stefan Wild and Jeff Larson 2019–2021
    Area: Learning to Optimize
    Next position: Senior Machine Learning Scientist, Biotech Startup, U.S.
15. Joceran Gouneau, Research Aide 2021–2022
    Area: Machine Learning
    Next position: PhD Student at Sorbonne Université

Supervision: Summer Students

16. Aiden Hamade (University of Kentucky), Undergraduate summer internship, 2024
17. Timothy Lang (College of Southern Nevada: Homepage), NNSA Minority Serving Institutions Internship Program, 2024
18. Thomas Randall (Clemson U.), Givens Fellow (2020, 2022)/DOE Omni Technology Alliance Internship Program 2024
19. Ololade Sowunmi (Florida State University), SRP-HPC Visiting student 2023
20. Gideon Idumah (Case Western Reserve U.), SRP-HPC Visiting student 2022
21. Bilas Talukdar (New Jersey Institute of Tech.), SRP-HPC Visiting student 2022
22. Akhil Akella (Northern Illinois U.), ALCF summer student, Givens Fellow 2020, 2022
23. Vincent Zhong (U. of Illinois Urbana-Champaign), Givens Fellow, with Tanwi Mallick 2022
24. Scott Emmons (U. of California, Berkeley), DOE CSGF Fellow 2021
25. Alec Linot (U. of Wisconsin Madison), Givens Fellow, with Romit Maulik 2021
26. Michael McCabe (U. of Colorado Boulder), Givens Fellow 2021
27. Felix Edward Perez (U. of Texas at Dallas), EERE HPC4Mfg Internship, with Romain Egele 2021
28. Edward Friesema (U. of Nevada), EERE Robotics Internship Program, with Sami Khairy 2021
29. Yixuan Sun (Purdue U.), Givens Fellow, with Tanwi Mallick 2020, 2021
30. Shengli Jiang (U. of Wisconsin Madison), Givens Fellow 2020
31. Yixuan Sun (Purdue U.), Givens Fellow, with Tanwi Mallick 2020
32. Andres Rodríguez Rey (U. of California, San Diego), NSF MSGI Fellow 2020
33. Sami Khairy (Illinois Institute of Tech.), ALCF Summer Student, Givens Fellow 2019, 2020
34. Peihong Jiang (Brown U.), NSF MSGI Fellow 2019
35. Tianchen (Eric) Zhao (U. of Michigan), Givens Fellow, with Sandeep Madireddy 2018
36. Nimish Awalgaonkar (Purdue U.), Givens Fellow 2018
37. Dipnil Chakraborty (U. of Texas at Dallas), NSF MSGI Fellow 2018
38. Dipendra Jha (Northwestern U.), Givens Fellow 2018
39. Akash Roy (U. of Texas at Dallas), Research Aide 2018
40. Andy Jin (U. of Chicago), Jeff Metcalf Undergraduate Fellow 2018
41. Tian Ma (U. of Chicago), Jeff Metcalf Undergraduate Fellow 2018
42. YiMing Yu (New Jersey Institute of Tech.), NSF MSGI Fellow, with Paul Hovland 2017

43. Prateek Agarwal (Illinois Institute of Tech.), Resident Associate, with Sven Leyffer 2017
44. Amal Fethi (ENS Paris, France), Resident Associate, with Sven Leyffer 2017
45. Salvador Aguinaga (U. of Notre Dame), DOE CGSR Fellow 2017
46. Juan Li (U. of Chicago), Resident Associate, with Sven Leyffer 2016
47. Amit Roy (U. of Utah), Givens Associate, with Paul Hovland 2015
48. Arnamoy Bhattacharyya (ETH Zurich, Switzerland), Givens Associate 2014

Other 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.

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 (2018–2023).
- **F.N.R.S. chargé de recherches fellowship**, from the Belgian Funds for Scientific Research from October 2010 to September 2013. This highly competitive postdoctoral fellowship was declined to accept a 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

- ACM Gordon Bell Prize Finalist (2024): Recognized for the development of ORBIT, a 113-billion-parameter vision transformer model for Earth system predictability, leveraging advanced parallelism techniques on the Frontier supercomputer to achieve 1.6 exaFLOPs of performance. This recognition highlights ORBIT's contribution to advancing AI-driven climate modeling and prediction accuracy.
- HPCWire 2024 Editors' Choice Awards – Top Supercomputing Achievement: Recognized for the development of ORBIT.
- Member of the Tennessee AI Advisory Council appointed by the Tennessee State Governor, August, 2024–Present.
- **R&D 100 Award Finalist 2024**: Selected as a finalist for the development of DeepHyper, an advanced software framework for scalable automated machine learning, enhancing optimization capabilities in scientific applications.
- **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.
- 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.

Program Development

Throughout a range of funded projects, I have secured over $20 million in direct funding support. The combined funding for all projects I have been involved in exceeds $200 million. This substantial financial support has played a pivotal role in advancing research and development across multiple scientific and technological domains.

Research proposals funded (External)

For the PI, the total funding for the project is listed. For projects where I was listed as a Co-I, the funding is reported for the institution. For the project where I was senior personnel, the funding listed was allocated to me.

Table 1: Summary of Externally Funded Research Proposals

Project Name	**Funding Agency**	**Role**	**Awarded Funding to Me**	**Total Funding**	**Duration**
SciGPT: Scalable Foundational Model for Scientific Machine Learning	DOE ASCR	Co-PI	$1.05M	$7.05M	09/24–08/27
DURBAN: Enhancing Performance Portability in HPC Software with Artificial Intelligence	DOE ASCR	Senior personnel	$200K	$7.05M	09/24–08/27
UT-ORII Convergent Research Initiative on Fusion Technology	UT-ORII	Co-PI	$500K	$10M	09/24–08/27
Exploring the Power of Distributed Intelligence for Resilient Scientific Workflows	DOE ASCR	Co-PI	$1.75M	$8.75M	09/23–08/28
Center for Steel Electrification by Electrosynthesis	DOE BES and ASCR	Co-PI	$1M	$15M	09/23–08/27
High-fidelity Digital Models for Fusion Pilot Plant Design	DOE FES and ASCR	Senior personnel	$800K	$12M	09/23–09/27
Development of High-Fidelity Simulation Capabilities for ELM-free Design Optimization	DOE FES and ASCR	Co-PI	$240K	$12M	09/23–08/27
Fundamental nuclear physics at the exascale and beyond	DOE Nuclear Physics and ASCR	Co-PI	$1.15M	$13M	10/22–09/27
Improving Projections of AMOC and its Collapse Through advanced Simulations (ImPACTS)	DOE BER and ASCR	Senior personnel	$800K	$15M	10/22–09/27
Community Research on Climate and Urban Science (CROCUS)	DOE BER	Senior personnel	$400K	$25M	10/22–09/27
Probabilistic Impact Scenarios for Extreme Weather Event Resilience	DOE OE and ASCR	Senior personnel	$275K	$700K	10/22–09/23
PosEiDon: Platform for Explainable Distributed Infrastructure	DOE ASCR	Co-PI	$725K	$3M	10/21–09/24
Inertial neural surrogates for stable dynamical prediction	DOE ASCR	Senior personnel	$300K	$3M	10/21–09/24
Probabilistic Machine Learning for Rapid Large-scale and High-rate Aerostructure Manufacturing	DOE EERE	PI	$550K	$2.43M	10/21–09/23

Continued on next page

Table 1 – *Continued from previous page*

Project Name	**Funding Agency**	**Role**	**Funding Support to Me**	**Total Funding**	**Duration**
RAPIDS2: A SciDAC Institute for Computer Science and Data	DOE ASCR	Senior personnel	$1.12M	$25M	10/20–09/25
Big Data and AI for Smart Mobility	DOE EERE	Senior personnel	$600K	$3M	10/20–09/22
DAIN: Dynamic architectures through introspection and neuromodulation	DARPA	Senior personnel	$200K	$2M	10/20–09/22
Scalable Deep Learning Framework for Optimal Control of Cascading Failures and Restoration in Power System	DOE OE	Senior personnel	$100K	$380K	04/20–03/22
Ab-initio Guided Design and Materials Informatics for AHSS	DOE High-Performance Computing for Manufacturing	Co-PI	$90K	$300K	10/20–09/21
PROTEAS-TUNE	DOE Exascale Computing Project	Senior personnel	$800K	$6M	10/19–09/23
Quantifying Energy Drivers in Chemical Separations	DOE BES	Senior personnel	$675K	$2.25M	09/19–09/22
Scalable Data-Efficient Learning for Scientific Domains	DOE ASCR	PI	$2.5M	$2.5M	09/18–08/23
Accelerating HEP Science: Inference and Machine Learning at Extreme Scales	DOE HEP and ASCR	Senior personnel	$500K	$8M	10/17–09/22
High Performance Computing and Big Data Solutions for Mobility Design and Planning	DOE EERE	Senior personnel	$600K	$3M	10/17–09/20
RAPIDS: A SciDAC Institute for Computer Science and Data	DOE ASCR	Senior personnel	$600K	$15M	10/17–09/20
CANDLE: Exascale Deep Learning and Simulation Enabled Precision Medicine for Cancer	DOE Exascale Computing Project	Senior personnel	$200K	$3M	10/16–09/20
Autotuning Compiler Technology for Cross-Architecture Transformation and Code Generation	DOE Exascale Computing Project	Senior personnel	$300K	$3M	10/16–09/19

Continued on next page

Table 1 – *Continued from previous page*

Project Name	**Funding Agency**	**Role**	**Funding Support to Me**	**Total Funding**	**Duration**
Automated Machine Learning Pipeline for City of Chicago's Smart Data Platform	Bloomberg/City of Chicago	Co-PI	$100K	$250K	04/16–10/17
XTUNE: Autotuning for Exascale	DOE ASCR	Senior personnel	$87K	$87K	01/15–12/16
SDAV/SUPER Supplement: Improving Computational Science Throughput via Model-Based I/O Optimization	DOE ASCR	Senior personnel	$200K	$320K	02/15–09/16
OrFPGA: Empirical Performance Tuning for FPGA Designs in Space Applications	NASA-SBIR	Senior personnel	$80K	$200K	02/15–09/16
RAMSES: Robust Analytical Models for Science at Extreme Scale	DOE ASCR	Senior personnel	$600K	$3M	07/14–06/17
Total			>$19M	>$200M	

Research proposals funded (Internal > $100K)

Table 2: Summary of Internally Funded Research Proposals

Project Name	**Funding Agency**	**Role**	**Total Funding**	**Duration**
Community and Infrastructure Adaptation to Climate Change: An AI-Driven Research Tool	ANL/LDRD (Climate and Energy Action)	Collaborator	$560K	10/21–9/23
Towards Neighborhood Scale Climate Simulations using AI and Accelerated GPUs	ANL/LDRD (Climate and Energy Action)	Collaborator	$508K	10/21–9/23
AI-Emulator Assisted Data-Assimilation	ANL/LDRD (Future Computing)	Co-PI	$1M	10/20–9/23
Automated Model Inference for Cosmological Structure Formation	ANL/LDRD (Future Computing)	Co-PI	$1M	10/20–9/23
Nuclear Quantum Monte Carlo with Machine Learning Techniques	ANL/LDRD (Future Computing)	Co-PI	$750K	10/20–9/23

Continued on next page

Table 2 – *Continued from previous page*

Project Name	**Funding Agency**	**Role**	**Total Funding**	**Duration**
A Framework for Device and Architecture Co-Design via Hyperparameter Optimization and Backpropagation through Hardware	ANL LDRD Prime - Microelectronics	Co-PI	$260K	1/20–9/20
Artificial Intelligence Assisted Safety Modeling and Analysis of Advanced Nuclear Reactors	ANL/LDRD	Co-PI	$505K	10/19–9/21
Towards the Creation of an Advanced Mobility Cybersecurity Testbed	ANL LDRD (Secure Energy and Critical Resources)	Co-PI	$360K	10/18–10/19
A.I C.D.I: Atomistically Informed Coherent Diffraction Imaging	ANL/LDRD	Co-PI	$180K	10/16–9/18
SLIK-D: Scalable Machine Learning Infrastructures for Knowledge Discovery	ANL/LDRD	Co-PI	$600K	10/16–9/18
End-to-End Genome Annotation and Phenotype Prediction with Deep Learning	ANL/LDRD	Co-PI	$300K	10/16–9/18

Complete List of Research Products

Citations: 5235; h-index: 38; i10-index: 91 (Source: Google Scholar; 02/24/2025)

Refereed Journal Articles and Book Chapters

1. L. Johnston, V. Patel, Y. Cui, and P. Balaprakash, "Revisiting the problem of learning long-term dependencies in recurrent neural networks," *Neural Networks*, vol. 106887, 2024.

2. W. F. Godoy, P. Valero-Lara, K. Teranishi, P. Balaprakash, and J. S. Vetter, "Large language model evaluation for high-performance computing software development," *Concurrency and Computation: Practice and Experience*, 2024.

3. X. Liu, M. Rüttgers, A. Quercia, R. Egele, E. Pfaehler, R. Shende, M. Aach, W. Schröder, P. Balaprakash, and A. Lintermann, "Refining computer tomography data with super-resolution networks to increase the accuracy of respiratory flow simulations," *Future Generation Computer Systems*, vol. 159, pp. 474–488, 2024.

4. O. Yildiz, K. Raghavan, H. Chan, M. J. Cherukara, P. Balaprakash, S. Sankaranarayanan, and T. Peterka, "Automated defect identification in coherent diffraction imaging with smart continual learning," *Neural Computing and Applications*, 2024.

5. S. Madireddy, C. Akçay, S. E. Kruger, T. B. Amara, X. Sun, J. McClenaghan, J. Koo, et al., "EFIT-Prime: Probabilistic and physics-constrained reduced-order neural network model for equilibrium reconstruction in DIII-D," *Physics of Plasmas*, vol. 31, no. 9, 2024.

6. T. Mallick, J. D. Bergerson, D. R. Verner, J. K. Hutchison, L.-A. Levy, and P. Balaprakash, "Understanding the impact of climate change on critical infrastructure through NLP analysis of scientific literature," *Sustainable and Resilient Infrastructure*, 2024.
7. S. Khairy and P. Balaprakash, "Multi-fidelity reinforcement learning with control variates," *Neurocomputing*, vol. 527, pp. 127963, 2024.
8. R. Egele, F. Mohr, T. Viering, and P. Balaprakash, "The unreasonable effectiveness of early discarding after one epoch in neural network hyperparameter optimization," *Neurocomputing*, vol. 527, pp. 127964, 2024.
9. Y. Sun, O. Sowunmi, R. Egele, S. H. K. Narayanan, L. Van Roekel, and P. Balaprakash, "Streamlining ocean dynamics modeling with Fourier neural operators: A multiobjective hyperparameter and architecture optimization approach," *Mathematics*, vol. 12, no. 10, pp. 1483, 2024.
10. A. Bhaduri, N. Ramachandra, S. K. Ravi, L. Luan, P. Pandita, P. Balaprakash, M. Anitescu, C. Sun, and L. Wang, "Efficient mapping between void shapes and stress fields using deep convolutional neural networks with sparse data," *Journal of Computing and Information Science in Engineering*, vol. 24, no. 5, pp. 051007, 2024.
11. T. Mallick, J. Macfarlane, and P. Balaprakash, "Uncertainty quantification for traffic forecasting using deep-ensemble-based spatiotemporal graph neural networks," *IEEE Transactions on Intelligent Transportation Systems*, 2024.
12. A. Samanta, P. Balaprakash, S. Aubry, and B. K. Lin, "Machine-learning-aided density functional theory calculations of stacking fault energies in steel," *Scripta Materialia*, vol. 241, pp. 115862, 2024.
13. X. Wu, J. R. Tramm, J. Larson, J.-L. Navarro, P. Balaprakash, B. Videau, M. Kruse, P. Hovland, V. Taylor, and M. Hall, "Integrating ytopt and libEnsemble to autotune OpenMC," *The International Journal of High Performance Computing Applications*, 2024.
14. S. Jiang, S. Qin, R. C. Van Lehn, P. Balaprakash, and V. M. Zavala, "Uncertainty quantification for molecular property predictions with graph neural architecture search," *Digital Discovery*, vol. 3, no. 8, pp. 1534–1553, 2024.
15. K. Raghavan, M. L. Avila, P. Balaprakash, H. Jayatissa, and D. Santiago-Gonzalez, "Classification of events from α-induced reactions in the MUSIC detector via statistical and ML methods," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1058, pp. 168786, 2024.
16. S. Khairy, P. Balaprakash, and L. X. Cai, "A Gradient-Aware Search Algorithm for Constrained Markov Decision Processes," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
17. J. Kim, S. Leyffer, and P. Balaprakash Learning symbolic expressions: Mixed-integer formulations, cuts, and heuristics, *INFORMS Journal on Computing*, INFORMS, 2023.
18. 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, 2023.
19. 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.

20. 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, Elsevier, 2023.

21. 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, Royal Society of Chemistry, 2023.

22. 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, 2023.

23. Y. Sun, K. Raghavan, and P. Balaprakash, "Introduction to Reinforcement Learning," in *Methods and Applications of Autonomous Experimentation*, Chapman and Hall/CRC, pp. 152–174, 2023.

24. 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.

25. 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.

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

27. 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.

28. 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.

29. 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.

30. 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.

31. 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.

32. 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.

33. 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)

34. 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.

35. 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.

36. 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.

37. 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.

38. 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)

39. 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)

40. 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.

41. 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.

42. 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.

43. 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.

44. 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.

45. 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.

46. 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.

47. 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.

48. 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.

49. 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.

50. 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.

51. 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.

52. 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.

53. 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.

54. 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.

55. 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.

56. 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.

57. 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

58. H. Jin, G. Papadimitriou, K. Raghavan, P. Zuk, P. Balaprakash, C. Wang, A. Mandal, and E. Deelman, Large Language Models for Anomaly Detection in Computational Workflows: From Supervised Fine-Tuning to In-Context Learning, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '24)*, IEEE Press, Article 92, pp. 1–17, 2024.

59. S. Dash, I. R. Lyngaas, J. Yin, X. Wang, R. Egele, J. A. Ellis, M. Maiterth, G. Cong, F. Wang, and P. Balaprakash, "Optimizing distributed training on Frontier for large language models," in *ISC High Performance 2024 Research Paper Proceedings (39th International Conference)*, Prometeus GmbH, pp. 1–11, 2024.

60. 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. Best paper award

61. A. Balaji and P. Balaprakash, "Application-Hardware Co-Optimization of Crossbar-Based Neuromorphic Systems," in *Proceedings of the 2023 International Conference on Machine Learning and Applications (ICMLA)*, IEEE, pp. 1253–1260, 2023.

62. W. Zhong, T. Mallick, J. Macfarlane, H. Meidani, and P. Balaprakash, "Graph Pyramid Autoformer for Long-Term Traffic Forecasting," in *Proceedings of the 2023 International Conference on Machine Learning and Applications (ICMLA)*, IEEE, pp. 384–391, 2023.

63. S. Madireddy, A. Yanguas-Gil, and P. Balaprakash, "Improving performance in continual learning tasks using bio-inspired architectures," in *Proceedings of the 2nd Conference on Lifelong Learning Agents (CoLLAs)*, PMLR, pp. 992–1008, 2023.

64. L. Luan, N. Ramachandra, S. K. Ravi, A. Bhaduri, P. Pandita, P. Balaprakash, M. Anitescu, C. Sun, and L. Wang, "Scalable Probabilistic Modeling and Machine Learning With Dimensionality Reduction for Expensive High-Dimensional Problems," in *Proceedings of the International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 87295, American Society of Mechanical Engineers, p. V002T02A011, 2023.

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

66. 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.

67. 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.

68. 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.

69. 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.

70. 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.

71. 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.

72. 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.

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

74. 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.

75. 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.

76. 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.

77. 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.

78. 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.

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

80. 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.

81. 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.

82. 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.

83. 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.

84. 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.

85. 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.

86. 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.

87. 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.

88. 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.

89. 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.

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

91. 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.

92. 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.

93. 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.

94. 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.

95. 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.

96. 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.

97. 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.

98. 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.

99. 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.

100. 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.

101. 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.

102. 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.

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

104. 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.

105. 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.

106. 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.

107. 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.

108. 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.

109. 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.

110. 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.

111. 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.

112. 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.

113. 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.

114. 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.

115. 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.

116. 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.

117. 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.

118. 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.

119. 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.

120. 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.

121. 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.

122. 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.

123. 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.

124. 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.

125. 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.

126. 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).

127. 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.

128. 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.

129. 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.

130. 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.

131. 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.

132. 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.

133. 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.

134. 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

135. 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.

136. 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.

137. 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.

138. 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.

139. 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.

140. 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.

141. 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.

142. 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.

143. 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.

144. 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.

145. 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

146. 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.

147. 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.

148. 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.

149. 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.

150. 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.

151. 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.

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

153. 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.

154. 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.

155. 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.

156. 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.

157. 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.

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

159. 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.

160. 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.

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

162. 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

163. C. Daniel, J. C. Gehin, K. Laurin-Kovitz, B. Morreale, R. Stevens, W. Tumas, M. Anitescu, A. Poczatek, A. Siegel, S. Som, R. Vilim, R. Grout, B. Kroposki, M. Yue, K. Rose, A. Al Rashdan, C. Ritter, P. Balaprakash, P. Jain, T. Kuruganti, M. A. Piette, T. Hong, C. Corley, R. Rallo, J. Grosh, B. Van Essen, M. Reno, H. Viswanathan, F. Alexander, and E. M. Dietrich, "Advanced Research Directions on AI for Energy," United States, 2024.

164. 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, United States, 2023.

165. 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)

166. 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)

167. 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).

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

169. 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.

170. 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.

171. 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

172. 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.

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

174. 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

175. 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.

176. 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.

177. 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.

178. 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.

179. 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. ORNL Artificial Intelligence Initiative, Keynote at 66th ACM Mid-Southeast Conference, Gatlinburg, TN, November 2024.

2. ORNL AI Advancements and Applications, Keynote at 8th Annual Commonwealth Computational Summit, University of Kentucky, October 2024.

3. From Prediction to Protection: Leveraging AI and HPC to Safeguard Lives and U.S. Infrastructure, Invited panelist at the USCAS HPC/AI Weather and Climate Panel for US Senate Staffers, September 2024.

4. Overview of Oak Ridge National Laboratory's AI Initiative: Advancing Secure, Trustworthy, and Energy-Efficient AI at Scale for Scientific Discovery Invited talk at the Smoky Mountains Computational Sciences & Engineering Conference, Knoxville, TN, September 2024.
5. Biological Computing: Are We Missing the Obvious? Invited talk at the 2024 Energy-Efficient Computing for Science Workshop, Bethesda, MD September 2024.
6. Reimagining Workflow and Resource Management Systems with Swarm Intelligence. Invited talk at 28th Annual IEEE High Performance Extreme Computing Conference. September 2024.
7. Advancing Secure, Trustworthy, and Energy-Efficient AI at Scale for Scientific Discovery, Plenary talk at the 2024 CNMS User Meeting, Knoxville, TN, August 2024.
8. Scalable Automated Deep Ensemble for Uncertainty Quantification in Scientific Machine Learning. Invited talk at Second USACM Thematic Conference on Uncertainty Quantification for Machine Learning Integrated Physics Modeling (UQ-MLIP 2024), Crystal City, Arlington, Virginia, August 2024.
9. Advancing Secure, Trustworthy, and Energy-Efficient AI to Accelerate Scientific Discovery and National Security, Invited talk at SCSP AI Expo for National Competitiveness, Washington, DC. May 2024.
10. ORNL Artificial Intelligence Programs, Keynote at the National Science Bowl, Washington, D.C., April 2024.
11. ORNL's AI Initiative: Advancing Secure, Trustworthy, and Energy-Efficient AI for Science and National Security, Invited talk at the U.S. Space & Rocket Center AI Symposium, Huntsville, AL February 2024.
12. ORNL's AI Initiative: Advancing Secure, Trustworthy, and Energy-Efficient AI for Science and National Security, Invited talk at LANL AI Forum, Los Alamos, NM, January 2024.
13. Advancing Safe and Trustworthy AI for Science: Understanding and Mitigating Threats to AI Data, Models, and Processes, Invited talk at the Monterey Data Conference, August 2023.
14. Extracting the impact of climate change from scientific literature using Snorkel-enabled NLP. Invited talk at the Future of Data-Centric AI, August 2022.
15. Automated machine learning with DeepHyper. Invited talk at the Center for Mathematics and Artificial Intelligence (CMAI) Colloquium, George Mason University, February 2022.
16. Scientific domain-informed machine learning. Invited talk at the Florida Georgia Louis Stokes Alliance for Minority Participation (FGLSAMP), Florida A&M University, November 2021.
17. Democratizing Deep Learning Development with DeepHyper. Invited talk at the Brookhaven National Laboratory, October 2021.
18. Artificial intelligence and machine learning. Invited talk at the AI for Sustainability Workshop, Northwestern University, April 2020.
19. Enabling ML approaches to HPC systems operations. Keynote and invited panelist at the Workshop on Monitoring and Analsis for High Performance Computing Sytems Plus Applications (HPCMASPA), September 2020.

20. The future of HPC systems in the presence of AI. Keynote and invited panelist at the Smoky Mountains Computational Sciences and Engineering Conference (SMC), August 2020.

21. Artificial intelligence to accelerate discovery and development. Invited talk at the Argonne Outloud Public Lecture, Argonne, May 2019.

22. Machine-learning-based performance modeling and tuning for high-performance computing. Keynote at the HPCaML 2019: The International Workshop on the Intersection of High Performance Computing and Machine Learning, Washington, DC, February 2019.

23. Machine-learning-based-search for automatic performance tuning. Keynote at the OMASE 2019: Optimization, Modeling, Analysis and Space Exploration Workshop, Washington, DC, February 2019.

24. Artificial intelligence for transportation and mobility. Invited talk at the Tech Hub, SAE World Congress 2017, Detroit, MI, May 2017.

Talks and Posters Presented at Conferences and Workshops

25. AI at the Nexus of Governance, Ethics, and Education, IEEE CIS Government Activities Committee YouTube Series, November 2024.

26. How the DOE Early Career Research Program has Advanced Scientific Discovery and Innovation, DOE Office of Science Webinar, October 2024.

27. Advancing Scientific Research with AI: Collaborative Efforts in Training Large Foundation Models on the Frontier Supercomputer Invited talk at the 26th Annual SOS conference, Cocoa Beach, FL April 2024.

28. ORNL's AI Initiative: Advancing Secure, Trustworthy, and Energy-Efficient AI for Science and National Security, Invited talk at Brown's National Labs Day, April 2024.

29. AI for Energy Storage: Advancing Secure, Trustworthy, and Energy-Efficient AI for Energy Storage, Invited talk at the Frontiers in Energy Storage: Next Generation AI Workshop, Berkeley, CA April, 2024.

30. Taming High-Performance Computing Platform Complexity with Machine Learning, Plenary Talk at CRNCH Summit 2024, Atlanta, GA February 2024.

31. Scalable, Automated Deep Neural Network Development for Complex Systems, Invited talk at Georgia Tech Workshop on Foundation of Scientific AI for Optimization of Complex Systems, January 2024.

32. Taming high-performance computing platform heterogeneity with machine learning. Sky Computing – Toward Efficient Computing on the Cloud Special, September 2022.

33. ORNL AI Initiative, Invited talk at the AIRES Workshop, May, 2024.

34. Advancing Science through AI: A Promising Pathway to Breakthroughs, Oppenheimer Science and Energy Leadership Program Fellows Visit, August 2023.

35. Democratizing deep learning development with DeepHyper. Invited talk at the Princeton-DataX Workshop Series, May 2022.

36. Graph neural network for anomalies detection in scientific workflows. ModSim 2002, August 2022.
37. 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.
38. DeepHyper: Scalable neural architecture and hyperparameter search for deep neural networks. DOE Booth Demo, November 2021.
39. AI for HPC: Everything that can be automated should be automated. JLESC workshop, AI for HPC panel, October 2021.
40. Scientific domain-informed machine learning. San Diego State University and Sustainable Horizon Institute, March 2021.
41. Graph-partitioning-based diffusion convolution recurrent neural network for large-scale traffic forecasting. DOE SMART All-Hands meeting (Invited talk), August 2021.
42. Large-scale traffic forecasting via spatial-temporal learning. DOE Workshop on Computational Mission Needs for Clean Energy and Advanced Manufacturing, June 2021.
43. Accelerating deep learning for science with SambaNova. The 2nd International Workshop on Machine Learning Hardware (Plenery talk), Co-located with ISC 2021, June 2021.
44. Neuromodulated neural architectures with local error signals for memory-constrained online continual learning. APS/CNM Users Meeting, May 2021.
45. Ytopt/SuRF: Machine-learning-based search for autotuning. 2021 ECP Annual Meeting, April 2021.
46. 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.
47. Machine-learning-based automatic performance tuning. Workshop on Program Synthesis for Scientific Computing, August 2020.
48. Automated machine learning for molecular chemistry. Artificial Intelligence for Water Workshop, Argonne, September 2020. (Invited talk)
49. Scientific domain-informed machine learning. Argonne Physics Division Colloquium, Argonne, May 2019.
50. 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.
51. 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.
52. Scientific domain-informed machine learning. Indian Institute of Science, Computational and Data Science Colloquium, India, December 2018.

53. DeepHyper: Asynchronous hyperparameter search for deep neural networks. 25th IEEE International Conference on High Performance Computing, Data, and Analytics, Bengaluru, India, December 2018.

54. Scientific domain-informed machine learning. Advanced Photon Source Colloquium, Argonne, December 2018.

55. Benchmarking machine learning methods for performance modeling of scientific applications. PMBS 2018: Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems, November 2018.

56. Spatial-temporal deep learning and hyperparameter search for traffic prediction. Workshop on Large-Scale Computing for Transportation Studies, Maui, HI, November 2018.

57. Reproducibility, portability and interpretability of deep learning. Workshop on Deep Learning for Multimessenger Astrophysics: Real-time Discovery at Scale, October 2018.

58. Scientific domain-informed machine learning. Department Seminar Series, Mechanical and Industrial Engineering, University of Illinois at Chicago, October 2018.

59. Scientific domain-informed machine learning. ALCF Simulation, Data, and Learning Workshop, Argonne, IL, October 2018.

60. Machine learning hardware. SciDAC Fusion Machine-Learning Workshop 2018, Princeton, NJ, June 2018.

61. Machine learning in high performance computing. Invited talk at the 7th Greater Chicago Area Systems Research Workshop, University of Chicago, April 2018.

62. Need for data locality in machine/deep learning. Fourth Workshop on Programming Abstractions for Data Locality (PADAL'17), Chicago, IL, August 2017.

63. Generative adversarial networks. US ATLAS Workshop, Argonne, July 2017.

64. Automatic multi-objective modeling with machine learning. Invited talk at the Workshop on Optimization and Machine Learning (ACNTW 17), Northwestern University, May 2017.

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

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

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

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

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

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

71. 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.

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

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

74. 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.

75. 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.

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

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

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

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

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

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

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

83. 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.

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

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

86. 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

87. AI — Friend or Foe? OACC Lunch & Learn, Haslam Business School, UTK, November 2024.
88. Responsible Deployment of AI in State Government Operations, Tennessee AI Advisory Council Engagement Subcommittee Meetings, October 2024.
89. AI for Neutrons: Capabilities and Opportunities, AI for Neutrons Workshop, ORNL, October 22, 2024.
90. AI Initiative at Oak Ridge National Laboratory Secure, Trustworthy, and Energy- Efficient AI to Accelerate Scientific Discoveries, CCSD Summer Series, Oak Ridge, TN, July 2024.
91. Advancements and Future Directions in AI Research, Keynote at 12th ORPA Research Symposium, July 2024.
92. Democratizing deep learning development with deephyper. AppliedAI Meetup, June 2022.
93. Overview of Argonne AI R&D. French Delegation visit at Argonne, April 2022.
94. DeepHyper: Automated deep learning at scale. Argonne/MCS-ARAMCO Seminar, March 2022.
95. Deephyper: Automated machine learning at scale. Argonne-IIT Collaboration Seminar, March 2022.
96. Deephyper for automated machine learning in VASP. Equation of State and Materials Theory group, March 2022.
97. DeepHyper: Scalable automated machine learning package. Argonne CPS Division seminar series, August 2020.
98. Neuromorphic acceleration for uncertainty quantification in deep neural networks. MCS All Hands Meeting, Argonne, IL, January 2020.
99. Artificial intelligence for science. Invited talk at the Transportation Research Board Executive Committee A.I. Policy Session, January 2019.
100. Scalable reinforcement-learning-based neural architecture search for scientific and engineering applications. LANS Seminar, Argonne, IL, February 2019.
101. Scientific domain-informed machine learning. Consortium for Computational Physics and Chemistry Meeting, Argonne, November 2018.
102. 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.
103. Multi objective optimization of HPC kernels for performance, power, and energy. 6th Joint Laboratory for Extreme Scale Computing (JLESC) Workshop, November, Japan, June 2016.
104. AutoMOMML: Automatic Multi-objective Modeling with Machine Learning. 5th Joint Laboratory for Extreme Scale Computing (JLESC) Workshop, Lyon, France, June 2016.
105. Predictive modeling for large scale vehicle simulations. Urban Data Analytics/City of Chicago SmartData Platform Workshop, Chicago, IL, May 2015.

106. Active-learning-based surrogate models for empirical performance tuning. Computation Institute Seminar, University of Chicago, Chicago, IL, February 2014.

107. Active-learning-based surrogate models for empirical performance tuning. 10th workshop of the INRIA-Illinois-ANL Joint Laboratory, Urbana-Champaign, IL, November 2013.

108. 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

109. Organizer and lead, ALCF DeepHyper Automated Machine Learning Workshop. ALCF, Argonne, July 2022. (Over 240+ registrants and 90+ online attendees)

110. Organizer and lead, Transfer learning and online tuning with PROTEAS-TUNE/ytopt. ECP AHM Tutorial, May 2022.

111. DeepHyper: A hyperparameter search package for deep neural networks. ALCF Simulation, Data, and Learning Workshop, Argonne, December 2020.

112. DeepHyper and hyperparameter optimization. ATPESC 2020: Argonne Training Program on Extreme-Scale Computing, Argonne, August 2020.

113. Reinforcement learning and applications. SciDAC TDS Meeting, March 2020.

114. AI tools. PSE AI in Science and Engineering Workshop, Argonne, September 2019.

115. Deep learning basics. ATPESC 2019: Argonne Training Program on Extreme-Scale Computing, St. Charles, August 2019.

116. Deep learning basics. Artificial Intelligence for Science Workshop, Argonne, August 2019.

117. Recurrent networks for time series data. Argonne Artificial Intelligence for Science Workshop, Argonne, August 2019.

118. Hyperparameter optimization and DeepHyper. ATPESC 2019: Argonne Training Program on Extreme-Scale Computing, St. Charles, August 2019.

119. Recurrent networks for time series data. Artificial Intelligence for Science Workshop for Summer Students, Argonne, July 2019.

120. Deep learning basics. Artificial Intelligence for Science Workshop for Summer Students, Argonne, July 2019.

121. Data parallel deep learning. CANDLE-EXALEARN ECP Workshop, Argonne, October 2018.

122. Data for machine learning. Argonne Geospatial Workshop, Argonne, August 2018.

123. Machine learning overview and applications. Argonne Environmental Science Division Retreat, Argonne, May 2018.

124. Deep learning basics. Argonne Deep Learning Workshop, Argonne, January 2018.

125. Recurrent neural networks. Argonne Deep Learning Workshop, Argonne, January 2018.

126. Automatic multi-objective modeling with machine learning. Argonne Training Program on Extreme-Scale Computing (ATPESC), St. Charles, IL, August 2017.

127. Overview of machine learning methods. Argonne Training Program on Extreme-Scale Computing (ATPESC), St. Charles, IL, August 2017.

128. Introduction to unsupervised and supervised learning in Python: Hands-on tutorial. Argonne Training Program on Extreme-Scale Computing (ATPESC), St. Charles, IL, August 2017.

129. Overview of machine learning methods. Argonne Machine Learning Workshop, Argonne, July 2017.

130. Automatic multi-objective modeling with machine learning. Argonne Machine Learning Workshop, Argonne, July 2017.

131. Introduction to unsupervised and supervised learning in python: Hands-on tutorial. Argonne Machine Learning Workshop, Argonne, July 2017.

132. Automatic multi-objective modeling with machine learning. Argonne Machine Learning Workshop, Argonne, April 2017.

133. Overview of machine learning methods. Argonne Machine Learning Workshop, Argonne, April 2017.

134. 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. 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

4. 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 Institutions

1. FY25 Technology Transfer Liaison, ORNL, 2025
2. DOE ASCR Reverse Site Visit, ORNL, 2024 *Role: Presented AI Initiative Highlights, detailing progress, impact, and strategic goals.*
3. CCSD New Employee Orientation, ORNL, 2024 *Role: Delivered AI Initiative Highlights to familiarize new employees with ORNL's cutting-edge AI advancements and strategic priorities for FY24.*
4. ASCR Observer at DOE SC AI Roundtables, Washington, D.C., October 28-30, 2024.
5. OSTP, Asad Ramzanali, ORNL, October 11, 2024 *Role: Presented ORNL's contributions to the national AI strategy and scientific initiatives.*
6. Canadian Cyber Security Centre, ORNL, March 27-28, 2024 *Role: Engaged with cybersecurity experts to share advancements and collaborate on mutual interests.*
7. Volkswagen Group and Porsche North America Leadership, ORNL, March 1, 2024 *Role: Showcased cutting-edge AI and computing capabilities relevant to the automotive industry.*
8. DOE FCIO/CIO Guests, ORNL, January 5, 2024 *Role: Facilitated discussions with DOE Chief Information Officers and leadership.*
9. Oppenheimer Science and Energy Leadership Program Fellows, ORNL, August 10, 2023 *Role: Hosted and presented AI-related work to a prestigious leadership program cohort.*
10. Michael Marn, Senator Blackburn Staffer, ORNL, August 16, 2023 *Role: Provided insights on AI initiatives and national laboratory research to a Congressional staffer.*
11. Member, Argonne AI Editorial Board, 2020–2023
12. Argonne session lead, AI for Science and Security Workshop series, 2022
13. Member, ANL Climate and Energy Action LDRD review committee, 2021, 2022 (Reviewed 16 proposals in total)
14. Co-lead, Argonne's AI for Science communications strategy, 2021
15. Postdoc mentor: Jordon Fox (NP, 2022–2023), Bryce Fore (NP, 2022–2023), Robert Underwood (MCS, 2022–2023), Zhen Xie (ALCF, 2022–2023), Hieu A. Doan (MSD, 2019–2021), Peco Myint (XSD, 2020–2023), Naveen Dandu (MSD, 2019–2022), Alp Dener (MCS, 2018–2021), Valentin Reis (MCS, 2019–2020), Andreas Victor Copan (CSE, 2019–2020)
16. Mentor, Early Career Research Proposals, 2019, 2020
17. Member, Red team review for research proposals, 2018–Present
18. Member, Wilkinson Fellowship Committee, 2016, 2018, 2020
19. Member, DOE Argonne Data Science Program review committee, 2017, 2018, 2019, 2020, 2021. (Reviewed approximately three proposals every year)
20. Member, Argonne's AI-X Task Force for Materials and Chemistry, 2018

21. Member, working group for ML strategy in Argonne Materials Science Division, 2018
22. Member, library committee for MCS, ANL, 2013–2015
23. Vice President, Postdoctoral Society of ANL, January 2012–December 2012

Services to Sponsors

24. Member, DOE ASCAC CSGF Fellowship Subcommittee, November 2024.
25. Member, DOE Computational Science Graduate Fellowship (CSGF) application screening committee, 2020, 2021, 2022, 2023, 2024.
26. Reviewer, DOE ASCR Early Career Research Proposals (ECRP) Review Panel, 2020, 2021, 2022, 2023.
27. AI/ML expert reviewer for DOE NP: SBIR/STTR Phase I, 2021, 2022, 2023, 2024.
28. Reviewer, DOE ASCR Express: 2022 Exploratory Research for Extreme-Scale Science, Explainable AI, 2022
29. Session lead, AI for Science and Security Workshop series, 2022
30. 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
31. DOE INCITE computational readiness review committee, 2016, 2017, 2018, 2019, 2020, 2021 (Reviewed two proposals every year)
32. 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
33. Session lead, Artificial Intelligence for Earth System Predictability Workshop, 2021
34. DOE AI Town Hall: Argonne co-lead for software environment and research and writing co-lead for AI for computer science, 2019
35. Co-lead for applied machine learning and member of the workshop organizing committee, DOE ASCR Smart Network Systems Workshop, Rockville, MD, December 2016
36. 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

37. Editorial Board, Frontiers in Systems Neuroscience, 2021–Present
38. Program Committee, 18th IEEE International Conference on eScience - eScience 2022
39. Program Committee, SIAM Conference on Parallel Processing for Scientific Computing (PP22)
40. Program Committee, NeurIPS 2021 Workshop on Machine Learning and the Physical Sciences
41. Ph.D. Thesis Committee, Albert Njoroge Kahira, Polytechnic University of Catalonia, 2021

42. Program Committee, ICPP: International Conference on Parallel Processing, 2021, 2022
43. Program Committee, IEEE/ACM International Conference on High Performance Computing, Networking, Storage, and Analysis (SC), 2020, 2021, 2022
44. Program Committee, ICONS: International Conference on Neuromorphic Systems, 2020, 2021, 2022
45. Program Committee, PMBS: International Workshop in Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems, 2015, 2016, 2017, 2018, 2019, 2020, 2021, 2022
46. Program Committee, tutorials, SC '20: IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis, 2020
47. Program Committee, Applications track, ICPP '20: 49th International Conference on Parallel Processing, 2020
48. Program Committee, AI4S20: Workshop on Artificial Intelligence and Machine Learning for Scientific Applications, 2020
49. Program Committee, ANTS 2016: Tenth International Conference on Swarm Intelligence, 2016
50. Program Committee, IPDRM 2016: First Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware, 2016
51. 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 Synergistic Activities

1. Mentor, TnAchieves – Guiding Tennessee high school seniors through the transition to post-secondary education, 2024–Present.
2. Panelist, After Trinity – A Screening of Oppenheimer After Trinity & a Student-Led Panel Discussion, Oak Ridge High School, July 2023.
3. Supervised 1 student under Sustainable Research Pathways for High-Performance Computing (SRP-HPC) 2023 Program.
4. Speaker at Sustainable Horizons Institute (SHI) seminar series, 2023.
5. Mentored two high school students from the Illinois Mathematics and Science Academy under the Student Inquiry and Research program, September 2022–Present
6. Supervised more than 30 summer graduate students from 20 different universities. May 2014–Present

7. Supervised two students under Sustainable Research Pathways for High-Performance Computing (SRP-HPC) Program, 2022–Present
8. Mentored 8 postdocs under Argonne mentoring program, 2019–Present
9. Was speaker at Sustainable Horizons Institute (SHI) seminar series, 2020–Present
10. Was speaker at the Florida Georgia Louis Stokes Alliance for Minority Participation (FGLSAMP), Florida A&M University, November 2021
11. 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)
12. Mentored ANL postdocs on preparing the elevator pitch for LDRD proposals, PSA LDRD Workshop and Networking Event, April 2015
13. 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
14. 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.