



Co-sponsored by Pacific Research Platform (PRP) and CineGrid.org
 Hosted by Calit2's Qualcomm Institute at UC San Diego
 Location: QI Auditorium, 1st Floor, Atkinson Hall, UCSD, La Jolla, California
 Dates: May 14-15, 2018

KEY

CHASE-CI Workshop Presenters are designated: C-CI-W

CHASE-CI Reception Demonstration Leads: C-CI-W/D

CineGrid Spotlight: AI Presenters and Panelists are designated: CS:AI

Ilkay Altintas

C-CI-W

Chief Data Science Officer at the San Diego Supercomputer Center (SDSC)
altintas@sdsc.edu

Ilkay Altintas specializes in workflows and computing for collaborative computational data science. Her research objective is delivering impactful results through making computational data science more reusable, programmable, scalable and reproducible. Her current focus is on dynamic data driven process management, and capturing these processes for practical data science, in a measurable and reproducible form. Her work has been applied to many scientific and societal domains, including: bioinformatics; geoinformatics; high-energy physics; multi-scale biomedical science; smart cities and smart manufacturing. Dr. Altintas enjoys teaching courses on big data and data science, and has taught hundreds of thousands of students through popular MOOCs.

David Arbuckle

C-CI-W & CS:AI

Director of IP Video Platform Engineering
 Comcast
david_arbuckle@comcast.com

David Arbuckle works in Comcast's downtown Denver office as director of IP Video Platform Engineering. There, he leads development and operations for the compute platform that powers IP linear video, video on demand, and cloud DVR in more than 20 datacenters. His team leverages a collection of open-source, proprietary, and internally-developed technologies to meet the aggressive scaling and performance demands of the IP video application stack. The Comcast platform supports over 1Tbit/s of egress traffic, 600k connections per second, and over 20,000 containers deployed daily.

Jessica Block

C-CI-W

Research Associate in Geospatial & Environmental Applications, Calit2
 Qualcomm Institute, UCSD
j.block@engr.ucsd.edu

Jessica Block is a spatial data scientist with the Qualcomm Institute at UC San Diego, and co-founded of their Big Pixel Initiative, applying big data techniques to satellite imagery analysis at scale. She is an interdisciplinary geologist and urban ecologist specializing in the use of sensor networks, remote sensing, and geospatial visualization tools for disaster response, public health, policy decision-making, and sustainability.

Jean Bolot

CS:AI

Vice President, Technicolor Research and Innovation, AI Lab
Technicolor, Inc.
jean.bolot@technicolor.com

Jean Bolot is Vice President of Research & Innovation at Technicolor where he leads the AI Lab. Before Technicolor he led the research lab of Sprint, was a founding team member of Ensim, a virtualization startup acquired by Ingram Micro, a visiting professor in Computer Science at UC Berkeley and a researcher at Inria in France. He holds a MS and PhD degree in Computer Science from the University of Maryland at College Park.

Daniel Crawl

C-CI-W

Associate Director, Workflows for Data Science Center, SDSC
crawl@sdsc.edu

Daniel Crawl specializes in scientific workflows for real-time and data-driven simulation and visualization. He is the lead architect for WIFIRE Firemap, a scalable data-driven monitoring, dynamic prediction and resilience cyberinfrastructure for wildfires. Other areas of research include [bioKepler](#): consisting of scientific workflow components to execute a set of bioinformatics tools using distributed execution patterns and the Kepler Scientific Workflow System. [Kepler](#) is an open-source scientific workflow system designed to help scientists, analysts, and computer programmers create, execute, and share models and analyses across a broad range of scientific and engineering disciplines.

Srinjoy Das

C-CI-W

PhD student, Electrical and Computing Engineering, Institute for Neural Computation, UCSD
S2das@ucsd.edu

Srinjoy Das' research focus is combining ideas from nonparametric statistics and generative machine learning to design algorithms for inference on time-series and image processing applications involving prediction, denoising and interpolation. The goal is to enable high performance and low-power implementation of these algorithms on state-of-the-art non Von-Neumann processing platforms.

Dimitris N. Politis

Distinguished Professor of Mathematics and Economics, UCSD
dpolitis@uscd.edu

Das and Politis' current research focuses on developing model-free prediction techniques for problems involving locally stationary time series and random fields thereby enabling accurate construction of both point prediction and interval estimates. Their work combines ideas from nonparametric statistics and generative deep learning to design algorithms for inference applications involving prediction, de-noising and interpolation. Enabling these algorithms to execute efficiently on state-of-the art non-Von-Neumann computing platforms is a focal point of the research.

Thomas A. DeFanti

C-CI-W & CS:AI

Research scientist, UCSD's Qualcomm Institute (QI)
Distinguished professor emeritus of Computer Science at the University of Illinois at Chicago
tdefanti@ucsd.edu

Thomas A. DeFanti, PhD, received the 1988 ACM Outstanding Contribution Award and became an ACM Fellow in 1994. He co-founded StarLight in Chicago with Joe Mambretti and Maxine Brown. He and Dan Sandin conceived the CAVE virtual reality theater in 1991.

Since coming to UCSD in 2004, he and Greg Dawe built the StarCAVE, NexCAVE, TourCAVE, WAVE, SunCAVE, and CAVEkiosk/4KAVE VR systems. With John Graham, Phil Papadopoulos and teams of colleagues, Tom designs

and builds 40/100Gbps state-of-the-art wide-area computer networked computing, storage, graphics processing and visualization facilities to serve advanced computing needs. He is co-PI of the NSF Pacific Research Platform project (prp.ucsd.edu, Larry Smarr, PI, and the CHASE-CI NSF community infrastructure project for machine learning. Tom is profiled in Wikipedia, at <http://www.evl.uic.edu/tom/>, and http://www.calit2.net/people/staff_detail.php?id=67

Shaw(feng) Dong

C-CI-W

Assistant Adjunct Professor, Department of Applied Mathematics and Statistics,
John Baskin School of Engineering, UCSC
shaw@ucsc.edu

Shaw Dong's research areas include: applied mathematics and statistics; high performance and distributed computing; cyberinfrastructure and cybersecurity; machine learning and deep learning focused on GPU-accelerated Gibbs Sampling and detecting damped Lyman-alpha absorbers with deep learning.

John Graham

C-CI-W

Chief Engineer for the NSF Pacific Research Platform (PRP), CHASE-CI and SDX awards
Qualcomm Institute, Calit2, UCSD
jjgraham@ucsd.edu

John Graham is the chief engineer for the NSF Pacific Research Platform (PRP), CHASE-CI, and SDX awards. John designs and implements dramatic new machine learning technologies changing the way we compute over networks. John is responsible for technical project coordination spanning across national and international partners. John has been training other engineers in hands-on workshops at Qualcomm Institute and elsewhere. In recognition of his empowering teamwork, John received two Innovations in Networking Awards for Experimental Applications at CENIC 2018: One for WIFIRE/HPWREN and one for Astronomy and Astrophysics: Big Data Using the Pacific Research Platform (see <https://cenic.org/NetworkingAwards>).

Jerry Gutierrez

CS-AI

Global HPC Sales Leader, IBM Cloud
IBM
jegutierrez@us.ibm.com

Jerry Gutierrez, Global HPC Sales Leader at IBM Cloud, has worked in the IT industry for over 20 years. In 2012, he joined SoftLayer, an early, successful global commercial cloud service provider, to assist businesses migrate their enterprise applications and infrastructure to Softlayer cloud services.

IBM acquired SoftLayer in 2014, and Jerry became responsible for working with IBM Cloud's HPC clients. More recently, Jerry led the IBM cloud team to architect, with NVIDIA and NetApp, the first global virtualized collaboration cloud service, designed for geographically dispersed creative teams to interact with their applications, without latency, and securely collaborate on graphics-intensive workloads. With NVIDIA GPUs becoming an essential part of machine/deep learning research, Jerry's role has expanded to work with IBM global cloud teams, across all industry sectors, to assist clients successfully utilize IBM Cloud infrastructure and Watson cognitive analytic services for their GPU-accelerated AI and Deep Learning initiatives.

Tomáš Ižo

CS:AI

Engineering Director, Machine Perception Group
Google Research
tomas.izo@gmail.com

Tomáš Ižo is an engineering director in the Machine Perception group at Google Research. He leads a team of engineers and scientists pursuing applied research in the areas of video understanding, video summarization, computational imaging, photography and videography, mobile perceptual computing, creative tools, and scalable, cross-platform media processing infrastructure. The team's work contributes to a wide range of products across Google and Alphabet, such as video understanding for YouTube, enabling content-based video search in Photos, imaging and video features in the Pixel camera, Cloud Platform APIs, video previews in Google Search, improving the quality of ads, and many more.

Recent work from the team includes improving YouTube thumbnails with deep neural nets, the YouTube-8M dataset for video understanding research, Motion Stills apps on Android and iOS for machine intelligence-enabled microvideo creation, and learned image super-resolution (RAISR).

Tomáš came to Google by way of the Massachusetts Institute of Technology, where he received a Ph.D. in computer science in 2007. His doctoral research focused on motion tracking, scene activity analysis and visual attention models. Prior to attending MIT, Tomas received a B.S. in computer science from Yale University, where he did undergraduate research in computational models of visual cortical processing.

Jeff Krichmar

C-CI-W

Jeffrey Krichmar, Professor, Cognitive Science & Computer Science Department,
University of California, Irvine (UCI)
jkrichma@uci.edu

Jeffrey Krichmar specializes in computational neuroscience and neuromorphic computing. Current research is on adaptive algorithms, creating neurobiologically plausible network simulations and constructing brain-based robots whose behavior is guided by neurobiologically inspired models, based on large-scale spiking neural networks.

Ken Kruetz-Delgado

C-CI-W

Director of the Calit2/QI Pattern Recognition Laboratory (QI-PRLab)
kruetz@eng.ucsd.edu

Ken Kruetz-Delgado, IEEE Fellow, is the Director of the Calit2 Qualcomm Institute Pattern Recognition Laboratory (QI-PRLab) and a Professor of Machine Learning, Statistical Signal Processing, and Robotics in the Electrical & Computer Engineering Department at UC San Diego. His current research includes the development of signal processing algorithms for EEG and fMRI brain imaging, the development of algorithms for signal processing on stochastic, spiking neuronal systems, and power-efficient implementation of machine learning algorithms on next-generation processors.

Falko Kuester

C-CI-W

Professor for Visualization and Virtual Reality, Calit2, UCSD

Dr. Falko serves as the director of the Cultural Heritage Engineering Initiative ([CHEI](#)), the Center of Interdisciplinary Science for Art, Architecture and Archaeology ([CISA3](#)), the Calit2 Center of Graphics, Visualization and Virtual Reality (GRAVITY) and the [DroneLab](#).

Prof. Kuester's research involves developing methodologies and techniques for cultural heritage diagnostics and preservation, including diagnostic and analytical imaging as well as visual and cultural analytics techniques that provide engineers, scientists, art historians and restorers, to intuitively and interactively explore historic artifacts. This research is creating the foundation for the development of digital surrogates of world cultural heritage sites and artifacts, providing a means for researchers, and the public, to study these artifacts and facilitate their preservation. His research interests also include tera-scale scientific data visualization, virtual reality and augmented reality, image-based modeling and rendering, robotics and layered manufacturing.

Anna Shcherbina for Anshul Kundaje

C-CI-W

PhD Student, Department of Computer Science, Computational Genomics, Stanford University
annashch@stanford.edu

Anna's research involves developing deep learning algorithms for identifying pathogenic variants in undiagnosed diseases. Her long-term goal is to contribute to precision medicine by integrating physical activity, medical history, and genetic information to build a more complete picture of patients' health.

Anshul Bharat Kundaje (see Anna Shcherbina)

C-CI-W

Assistant Professor of Genetics and Computer Science at Stanford University
akundaje@stanford.edu

Professor Kundaje's primary research area is large-scale computational regulatory genomics. He specializes in developing statistical and machine learning methods for large-scale integrative analysis of heterogeneous, high-throughput functional genomic and genetic data to decipher regulatory elements and long-range regulatory interactions, learn predictive regulatory network models across individuals, cell-types and species and improve detection and interpretation of natural and disease-associated genetic variation.

Alexandr Kuznetsov

C-CI-W

PhD Student, Computer Graphics and Vision, Computer Science and Engineering, UCSD

Ravi Ramamoorthi (see Ramamoorthi)

CS:AI

Ronald L. Graham Professor of Computer Science and Engineering
Director UCSD Center for Visual Computing

Kuznetsov's research interests focus on realistic computer graphics, and how to denoise Monte Carlo renderings, to speed up the rendering process and its potential applications in virtual reality. Additional work involves rendering wave optics phenomena such as light diffraction and iridescence. One of the joint projects underway by [Kuznetsov and Ramamoorthi](#), involves modifying Kubernetes for FIONA8s for computer vision and computer graphics research. Research includes: physically based rendering, GPU based ray tracers, sampling and reconstruction and deep learning.

Louise Ledeen

CS:AI

Director of Strategic Development, CineGrid
Moderator and Organizer, CineGrid Spotlight: AI
louise@cinagrid.org

Louise Ledeen's career has gravitated to the intersection of technology and the creative process. As Louise Etra, she worked as a video artist and curator of technology-based art, introducing video art, computer graphics and electronic music to the public at museums and galleries around the world. Louise was a co-founder of Diaquest, a company which built video animation controllers, for scientists, engineers and artists to record computer graphics to video from PC, Apple and SGI computers. At Silicon Graphics, Louise Ledeen worked with engineers and strategic partners, to deploy HPC computers, graphics, shared storage and emerging network technologies, to change the landscape of content creation software and workflows, by enabling early digital film scanners to reach 24fps realtime performance, which was the impetus for interactive color grading and finishing.

More recently, on behalf of NetApp, Louise worked with IBM Cloud and NVIDIA, to create the first commercial, secure virtualized cloud implementation, for geographically dispersed creative and engineering teams around the world, to interact with their applications, without latency, and collaborate in realtime on graphics-intensive workloads. Louise is now focusing her efforts on researching machine learning applications to accelerate the restoration of scientific and cultural heritage films, with the goal of implementing advanced, prescriptive methodologies for collaborative restoration efforts, at scale.

Eric Lo (see Falko Kuester)

C-CI-W/D

Design and Development Engineer, Cultural Heritage Engineering Initiative (CHEI), UCSD

Eric is focused on hardware and software development and test procedures in robotics and embedded systems for Engineers for Exploration projects, at CHEI, UCSD. His expertise includes unifying the hardware and simulation test designs to accelerate development, simulation tests and integration.

Robert Lockwood Logan IV (with Padhraic Smyth, see: Smyth)

C-CI-W

PhD Student, School of Information & Computer Science, UCI

rlogan@uci.edu

Robert' current research focuses on using multimodal data to solve problems in information extraction. Logan and Smyth's research work is in the fields of biomedicine and climate science, with a focus on machine learning, intelligent systems, pattern recognition, algorithms for modeling climate science and predicting with dynamical data, using Markov and latent variable models.

Chris McFarland

C-CI-W/D

Programmer, Analyst, Qualcomm Institute, UCSD

Chris McFarland is the programming lead on a library-based online database for 'big' 3D data assets like high resolution stereo photographs and laser scan files of archaeological sites. The "CAVEbase" project was developed by Chris McFarland in collaboration with the Center for Cyber-Archaeology and Sustainability and the Cultural Heritage Engineering Initiative (CHEI)

cmcfarland@ucsd.edu

Dmitry (Dima) Mishin

C-CI-W

Applications Developer at San Diego Supercomputer Center SDSC) and PRP Researcher

dmishin@ucsd.edu

Dima's research is based on hardware and software system integration of complex computing, data and visualization workflows. His recent research is focused on the development and deployment of the large highly distributed kubernetes cluster based on the PRP network, highly distributed Kubernetes cluster based on the PRP network, Other activities include developing components for the Comet supercomputer, one of which is the Virtual Clusters infrastructure, His expertise extends to the GPU-specific uses of Kubernetes, as part of the CHASE-CI initiative. Traceroute tool for visualization of the Perfsonar r is one of the many recent development projects.

Michelle Munson

CS:AI

CEO and Co-Founder

Eluvio

michelle@eluv.io

Michelle Munson is the co-founder and CEO of Eluvio, a new Berkeley-based start-up creating new software technologies for a content-centric Internet. She previously founded East Bay software company Aspera in 2004 and led the company as CEO until May 2017, including through acquisition by IBM in 2014. She and co-founder Serban Simu created the Aspera FASP™ transport technology, an Emmy-award winning technology used throughout the digital media supply chain for high-speed low cost secure digital content transport, replacing satellite and traditional tape based delivery technologies. Michelle holds several patents and is a frequent speaker in the areas of content networking innovation including high performance delivery, machine learning, blockchain security and cloud.

Ms. Munson was the 2016 Charles Swartz awardee and is a SMPTE Fellow. Michelle has dual B.Sc. degrees in Electrical Engineering and in Physics from Kansas State University and was a Goldwater Scholar for achievement in Science and Mathematics, and later a Fulbright Scholar at Cambridge University where she received a postgraduate Diploma in Computer Science.

Emri Neftci

C-CI-W

Assistant Professor, Cognitive Sciences, School of Social Sciences, UCI

eneftci@uci.edu

Dr. Neftci's current research focuses on theoretical and computational modeling of learning in neural systems that exploit the characteristics of neuromorphic hardware. The focus of research efforts by Dr. Emre Neftci and his

collaborators, is on embedded learning to create computing systems that work more like the brain, with the long-term research goal to emulate high-level, computational substrates implementing a class of machine learning algorithms, with the long-term goal of designing brain-inspired (neuromorphic) architectures that efficiently implement these algorithms. The anticipated outcomes from this research range from brain-computer interfaces to goal-directed and adapting robotic systems

Mai Nguyen

C-CI-W

Lead Data Scientist, San Diego Supercomputer Center (SDSC)

mhnguyen@ucsd.edu

As the lead data scientist at SDSC, Dr. Nguyen's conducts research on applying the scalable, distributed application of machine learning algorithms to big data problems including digital medicine, satellite image analysis, geospatial applications, and wildfire prevention.

This SDSC and UCSD combined research team works on applying machine learning, geospatial data, and distributed platforms to interdisciplinary problems. Research expertise encompasses machine learning, deep learning, geology, remote sensing, scientific workflows, and distributed computing. Current projects include wildfire modeling and management, data integration and visualization for disaster response, and satellite image processing for demographics analysis and land cover mapping.

Vid Petrovic

C-CI-W/D

PhD Student, Department of Computer Science and Engineering, UCSD Calit2

vipetrov@ucsd.edu

Petrovic is a development engineer, researching techniques for providing virtual access to remote study sites via the interactive visual exploration, manipulation, analysis, and annotation of massive heterogeneous datasets, such as: point clouds, images and mesh/vector models, with a current focus on applications in digital archaeology/cultural heritage and marine ecology.

Joel Polizzi

C-CI-W/D

Visualization Engineering Technician at UCSD

jpolizzi@eng.ucsd.edu

Joel works with researchers, faculty and staff for all activities pertaining to media and media production, at Calit2 Qualcomm Institute (QI). This includes the operation and maintenance of a 200-seat auditorium, 32 screen visualization wall, CAVE VR environments, performance technical design and technical aspects of conferences and an on-site gallery.

Ravi Ramamoorthi

CS:AI & C-CI-W

Ronald L. Graham Professor of Computer Science and Engineering and Director UCSD Center for Visual Computing

ravir@cs.ucsd.edu

Ramamoorthi's personal research work on spherical harmonic lighting and irradiance environment maps is now widely adopted by motion picture special effects companies, such as Pixar and Weta Digital, and used by popular gaming titles, such as the Halo series. As the Founding Director of UCSD Center for Visual Computing, Ravamoorthi's research group develops the theoretical foundations, mathematical representations and computational models for the visual appearance of objects, digitally recreating or rendering the complexity of natural appearance. The group's research program cuts across computer graphics, computer vision and signal processing with applications in sparse reconstruction and frequency analysis, interactive photorealistic rendering, acquisition and representation of data-driven appearance, image and video editing, light-field cameras, physics-based vision and lighting-insensitive recognition.

Anna Shcherbina

C-CI-W

PhD Student, Department of Computer Science, Computational Genomics, Stanford University
annashch@stanford.edu

Anna Shcherbina is interested in developing algorithms that utilize machine learning and data mining approaches to derive medically relevant conclusions from multi-layer genomics data. Research is focused on developing deep learning algorithms for identifying pathogenic variants in undiagnosed diseases, and mining big data resources such as the “UK Biobank” and the “MyHeart Counts” mobile health data for meaningful associations between physical activity and health.

Anshul Kundaje and Anna Shcherbina are working together to develop machine learning methods for interpretive, integrative analysis of functional genomic data. Long-term research involves the development of statistical and machine learning methods for integrative analysis of heterogeneous, high-throughput biological data to obtain a systems-level view of cellular mechanisms.

Jürgen Schulze

C-CI-W

Adjunct Professor, Computer Science and Engineering, UCSD and Research Scientist, QI, Calit2, UCSD
jschulze@ucsd.edu

Schulze’s research focuses on scientific visualization, visual data analytics and volume rendering, and includes: machine learning for interactive algorithms for augmented and virtual reality applications, human-computer interaction, such as 3D user interface design/evaluation, and mobile devices for graphics application.

Larry Smarr

C-CI-W & CS:AI

Director, California Institute for Telecommunications and Information Technologies
ismarr@ucsd.edu

Dr. Larry Smarr, Computer Science and Engineering, UCSD is PI of the NSF Cognitive Hardware and Software Ecosystem Community Infrastructure (CHASE-CI) grant, which is providing high-speed network access over the NSF-funded Pacific Research Platform to hundreds of affordable GPUs, networked together into a High Throughput Computing purpose-built cloud, and to a variety of non von Neumann (NvN) architectures, all running a wide array of machine learning algorithms, optimized for performance on both GPUs and NvN chips. In addition, CHASE-CI will enable detailed measurements of component and system performance and energy efficiency, thereby providing fast throughput computation with the lowest power consumption.

Padhraic Smyth

C-CI-W

Professor, Department of Computer Science, UCI and Director of UCI Data Science Initiative
smyth@ics.uci.edu

Padhraic Smyth is a Professor at UCI, Department of Computer Science with a joint appointment in Statistics, He is the Director of the Center for Machine Learning and Intelligent Systems. His research interests include machine learning, data mining, pattern recognition, and applied statistics.

Adam Tilghman

C-CI-W

IT Services/Academic Technology Services, UCSD
agt@ucsd.edu

Adam Tilghman is an IT systems architect, specializing in optimizing applications and systems architecture, with a focus on instructional use of HPC systems. Tilghman is responsible for implementing FIONA8 systems with Kubernetes across all UCSD instructional laboratories.

Director of Cyberinfrastructure and Research Computing
University of California Merced, Office of Information Technology
jweekley@ucmerced.edu

Mr. Weekley works closely with the faculty to develop scalable, sustainable approaches to high performance computing, high-speed networking (Science DMZs) and large-scale visualization through the UC Merced Wide Area Visualization Environment (WAVE) - similar in to the UCSD SunCAVE. He is active in collaborative efforts such as the Pacific Research Platform and CHASE-CI.

After over fifteen years as a member of the Research Faculty at the MOVES Institute at the Naval Postgraduate School in Monterey, CA. Mr. Weekley spent a year in New Zealand, as the Science and Education Manager for the New Zealand national network REANNZ. His work as applied to science and engineering programs, includes 3D graphics, motion capture systems, high-performance computing architectures and tiled-display visualization environments. Jeffrey's research in infrastructures design, development of cloud services and advanced media workflows, include network-distributed 4K digital video and film restoration applications. These groundbreaking projects, include the world's first dual live-streaming 4K from California to Japan; deploying a live-streaming 4K camera aboard the USN Airship MZ-3A; deploying HD and 4K video in autonomous systems; and developing video cloud services for the Navy. Mr. Weekley is well known in the CineGrid community for asking the first question.