2024

Christine Mirzayan Science and Technology Policy Graduate Fellowship Program

Statement of Interest

Question:
In 500 words or fewer, please answer the following questions:

- Why are you interested in being a Mirzayan Fellow?
- How do you hope this experience will contribute to your career goals?
- What unique perspectives and/or personal experiences can you bring to discussions during the Mirzayan Fellowship?
- Is there anything else you would like us to know about you?

Answer:
The Christine Mirzayan Science and Technology Policy Graduate Fellowship presents a unique opportunity for me to blend my expertise in high-performance computing (HPC) and artificial intelligence (AI) with science and technology policy, a critical combination as we transition toward an increasingly AI-driven world. My research focuses on developing AI-driven solutions for efficient resource allocation in HPC environments, explicitly exploring how AI can optimize computing infrastructure, reduce "time-to-science," and enhance the productivity of computational scientists. An NSF AI institute funds my work. It aims to create an AI-driven cyberinfrastructure that democratizes AI by making it accessible to a broader range of users, including academic institutions, industries, and communities with limited resources for independent AI development. This initiative supports research by providing access to advanced computing resources, datasets, and tools, ensuring that AI technologies are available to those who may otherwise lack the means to develop and utilize AI solutions.

While technical advancements in AI and HPC are essential, well-informed policies are equally important to ensure these benefits are distributed equitably and sustainably across society. Through this fellowship, I hope to gain valuable insights into how policies are crafted and implemented to govern emerging technologies like AI. This includes ensuring equitable access to HPC resources and supporting the growth of AI-enabled ecosystems, such as intelligent edge devices like phones, cameras, and drones. Additionally, I am keen to explore how the shift to AI-driven solutions will affect accountability and sustainability policies, especially given the increasing energy and hardware demands for training and deploying AI models.

One of the unique perspectives I bring to this fellowship is my firsthand experience with the disparities in educational and research resources between developing and developed countries. During a summer internship, I assessed the AI readiness of an Indian university—my undergraduate institution—and found that the lack of computational resources significantly limits students' ability to apply their theoretical knowledge to practical AI model development. This resource gap places students in countries like India at a disadvantage compared to their peers in the U.S., where access to advanced computational resources, such as those available at the Texas Advanced Computing Center (TACC) and the Ohio Supercomputer Center (OSC), is much more common. This disparity hampers education and limits global contributions to AI advancements.

This issue is not just confined to developing countries; it extends to the growing divide between academia and industry. Industry-affiliated researchers dominate major AI conferences such as NeurIPS and ICML, benefiting from vast computational resources provided by tech giants like Google, Microsoft, and DeepMind. At NeurIPS 2020 and 2021, 40-50% of accepted papers were co-authored by industry researchers, underscoring the widening gap between those with access to advanced resources and those without.

The increasing influence of industry in academic research highlights the urgent need for policies that ensure equitable access to computational resources. Through the Mirzayan Fellowship, I plan to explore how science and technology policies can address these challenges and facilitate a more balanced and inclusive research ecosystem.

Furthermore, the rapid development of AI models such as GPT-3.5 and GPT-4, which require immense computational resources, underscores the growing divide between institutions with access to these resources and those without. I plan to study the policy implications of this demand and explore solutions that promote a more equitable distribution of AI technologies across institutions, ensuring that the benefits of AI development are shared more broadly.

Question:
In 300 words or fewer, please tell us about a policy issue you are interested in or passionate about and why you think it is an important issue. The topic does not need to align with your professional or educational experience.

Answer:
The rapid integration of artificial intelligence (AI) across various sectors is enabling scientists and industries to democratize AI while creating the need for clear operational guidelines for these models. For example, products come with user manuals detailing product descriptions, usability guidelines, safety, regulatory compliance, and environmental impact—information users need before purchase or operation. Similarly, any AI-enabled software, service, or device (like smart devices) should include comprehensive safety, usability, and privacy documentation, enabling users to make informed decisions.
Initiatives such as the European Union (EU) AI Act, Institute of Electrical and Electronics Engineers (IEEE) Global Initiative on Ethics of Autonomous and Intelligent Systems, and Organization for Economic Co-operation and Development (OECD) AI Principles primarily focus on trustworthiness, transparency, explainability, data, and user privacy concerns. However, minimal research addresses the necessity of the cyberinfrastructure needed to develop and deploy these models, which would impact the broader ecosystem. This is a policy issue I am deeply passionate about.
Despite the immense potential benefits of the AI revolution, such as improved medical diagnostics and public services, it raises critical concerns that necessitate careful policy consideration at multiple levels. One of the most pressing issues is the environmental impact of AI.

- The energy consumption required to train and deploy large AI models is substantial. For example, training a single large language model like ChatGPT (version 3) consumed approximately 5.4 GWh of energy, enough to power 504 average American households annually. This energy-intensive process involves using thousands of powerful GPUs for extended periods.
- Transitioning to an AI-enabled ecosystem necessitates deploying vast cyberinfrastructures with newer architectures capable of training and deploying large AI models like LLMs on edge devices. This increased demand for specialized hardware will likely strain raw material mining and manufacturing processes.
- As AI becomes more integrated into everyday life, from smartphones to smart homes, energy consumption will rise significantly. While training and fine-tuning

AI models are energy-intensive, deploying them on edge devices creates an even more significant challenge due to the need for widespread manufacturing and use. This growing reliance on AI raises urgent sustainability concerns.

This raises several concerns about the development and usage of AI models and their environmental impact. The problems posed by the AI revolution echo the establishment of standards by the World Wide Web Consortium (W3C) for web services, a crucial part of the digital revolution. It was founded to establish web standards focused on principles like openness and consistency, promoting unified development and deployment across platforms. Policymakers must address similar challenges in the AI domain by:

1. Implementing energy efficiency standards: Encourage the reuse of existing infrastructures until they reach the end of their useful life, gradually transitioning to more energy-efficient AI systems.
2. Establishing evaluation guidelines: Develop criteria for assessing the need for AI-enabled alternatives compared to current services.
3. Ensuring data protection: Implementing clear and comprehensive guidelines for data protection and usage in AI systems is imperative, fostering transparency and trust in the AI industry.
4. Promoting AI literacy: Create educational programs to enhance public understanding of AI and empower individuals to make informed choices about its use and how it impacts their lifestyle.
5. Establishing accountability: Develop accountability frameworks for AI-driven decisions in critical sectors.

2025

Statement of Interest

Statement of Interest

In 500 words or fewer, please answer the following questions:

- Why are you interested in being a Mirzayan Fellow?
- How do you hope this experience will contribute to your career goals?
- What unique perspectives and/or personal experiences can you bring to discussions during the Mirzayan Fellowship?
- Is there anything else you would like us to know about you?

FULL

As a Ph.D. candidate at The Ohio State University, my research is at the intersection of AI, high-performance computing (HPC), and next-generation cyberinfrastructure. I am part of the NSF-funded AI Institute for Intelligent Cyberinfrastructure with Computational Learning in the Environment (ICICLE), where my work is driven by a desire to build systems that are not only powerful but also sustainable and equitable. I am deeply interested in the Christine Mirzayan Science & Technology Policy Graduate Fellowship because it offers a unique and critical opportunity to move beyond the technical confines of the lab and engage directly with the policy landscape. My hope is to contribute my expertise in computational efficiency and AI governance to the national discourse and to learn how to translate complex scientific challenges into actionable, evidence-based policy.

My career has been shaped by a central question: “In what ways do HPC resources accelerate computationally driven science, particularly in achieving speedups that reduce time-to-science?” HPC centers at the Department of Energy's national labs have already transformed entire domains of science. In nuclear physics, for instance, large-scale hydrodynamics and neutron transport simulations have replaced underground testing, ensuring stockpile reliability entirely through computation. In climate science, the Energy Exascale Earth System Model (E3SM) integrated atmosphere, ocean, and ice dynamics at unprecedented resolution, producing projections that now directly inform global climate policy and IPCC assessments. And in health, HPC made the Human Genome Project possible and enabled rapid pandemic responses, with supercomputers screening billions of compounds against SARS-CoV-2 proteins in a matter of days—a task that would have taken years experimentally. These breakthroughs demonstrate how advanced computing has become a critical driver of discovery, accelerating progress across disciplines and fundamentally redefining how science is conducted.

Building on these foundations, my research explores a new set of questions as AI increasingly augments and replaces traditional computational methods. I ask: “How does the efficiency of a scientific workload traditionally powered by native algorithms change when restructured with AI backends?” AI is now transforming the pace and accessibility of discovery across domains. In nuclear security, surrogate AI models accelerate hydrodynamics and radiation transport, while anomaly detection systems process massive simulation outputs more efficiently. In climate science, deep learning downscales global models into high-resolution local forecasts, and AI parameterization captures processes that traditional codes cannot resolve. In genomics, AlphaFold redefined protein structure prediction, while ML-based docking and generative chemistry narrowed billions of compounds into promising therapeutics during the COVID-19 response. In engineering, reduced-order ML models replace costly fluid dynamics runs, while smart sensors enable real-time monitoring of structural health. Extending beyond laboratories, AI is also being used to predict natural disasters—from detecting earthquake precursors to forecasting floods and hurricanes—while in construction, it supports earthquake-resilient design and robotic precision building.

At the same time, I am compelled to ask: “If science is being accelerated by computational resources that are rare, expensive, and resource-intensive, how efficiently are they being used?” My analysis of a local Ohio Supercomputer Center cluster revealed striking inefficiencies: GPU utilization rarely exceeds 50% and averages around 23%, while GPU memory averages only 12%. Furthermore, inefficient provisioning—such as using powerful A100 GPUs for jobs that could run on V100s or reserving entire nodes while leaving most CPUs idle—further exacerbates this waste. These findings highlight a critical policy gap and underscore the urgent need for intelligent resource prediction and policy interventions to ensure that our nation’s most valuable computational assets are used responsibly, ethically, and equitably. My hands-on experience with this "efficiency problem" offers a unique, data-driven perspective to the fellowship, providing concrete evidence of how federal resources are being used—and misused—on the ground.

My broader motivation extends beyond resource efficiency to the societal implications of AI. I am driven by the question: “What frameworks should be designed to safely and fairly integrate AI into both scientific discovery and everyday life?” AI is now embedded in every digital realm, with every individual directly or indirectly interacting with AI systems. Human-Computer Interaction is no longer limited to robots engaging with people, but encompasses how all forms of AI—including generative and agentic systems—shape society. Like the World Wide Web Consortium (W3C), which ensured the internet remained a democratic medium through open standards, accessibility, and inclusivity, I believe AI policy must uphold the same pillars: openness, accessibility, equity, and trust.

Within ICICLE, I have realized that creating equitable frameworks alone is not enough; broader reform in education is essential. I often reflect: “If I am contributing to the advancement of AI—even with the goal of democratization—how can I do so fairly, without reinforcing harmful power dynamics?” Those with technical expertise can manipulate AI systems for gain, while those without an understanding are left vulnerable. By teaching not only scientists but also the public how AI works, we can reduce fear of the unknown, enable meaningful feedback, and create safeguards such as “no-AI modes” for critical systems.

As someone who aspires to academia, I see my responsibility as twofold: to advance technical innovation and to prepare the next generation of scientists and engineers to use AI responsibly. The Mirzayan Fellowship would be a pivotal step in this journey, providing me with the essential skills to bridge the gap between the lab and the halls of government. Working across the computing continuum—from edge devices and high-speed networks to large-scale supercomputers—I want to instill not only technical expertise but also a profound commitment to fairness, accessibility, and inclusivity. Only by combining innovation with ethical responsibility can we ensure that AI and HPC serve society broadly, rather than reinforcing inequality.

Short

As a Ph.D. candidate at The Ohio State University, my research within the NSF-funded **AI Institute for Intelligent Cyberinfrastructure (ICICLE)** is guided by a central question: *How can we design next-generation computing systems that make the pursuit of AI-driven scientific discovery both powerful and responsible?* As a computational scientist, I focus on building models and frameworks that improve efficiency and accuracy in scientific workflows. While this technical work drives measurable results, it often leaves little space to engage with the broader policies that determine how these technologies are shared, governed, and applied. The Mirzayan Fellowship offers a unique opportunity to bridge that gap by translating my technical expertise into actionable policy and deepening my understanding of the critical role scientists play in advising the nation.

My research is shaped by the transformative power of advanced computing and its convergence with AI. I ask questions such as: *How can AI, running across the computing continuum from lightweight edge devices like Raspberry Pis to cutting-edge accelerators such as Intel Broadwell, NVIDIA Hopper, or Grace Hopper chips, expand access to discovery? What tools must we create so domain scientists can fully benefit from this continuum while ensuring that AI remains democratized—open, inclusive, accessible, accountable, and fair?* HPC centers have already enabled breakthroughs that were unthinkable just a decade ago, from nuclear security to climate modeling and pandemic response. Building on these foundations, I explore how the efficiency of these workflows changes when restructured with AI. Yet, while AI accelerates science, it remains dependent on rare, expensive, and resource-intensive infrastructure. My analysis of an Ohio Supercomputer Center cluster revealed striking inefficiencies: GPU utilization rarely exceeds 50%, GPU memory averages only 12%, and CPUs are often left idle. These findings highlight a pressing policy gap: without intelligent resource management and governance frameworks, critical national infrastructure risks being underutilized and inequitably shared.

I believe my hands-on experience with this “efficiency problem” offers a unique perspective to the Mirzayan Fellowship. I can bring a data-driven understanding of how federal resources are being used—and misused—on the ground. Furthermore, my broader motivation extends to the societal implications of AI. I often reflect on the power dynamics created by these technologies, where those with technical expertise can shape systems that impact everyone, while the public is left vulnerable. This perspective has led me to a core belief that AI policy should be guided by the same principles that made the World Wide Web a democratic medium: **openness, accessibility, equity, and trust**. To achieve this, we must not only create ethical frameworks but also reform education to demystify AI for the public. A scientifically literate citizenry is essential for creating meaningful feedback loops and safeguards for AI systems.

Beyond research, I bring a deep commitment to education and mentoring. I have mentored more than twenty undergraduate, graduate, and doctoral students, guiding them through theses and projects that shaped their academic and professional trajectories. As the ICICLE NextGen representative, I also work to build a pipeline of future researchers by engaging with middle and high school students, introducing them to the potential benefits of AI and helping to remove fear of this technology. I believe that *fear of a science prevents its advancement by making it less* *accessible*, and my efforts focus on making AI approachable, inspiring curiosity, and creating trust in its responsible use.

I bring over two decades of experience in computing—from early software development in high school to nearly a decade of computational science research. I have designed and prototyped AI-centric frameworks, presented research at international HPC conferences such as **SC (Supercomputing), PEARC (Practice and Experience in Advanced Research Computing), and Science Gateways**, and collaborated with researchers from national laboratories and global summer schools. My work has been recognized with several **best paper awards**, specifically for insights into the *applicability and bottlenecks of AI from the perspective of generalizability, feasibility, explainability, and usability*. These experiences have given me both technical expertise and a strong appreciation for the societal conversations that must shape AI's future.

Through the Mirzayan Fellowship, I also hope to complement this bottom-up, researcher-driven approach with top-down perspectives on how policy is formally designed and implemented. Just as the **World Wide Web Consortium (W3C)** preserved the openness and accessibility of the internet through shared governance and transparent standards, I want to learn how policies for AI and HPC can achieve similar democratization, equity, and trust. My aspiration is to pursue an academic career where I not only advance AI research but also train the next generation of scientists and engineers to use these systems responsibly. The Fellowship will provide the policy foundation to align my technical work with societal needs, ensuring that innovation and fairness progress hand in hand.

Short Strong

As a Ph.D. candidate at The Ohio State University, my research within the NSF-funded AI Institute for Intelligent Cyberinfrastructure (ICICLE) is guided by a central question: How can we design next-generation computing systems that make the pursuit of AI-driven scientific discovery both powerful and responsible? As a computational scientist, I focus on building models and frameworks that improve efficiency and accuracy in scientific workflows. The Mirzayan Fellowship offers a unique opportunity to bridge the gap between this technical work and the broader policies that determine how these technologies are shared, governed, and applied, while deepening my understanding of the critical role scientists play in advising the nation.

My research is shaped by the growth of advanced computing and its overlap with AI. I ask questions such as: How can AI, running across the computing continuum from small devices like Raspberry Pis to accelerators such as Intel Broadwell, NVIDIA Hopper, or Grace Hopper chips, expand access to discovery? What tools must we create so domain scientists can fully benefit from this continuum while ensuring that AI remains open, inclusive, accessible, accountable, and fair? HPC centers have already enabled breakthroughs from nuclear security to climate modeling and pandemic response. Building on these foundations, I study how the efficiency of these workflows changes when restructured with AI. Yet, while AI speeds up science, it relies on rare, expensive, and resource-heavy infrastructure. My analysis of an Ohio Supercomputer Center cluster showed clear inefficiencies: GPU use rarely passed 50%, GPU memory averaged only 12%, and CPUs were often left idle. These findings show a real policy gap: without smarter resource management and governance, national infrastructure risks being wasted and unfairly shared.

I believe my hands-on experience with this “efficiency problem” gives me a useful perspective for the Mirzayan Fellowship. I can bring a data-driven view of how federal resources are being used—and misused—on the ground. This connects to my broader interest in the social impact of AI. I often think about how those with technical expertise can shape systems that affect everyone, while the public is left with little say. I believe AI policy should follow the same values that made the World Wide Web work for all: openness, accessibility, equity, and trust. To reach this, we need both clear ethical rules and better education that removes fear of AI. A public that understands science is key to building trust and strong safeguards.

Beyond research, I bring a strong focus on teaching and mentoring. I have mentored more than twenty undergraduate, graduate, and doctoral students, and as the ICICLE NextGen representative, I also work with middle and high school students to introduce the benefits of AI and reduce the fear of this field.

I have designed AI frameworks, presented research at SC, PEARC, and Science Gateways, and worked with national labs and global summer schools. My work has won several best paper awards for insights into the applicability and limits of AI in terms of generalizability, feasibility, explainability, and usability.

Through the Mirzayan Fellowship, I hope to add top-down policy experience to my bottom-up research work. My goal is to pursue an academic career where I both advance AI research and train the next generation to use these systems responsibly. The Fellowship will give me the policy background to align technical progress with fairness and shared national needs.

FInal

As a Ph.D. candidate at The Ohio State University, I am driven by a central commitment: to shape the next generation of AI-enabled computing systems so that scientific discovery is both powerful and responsible. My research at the NSF-funded AI Institute for Intelligent Cyberinfrastructure (ICICLE) focuses on developing software frameworks that enhance efficiency and accuracy in scientific workflows, with a particular emphasis on aligning technical progress with responsible innovation. The Mirzayan Fellowship is the bridge I seek to connect my technical expertise to the policy frameworks that ensure these technologies serve society's broader needs, deepening my understanding of how scientists can guide national priorities.

My research is shaped by the growth of advanced computing and its overlap with AI. I ask: How can AI, running from low-power edge devices like Raspberry Pis to high-power center architectures such as Intel Broadwell CPUs or NVIDIA Hopper GPUs, expand access to discovery? What tools will let domain scientists fully benefit from this continuum while keeping AI open, inclusive, and accountable? HPC centers have driven breakthroughs in nuclear security, climate modeling, and pandemic response. Building on this, I study how workflow efficiency changes when restructured with AI. Yet, AI depends on scarce and costly infrastructure. My review of an Ohio Supercomputer Center cluster revealed inefficiencies: GPU use rarely topped 50%, memory averaged 12%, and CPUs sat idle. These findings show a policy gap: without smarter resource management and governance, national infrastructure risks are being wasted and unfairly shared.

My hands-on experience with this efficiency issue offers a valuable perspective for the Mirzayan Fellowship. I bring a data-driven view of federal resource use that connects to my interest in AI's social impact. I consider how technical experts shape systems that affect everyone, while the public has little influence. AI policy should follow the values that made the World Wide Web broadly accessible: openness, accessibility, equity, and trust. Achieving this requires clear ethical guidelines and better education to reduce fear of AI. A public that understands science helps build trust and safeguards.

Beyond research, I focus on teaching and mentoring. I have mentored over twenty undergraduate, graduate, and doctoral students. As ICICLE NextGen representative, I also work with middle and high school students, introducing AI's benefits and reducing their fear of the field.

I have designed AI frameworks, presented at several conferences and workshops, like Supercomputing and Science Gateways, and collaborated with national labs and global summer schools. My work has earned Best Paper awards for its insights into AI's applicability and limits, especially generalizability, feasibility, explainability, and usability.

Through the Mirzayan Fellowship, I will integrate my technical research with policy engagement to advance AI in a way that aligns with fairness, shared national needs, and responsibility. My ultimate goal is to pursue an academic career that not only advances AI research but also prepares future students, researchers, and practitioners to ensure that these systems serve the public interest. The Mirzayan Fellowship will equip me to connect rigorous research with policy for maximal societal benefit.

Description of Policy Issue Interest

Description of Policy Issue Interest

In 300 words or fewer, please tell us about a policy issue you are interested in or passionate about and why you think it is an important issue. The topic does not need to align with your professional or educational experience.*

Tab 10

Policy Issue Essay (Integrated CI + AI Democratization)

The rapid integration of artificial intelligence (AI) across sectors is transforming science, industry, and society. Yet, this growth has outpaced the creation of clear operational guidelines and sustainability safeguards. I see two intertwined policy gaps: one around the efficiency and sustainability of the cyberinfrastructure (CI) that enables AI, and another around democratization and documentation that ensures the public can engage with AI responsibly.

From the CI perspective, my own research on high-performance computing (HPC) reveals inefficiencies that translate directly into wasted energy and public resources. At the Ohio Supercomputer Center, GPU utilization often falls below 50%, memory use averages only 12%, and CPUs are left idle. At the same time, training a large model like GPT-3 required an estimated 5.4 GWh of electricity—enough to power more than 500 American households annually. As AI adoption spreads to edge devices in homes, schools, and hospitals, the combined energy demand will expand exponentially. This is not only a technical issue but also a national infrastructure and sustainability concern that requires thoughtful policy.

From the democratization perspective, I believe AI-enabled systems need the equivalent of a “user manual.” Just as products come with documentation outlining usability, safety, compliance, and environmental impact, AI services and devices should provide clear information on their operation, risks, and resource use. Initiatives like the EU AI Act, IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, and OECD AI Principles have advanced conversations on trustworthiness, transparency, and privacy. But far less attention has been paid to connecting these values to the infrastructures that make AI possible, or to the basic documentation that would empower everyday users to make informed choices.

The precedent of the World Wide Web is instructive. The Web flourished because the World Wide Web Consortium (W3C) established standards rooted in openness and consistency. AI needs a similar framework—one that ties together ethical use, sustainable infrastructure, and public accessibility. Without this, the benefits of AI will remain concentrated among a few, while costs—energy consumption, e-waste, raw material mining—are borne by society at large.

Policymakers could address these challenges by:

- **Implementing energy efficiency standards** for compute infrastructure, encouraging the reuse of existing resources and incentivizing transitions to more efficient systems.

- **Establishing documentation requirements** so all AI-enabled products disclose safety, usability, privacy, and environmental impact.

- **Developing evaluation guidelines** to determine when AI-enabled alternatives are truly needed over existing solutions.

- **Ensuring accountability frameworks** for AI-driven decisions in critical sectors like health and education.
- **Promoting AI literacy** to help citizens engage with AI without fear and shape its responsible use.

This issue matters to me because I sit at the intersection of these challenges. My work builds AI frameworks that address inefficiencies in HPC, while my outreach through the ICICLE NextGen program introduces students to AI in ways that demystify the technology and build trust. I see both the promise of AI in accelerating discovery and the risks of unchecked growth. For me, the core question is not whether AI will reshape society—it already is—but whether our policies will ensure that transformation is sustainable, transparent, and inclusive.

Final

The policy issue I care most about is building transparent, sustainable AI by democratizing access to compute. By 'democratized access,' I mean ensuring that more individuals and organizations can access AI technologies via broadening eligibility for education programs, creating open platforms, or establishing equitable resource-sharing policies. AI is reshaping science, industry, and society, but its growth has outpaced policies that ensure sustainability and public trust. I see two urgent gaps: efficient national cyberinfrastructure and democratized AI access.

Training large-scale AI models consumes staggering amounts of energy, and demand is projected to triple by 2030. Yet, powerful computing centers often run far below capacity—GPUs are underutilized, memory use remains low, and CPUs are frequently idle. These inefficiencies waste taxpayer-funded infrastructure, drain electricity, and contribute to rising carbon emissions. Without stronger policies, AI's rapid expansion could strain national energy systems and worsen climate burdens.

Accessibility is just as critical. Today, only well-funded corporations and top institutions can test, validate, or challenge AI claims, leaving most researchers and the public excluded. This digital divide limits both innovation and accountability. If cyberinfrastructure resources were shared more equitably, ecologists could run biodiversity models, farmers could simulate precision agriculture outcomes, and students could gain direct experience with AI in classrooms. Public representatives and watchdog groups could also independently test AI models in sensitive areas like credit scoring, hiring, or health diagnostics, ensuring that claims of fairness and accuracy withstand scrutiny. In this way, democratized access to computational power would not only broaden participation in scientific discovery but also create essential pathways for transparency, oversight, and public trust in AI.

I believe policy should address both sides: first, by setting clear efficiency and carbon-reduction standards for compute infrastructure; second, by democratizing access to ensure AI drives collective progress and accountability.

Professional Profile

Professional Profile

Short Biography:*

Tab 12

1. I wrote my first program at the age of four, and since then computing has been both my language and my lens on the world. From building small applications in high school to designing AI-driven frameworks for national supercomputing centers, I have spent over two decades exploring how computation can accelerate discovery. Today, as a Ph.D. candidate in Computer Science and Engineering at The Ohio State University, my research focuses on making high-performance computing (HPC) and artificial intelligence (AI) more efficient, accessible, and responsible. Over the past ten years, I have developed frameworks like the HPC Application Resource Predictor (HARP) and AI-enabled schedulers that improve resource use in large-scale systems. I have published peer-reviewed papers and presented numerous posters at international conferences such as SC (Supercomputing), PEARC, and Science Gateways. My work has been recognized with best paper awards and invitations to global summer schools, where I collaborated with researchers across disciplines and cultures. Beyond research, I am deeply committed to education and community leadership. As the ICICLE NextGen representative, I have mentored more than twenty undergraduate, graduate, and doctoral students. I also created a dedicated NextGen session where students present directly to the research community, ensuring that their voices are heard in shaping the future of cyberinfrastructure. My outreach extends to middle and high school students, where I demystify AI and help them see it as a tool for creativity rather than fear.i also contribute regularly to ICICLE's AVIO newsletters, where I highlight student perspectives, project milestones, and community stories. These efforts reflect my belief that science moves forward not just through innovation, but through inclusivity, transparency, and trust. My journey weaves together research, mentorship, and global collaboration with one goal: to build AI and HPC systems that serve both science and society.
2. I am Swathi Vallabhajosyula, a Graduate Research Assistant and a 5th year Ph.D. student in Computer Science and Engineering. I am pursuing a Doctor of Philosophy degree at The Ohio State University, where my passion for academic excellence in computer science continues to thrive. Prior to embarking on my doctoral studies, I obtained a Master of Science degree in Computer Science and Engineering from The University of Minnesota, Duluth. My undergraduate education was completed at Vallurupalli Nageswara Rao Vignana Jyothi Institute of Engineering & Technology (VNRVJIET), affiliated with Jawaharlal Nehru Technological University Hyderabad (JNTU-Hyd), where I earned a Bachelor's in Technology degree in Computer Science and Engineering. This educational journey has taken me from Hyderabad, Telangana, India, to the vibrant city of Columbus, Ohio. I am gaining valuable experience as an SGX3 Summer Intern at the Texas Advanced Computing Center. In the past, I have had the privilege of serving as a Teaching Research Fellow at Vallurupalli Nageswara Rao Vignana Jyothi Institute of Engineering & Technology (VNRVJIET) in India and as a Research Intern at Oak Ridge National Laboratory, where I worked remotely. Additionally, I have had the opportunity to contribute my skills as a Web Application Developer at Spherexx.com, an IT Intern at CHART Industries, and a Software Engineer

at CA Technologies in Hyderabad, India. My diverse research, teaching, and software development background helped me build a wide range of skills and a strong work ethic. I am dedicated to my pursuits, adaptable to new challenges, and always looking for growth opportunities. Continuous learning and personal development are key values that drive me to excel in my academic and professional endeavors. During my leisure time, I enjoy immersing myself in the captivating realms of scientific fiction, thrilling mysteries, and engaging novels. As a devoted Potterhead, I cherish the magical world crafted by J.K. Rowling. To unwind and relax, I find solace in assembling jigsaw puzzles, experiencing the satisfaction of seeing the pieces fit together flawlessly. Additionally, I thoroughly enjoy the thrill of solving puzzles and unraveling mysteries in engaging escape room games.

3. Collaboration has been central to my research philosophy, helping me bridge AI and HPC innovations with domain sciences. During my Ph.D., I worked with Dr. Bryan Carstens (Dept. of Evolution, Ecology, and Organismal Biology) [https://eeob.osu.edu/people/carstens.12] to improve HPC workflows and resolve genome sequencing bottlenecks. Their lab relied on outdated allocation scripts, leading to poor HPC utilization. Building on this, I applied my tools to another domain, collaborating with the Imageomics Institute [https://imageomics.osu.edu/] and the NSF AI Institute ICICLE at Ohio State. There, I adapted HPC frameworks for computational biology workflows, AI training, and animal ecology inference. At ICICLE, I also worked with Dr. Joe Stubbs [https://tacc.utexas.edu/about/staff-directory/joe-stubbs/] at TACC on edge-to-center workflows, integrating AI pipelines for efficient field sensor data processing. My involvement with the Science Gateways Community Institute [https://sciencegateways.org/internships] further deepened my expertise in designing user-focused middleware tools to support domain scientists.These experiences taught me to listen, translate computational advances into practical tools, and build trust across disciplines. In future work, I will co-design solutions with stakeholders and foster inclusive spaces where diverse expertise thrives. I view collaboration as key to democratizing computational resources and accelerating scientific discovery.
4. During my Ph.D., one of my most impactful contributions has been the development of HARP (HPC Application Resource Predictor), an AI-driven framework designed to address the persistent challenge of poor resource utilization in high-performance computing (HPC) environments. This work originated from a simple yet critical question: “Why do we need AI to optimize resource allocation when experienced users can fine-tune it themselves?” Despite early skepticism—since resource estimation is traditionally guided by rule-of-thumb methods and considered NP-Hard—my findings on inefficiencies in HPC job behavior motivated me to explore AI-based solutions for smarter, fairer, and more efficient allocation. My solo-authored short paper on HARP received significant recognition at PEARC23 and PERAC25, winning: Phil Andrews Award for Best Paper Overall Best Paper – Systems & Systems Software Track (Overall and Student Category) Alongside publications and over a dozen poster presentations, I was awarded the Science Gateways Graduate Fellowship for three consecutive years, hosted by the Texas Advanced Computing Center (TACC) at The University of Texas at Austin—one of the leading supercomputing centers in the U.S. As part of the

NSF-funded ICICLE Institute, my research supports the development of open-source middleware that enables domain scientists to run AI-driven workflows through TAPIS-powered gateways, without requiring deep expertise in HPC or AI. This work also earned me the 2024 The OSU Graduate Student Research Award (Honorable Mention), runner-up for Best Poster at Science Gateways 2023, and Honorable Mention at the 2024 CSE Annual Graduate Poster Competition. I also serve as the ICICLE NextGen Student Group Manager, coordinating student-led sessions and outreach. Representing The OSU and ICICLE, I collaborate with OHI/O to present my work at AI exploration events in local schools, inspiring students to embrace AI with curiosity rather than fear. I often tell students: “Read, survey, and apply—use your coding skills to test your hypothesis through experiments.” As part of The OSU's research capstone courses, I mentored 10 Master's students through MS projects and 4 undergraduates in senior projects. Through this experience, I launched a summer research program in India for nearly 200 students, guiding hands-on projects and encouraging inclusive, accessible engagement with research.