CASC

Coalition for Academic Scientific Computation

Safe, Smart, Secure.

Research Computing to Inspire Innovation

Contents

4 Letter from the Chair

6 Leveraging Technology to Promote Public Health

8 Creating a Smarter, Safer Internet

10 Building Resilience in a Changing Climate

12 Inspiring Innovation in Energy and Materials

14 Fascinating Frontiers

16 Navigating Uncertain Ground

17 Mobilizing a New Generation

18 CASC Members

Our Mission

Advocate for public and private investment in research computing and data services.

Advise federal agencies on relevant funding programs.

Engage in policy discussions on research computing and data services.

Foster a diverse community of leaders in this field.

Provide a forum for sharing strategic ideas and best practices.

About the Cover

The image shows a visualization of blood flowing through a brain aneurysm. Aneurysms, which are weakened bulges in blood vessel walls, can become life-threatening if they rupture. A better understanding of the blood flow dynamics within aneurysms could enable early detection and treatment, preventing serious complications.

For this work, Trung Le and Tam Nguyen from North Dakota State University used a data-driven machine-learning technique called dynamic mode decomposition to quantify the impact of the inflow jet — a strong stream of blood that flows into the aneurysm — in 3D computer models of aneurysms based on real patient data. For each aneurysm modeled, they observed signs that the blood flow transitioned from laminar to turbulent, a significant finding since long-term turbulence is a key contributor to rupture. Additionally, the researchers found that dynamic mode decomposition could potentially be applied to medical imaging techniques such as MRI or CT scans, suggesting the approach could be used for clinical applications in the future. The simulations, visualizations, and data analyses for this research were performed using the computational resources at the North Dakota State University Center for Computationally Assisted Science and Technology.

Trung Le, North Dakota State University
Copyright Trung Le, North Dakota State University

Tam Nguyen, North Dakota State University
Copyright Tam Nguyen, North Dakota State University

Dear Colleagues and Friends of CASC,

As I reflect on the past year, I am filled with gratitude for the incredible strides we've made as a community. Gratitude because this progress would not be possible without the efforts of our members, the Executive Committee, and our Executive Director and Operations Manager. Since its inception in 1989, the Coalition for Academic Scientific Computation (CASC) has convened leaders, supporters, and practitioners of research computing to address emerging challenges and seize opportunities. Today our network spans over 105 member institutions, and we are leveraging our individual and collective efforts to build a brighter future through scientific discovery, scholarly research, innovation, and collaboration.

For those new to CASC, our annual brochure highlights the impactful work happening across our member institutions. These pages capture the transformative role research computing and data services play in advancing science and technology. Through our collective efforts, we've made significant contributions to the research ecosystem, helping to shape the future of scientific inquiry.

This past year, CASC and our members have reached some remarkable milestones:

Expanding Access to Cutting-Edge Resources: CASC has successfully partnered with federal agencies and private sector leaders to secure funding and infrastructure, significantly enhancing access to state-of-the-art computing resources. These advancements have empowered researchers to make breakthroughs in critical areas such as climate modeling and genomics. Notably, 14 of the 25 National Artificial Intelligence Research Institutes funded by the National Science Foundation (NSF) are led by CASC members, and several members are providing shared resources and applications support for the National Artificial Intelligence Research Resource (NAIRR) pilot project.

Influencing National Policy and Strategy: CASC has been instrumental in shaping national discussions on the future of scientific computation. Our members serve on key advisory boards, contributing to policy recommendations that underscore the importance of sustained investment in advanced computing. Highlights include participating in the White House Office of Science and Technology Policy forum on the "Power of Federal Research and Development" and contributing to NSF's Sustainable Computing for Sustainability workshop.

Strengthening Collaboration and Knowledge Sharing: Through a series of meetings, working groups, and presentations, CASC has fostered deeper collaboration among members, driving interdisciplinary research and sharing best practices in areas like artificial intelligence (AI), energy efficiency, and data center management.

Promoting Workforce Development and Diversity: Recognizing the urgent need for skilled professionals in scientific computing, CASC has supported initiatives such as the Cyberinfrastructure Leadership Academy to develop the next generation of leaders. Our Democratization & Inclusivity working group is actively charting a path to increase participation from underrepresented groups, ensuring a more diverse and vibrant community.

Looking ahead, our priorities for the coming year are clear and ambitious:

Advancing AI and Research Computing Integration: We will continue to champion the integration of AI and machine learning with advanced scientific computing, accelerating innovation and discovery. CASC's Dynamic State of AI working group is at the forefront of this effort, alongside groups exploring AI-driven energy-efficient computing and advanced research infrastructure management.

Enhancing Cyberinfrastructure Security: As cyber threats become more sophisticated, CASC is committed to promoting the development and implementation of advanced regulated research practices to safeguard our computational resources and the invaluable data they manage.

Expanding Outreach and Education: We aim to broaden our outreach to engage policymakers, educators, and corporate stakeholders, fostering a deeper public understanding of the vital role of research computing in advancing science.

Fostering Collaboration with Strategic Partners: We are increasing CASC's presence at key conferences and meetings, such as Practice and Experience in Advanced Research Computing (PEARC), EDUCAUSE, SC, Research Computing at Smaller Institutions (RCSI), and Minority Serving-Cyberinfrastructure Consortium (MS-CC) to strengthen our partnerships and ensure our members have access to the latest technologies and methodologies.

As we embrace the challenges and opportunities ahead, I am confident that our collective efforts will continue to push the boundaries of scientific computation. The work of CASC members as outlined in this annual brochure — whether it's improving public health, advancing AI, building climate resilience, protecting privacy, or driving new discoveries — will inspire us all.

Warm regards,

James "Jim" Wilgenbusch
2023-2024 CASC Chair

CASC members enjoyed a lively exchange of ideas at the Spring 2024 Member Meeting

CASC Executive Director Kathryn Kelley (left) and Operations Manager Carolyn Casler
Photo by Sarah Miller, Mississippi State University

CASC reception at PEARC24

Leveraging Technology to Promote Public Health

Supporting a Safer Food Supply

Farming at the scale needed to keep America's grocery shelves stocked requires sophisticated systems. One example is the poultry incubators that farmers use to nurture hundreds or thousands of eggs simultaneously. In these high-tech hatchers, fluctuations in temperature and humidity can influence chick growth and development — and the spread of pathogens like *E. coli.*

To better understand the risks and inform healthier hatchery practices, an interdisciplinary team led by Shanti Bhushan at Mississippi State University (MSU) developed a "digital twin" of a poultry incubator. With a computational grid consisting of 15 million individual cells executed using 120 processors, their simulations provided detailed insights into the thermo-fluid dynamics and particle transport within an incubator. It took 150 hours of computing runtime on the Orion supercomputer at MSU's High Performance Computing Collaboratory to create just 60 seconds of simulation time.

The image shows the spread of *E. coli* from a few infected eggs in the center of the incubator. This research was performed in association with MSU's Center for Advanced Vehicular Systems with funding from the U.S. Department of Agriculture.

CASC Priority Area: Team Science

Solving complex problems increasingly requires expertise in multiple fields — in this example, bridging computational fluid dynamics to biology and agricultural engineering. Research computing and data organizations are a critical nexus for multidisciplinary collaborations. CASC advocates for funding and policies that support team science and helps our members cultivate best practices for successful collaborations.

Shanti Bhushan, Mississippi State University
Copyright Shanti Bhushan, Mississippi State University

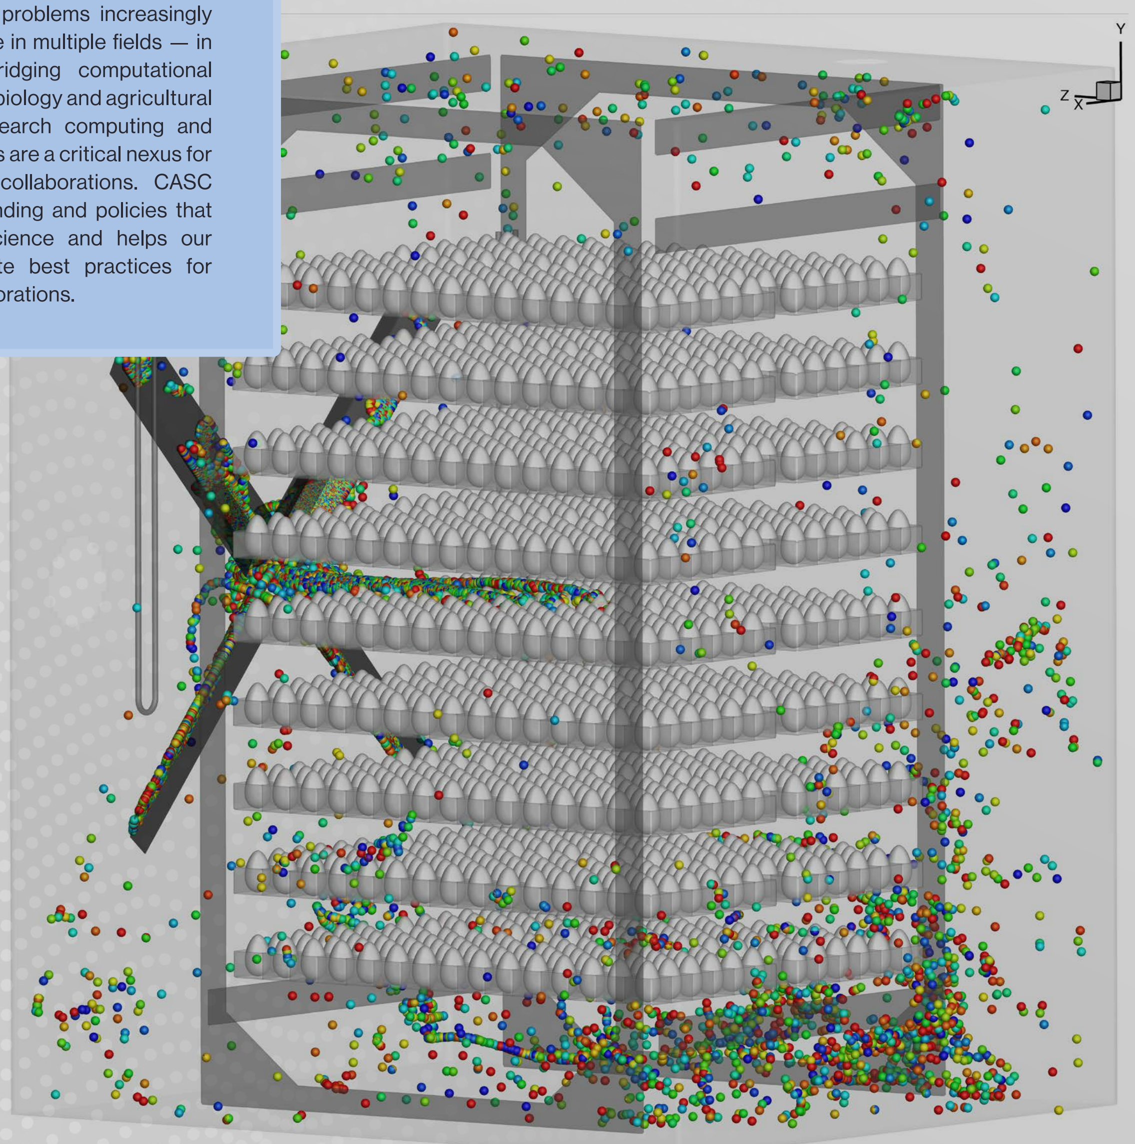

Copyright Shanti Bhushan, Mississippi State University

Addressing the Hidden Dangers of Street Drugs

Today's street drugs can be deadly, especially if they contain harmful or unexpected substances that increase the risk of overdose or severe reactions. Nabarun Dasgupta and colleagues at the University of North Carolina at Chapel Hill (UNC) are working to save lives by understanding what is in the ever-changing drug supply.

Here's how it works: Any public health organization can anonymously mail a drug sample — taken, for example, by swabbing an empty baggie — to the UNC Street Drug Analysis Lab. The lab uses a chemical analysis method called high-resolution gas chromatography-mass spectrometry to determine what's in the sample and then posts the results online, typically within a few days. Using this system, the lab monitors the street drug supply in partnership with over 160 front-line public health programs in 35 states. To streamline data sharing, the team uses Deepnote, a cloud-hosted Jupyter-like notebook with AI features that automate routine tasks, enabling researchers to focus on effective computational science and data presentation.

So far, the researchers have detected over 325 unique substances in more than 7,000 street drug samples. Their work has helped solve medical mysteries in hospitals, allowed EMS teams to more accurately respond to overdoses, and enabled drug treatment regimens that are tailored to actual substance use. Last year, Dasgupta was recognized as a TIME100 Next emerging leader for his work combating the national opioid overdose epidemic.

Nabarun Dasgupta, University of North Carolina at Chapel Hill
Copyright Pearson Ripley for University of North Carolina

Copyright Pearson Ripley for University of North Carolina

Strategies to Beat the Heat

With over half the world's population now living in urban areas — and many cities facing increasingly long and intense heat waves — lowering the burden of heat on urban populations is becoming a serious public health challenge. Cities can help to reduce the root causes of extreme heat by lowering emissions of greenhouse gases. They can also adapt to the higher temperatures through measures like planting trees or installing cool roofs. But what happens when both adaptation and mitigation are applied in tandem?

To help city managers know where to focus their efforts, Matei Georgescu, Director of the Urban Climate Research Center at Arizona State University, led a project modeling the impacts of different heat mitigation and adaptation strategies in 48 of the largest U.S. cities. Drawing on regional climate modeling simulations conducted by Scott Krayenhoff using the university's supercomputer resources, the team, including Ashley Broadbent, incorporated projections of urban expansion, greenhouse gas emissions, and population migration to estimate the cooling potential of various measures.

Overall, they found that combining adaptation and mitigation is likely to reduce exposure to extreme heat by the last decade of this century to a greater extent than if these strategies were applied separately. However, the work also revealed important differences in the relative impacts of various measures in different regions, providing a basis to help cities make decisions that reflect what is likely to work best for them. This work was published as the inaugural paper in the new Nature journal, *Nature Cities*.

Matei Georgescu, Arizona State University
Copyright Matei Georgescu, Arizona State University

Creating a Smarter, Safer Internet

Sniffing Out Scams

Romance fraud — which uses the pretense of a romantic interest to lure people to hand over money or personal information — is among the fastest-growing online scams. While younger people are the most frequent victims, older adults are the most likely to incur a big financial hit. To address this growing issue, Anita Nikolich, a researcher at the University of Illinois Urbana Champaign School of Information Science, worked on an NSF-funded multi-institutional team to help older adults recognize romance scams before they become victims.

Since people are often embarrassed to talk about being scammed, Nikolich decided a smartphone game — akin to Candy Crush, Wordle, or other games popular with the over-55 set — could be a good way to help people learn about scams in a private and engaging way. To create the game, she consulted with cybersecurity experts at the University of Illinois at Urbana Champaign's National Center for Supercomputing Applications and experts in hacking and cybersecurity from institutions around the country. In the game, called DeepCover, players create an agent and embark on a mission to help stop a scam. As they work through puzzles and earn prizes, they're provided with tips on how to identify scams. Since scammers are exceptionally adaptable, Nikolich's team views the gaming app as a dynamic, ongoing project, which they aim to expand and customize for specific countries.

Anita Nikolich, University of Illinois Urbana Champaign
Copyright Anita Nikolich, University of Illinois

A screenshot from the DeepCover app
Copyright DART Collective, dartcollective.net

CASC Priority Area:
Cybersecurity

Cybersecurity requires diligence at all levels, from individuals to institutions. To advance cybersecurity best practices, CASC works with our member institutions to gain institutional buy-in and inform clear policies to support sustainable, productive, and secure systems for protecting research and data. We also advocate for federal and state government actions to support compliance and enhance internet safety for everyone.

Stopping the Spread of Misinformation

The spread of misinformation online leads people to believe lies and undermines trust in crucial institutions, with real-world consequences. By modeling the propagation of misinformation through online social networks, David Butts, then a graduate student at Michigan State University (MSU), set out to find the best ways to stop the spread.

Butts, along with colleagues Sam Bollman and Michael Murillo, used a technique called agent-based modeling. While most analytical models work from the top down and often include sweeping generalizations, agent-based models work from the bottom up, accounting for detailed individual differences and specific interactions. The extra time it takes to generate these models is worth it because they allow researchers to recreate specific environments that can provide more precise results.

According to the models, the most promising ways to decrease the spread of misinformation are education-based policies, which increase skepticism, followed by counter-campaign policies, which directly address false narratives with corrective advertising. Thanks to the supercomputer at MSU's Institute for Cyber-Enabled Research, Butts was able to perform the simulations and data analysis in about a week — a process that would have taken about two years on a standard laptop.

David Butts, Michigan State University; *the screen shows Butts' simulation of misinformation spread among more than 7,000 users in a social network*
Copyright Michelle David, Michigan State University

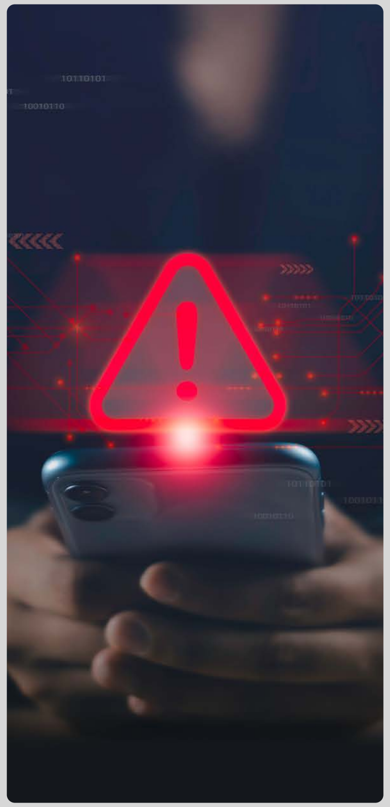

Crisis-Proof Connectivity

On Christmas morning 2020, a bomb ripped through downtown Nashville, Tennessee, killing the bomber, injuring eight others, and damaging buildings and infrastructure. The explosion also knocked out part of AT&T's regional access network, which connects local internet networks to regional backbone networks. For days afterward, communities across three states grappled with cellular and internet outages that disrupted critical services including 911 emergency dispatch, air traffic control, hospital networks, and credit card processing.

The incident prompted an urgent question: How can we better prevent internet outages resulting from intentional attacks or natural disasters? To explore this question, researchers led by Kimberly Claffy from the University of California San Diego Center for Applied Internet Data Analysis (CAIDA) at the San Diego Supercomputer Center conducted a year-long study of real-world access network failures across the United States. The lead author of the study was Alexander Marder, then a research scientist at CAIDA and now an assistant professor at Johns Hopkins University. The researchers combined new techniques for analyzing access-network infrastructure deployments with measurements of large-scale outages to demonstrate the feasibility of targeted attacks and quantify their potential impacts. Their study yielded insights into physical attack surfaces and resiliency limits of regional access networks, as well as suggested strategies to mitigate risks.

Kimberly Claffy, University of California San Diego
Copyright Timothy Dunnigan

Alexander Marder, Johns Hopkins University
Copyright Alexander Marder, Johns Hopkins University

Building Resilience in a Changing Climate

Broadening Access to Climate Data

Hawai'i faces unique challenges from the effects of climate change, impacting its resources, ecosystems, economic vitality, and human health. The Change Hawai'i project, part of the National Science Foundation-funded University of Hawai'i Established Program to Stimulate Competitive Research (Hawai'i EPSCoR), is a multidisciplinary research effort to help the state address these issues by developing new climate data products and tools and sharing knowledge with the community.

As part of this project, Ryan Longman from the East-West Center and the University of Hawai'i and a team of about 20 climate and data scientists from across the university developed a comprehensive hub for Hawai'i climate data known as the Hawai'i Climate Data Portal (HCDP). HCDP, which took over a decade to develop, provides streamlined access to high-quality, reliable information crucial for local planning and decision-making.

Ryan Longman, University of Hawai'i
Copyright Ryan Longman, East-West Center, Honolulu, Hawai'i

Portal users can access more than 100 years of monthly rainfall maps, 30 years of daily temperature maps, and a variety of other data products, most of which are updated in near-real-time. HCDP also features map and data visualizations, a library of related scientific publications, climate research highlights, traditional ecological knowledge and climate perspectives, and several decision-support tools to help stakeholders tap into the power of data as they respond to climate challenges in their communities.

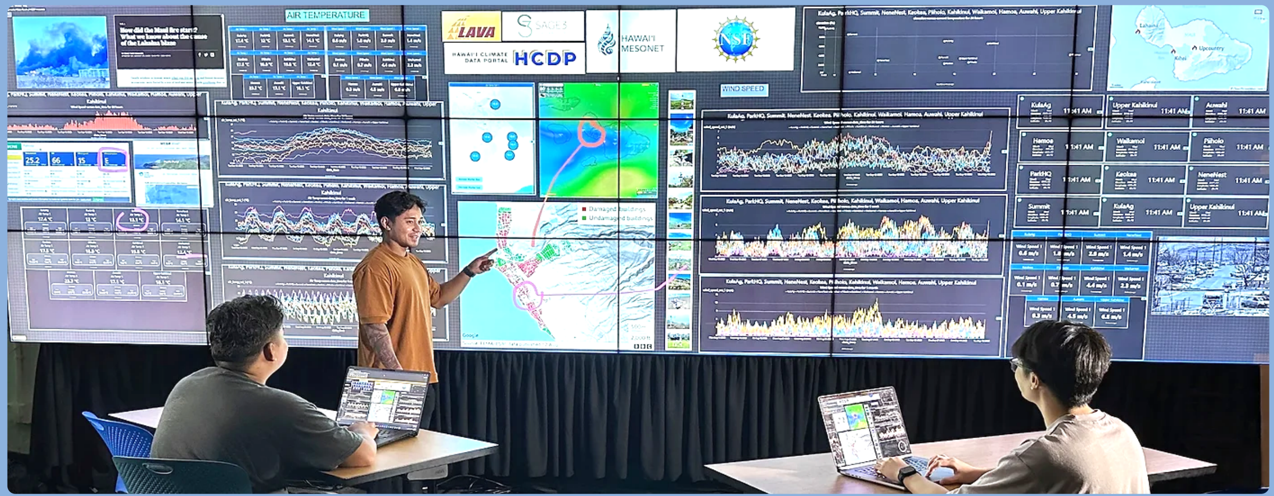

Copyright Laboratory for Advanced Visualizations and Applications (LAVA), University of Hawai'i at Manoa

CASC Priority Area:
Alignment

Big challenges like climate change touch every sector of society, and an effective response requires aligning community priorities, government policies, industry interests, and research insights and resources. CASC promotes alignment among these diverse stakeholders by bolstering communication channels within institutions and facilitating productive dialogue among CASC members, federal agencies, and industry.

Looking Back to Look Ahead

To understand and predict rare events like extreme floods or long-term events like droughts, scientists require not just years but decades of high-resolution data. The daunting task of providing all this data is now much easier thanks to a new weather simulation dataset from researchers led by Roy Rasmussen from the NSF National Center for Atmospheric Research, in partnership with the U.S. Geological Survey.

Roy Rasmussen, National Center for Atmospheric Research
Copyright Roy Rasmussen, National Center for Atmospheric Research

The nearly one-petabyte dataset, called CONUS404, covers the contiguous United States for more than 40 years (1980–2021) with a 4-kilometer grid spacing. Researchers created it by downscaling ERA5, one of the most widely used global reanalysis datasets. Creating a dataset encompassing such a vast area for so many years at such high resolution required nearly a year's worth of supercomputing time and was only possible due to recent advancements in supercomputing and weather models.

Scientists are already using the data to explore ways to improve long-range forecasting, better plan for water resource allocation, and understand the causes and impacts of extreme and rare weather events. The data also reveals some ways that weather patterns may have already changed due to climate warming.

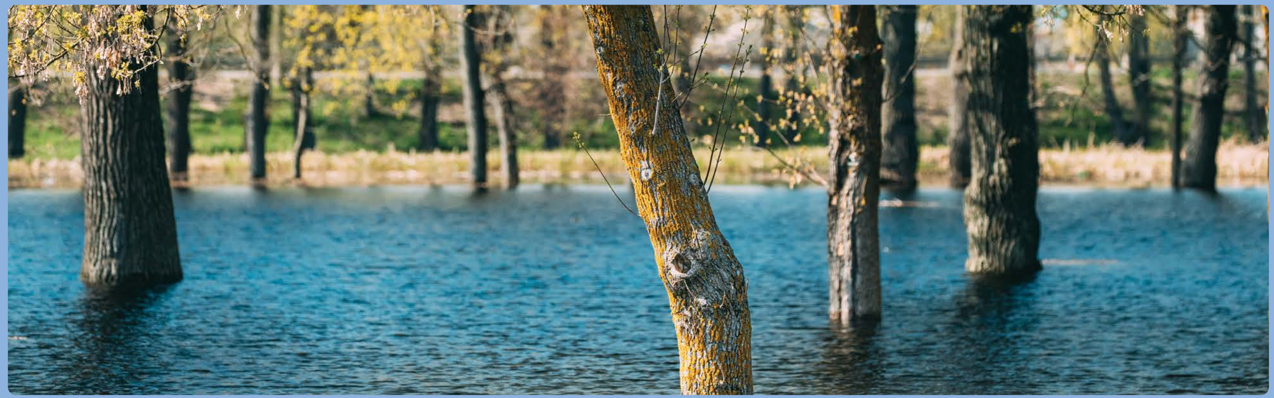

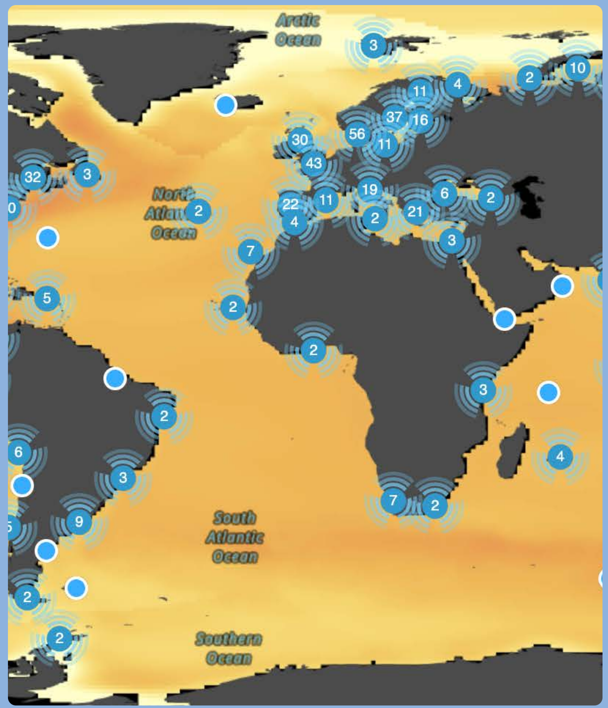

Copyright Sea Level Projection Tool, NASA

Informing Action Amid Uncertainty

There are many ways to make coastal communities more resilient in the face of climate change, but deciding which path to take often comes down to one key question: How much will sea levels rise? While planners want clear answers, even the best sea-level projection models have many sources of uncertainty. To better represent and characterize these uncertainties, Robert E. Kopp, Praveen Kumar, and Alex Reedy from the Rutgers University Earth System Science & Policy Lab are developing new tools to provide precise sea-level projections with quantified uncertainty estimates to inform decisions about coastal planning and investments.

A key tool in this effort is the Framework for Assessing Changes to Sea-level (FACTS), which is designed to help users understand how sea levels may respond to future climate variability. Used in the Intergovernmental Panel on Climate Change Sixth Assessment Report, FACTS integrates data from ice sheets, glaciers, ocean models, and population dynamics, while accounting for emissions and other factors that influence sea-level change. Sea-level projections from FACTS are intended to empower coastal communities to navigate the challenges posed by sea-level rise with greater confidence and resilience. The image shows a screenshot of the FACTS web interface, from which anyone can access global and regional (blue dots) sea-level projections under different scenarios.

Inspiring Innovation in Energy and Materials

Turning Trash into Treasure

While we like to think that most of what is recyclable gets recycled, the truth is that today's plastic recycling methods are often not efficient enough to make them economically worthwhile, especially for items that contain multiple materials. As a result, many recyclable products simply wind up in landfills.

One promising solution to this issue is a chemical recycling method known as pyrolysis. It can be used to extract valuable chemical components from unsorted, unwashed materials through decomposition in an oxygen-free, high-temperature environment. However, a major challenge for pyrolysis is identifying the ideal temperature, pressure, and reactor type, a process that currently relies on trial and error.

To optimize pyrolysis conditions so that plastics can be more efficiently converted into valuable chemicals, University of Pittsburgh Professor Giannis Mpourmpakis and postdoctoral researcher Hyunguk Kwon developed a new computational method to calculate temperature-dependent thermochemistry for the large and flexible molecules that make up plastics. Using computing resources from the University of Pittsburgh Center for Research Computing, the researchers showed that their approach accurately predicts how octadecane, a model compound for polyethylene, breaks down at different temperatures. The image illustrates how the thermal decomposition of large molecules produces small depolymerization products at elevated temperatures.

Giannis Mpourmpakis, University of Pittsburgh
Copyright Giannis Mpourmpakis, University of Pittsburgh

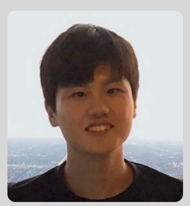

Hyunguk Kwon, now at Seoul National University of Science and Technology
Copyright Hyunguk Kwon, Seoul National University of Science and Technology

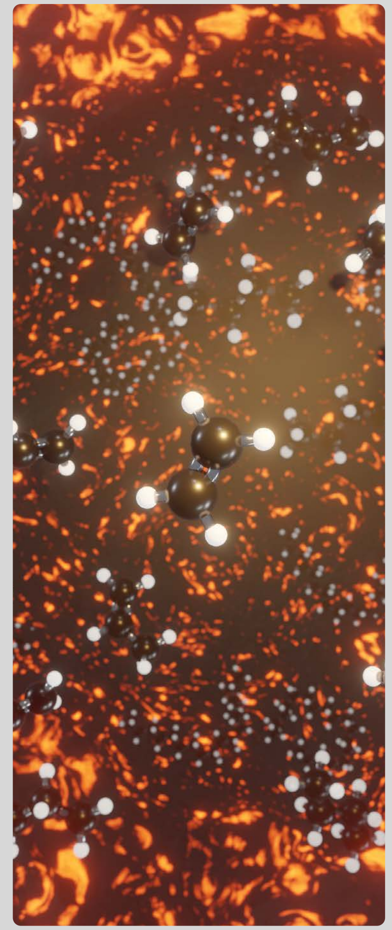

Copyright Giannis Mpourmpakis, University of Pittsburgh

A Step Forward for Fusion

Magnetic fusion energy tokamaks are experimental fusion reactors designed to use the same process that powers stars to potentially produce clean and virtually limitless energy. While these reactors are promising, they involve the confinement and control of extremely hot plasma using powerful magnetic fields, making monitoring, optimizing, and predicting plasma behavior incredibly complex.

To make fusion experiments more efficient and successful, Eliot Feibush and colleagues from the Princeton University Plasma Physics Laboratory, in collaboration with NVIDIA and General Atomics, developed a high-performance "digital twin" of the experimental DIII-D tokamak in San Diego. The simulation integrates AI-enabled design, experimentally validated theories, and advanced computation with sophisticated visualization techniques, allowing researchers to simulate and predict plasma dynamics, detect potential instabilities, and optimize tokamak performance before running physical experiments. An innovative machine-learning-based gyrokinetic simulation is trained and executed on the Stellar computing cluster at Princeton and helps accelerate workflows, producing a realistic and near-real-time model of a tokamak that could pave the way for clean fusion power. The image shows a cut-away view of the DIII-D vacuum vessel containing a visualization of the fluctuation of electrostatic potential in the plasma.

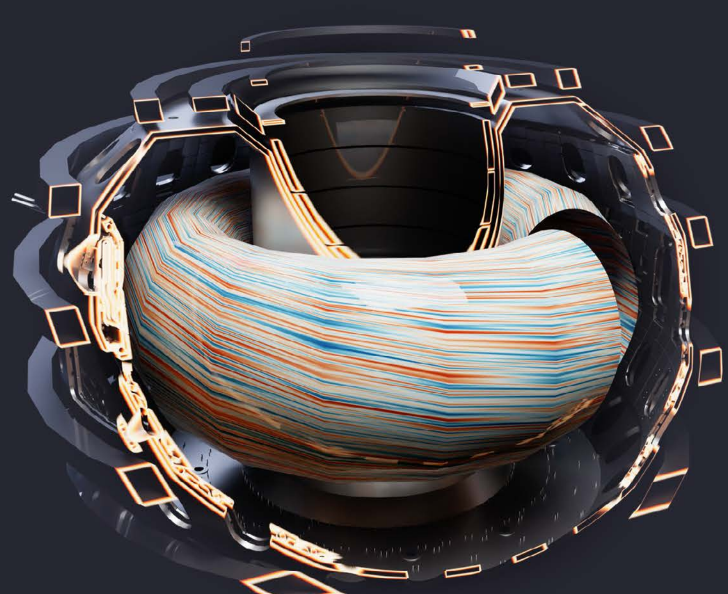

Copyright Eliot Feibush, Princeton University; John Stone, Nvidia Corporation

Eliot Feibush, Princeton University
Copyright Eliot Feibush, Princeton University

Supercharging Battery Technology

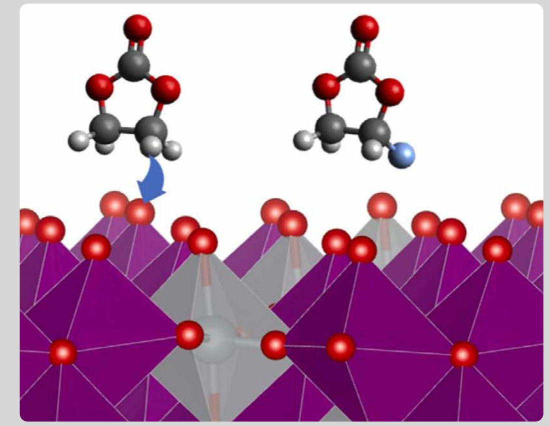

Copyright Next Energy, volume 4, July 2024

Marco Olguin, University of Southern California
Copyright Marco Olguin, University of Southern California

Lithium nickel manganese oxide (LNMO) is a promising cathode material for lithium-ion batteries due to its high energy density, safety, and cost-effectiveness. However, it is not widely used because the material tends to lose capacity at high voltages and isn't stable at high temperatures. Marco Olguin from the University of Southern California Center for Advanced Research and Computing (CARC) and colleagues from the University of California San Diego, the University of Chicago, and the Army Research Laboratory are using computation-guided experimental chemical synthesis to find the best way to make LNMO more practical for lithium-ion batteries, including those used in electric vehicles and energy storage systems.

Given the critical role of the battery's electrolyte in allowing the flow of ions between the anode and cathode during charging and discharging, the researchers sought to find an electrolyte that was best suited for use with LNMO. With CARC computing resources, they ran interface calculations for several potential electrolyte molecules (one of which is depicted in the visualization) and predicted that fluorinated carbonate would be the most stable in terms of preventing decomposition at the cathode. They then synthesized fluorinated carbonate and found that it indeed was stable at both high temperatures and high voltages. This advanced electrolyte, which is also nonflammable, could help extend the lifespan of batteries by allowing them to retain more than 30% of their capacity, even when stored at high temperatures.

Particle Detector Gets a Software Upgrade

Peter Wittich of Cornell University and Steve Lantz from the Cornell Center for Advanced Computing (CAC), in collaboration with other U.S. institutions, are advancing the software infrastructure for collecting, managing and analyzing data from the Large Hadron Collider (LHC). The LHC is essential for studying the fundamental particles and forces of the universe, and efficient data analysis is crucial for making groundbreaking discoveries, such as the Higgs boson in 2012 that explained why particles have mass. However, as instrumentation improves and data volumes increase, analyzing this data becomes increasingly complex.

The current project, which is supported by the NSF Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP), focuses on enhancing the software for the Compact Muon Solenoid (CMS), one of the two major detectors at the LHC. The CMS generates 20 petabytes of data annually, a figure expected to increase tenfold following an upgrade to boost the detector's luminosity, scheduled for completion in 2029. To manage this data surge, the researchers will upgrade particle-tracking software and develop a new algorithm tailored for massively parallel hardware. This will allow physicists to process the massive amounts of data generated by the CMS.

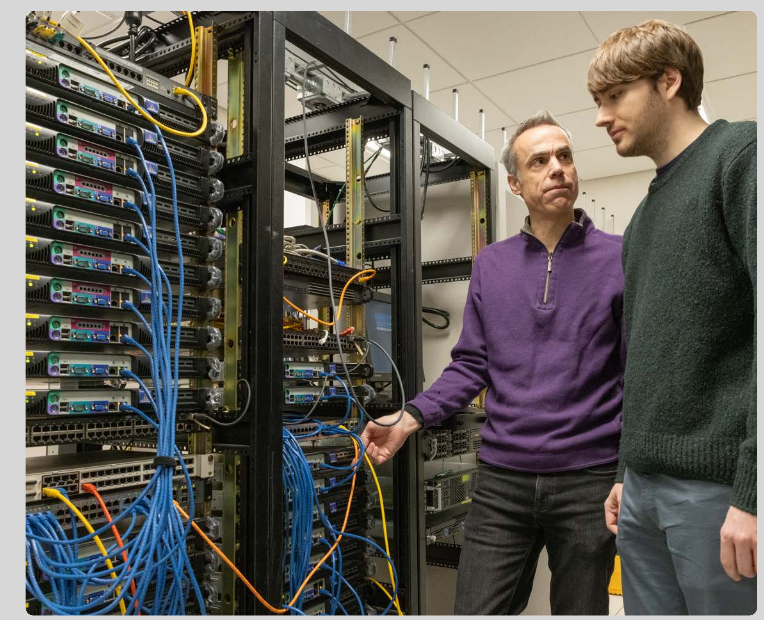

Peter Wittich (left) and doctoral student Gavin Niendorf inspecting a server rack that runs hardware used for analyzing data from the CMS detector. Copyright Cornell University.

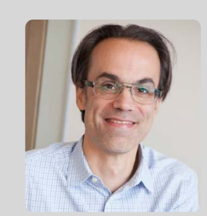

Peter Wittich, Cornell University
Copyright Cornell University

Steve Lantz, Cornell CAC
Copyright Cornell University

Fascinating Frontiers

Animals Offer Lessons on Learning to Speak

Despite our vastly different vocalizations, humans, bats, whales, and seals show surprising similarities in the brain regions used for vocal learning. A multi-institutional team, including Andreas Pfenning from Carnegie Mellon University and Michael Yartsev from the University of California, Berkeley, combined machine learning with studies of bat brains to reveal the genomic factors behind how we learn to speak. The research sheds light on our evolutionary relationships with other animals and even offers new clues about the drivers of autism spectrum disorder.

The researchers examined genes and the regulatory elements influencing gene expression across hundreds of genomes, a task made possible by new AI methods that analyze tissue-specific expression patterns of genes. Running these analyses on the Pittsburgh Supercomputing Center's Bridges-2 system, they identified 50 gene regulatory elements from the brains of humans, bats, and marine mammals that are closely linked with vocalization. They also found that these regulatory elements tend to reside in parts of the genome related to autism spectrum disorder. This suggests that shared regulatory networks for vocal and social behavior exhibit similar patterns of evolutionary adaptation.

Andreas Pfenning, Carnegie Mellon University
Copyright Department of Computational Biology, Carnegie Mellon University

Michael Yartsev, University of California, Berkeley
Copyright Alison Yin

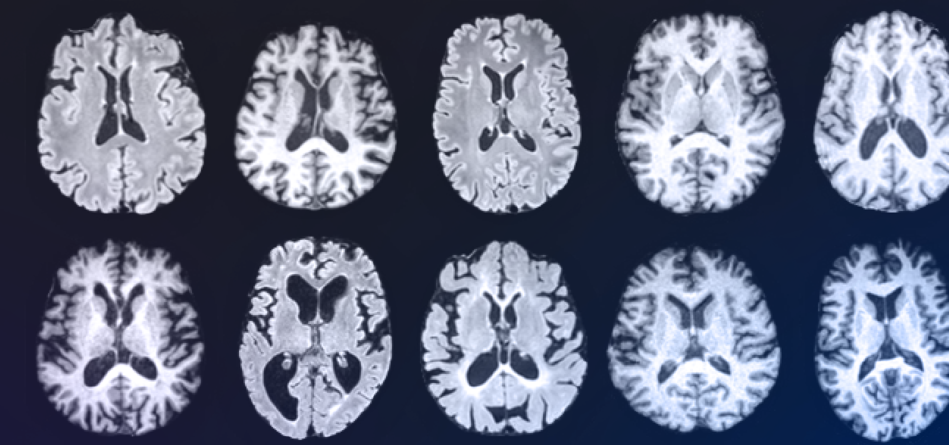

Copyright Vijaya B. Kolachalama, Boston University

AI Sharpens Dementia Diagnosis

Doctors can often diagnose dementia with confidence, but identifying the exact type can be challenging. With more accurate diagnoses, doctors could use targeted treatments, improving the outlook for patients and their families. Toward this goal, an interdisciplinary research team led by Vijaya Kolachalama at Boston University developed an AI tool that can distinguish between 10 types of dementia, including Alzheimer's disease and mixed dementia.

The researchers trained their machine learning model using routinely collected clinical data such as demographic information, medical history, medication use, neurological and neuropsychological exam scores, and neuroimaging data from nine independent, geographically diverse datasets of over 50,000 participants. To interpret such a wide range of diagnostic parameters, they used a transformer architecture, which processes and understands data by focusing on the most important parts of the input. This approach enhanced the model's robustness and predictive power, enabling it to handle real-world scenarios with incomplete data. During testing, the AI tool accurately distinguished 10 types of dementia, scoring 0.96 out of 1.0 on a performance scale. When the AI tool was combined with neurologist assessments, it boosted the accuracy of diagnoses by over 26% across all dementia types.

Vijaya B. Kolachalama, Boston University
Copyright Vijaya B. Kolachalama, Boston University

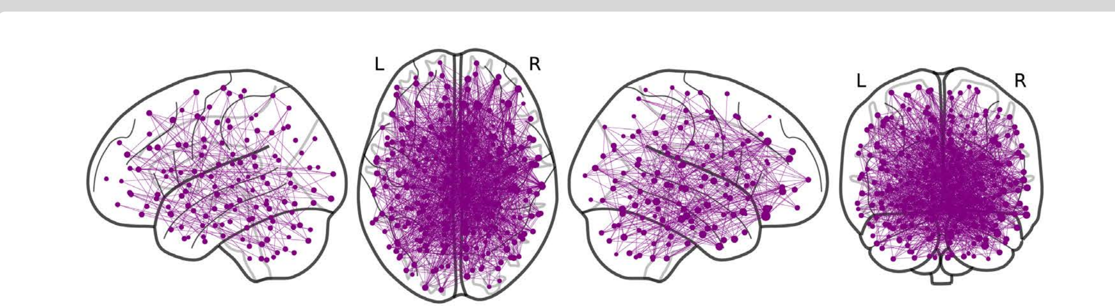

Copyright Ziwei Zhang and Monica Rosenberg, University of Chicago

The Science of Surprise

There's nothing quite like experiencing a sudden jolt of surprise. Decoding exactly what happens in the brain during such moments could help scientists enhance AI systems, identify optimal learning conditions, and aid in disaster planning or even political polling by predicting how people may respond to unexpected events.

Monica Rosenberg and Ziwei Zhang from the University of Chicago used thousands of brain scan images from dozens of people to train a machine learning model to predict the feeling of surprise during various situations. The findings revealed that our brains process surprises through a complex interplay of brain networks, with the relationships between activity in different brain regions being key predictors of surprise. The image shows a brain network with high connectivity among regions, reflecting high surprise.

To deal with the large amounts of imaging data involved in this project, the researchers relied on high performance computing resources like the Midway cluster at the University of Chicago's Research Computing Center. This significantly accelerated the data processing and analysis and enabled the conversion of complex brain images into valuable statistical data.

Ziwei Zhang, University of Chicago
Copyright Ziwei Zhang, University of Chicago

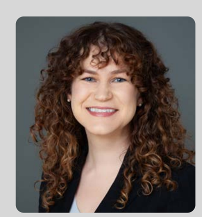

Monica Rosenberg, University of Chicago
Copyright Hummingbird Studios

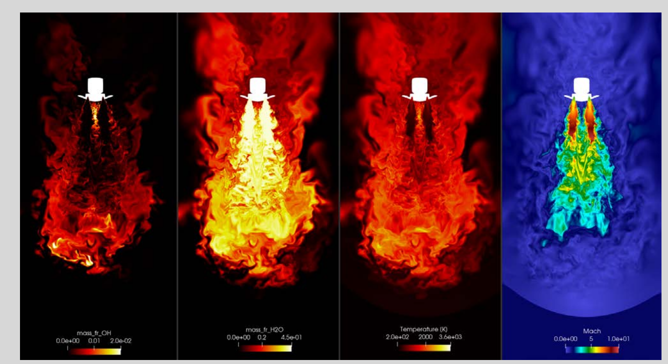

Mars lander retropropulsion simulation
Copyright Gabriel Nastac, NASA Langley Research Center

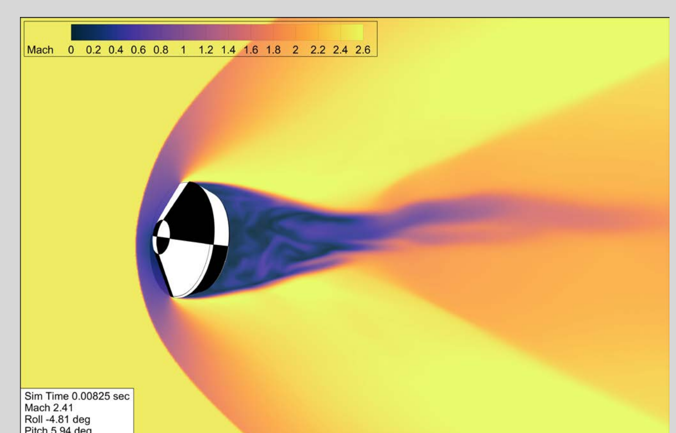

Blunt-body spacecraft simulation
Copyright Brenton J. Willier, Georgia Tech

Simulating Spaceflight

Since it's impossible to fully replicate the unique conditions of space here on Earth, scientists use computational fluid dynamics (CFD) simulations to model the aerodynamic forces, heat transfer, and other factors encountered by spacecraft. These simulations inform designs that will make space missions safer and more efficient.

Led by Eric Nielsen from NASA's Langley Research Center, scientists are using CFD tools to run simulations of a human-scale Mars lander on Oak Ridge Leadership Computing Facility (OLCF) supercomputers. Unlike past missions to Mars, landing a vehicle with humans onboard will most likely use retropropulsion, in which rockets are fired downward to slow the vehicle's descent. Since this hasn't been done before on Mars, advanced CFD simulations are critical for determining if it can be done safely. After five years of running increasingly sophisticated simulations, the researchers completed a 35-second autonomous test flight simulating the vehicle's final descent. The groundbreaking results were made possible by OLCF's new Frontier supercomputer, the world's first computer to achieve an extremely high-speed type of parallel computing known as exascale.

CFD simulations are also key to research led by Dimitri Mavris at the Georgia Institute of Technology School of Aerospace Engineering's Aerospace Systems Design Laboratory as part of a NASA Early Stage Innovation grant. His team developed a simulation to mimic blunt-body entry vehicles — like those from the Apollo missions — in free flight. High-fidelity CFD simulations performed using compute nodes at Georgia Tech's Partnership for Advanced Computing Environment (PACE) will support the design of future spacecraft by creating advanced models that can predict the pressure and forces on a spacecraft during reentry.

Navigating Uncertain Ground

Unzipping an Earthquake Mystery

The 2019 Ridgecrest earthquakes, a series of quakes in southern California with magnitudes of 6.4 and 7.1, were not only the largest the state had experienced in 20 years but also the most baffling. Scientists did not understand why the quakes cascaded across areas that were not previously known to contain fault lines, nor why there was an almost 34-hour gap between the foreshock and the larger mainshock.

Recognizing that a better understanding of the underlying processes involved in these events would help to improve earthquake preparedness and hazard assessment in the region, Alice-Agnes Gabriel from the University of California San Diego together with Gregory Abram and Francesca Samsel from the Texas Advanced Computing Center at the University of Texas at Austin sought to uncover what might have caused this complex earthquake rupture sequence. They used high performance computing to model the known faults and the subsurface and then applied visualization techniques to analyze the simulation results in a realistic 3D Earth structure. The resulting visualizations (pictured) revealed the seismic waves as they ripple through the Earth, generated by the "unzipping" of fault lines (including previously unknown secondary faults) during the Ridgecrest earthquakes. The visualizations also revealed how multiple faults can rupture in unison and in complementary ways, defying prior assumptions of earthquake behavior.

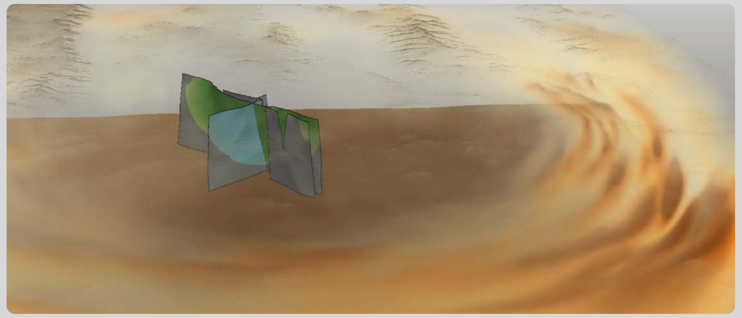

Copyright Alice-Agnes Gabriel, University of California San Diego and Gregory Abram and Francesca Samsel, University of Texas at Austin

Alice-Agnes Gabriel, University of California San Diego
Copyright Prof. Dr. habil. Alice-Agnes Gabriel, UC San Diego and LMU Munich, Germany

Gregory Abram, University of Texas at Austin
Copyright Gregory Abram, University of Texas at Austin

Francesca Samsel, University of Texas at Austin
Copyright Francesca Samsel, University of Texas at Austin

When a Landslide Triggers a Tsunami

As temperatures rise and glaciers shrink, hills and mountains once supported by ice are left exposed, making them vulnerable to landslides. In places like coastal Alaska, this phenomenon is causing large land masses to spill into water, which can trigger dangerous tsunamis.

To provide early warnings to communities, University of Alaska Fairbanks Geophysical Institute researchers Ezgi Karasözen and Michael West developed a new method for remotely detecting large landslides just minutes after they occur and then quickly determining if they present a tsunami hazard. Their method uses an algorithm that continually scans data from multiple seismic stations to look for a landslide wave signature and then estimates the volume of the landslide. They have installed a prototype detection system close to the Barry Arm of Prince William Sound, a vulnerable area that has already experienced several landslides. Using data from seismic stations already in the Alaska network, the researchers showed that their system can evaluate landslides within three minutes and generate rapid volume estimates directly from seismic amplitudes.

The debris from a 2015 landslide (left) extends over the bottom of Alaska's Tyndall Glacier. Copyright Chris Larsen, UAF Geophysical Institute.

Ezgi Karasözen, University of Alaska Fairbanks
Copyright University of Alaska Fairbanks Geophysical Institute

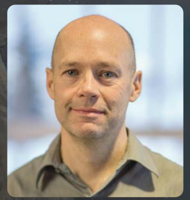

Michael West, University of Alaska Fairbanks
Copyright University of Alaska Fairbanks Geophysical Institute

Mobilizing a New Generation

A still image from a digital video by Ohio University art student Alyssa Vandale; created using Adobe Premiere, Illustrator, Photoshop, After Effects, Fresco, Audition, and batch processing on Stable Diffusion to create a sequencing video. Copyright Alyssa Vandale, Digital Art + Technology, Ohio University.

Graduate student Noufel Sharif with Professor Basil Masri Zada (right) experimenting with lights, shadows, and projections. Copyright Basil Masri Zada, Digital Art + Technology, Ohio University.

The Creative Side of AI

Although AI is gaining acceptance in the art world, it is still often met with a degree of reservation. Artists worry about issues such as authenticity, copyright, and bias that can come with applying AI to their work. Basil Masri Zada from Ohio University (OU) is facing these concerns head-on by incorporating AI into coursework he developed for the OU School of Art and Design's new Bachelor of Fine Arts concentration in digital arts and technology, which focuses on the evolving relationship between technology and art.

Two courses within the curriculum emphasize using AI as a creative tool to support artistic expression. Before launching these courses, Masri Zada worked with the Ohio Supercomputer Center (OSC) to host a version of Stable Diffusion, a popular open-source generative AI software. Unlike most text-to-image AI software, Stable Diffusion allows precise editing within the generated image. This allows artists to seamlessly integrate it into their workflow, rather than the software being the sole generator of the creative product. Hosting the software on OSC's remote access platform provided users with greater privacy and control over the data generated, and also made it possible to examine the influence of training materials (and possible sources of bias) on AI behavior. This setup empowers students to leverage AI tools while remaining firmly in control of their creative expression.

CASC Priority Area:
Workforce

Even as machines perform an even broader array of tasks, people remain at the heart of the computing and data services enterprise. CASC advocates for policies and practices that foster a robust, more inclusive pipeline of skilled experts, creative thinkers, and capable collaborators who will become tomorrow's research computing and data workforce.

Expanding Access to Cutting-Edge Computing

Undergraduate-focused institutions are a critical part of the STEM workforce pipeline, but students and faculty at these institutions often lack access to high performance computing resources. Partnerships can help to bridge the gap. Jonathan Lyon from Murray State University, a primarily undergraduate institution, provides his chemistry students experience with top-of-the-line capabilities through a remote computing resource provided by the San Diego Supercomputer Center (SDSC) at the University of California San Diego.

Lyon and his students tap into SDSC's Expanse supercomputer through the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which includes remote login, remote computation, data movement, science workflow support, and science gateway support toolkits. With SDSC computational resource allocations, the group studies structure-property relationships in strongly bound atomic clusters as their size increases from individual atoms to bulk quantities. This work, which earned the team a spot to present at an international symposium, could help identify opportunities to alleviate the hazardous conditions that sometimes occur with the high pressures and temperatures required for the industrial synthesis of certain chemicals. Ten CASC member institutions provide resources and support through the ACCESS program.

(from left) Murray State University students Joey Quilliam, Madison Winkeler, Ciara N. Richardson, and Jenna Kesselring with their professor, Jonathan Lyon. Copyright Jonathan T. Lyon, Murray State University.

CASC Members

Albert Einstein College of Medicine Department of Information Technology

Arizona State University Research Computing

Boston University

Brown University Center for Computation and Visualization

Carnegie Mellon University

Case Western Reserve University Core Facility Advanced Research Computing

Chan Zuckerberg Biohub CZ BioHub Scientific Computing

City University of New York High Performance Computing Center

Clemson University Computing and Information Technology (CCIT)

Columbia University

Cornell University Center for Advanced Computing

Dartmouth College

Flatiron Institute

George Mason University

George Washington University

Georgetown University UIS

Georgia Institute of Technology PACE

Harvard University

Icahn School of Medicine at Mount Sinai

Indiana University Pervasive Technology Institute

Johns Hopkins University

Lawrence Berkeley National Laboratory

Louisiana State University Center for Computation & Technology (CCT)

Massachusetts Green High Performance Computing Center

Michigan State University High Performance Computing Center

Mississippi State University High Performance Computing Collaboratory (HPC2)

National Center for Atmospheric Research (NCAR)

New Jersey Institute of Technology

New York Genome Center

New York Structural Biology Center Simons Electron Microscopy Center

New York University

North Carolina State University

North Dakota University System

Northwestern University

NYU Langone Hospitals

Oak Ridge National Laboratory (ORNL) National Center for Computational Sciences

Ohio Supercomputer Center (OSC)

Old Dominion University Information Technology Services

Oregon State University

Pennsylvania State University

Pittsburgh Supercomputing Center

Princeton University

Purdue University Research Computing

Rensselaer Polytechnic Institute

Rice University Center for Research Computing

Roswell Park Comprehensive Cancer Center

Rutgers University

Southern Methodist University

Stanford University

Stony Brook University Research Technologies

Swarthmore College

Tennessee Technological University

Texas A&M University High Performance Research Computing

Texas Advanced Computing Center, The University of Texas at Austin

Texas Tech University High Performance Computing Center

University at Buffalo, State University of New York Center for Computational Research

University of Alabama

University of Alabama at Birmingham IT-Research Computing

University of Alaska Fairbanks Research Computing Systems

University of Arizona Research Computing

University of Arkansas High Performance Computing Center

University of California, Berkeley Berkeley Research Computing

University of California, Davis HPC Core Facility

University of California, Irvine Research CyberInfrastructure Center

University of California, Los Angeles Institute for Digital Research and Education

University of California, Merced Cyberinfrastructure and Research Technologies

University of California, San Diego San Diego Supercomputer Center (SDSC)

University of Chicago & Argonne National Laboratory Research Computing Center

University of Cincinnati Research Technologies

University of Colorado Boulder

University of Connecticut

University of Florida

University of Georgia Advanced Computing Resource Center (GACRC)

University of Hawaii Information Technology Services

University of Illinois at Chicago Advanced Cyberinfrastructure for Education and Research

University of Illinois at Urbana-Champaign National Center for Supercomputing Applications (NCSA)

University of Iowa

University of Kentucky Center for Computational Sciences

University of Louisville

University of Maryland Division of Information Technology

University of Massachusetts

University of Miami Institute for Data Science and Computing

University of Michigan Office of Research

University of Minnesota Minnesota Supercomputing Institute for Advanced Computational Research

University of Nebraska Holland Computing Center

University of Nevada, Las Vegas National Supercomputing Institute (NSI)

University of Nevada, Reno Research Computing

University of New Hampshire Research Computing Center

University of New Mexico Center for Advanced Research Computing

University of North Carolina at Chapel Hill

University of North Carolina at Chapel Hill Renaissance Computing Institute (RENCI)

University of Notre Dame Center for Research Computing

University of Oklahoma Supercomputing Center for Education and Research

University of Oregon Research Advanced Computing Services (RACS)

University of Pittsburgh Center for Research Computing

University of Rhode Island

University of Southern California Center for Advanced Research Computing

University of Tennessee National Institute for Computational Sciences (NICS)

University of Utah Center for High Performance Computing

University of Virginia Research Computing (RC)

University of Wyoming Advanced Research Computing Center (ARCC)

Vanderbilt University Advanced Computing Center for Research and Education

Virginia Tech Advanced Research Computing

Washington University in St. Louis

West Virginia University

Yale University Yale Center for Research Computing (YCRC)

CASC | Coalition for Academic Scientific Computation

The Coalition for Academic Scientific Computation (CASC) is a nonprofit organization dedicated to the use of advanced computing technology to accelerate discovery.

For over 30 years, CASC has been the leading voice for many of the nation's most forward-thinking research computing and data services organizations. Today, our vibrant community includes over 100 member organizations. We champion our members by fostering idea and resource exchange, advocating for critical funding and policies, and driving our vision for a robust, sustainable research ecosystem. This ecosystem supports national competitiveness, global security, economic growth, and a diverse and expertly trained 21st-century workforce.

1717 Pennsylvania Ave., NW, Ste. 1025
Washington, DC 20006

Learn More @ casc.org