Intro to HPC Bootcamp: Engaging New Communities Through Energy Justice Projects

Mary Ann Leung
Sustainable Horizons Institute
mleung@shinstitute.org

Katharine Cahill
Ohio Supercomputer Center
kcahill@osc.edu

Rebecca Hartman-Baker
Lawrence Berkeley National Lab
rjhartmanbaker@lbl.gov

Paige Kinsley
Argonne National Lab
pkinsley@anl.gov

Lois Curfman McInnes
Argonne National Lab
mcinnes@anl.gov

Suzanne Parete-Koon
Oak Ridge National Lab
paretekoonst@ornl.gov

Sreeranjani "Jini" Ramprakash
Argonne National Lab
ramprakash@anl.gov

Subil Abraham
Oak Ridge National Lab
abrahams@ornl.gov

Lacy Beach Barrier
Sustainable Horizons Institute
lbeach@shinstitute.org

Gladys Chen
Oral Roberts University
gladyschenyx@gmail.com

Lizanne DeStefano
Georgia Institute of Technology
ldestefano6@gatech.edu

Scott Feister
California State University Channel Islands
scott.feister@csuci.edu

Sam Foreman
Argonne National Lab
foremans@anl.gov

Daniel Fulton
Lawrence Berkeley National Lab
DFulton@lbl.gov

Lipi Gupta
Lawrence Berkeley National Lab
lipigupta@lbl.gov

Yun (Helen) He
Lawrence Berkeley National Lab
yhe@lbl.gov

Anjuli Jain Figueroa
DOE Office of Economic Impact and Diversity
anjuli.jainfigueroa@hq.doe.gov

Murat Keceli
Argonne National Lab
kecelim@anl.gov

Talia Capozzoli Kessler
Georgia Institute of Technology
talia.capozzoli@ceismc.gatech.edu

Kellen Leland
Oak Ridge National Lab
lelandkr@ornl.gov

Charles Lively
Lawrence Berkeley National Lab
CharlesLively@lbl.gov

Keisha Moore
Sustainable Horizons Institute
kmoore@shinstitute.org

Wilbur Ouma
Ohio Supercomputer Center
wouma@osc.edu

Michael Sandoval
Oak Ridge National Lab
sandovalma@ornl.gov

Rollin Thomas
Lawrence Berkeley National Lab
rcthomas@lbl.gov

Alvaro Vazquez-Mayagoitia
Argonne National Lab
vama@alcf.anl.gov

ABSTRACT

The U.S. Department of Energy (DOE) is a long-standing leader in research and development of high-performance computing (HPC) in the pursuit of science. However, we face daunting challenges in fostering a robust and diverse HPC workforce. Basic HPC is not typically taught at early stages of students' academic careers, and the capacity and knowledge of HPC at many institutions are limited. Even so, such topics are prerequisites for advanced training programs, internships, graduate school, and ultimately for careers in HPC. To help address this challenge, as part of the DOE Exascale Computing Project's Broadening Participation Initiative, we recently launched the *Introduction to HPC Training and Workforce Pipeline Program* to provide accessible introductory material on HPC, scalable AI, and analytics.

We describe the *Intro to HPC Bootcamp*, an immersive program designed to engage students from underrepresented groups as they learn foundational HPC skills. The program takes a novel approach

---

https://doi.org/10.22369/issn.2153-4136/5/1/10

to HPC training by turning the traditional curriculum upside down. Instead of focusing on technology and its applications, the bootcamp focuses on energy justice to motivate the training of HPC skills through project-based pedagogy and real-life science stories. Additionally, the bootcamp prepares students for internships and future careers at DOE labs. The first bootcamp, hosted by the advanced computing facilities at Argonne, Lawrence Berkeley, and Oak Ridge National Labs and organized by Sustainable Horizons Institute, took place in August 2023.

KEYWORDS

high-performance computing, training, bootcamp, energy justice, diversity, workforce development

1 WHY AN INTRO TO HPC BOOTCAMP?

The U.S. Department of Energy (DOE) employs a mission-driven team science approach to basic and applied research, and the DOE national laboratories are renowned leaders in research and development (R&D) in high-performance computing (HPC). DOE's investments have pushed the growth of computational and data-enabled science and engineering as essential drivers of scientific and technological progress, in conjunction with theory and experiment. And yet, the DOE national laboratories remain somewhat of a hidden gem; they are not well-known among the emerging talent pool. This limited visibility, coupled with urgent workforce development and training challenges, call for interventions aimed at raising awareness and engaging a broader set of people in the mission-driven team science of DOE laboratories.

1.1 DOE HPC Workforce Challenges

The DOE labs face critical HPC workforce challenges, similar to those in science and engineering generally [28], but exacerbated due to the inter- and multidisciplinary nature of the work and the reliance on an understanding of advanced and high-performance computing [12]. As stated by the DOE Advanced Scientific Computing Advisory Committee (ASCAC) Workforce Subcommittee [5],

> "All large DOE national laboratories face workforce recruitment and retention challenges in the fields within Computing Sciences that are relevant to their mission. ... Future projections indicate an increasing workforce gap and a continued underrepresentation of minorities and females in the workforce unless there is an intervention."

Addressing these workforce challenges requires broad community collaboration to change computational science's culture and profile to match the changing demographics of the future workforce. Impactful DOE-wide programs, lab-specific regional initiatives, and activities in the wider computing community are making headway. Additional work is underway within the Broadening Participation Initiative [2, 11, 30] launched by the DOE Exascale Computing Project (ECP) [9, 26], which includes three complementary thrusts: (1) Establishing an *HPC Workforce Development and Retention (HPC-WDR) Action Group*, to foster a supportive and inclusive culture in DOE labs and communities; (2) expanding the *Sustainable Research Pathways (SRP)* [27] internship and workforce development program as a multi-lab cohort of students from underrepresented groups (and faculty working with them), who collaborate with DOE lab staff on world-class HPC projects; and (3) building an *Introduction to HPC Training and Workforce Pipeline Program* to provide accessible introductory material on HPC, scalable AI, and analytics to a broader audience. The Intro to HPC Bootcamp is one facet of the third thrust.

1.2 HPC Training Challenges

The idea of engaging a broader group of people earlier in their academic careers compelled us to rethink the existing DOE HPC training curricula. Broadening participation in this context translates to several dimensions of diversity: demographic, academic discipline, generational, and others. The challenge is in reaching new communities who have little exposure to HPC and may not know why or how HPC can address issues they care about. The future workforce will be comprised of Millennials, Gen Z, women, people of color, and others, who strongly value having a positive social impact on their communities. Further complicating the training challenge, the basics of HPC are not typically taught at the early stages of students' careers, and the capacity and knowledge of HPC at many institutions are limited [13, 15, 32]. Even so, such topics are prerequisites for advanced opportunities such as internships, the Argonne Training Program for Extreme-Scale Computing [1], the DOE Computational Science Graduate Fellowship Program [8], and ultimately for careers in HPC.

1.3 Raising Awareness and Engaging a Broader Set of People in Mission-driven Team Science

Thus, introductory HPC training must employ a fundamentally different approach that addresses the values and needs of this broader audience. The Intro to HPC Bootcamp's novel approach engages students in energy justice projects, where they learn HPC fundamentals using culturally relevant pedagogy, while gaining exposure to exciting career opportunities in computing sciences at DOE national labs.

We envision the Intro to HPC Bootcamp as an accessible on-ramp for students early in their academic careers, who will subsequently be prepared to complete internships and other HPC training offerings at the DOE labs and continue to graduate school for eventual careers in the computing sciences. As illustrated in Figure 1, the following sections describe our goals and approach, diversifying team science, bootcamp innovations, and related topics.

2 GOALS AND APPROACH

The first Intro to HPC Bootcamp [19], hosted by the advanced computing facilities at Argonne, Lawrence Berkeley, and Oak Ridge National Labs and organized by Sustainable Horizons Institute, was held August 7–11, 2023 at Lawrence Berkeley National Lab. Sixty students [20] worked in groups supported by fourteen trainers [22] and ten peer mentors [21] on seven energy justice projects [23] that explored issues related to the social impact of climate risk and resilience, solar power, sustainable cities, and energy usage. The bootcamp goals focused on engaging new communities in HPC in order to:

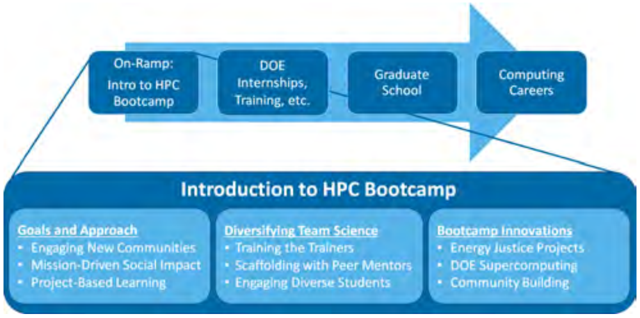

Figure 1: Overview of the Intro to HPC Bootcamp.

- Raise general awareness of the value, benefits, and rewards of work in HPC, especially at DOE labs;
- Engage more people from historically underrepresented groups, including those with no (or limited) HPC exposure;
- Include a broader set of disciplines in the HPC community; and
- Increase the available talent to work in HPC and computing sciences, where we have urgent needs for skilled researchers, research software engineers, systems administrators, and more.

Thus, as shown in Figure 1, the bootcamp can be considered as an accessible on-ramp to HPC training, internships, and ultimately careers in the computing sciences at DOE national labs.

2.1 Engaging New Communities

Changing the HPC workforce culture and profile to match the evolving demographics of the emerging workforce requires intentional approaches to engaging new communities and addressing the current state of very low numbers of people from underrepresented groups. In addition, many people in the developing workforce are driven by a desire to have a positive social impact on their communities [7]. The combination of these factors suggests employing culturally relevant pedagogy, including: creating a safe learning space, fostering a sense of belonging, challenging and engaging learners in problem solving, providing checkpoints and proctors, and enabling students to see where they are going throughout the learning process. This approach requires trainers to be intentional about building trust with learners—sharing their experiences and providing examples of breakthrough science accomplishments and projects that are led or contributed to by researchers with diverse backgrounds and demographics.

We therefore designed the Intro to HPC Bootcamp to provide a culturally relevant curriculum through social impact projects that are driven by DOE's mission and help to build foundational skills in HPC, scalable AI, and analytics. Accomplishing this required turning the traditional HPC curriculum upside down. Instead of teaching technology and its applications, the bootcamp uses project-based learning to engage students to answer social impact questions through energy justice projects while gaining hands-on experience using state-of-the-art computational and data science tools and techniques.

As shown in Figure 2, the five-day bootcamp framework begins by building community and establishing a friendly learning environment, motivating learners through challenging problems in science and society, and then introducing foundations of computation, HPC, and mathematics, while exposing students to the exciting R&D in high-performance computing at DOE labs. Capstone projects are introduced on the first day, and learners work throughout the week in small groups toward a final presentation at the conclusion of the bootcamp. The bootcamp agenda and plenary lectures are available via the bootcamp website [19]. This approach is motivated by the programs Advanced Computing for Social Change [14] and Computing4Change [6].

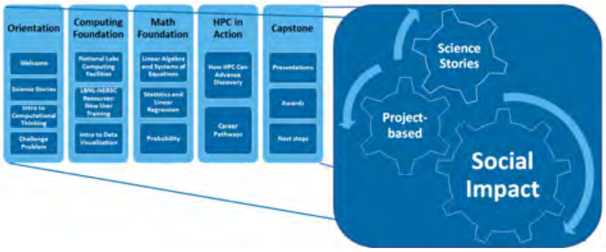

Figure 2: Framework for engaging new communities in high-performance computing: Solving problems with mission-driven social impact.

2.2 Mission-driven Social Impact

Research and development throughout DOE national laboratories are driven by the mission of ensuring America's security and prosperity by addressing its energy, environmental, and nuclear challenges through transformative science and technology solutions [39]. The topic of *energy justice* is an umbrella for a wide range of work with compelling social impact. The topics of projects in the Intro to HPC bootcamp were inspired by the Justice40 Initiative [25], whose implementation at DOE is led by the DOE Office of Economic Impact and Diversity [10]. By applying the tenets of justice (procedural, distributional, recognition, and restorative justice), energy justice [24] can promote policies that center community needs like reducing energy burden, avoiding disproportionate environmental impacts, ensuring equitable distribution of benefits of energy generation, creating opportunity for reliable access to clean energy, and encouraging democratic participation and enterprise creation in the energy system. DOE's Justice40 framework outlines eight such policy priorities that can better address the energy systems needs of overburdened communities.

Energy justice recognizes that the benefits of energy technologies have not been equally distributed across all Americans, often leaving out Black, Brown, Indigenous, and low-income communities. Underlying structural inequalities have resulted in development projects with higher rates of pollution, lower adoption of renewables, negative health impacts, and a higher energy burden in these communities. Energy justice is a complex issue with economic, racial, geographic, and social implications covering issues from energy affordability and access to infrastructure development. As we embark on an energy transition demanded by climate change concerns and fairness, we need to incorporate energy justice in the earliest stages of R&D and workforce development to enable a more just technology development and deployment.

2.3 Project-based Learning

Project-based learning in computational science [34] engages students through an active educational approach focused on real-world problems. It involves collaborative project work where students utilize computational tools to address complex challenges, fostering a deeper understanding of both subject matter and practical applications. This method introduces authentic issues mirroring research, industry, and societal contexts, motivating students by demonstrating the tangible impact of their skills. Project-based learning emphasizes critical thinking, problem-solving, and active participation—empowering students to drive their own investigations, make decisions, and take ownership of projects, thereby fostering a sense of agency and responsibility. Collaborative teamwork is highlighted, mirroring the collaborative nature of scientific research. Students gain proficiency in using computational tools through hands-on projects, with the iterative process fostering skills in inquiry, reflection, and resilience. Instructors and mentors offer personalized guidance, enhancing the learning experience, while presentations and communication of project outcomes refine students' ability to convey technical concepts. Overall, project-based learning in computational science equips students with practical skills, nurturing curiosity, critical thinking, and readiness to contribute to scientific and technological advancements.

3 DIVERSIFYING TEAM SCIENCE

Numerous studies have shown that diverse organizations, teams, and communities perform more creatively and effectively—and thus are demonstrably more innovative and productive [18, 33]. In the case of the Intro to HPC Bootcamp, diversity of team science took on a variety of forms. We employed a diverse set of trainers, organizers, peer mentors, and students on several dimensions of diversity, as described below.

Collaboration on the Intro to HPC Bootcamp is a partnership among experts in advanced computing, computational science, energy justice, workforce development, education, training, social science, and program evaluation. With collaborators from multiple DOE national laboratories, Sustainable Horizons Institute, the DOE Office of Economic Impact and Diversity, and academic partners, the team's breadth of experience has been essential for devising this first-of-a-kind program to introduce HPC at DOE labs in the context of mission-centered energy justice projects, with emphasis on engaging early-career students from underrepresented groups. The diversity of the bootcamp trainers, project leaders, and organizers is another facet of the diversity of this effort.

3.1 Training the Trainers

A pivotal step in developing the bootcamp program was hosting a "Train the Trainers" workshop in spring 2023, which served as a platform for promoting effective team collaboration. Through two days of in-person interactions, the workshop fostered strong interrelationships among the bootcamp leadership team and introduced the bootcamp concept to project leaders and trainers.

Workshop participants discussed the foundational principles of energy justice and inclusivity, enriched by a panel of organizers and mentors with extensive experience in social justice-focused computational science workshops. A key focus of the workshop was training the trainers in the art of modifying existing HPC training materials to suit the unique requirements of the bootcamp pedagogy, while also imparting effective coaching techniques to ensure students' successful engagement with the bootcamp's social challenge projects. The forum provided an opportunity for the team to deliberate on anticipated outcomes of the bootcamp, delineating the essential takeaways for students, along with post-bootcamp opportunities for their continued growth. Notably, the workshop stimulated the generation of innovative HPC project ideas aligned with the theme of energy justice, laying the foundation for a series of compelling energy justice projects to be explored in the bootcamp.

3.2 Scaffolding with Peer Mentors

Inspired by the approach used in the Advanced Computing for Social Change workshops, we engaged peer mentors [21] to provide support to students during the bootcamp. Peer mentors addressed technical questions, while also providing guidance on collaboration, presentations, and workshop expectations. Their contributions to the bootcamp were considerable, from providing a manual for mentors to use as a guide, to being a conduit of information between the students and the organizers. For example, peer mentors helped to foster a supportive and responsive environment by conveying to the project leaders and trainers questions that arose in conversations with students. The peer mentors facilitated discussions and tutorial sessions, while working closely with trainers and project groups. Most peer mentors were invited to apply to be a mentor, having had previous mentoring experiences; others were invited from the student applicants.

3.3 Engaging Diverse Students

Recruiting a diverse set of bootcamp participants early in their academic careers was central to the Intro to HPC Bootcamp success. Following is a description of the student profiles and background as well as our methods for recruiting and engaging them.

Demographics. The bootcamp partcipants were comprised of a highly diverse group of students on several dimensions. As shown in Figure 3, almost sixty percent of the sixty bootcamp participants were female, nearly a third were Hispanic/Latinx, thirteen percent had a disability, and nearly half were first-generation scholars (the first in their family to attend college). The racial distribution of students, also shown in Figure 3, illustrates an unusual representation for HPC—with Black or African Americans at twenty-eight percent (the largest racial group), followed by twenty percent Asian, twenty percent multiple race, and eighteen percent Caucasian.

As shown in Figure 4, the student academic profile included nearly eighty percent undergraduates, fourteen percent masters students, five percent students from community colleges, and just three percent doctoral students. Institution types ranged from eight Historically Black Colleges and Universities (HBCU), twenty Hispanic Serving Institutions (HSI), ten Asian American Native American Pacific Islander Serving Institutions (AANAPISI), three community colleges, thirteen liberal arts colleges, twenty-nine public/state universities, and thirteen high research institutions. Note that some institutions are counted in multiple types. In terms of academic focus, while approximately forty percent of students are studying

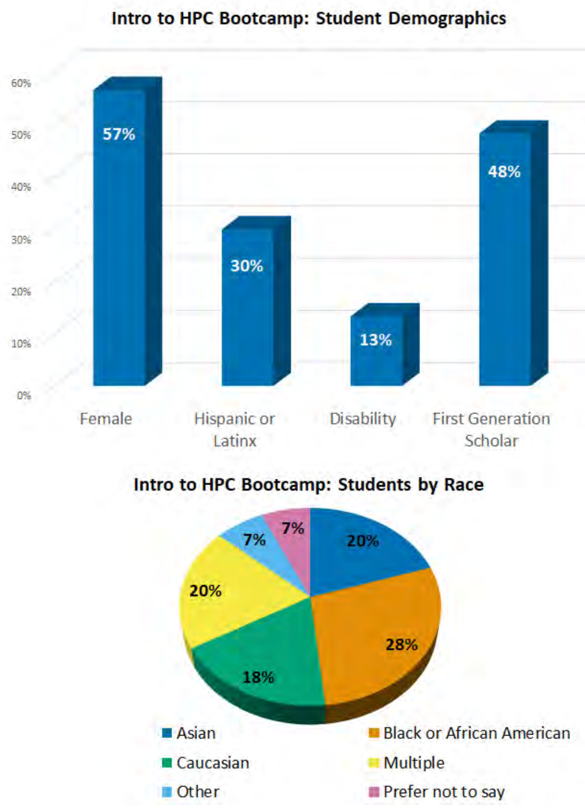

Figure 3: Demographics of bootcamp participants.

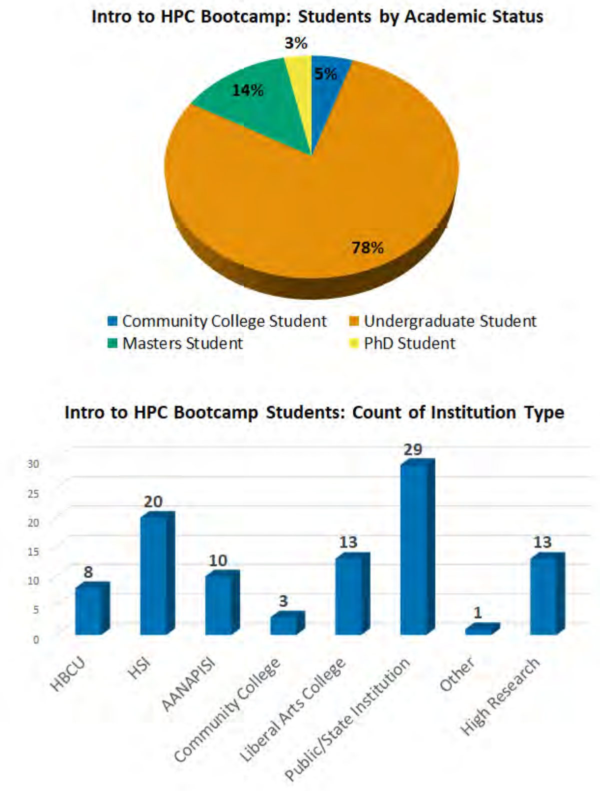

Figure 4: Academic status and institution type of bootcamp participants.

computer science, Figure 5 shows a wide distribution of domains in science, engineering, and mathematics.

Student recruitment. We achieved this demographic profile by broadly recruiting students from underrepresented groups to apply to the bootcamp, emphasizing the unique opportunity to advance understanding of *both* energy justice and HPC, while learning about career opportunities at DOE national labs. Although applications did not require CVs or letters of recommendation (with the goal of removing barriers for students to apply), applicants were asked to provide information about their motivation for participating in the bootcamp and their experience in computing, including R, Python, and HPC. This approach appears to have strongly resonated with the target community, as we received applications from several hundred students at US-based institutions, even though we could accommodate only sixty learners in this pilot session.

HPC skills. Typical skills needed by students to become interested and involved in HPC-based research include subsets of any number of the following: experience working with a programming language, exposure to parallel programming models and parallel thinking, exposure to techniques for visualizing and managing large data sets, experience with machine learning, and exposure to computational science. Requirements for the bootcamp were some experience in computing *and* an interest in energy justice topics. Because the bootcamp focused on introducing HPC, we did not expect applicants to have significant prior HPC exposure. We were able to engage a few applicants whose HPC skills exceeded the program scope as *peer mentors* in the bootcamp.

Student travel and stipend. The bootcamp fully supported travel, lodging, and meals for participants in the 5-day program. Participants received a modest stipend of $500 after completion of the bootcamp. We consider a stipend to be essential—making it possible for students who truly need the earnings from their usual jobs to sacrifice a week's employment in order to attend the bootcamp; without the stipend, some students could not participate.

4 BOOTCAMP INNOVATIONS

The centerpiece of the Intro to HPC Bootcamp was tackling energy justice projects using DOE supercomputing, while building community among the students, peer mentors, and lab staff.

4.1 Energy Justice Projects

Throughout the bootcamp, students worked in twelve small groups on projects, led by trainers and peer mentors. Bootcamp projects [23] were developed to teach HPC and AI tools in the context of wide-reaching energy justice problems, an example of our intentional design of melding social impact with technical skills development.

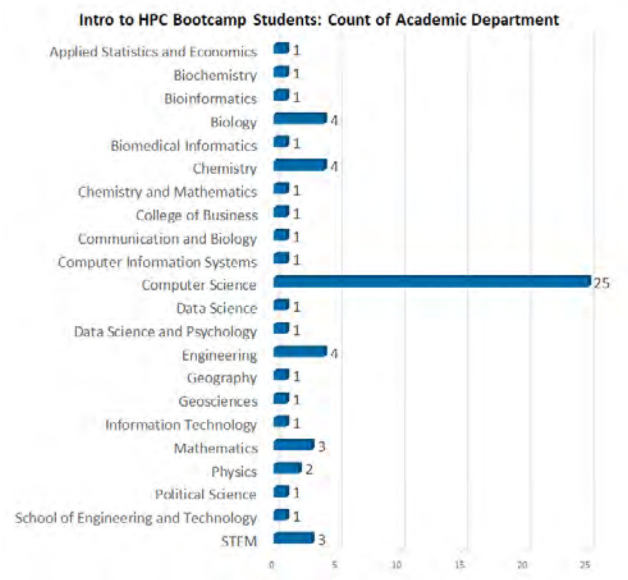

Figure 5: Academic departments of bootcamp participants.

The following are key themes in social impact and HPC technologies for the seven bootcamp projects.

- **AI-powered equity analysis of renewable energy**
  - *Social*: Equitable energy law, accessible energy justice info
  - *Tech*: Machine learning, large language models, web scraping
- **Energy justice analysis of climate data**
  - *Social*: Geographic climate disparities, variable correlations
  - *Tech*: Data visualization, big data wrangling
- **Solar power for affordable housing through computational design of low-cost/high-efficiency solar cells**
  - *Social*: Sustainable solar cells, renewable energy optimization
  - *Tech*: Machine learning, big data analysis
- **Energy cost for disadvantaged populations and methods of energy efficiency and optimization in computing systems**
  - *Social*: Power sustainability in HPC, equitable energy affordability and distribution
  - *Tech*: Top500 and Green500 energy efficiency ratings, HPC resource utilization
- **Understanding the impact of HPC center energy usage on low-income and minority populations**
  - *Social*: Power sustainability in HPC, equitable energy affordability and distribution
  - *Tech*: Code optimization, Top500 energy efficiency ratings, strategies for HPC resource utilization and allocation
- **Power outages and inequities in energy access for medically vulnerable populations**
  - *Social*: Weather impact on energy access, energy access for medically vulnerable populations
  - *Tech*: Data visualization, parallel wrangling of large data sets
- **Socioeconomics of power outages and heatwaves**
  - *Social*: Power outage data, weather impact on energy access, discrepancies in weather event impact across populations
  - *Tech*: Data visualization, statistical analysis of large data sets, MPI for python

To address the energy justice questions posited by the projects, bootcamp participants employed publicly available data sets, including climate data for national lab projects such as ClimRR [4] and EAGLE-I [35], HPC Top500 and Green500 power consumption statistics [16, 37], a dye-sensitized solar-cell device database [3], the HHS emPOWER Program [17], and U.S. census data [38].

4.2 DOE Supercomputing

Participants received access to the NERSC supercomputer Perlmutter [29], ranked 8th on the June 2023 Top500 List [36], providing them a suite of powerful computational tools to work on their energy justice projects. In contrast to training accounts usually provided for user training events, NERSC created a regular allocation project using the Director's Reserve Pool, so that students could continue Perlmutter access through the end of the current allocation year (until mid-January 2024). Students received detailed instructions about applying for a Perlmutter account and setting up multi-factor authentication. Multiple office hours were held prior to the bootcamp to help resolve any account and login issues, so that students were ready to work on projects immediately during the bootcamp. The students also had the opportunity to view Perlmutter in action during a tour of NERSC. Due to the large size of the student group and the noise level inside the machine room, a presentation about the NERSC facility and machine room took place before the tour, and an engaging Q&A session took place after the tour.

Students largely employed Jupyter Notebooks for bootcamp projects. NERSC staff prepared a bootcamp-specific kernel for use with the Notebooks, thereby streamlining access to necessary Python and other data analytics and plotting packages. GPU compute node reservations were made in advance; usage in a shared mode by multiple students was integrated smoothly upon launching the Notebooks.

A single GitHub repository provided easy access to all bootcamp project materials and presentations. The data sets and Jupyter Notebooks for bootcamp projects also were made available on the Perlmutter file system in each project's workspace.

4.3 Community Building

Community building served as an important theme and thread throughout the bootcamp, with the aim of supporting the goals of culturally relevant pedagogy: providing a safe learning environment, fostering a sense of belonging, and helping participants envision themselves not just "fitting in" but, also developing productive, rewarding careers.

Staff and leaders of DOE national labs spoke to the students about their career paths and motivations, oftentimes describing some of their challenges and nonlinear pathways. During a DOE staff panel, students learned about research applications of HPC and engaged in questions/answers about HPC careers.

Networking and relationship-building opportunities were built in formally and informally throughout the bootcamp, giving participants opportunities to connect with others sharing common interests in computing, energy justice, career paths, and other topics. Students met with other participants, with peer mentors who are a little further along in their academic journeys, and with DOE lab staff who have careers in HPC. During lunch each day, students and invited speakers discussed topics ranging from energy justice to whether to pursue a Ph.D., to work-life balance.

On the final bootcamp day, each of the twelve project groups presented an overview of the team, project goals, approach, insights gained during the bootcamp, and ideas for future work. Each bootcamp participant provided feedback to other groups on their presentation, and lively questions and discussion followed each presentation.

5 EVALUATION INSIGHTS

Bootcamp participants were surveyed both before and after the bootcamp by an external evaluator, employing instruments developed with a utilization-focused evaluation design [31]. The pre-survey, a formative tool to assist in pre-bootcamp planning, was distributed electronically to 60 participants, who all responded (100 percent). The post-survey included both qualitative and quantitative items related to participants' satisfaction with the bootcamp, as well as the perceived impact on participants' technical skills, understanding of energy justice, and future research or career attainment. The post-survey was distributed electronically to 60 participants at the conclusion of the bootcamp, and 54 responded (90 percent).

Preliminary feedback indicates that by focusing on project-based HPC for energy justice topics, the bootcamp captured the interest of a wide range of students from underrepresented groups, exposing them to the power of applying HPC to challenges in science and society. Further, participants indicated interest in pursuing a career in HPC or energy justice (Table 1). The bootcamp also introduced students to the impactful HPC research underway at DOE labs, providing pathways for future internships, education, networking, and more; at the conclusion of the bootcamp, participants indicated interest in pursuing a career at a DOE lab (Table 1).

Table 1: Interest in Careers

Key: 1: Not interested at all / 2: Somewhat interested / 3: Neutral
4: Interested / 5: Very interested

How interested are you in ...	1	2	3	4	5	Mean	SD
continuing to pursue HPC in a future career	1 (2%)	3 (6%)	4 (7%)	21 (39%)	25 (46%)	4.22	0.95
energy justice in a future career	1 (2%)	5 (9%)	8 (8%)	20 (37%)	20 (37%)	3.98	1.04
a career at a DOE lab	1 (2%)	1 (4%)	2 (4%)	16 (30%)	33 (61%)	4.44	0.88

In preliminary post-survey feedback, participants referenced the wide range of participant proficiency in technical fields such as coding, and some indicated it presented a challenge during groupwork. In order to balance the strengths and weaknesses of each group, participants suggested that groups should be intentionally comprised of members that capture a variety of skillsets. Additionally, participants identified time constraints as a challenge to the bootcamp, sharing that more time would have been preferable to complete their project. As some participants suggested, condensing the bootcamp lectures or expanding the length of the bootcamp, in general, would allow for more extensive groupwork.

6 LESSONS LEARNED AND NEXT STEPS

This pilot *Intro to HPC Bootcamp* represents our first iteration of an innovative approach to HPC training—emphasizing mission-driven social impact using project-based pedagogy and real-life science stories to prepare students for internships and future careers at DOE labs. Others could adapt and customize this model according to their specific needs and communities.

Looking ahead, the team is working to refine the bootcamp model and expand its reach through other modalities such as academic curricula and online learning, especially targeting students from underrepresented communities. At a high level, these measures could include, but are not limited to:

- **Partnering with faculty members:** To develop course modules for flexible use in campus settings, including through integration with subject-matter courses in the physical and social sciences
- **Creating an asynchronous online learning component:** To accommodate various schedules and learning styles and supplement the synchronous team and project-based components
- **Providing multiple offerings:** Working toward offering the bootcamp multiple times per year, considering regional emphasis
- **Broadening involvement of DOE lab staff:** To enhance the depth and breadth of expertise available to participants.
- **Offering training to progressively build HPC skills:** To bridge the gap for beginners (such as bootcamp alumni) and prepare them for more advanced HPC training and internships

Partnerships with faculty at colleges and universities, especially minority-serving institutions without on-site HPC research staff, offer a scalable avenue to connect with diverse students. We are piloting this strategy, and plans to expand are under way.

Acknowledging the diverse commitments and schedules of potential participants, we are considering a hybrid approach including an asynchronous learning component to complement the synchronous team-oriented core activities and prepare students who have less experience with programming. While asynchronous offerings on their own do not naturally support the community building aspect of the bootcamp, we believe that a a hybrid approach could enable participants to access introductory content at their own pace, while also engaging synchronously for group work on HPC projects and making personal connections with peers and lab staff.

To meet the urgent demand for HPC training and engagement, we are working toward offering the bootcamp multiple times per year, considering enhancements such as regional emphasis and avenues for more local lab involvement. Also, development of additional training modules would enable pathways for bootcamp alumni and others to progressively build HPC skills. Goals include reaching more participants and building a sustainable pipeline of talent for DOE national labs.

As we look ahead, these strategic directions will guide us in furthering the impact and accessibility of the Intro to HPC Bootcamp, as a core pillar of the ECP Broadening Participation Initiative. By embracing innovative approaches, emphasizing flexibility, and increasing our offerings, we aspire to cultivate a workforce that reflects the diversity of our society and leverages the power of HPC to address critical challenges in energy justice and beyond.

ACKNOWLEDGMENTS

We sincerely thank Linda Akli for insightful discussion about training strategy. We thank all Intro to HPC Bootcamp participants and peer mentors for their contributions to the program.

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC award DDR-ERCAP0025729.

REFERENCES

[1] Argonne Training Program on Extreme-Scale Computing (ATPESC) 2023. https://extremecomputingtraining.anl.gov.

[2] A. Barker, D. Martin, M. A. Leung, L. C. McInnes, J. Ramprakash, J. White, J. Ahrens, E. Draeger, S. Fomby, Y. Ghadar, R. Gupta, M. Halappanavar, M. Heroux, C. Kelly, M. Miller, H. A. Nam, S. Peles, D. Rouson, D. Turner, T. Turton, D. Brown, and V. Taylor. 2021. A multipronged approach to building a diverse workforce and cultivating an inclusive professional environment for DOE high-performance computing. Response to the Stewardship of Software for Scientific and High-Performance Computing RFI, https://doi.org/10.6084/m9.figshare.17192492.

[3] Edward Beard and Jacqueline Cole. 2022. Dye-sensitized Solar Cell Database. (2022). https://doi.org/10.6084/m9.figshare.13516220.v1

[4] Center for Climate Resilience and Decision Science (ClimRR) 2023. ClimRR Data Catalog. https://disgeoportal.egs.anl.gov/ClimRR/?page=Data-Catalog.

[5] B. Chapman, H. Calandra, S. Crivelli, J. Dongarra, J. Hittinger, S. Lathrop, V. Sarkar, E. Stahlberg, J. Vetter, and D. Williams. 2014. DOE Advanced Scientific Advisory Committee (ASCAC): Workforce Subcommittee Letter. https://dx.doi.org/10.2172/1222711.

[6] Computing 4 Change 2021. https://www.sighpc.org/for-our-community/computing4change.

[7] Deloitte 2021. The Deloitte Global 2021 Millenial and Gen Z Survey. https://www2.deloitte.com/content/dam/Deloitte/mk/Documents/about-deloitte/2021-deloitte-global-millennial-survey-report.pdf.

[8] DOE Computational Science Graduate Fellowship (CSGF) Program 2023. https://www.krellinst.org/csgf.

[9] DOE Exascale Computing Project (ECP) 2023. https://www.exascaleproject.org.

[10] DOE Office of Economic Impact and Diversity 2023. https://www.energy.gov/diversity/office-economic-impact-and-diversity.

[11] ECP Broadening Participation Initiative 2023. https://www.exascaleproject.org/hpc-workforce.

[12] Roscoe Giles et al. 2020. Transforming ASCR after ECP. https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/202004/Transition_Report_202004-ASCAC.pdf.

[13] Scott Feister and Elizabeth Blackwood. 2022. HPC Workforce Development of Undergraduates Outside the R1. *The Journal of Computational Science Education* 13 (Dec. 2022), 8–11. Issue 2. https://doi.org/10.22369/issn.2153-4136/13/2/2

[14] K. Gaither, R. Gomez, L. Akli, R. Mendenhall, M. Bland, S. Fratkin, L. Rivera, and L. DeStefano. 2017. Advanced Computing for Social Change: Educating and Engaging Our Students to Compete in a Changing Workforce. In *PEARC17: Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact*. https://doi.org/10.1145/3093338.3093391

[15] Steven I. Gordon and Katharine Cahill. 2020. The State of Undergraduate Computational Science Programs. *The Journal of Computational Science Education* 11 (April 2020), 7–11. Issue 2. https://doi.org/10.22369/issn.2153-4136/11/2/2

[16] Green500 List 2023. https://www.top500.org/lists/green500/2023/06/.

[17] HHS emPOWER Map 2023. Medicare At-Risk Populations by Geography. https://empowerprogram.hhs.gov/empowermap/.

[18] Dame Vivian Hunt, Dennis Layton, and Sara Prince. 2015. *Why diversity matters*. Technical Report. McKinsey and Company.

[19] Intro to HPC Bootcamp 2023. https://shinstitute.org/introduction-to-high-performance-computing-bootcamp.

[20] Intro to HPC Bootcamp: Participants 2023. https://shinstitute.org/intro-to-hpc-participants.

[21] Intro to HPC Bootcamp: Peer mentors 2023. https://shinstitute.org/intro-to-hpc-peer-mentors.

[22] Intro to HPC Bootcamp: Project leaders and trainers 2023. https://shinstitute.org/intro-to-hpc-project-leaders-and-trainers.

[23] Intro to HPC Bootcamp: Projects 2023. https://shinstitute.org/intro-to-hpc-energy-justice-projects.

[24] Kirsten Jenkins, Darren McCauley, Raphael Heffron, Hannes Stephan, and Robert Rehner. 2016. Energy justice: A conceptual review. *Energy Research and Social Science* 11 (2016), 174–182. https://doi.org/10.1016/j.erss.2015.10.004

[25] Justice 40 Program, DOE Office of Economic Impact and Diversity 2023. https://www.energy.gov/diversity/justice40-initiative.

[26] D. Kothe, S. Lee, and I. Qualters. 2019. Exascale Computing in the United States. *Computing in Science and Engineering* 21, 1 (2019), 17 – 29. https://doi.org/10.1109/MCSE.2018.2875366.

[27] M. A. Leung, S. Crivelli, and D. Brown. 2019. Sustainable Research Pathways: Building Connections across Communities to Diversify the National Laboratory Workforce. in Collaborative Network for Engineering and Computing Diversity (CoNECD), Washington, D.C..

[28] National Center for Science and Engineering Statistics (NCSES) Directorate for Social, Behavioral and Economic Sciences, National Science Foundation. 2023. Diversity and STEM: Women, Minorities, and Persons with Disabilities. https://ncses.nsf.gov/wmpd.

[29] NERSC Perlmutter System 2023. https://www.nersc.gov/systems/perlmutter/.

[30] S. Parete-Koon, M.A. Leung, J. Ramprakash, and L.C. McInnes. 2022. Exascale Computing Project's Broadening Participation Initiative. In *Ninth SC Workshop on Best Practices for HPC Training and Education (with SC22)*.

[31] M. Q. Patton. 2000. Utilization-focused Evaluation. 49 (2000). https://doi.org/10.1007/0-306-47559-6_23

[32] Rajendra K. Raj, Carol J. Romanowski, John Impagliazzo, Sherif G. Aly, Brett A. Becker, Juan Chen, Sheikh Ghafoor, Nasser Giacaman, Steven I. Gordon, Cruz Izu, Shahram Rahimi, Michael P. Robson, and Neena Thota. 2020. High Performance Computing Education: Current Challenges and Future Directions. In *Proceedings of the Working Group Reports on Innovation and Technology in Computer Science Education (ITiCSE-WGR '20)*. Association for Computing Machinery, New York, NY, USA, 51–74. https://doi.org/10.1145/3437800.3439203

[33] David Rock and Heidi Grant. 2016. Why Diverse Teams Are Smarter. *Harvard Business Review* (2016).

[34] Namsoo Shin, Jonathan Bowers, Joseph Krajcik, and Daniel Damelin. 2021. Promoting computational thinking through project-based learning. *Discip Interdscip Sci Educ Res* 3, 7 (2021). https://doi.org/10.1186/s43031-021-00033-y

[35] Varisara Tansakul, Aaron Myers, Sarah Tennille, Matthew Denman, Alec Hamaker, Jonathan Huihui, Kyle Medlen, Karl Allen, Daniel Redmon, Supriya Chinthavali, Mark Coletti, Joshua Grant, Matt Lee, Dakotah Maguire, Scott Newby, Chelsey Stahl, Budhu Bhaduri, and Jibo Sanyal. 2023. EAGLE-I Power Outage Data 2014 - 2022. https://doi.org/10.13139/ORNLNCCS/1975202

[36] Top500 List 2023. https://www.top500.org/lists/top500/2023/06.

[37] Top500 Statistics 2023. https://www.top500.org/statistics/.

[38] United States Census 2019. County Population Totals: 2010-2019. https://www.census.gov/data/datasets/time-series/demo/popest/2010s-counties-total.html.

[39] U.S. Department of Energy 2023. https://www.energy.gov/about-us.