Dynamic Capability Composition at the Digital Continuum from Edge to HPC

Keynote presentation for the Chesapeake Large-Scale Analytics Conference 2022
October 26, 2022 – Annapolis, MD

İlkay ALTINTAŞ, Ph.D.

Chief Data Science Officer & Division Director of Cyberinfrastructure and Convergence Research and Development, **San Diego Supercomputer Center**
Founding Fellow, **Halıcıoğlu Data Science Institute**
Founding Director, **Workflows for Data Science Center of Excellence**
Founding Director, **WIFIRE Lab**

University of California, San Diego

Joint Faculty Appointee, ***Los Alamos National Laboratory***

UC San Diego

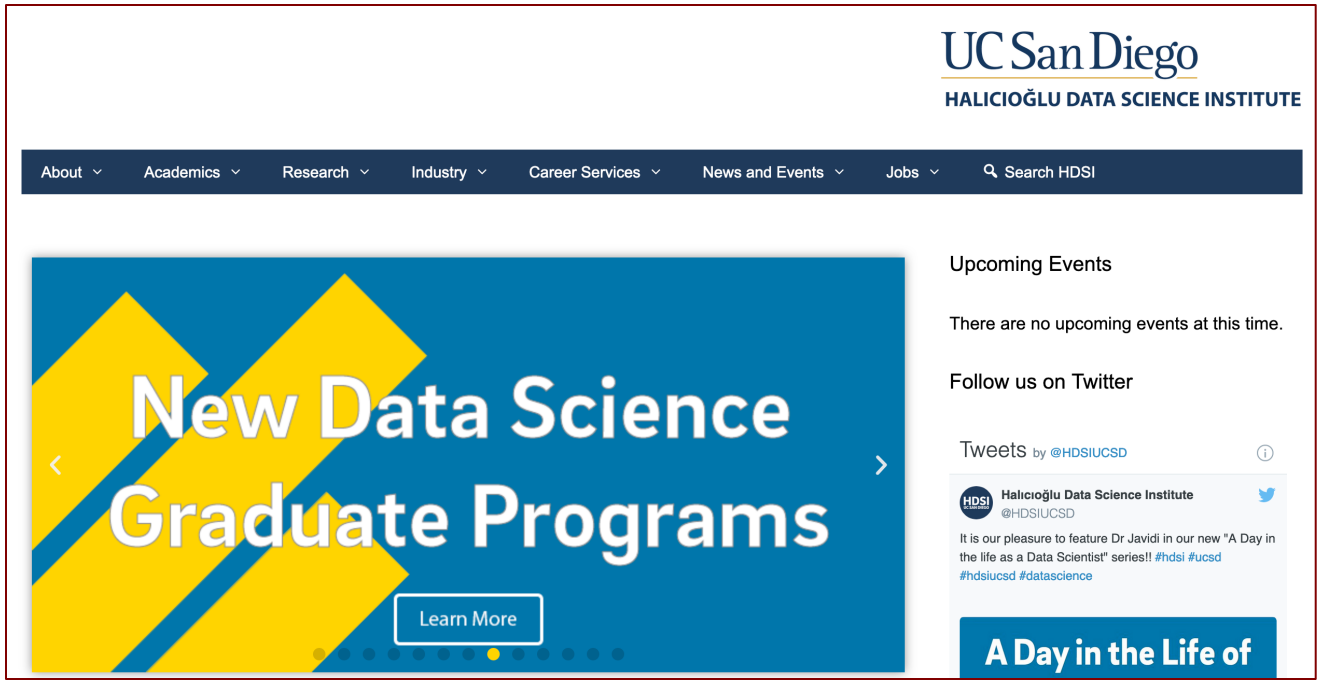

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Cyberinfrastructure and Convergence Research @SDSC

Translating cyberinfrastructure research for impact at scale

Convergence research is driven by a specific and compelling societal problem and works towards integrating innovative and sustainable solutions into society.

Why is data cyberinfrastructure and ML/AI critical?

- *Number of information sources*
- *Complexity of the data*
- *Number and needs of different users*

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

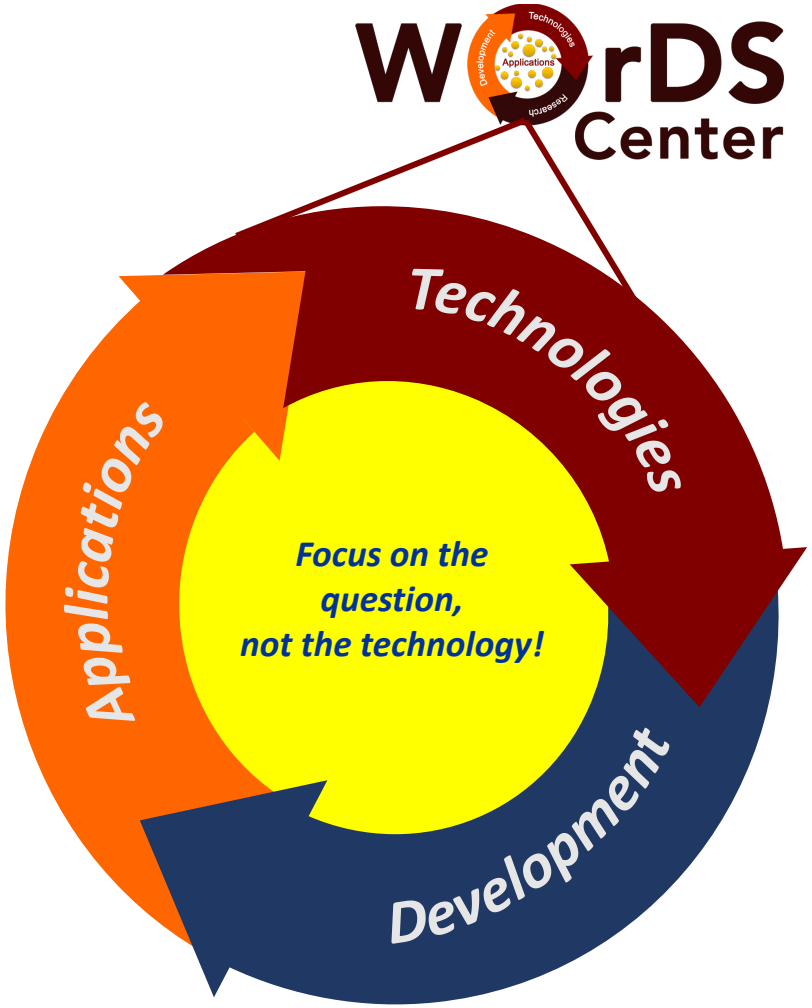

Workflows for Data Science Center of Excellence at SDSC

http://WorDS.sdsc.edu

Research and Development Mission:

- *Methodology and tool innovation to enable collaborative workflow-driven science*
- *Create solutions on top of big data and advanced computing platforms.*

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Evolving advanced computing and workflow ecosystem…

“Big” Data, Computational Science, Data Science, Cyberinfrastructure, and Their Applications

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Application Context # 1

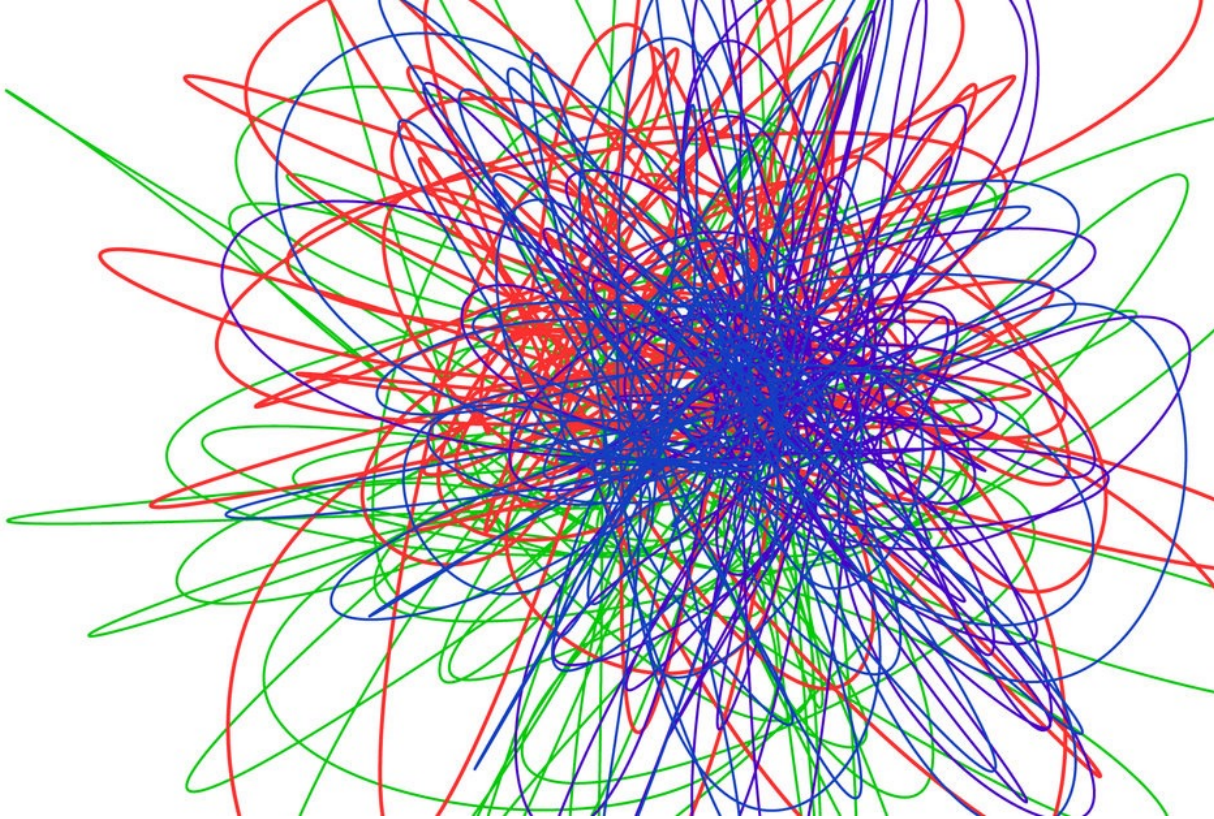

We are in the age of complexity!

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

https://www.wcrp-climate.org/learn-grand-challenges

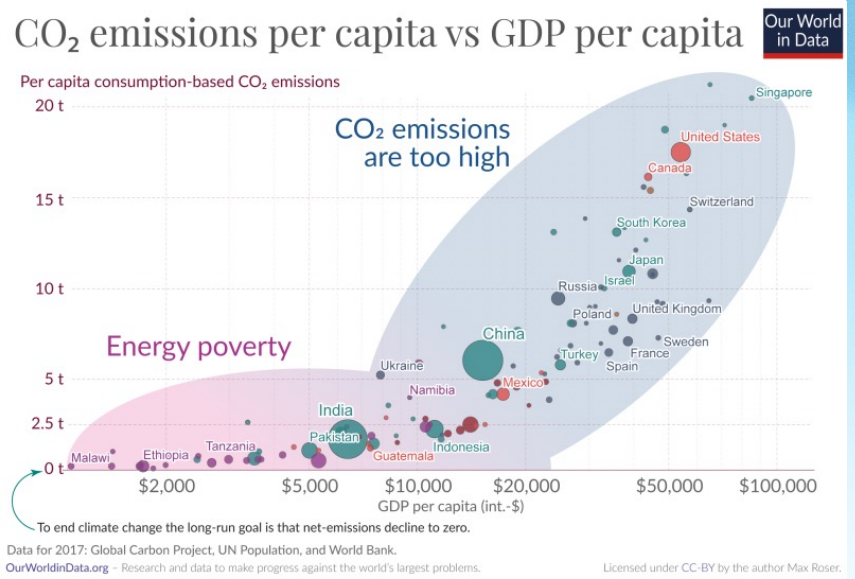

https://ourworldindata.org/worlds-energy-problem

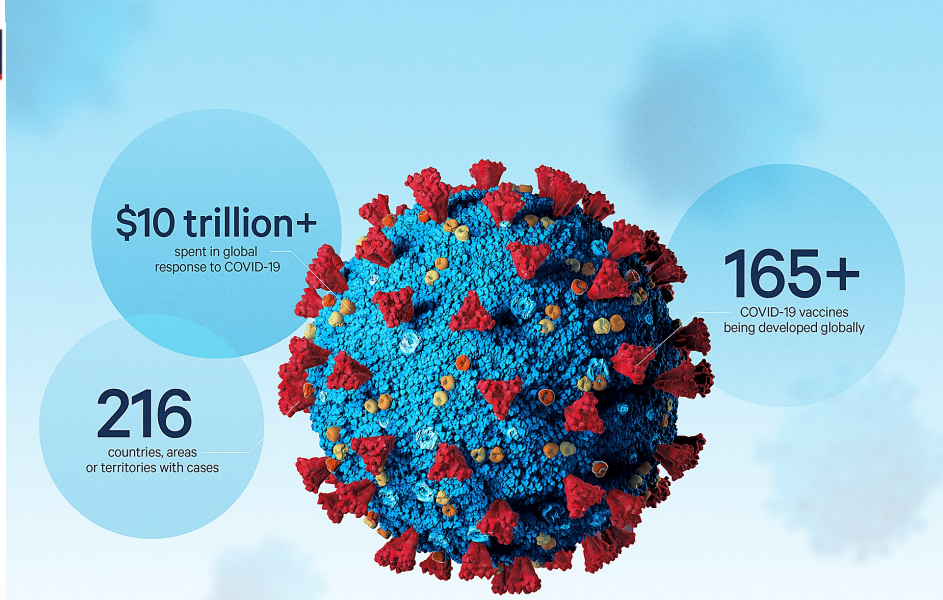

https://www.10xgenomics.com/research-areas/infectious-disease

The biggest challenges of our time are too difficult to solve alone!

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Advanced computing and data technologies can help, and… …we need to collaborate effectively.

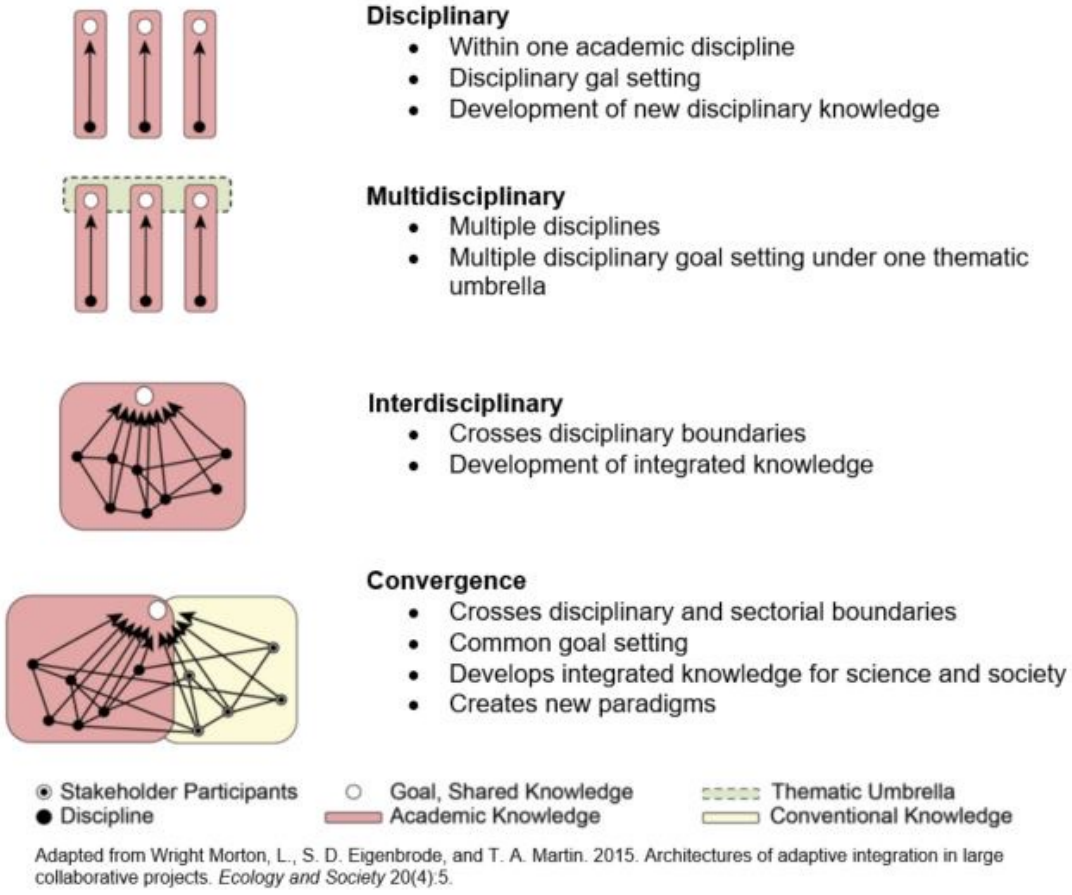

Disciplinary
- Within one academic discipline
- Disciplinary gal setting
- Development of new disciplinary knowledge

Multidisciplinary
- Multiple disciplines
- Multiple disciplinary goal setting under one thematic umbrella

Interdisciplinary
- Crosses disciplinary boundaries
- Development of integrated knowledge

Convergence
- Crosses disciplinary and sectorial boundaries
- Common goal setting
- Develops integrated knowledge for science and society
- Creates new paradigms

Adapted from Wright Morton, L., S. D. Eigenbrode, and T. A. Martin. 2015. Architectures of adaptive integration in large collaborative projects. *Ecology and Society* 20(4):5.

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Building Cyberinfrastructure Systems for Impact

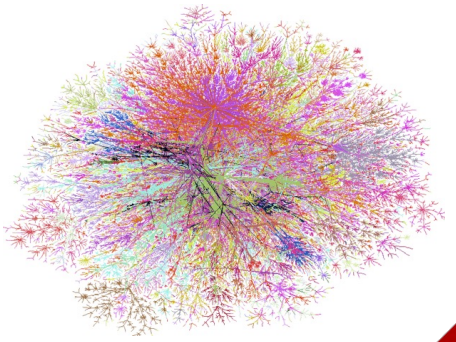

Data

Cyberinfrastructure

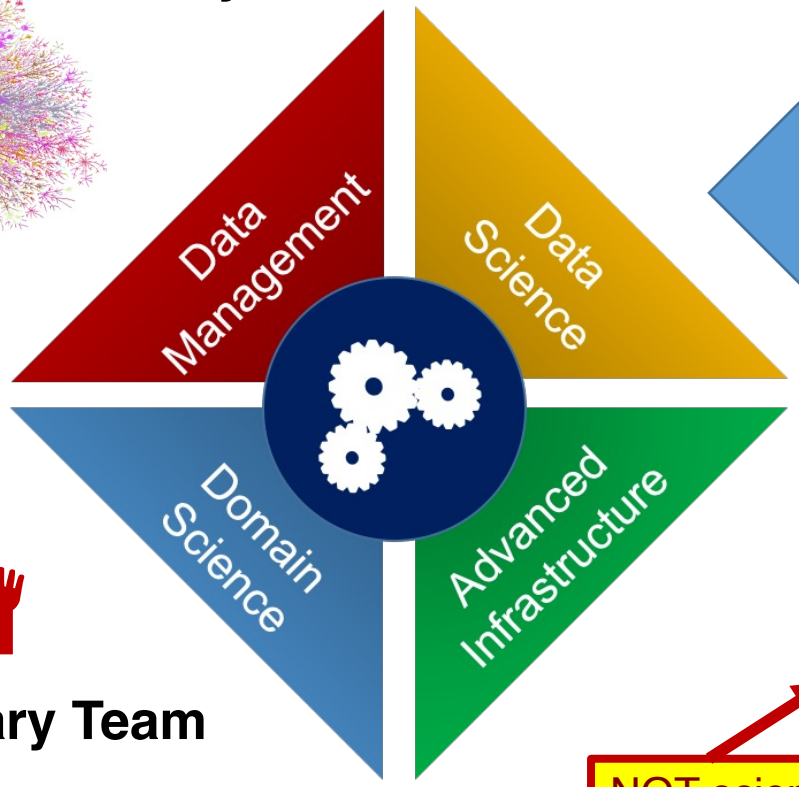

Crossdisciplinary Team

Convergence Research

Translational Innovation

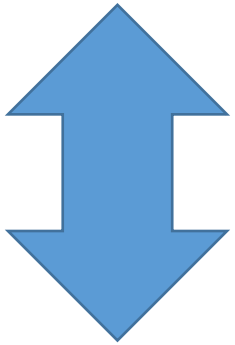

Scientific, Societal and Educational Impact

NOT science and tech to practice, it is practice to science and tech!

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego HALICIOĞLU DATA SCIENCE INSTITUTE

Translational Computer Science

https://www.computer.org/csdl/magazine/cs/2021/05/09547673/1x9TDQr5c2c

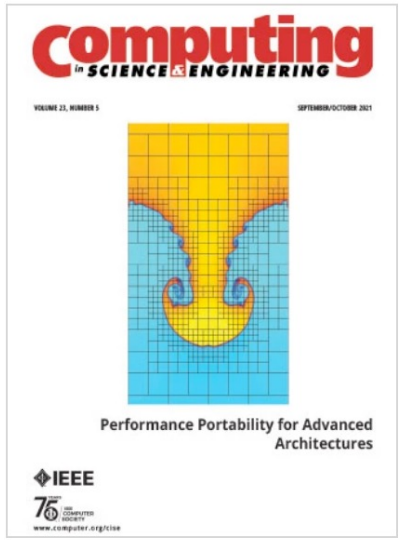

Computing in Science & Engineering

Translational Computer Science for Science and Engineering

Sept.-Oct. 2021, pp. 5-6, vol. 23
DOI Bookmark: 10.1109/MCSE.2021.3109962

Authors

Manish Parashar, University of Utah, Salt Lake City, UT, USA
David Abramson, University of Queensland, Brisbane, QLD, Australia

Systems should enable collaborative teams and thinking!

Think ”teamflows” along with individual workflows.

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Application Context # 2

We are in the age of big data and AI!

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

AI-Integrated Applications at the Digital Continuum

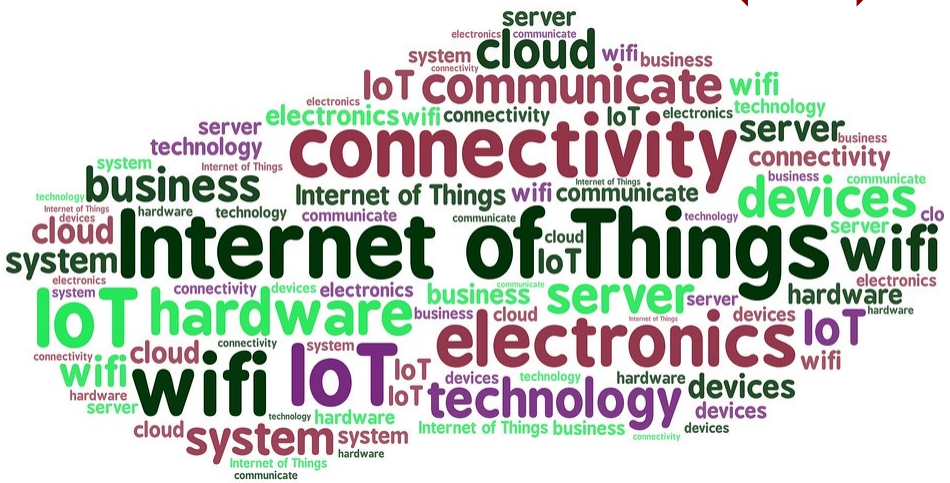

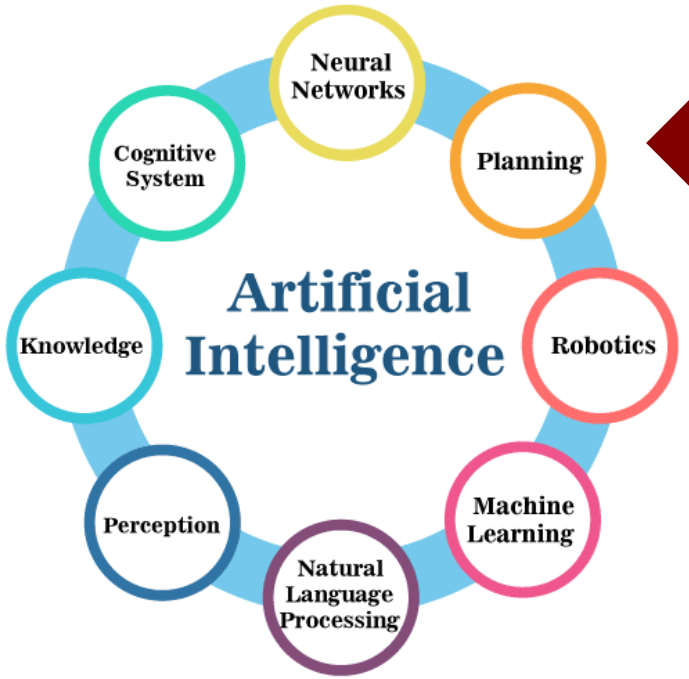

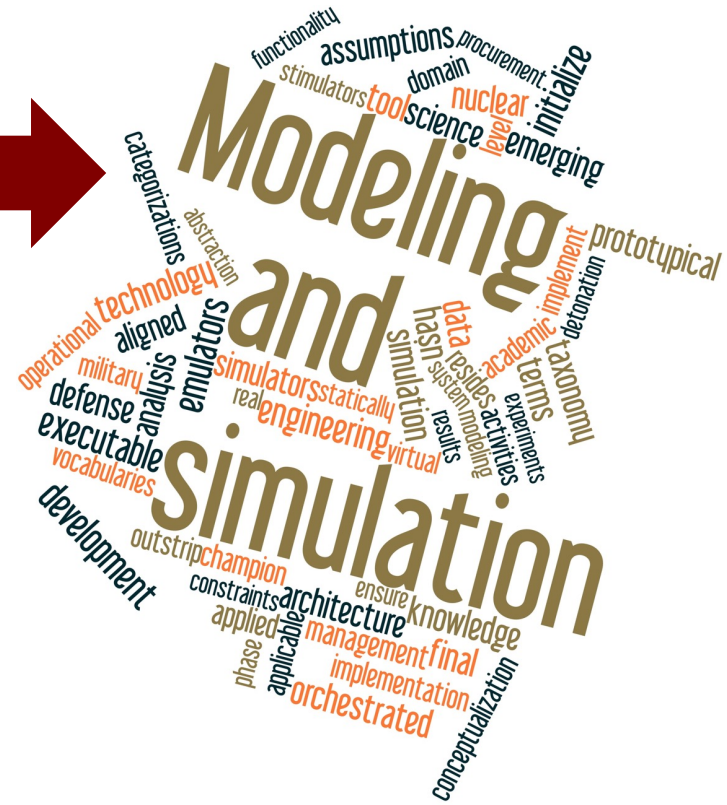

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Big Data and IoT ⟷ **Artificial Intelligence** ⟷ **Modeling and Simulation**

Capability ⟷ Capacity

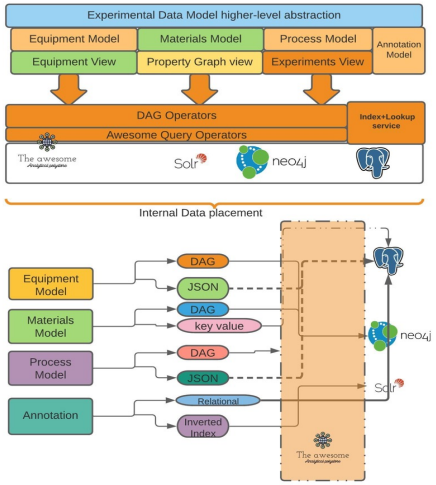

Big Data

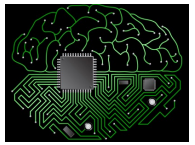

xPU → GPU, CPU, TPU, IPU, QPU,…

Edge

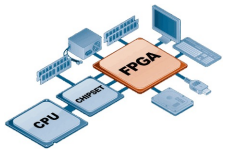

FPGA

Cloud, HPC, Storage

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Evolution of Applications with AI

- Thinking in workflows moves AI closer to solutions in integrated applications.
- Teamwork is key to these workflows -- think workflows and teamflows!
- Enabling science and solutions requires open and transparent sharing and use of big data.
- Continuous integration of computing, data and AI into workflows needs dynamic systems capacity.

Release Date: May 10, 2022

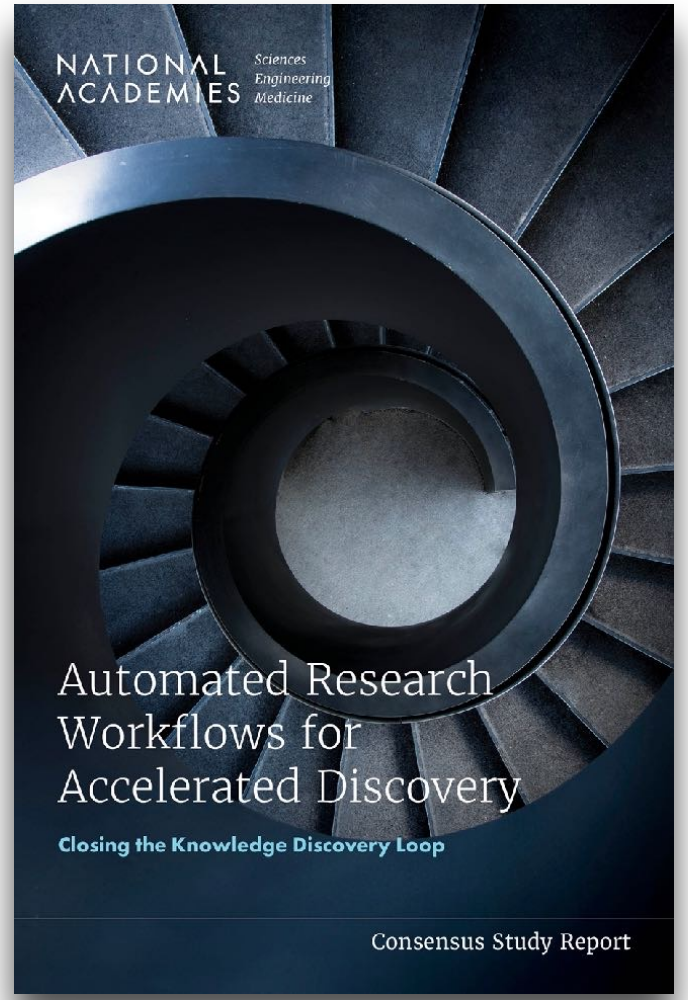

https://www.nationalacademies.org/our-work/realizing-opportunities-for-advanced-and-automated-workflows-in-scientific-research

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Systems should enable seamless integration of AI-integrated application workflows!

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Application Context # 3

We are in the age of many systems!

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Some recent news…

HPC wire

NSF Funds Five New XSEDE-Allocated Systems

August 7, 2020

Aug. 7, 2020 — This summer, five new National Science Foundation (NSF)-
ded advanced computing systems have been awarded to partner
itutions across the country, all with their own unique specialties and
lications. Once deployed, all of these systems will be primarily allocated
ugh XSEDE to help connect researchers, regardless of physical location
he United States, to the system that best suits their research needs.

e recent
ouncements
wards by
NSF is a
ar indicator
he NSF's re-

SDSC SAN DIEGO SUPERCOMPUTER CENTER

ILLINOIS NCSA | National Center for Supercomputing Applications

TACC TEXAS ADVANCED COMPUTING CENTER

PITTSBURGH SUPERCOMPUTING CENTER

PURDUE UNIVERSITY

UNIVERSITY INFORMATION

Jetstream 2 (Indiana University / Texas Advan
University of Arizona / Arizona State University
University Corporation for Atmospheric Resear
University of Hawaii)

The NSF has awarded a $10 million gran
distributed cloud computing system to su
artificial intelligence, and enhanced large
nation. The project is led by the Pervasiv
Indiana University. Jetstream 2 is a follow
was funded in 2014 as the NSF's first pro
engineering research cloud system for thi

Delta (National Center for Supercomputing App

NCSA will integrate Delta into the nationa
through XSEDE. Integration into XSEDE
substantial portfolio of services and suppo
deliver unprecedented advances in resea
collaboration will promote synergy among
include campus, national, and commercia

Anvil (Purdue University)

Anvil, which is funded by a $10 million aw
significantly increase the capacity availab
has been a partner for the past nine year
2021 and will serve researchers for five y
the NSF will support Anvil's operations ar

Neocortex (Pittsburgh Supercomputing Center

A $5 million NSF award will allow PSC to
performance artificial intelligence (AI) sys
fundamentally new hardware to greatly sp
research organization of Carnegie Mellon
of Pittsburgh, will build the new supercom
with Cerebras Systems and Hewlett Pack

Voyager (San Diego Supercomputing Center)

The NSFhas awarded DSC at UC San Di
develop a high-performance resource for
research across a wide swath of science
Called Voyager, the system will be the firs
NSF resource portfolio. In addition to the
equivalent amount of funding is expected

HPC wire

NCSA Gives Update on FABRIC Research Infrastructure Project

May 16, 2022

May 16, 2022 — When the internet launched in the last century it was an amazing feat of science and technology. But it was not designed for the massive data sets, machine-learning tools, advanced sensors and Internet of Things devices that have become central to many research and business endeavors and our homes.

Funded with a $20 million grant from the National Science Foundation, FABRIC (Adaptive Programmable Research Infrastructure for Computer Science and Science Applications) is exploring ways to replace an internet infrastructure that's been showing its age for the last 20 years.

NCSA is one of 13 collaborating institutions helping create a platform for testing novel internet architectures that could enable a faster, more secure internet better suited for today's users and future needs. One that's also able to do things not possible now. The FABRIC project is led by the Network Research and Infrastructure Group at RENCI at the University of North Carolina at Chapel Hill.

NCSA: First with the Gigs

Last fall, NCSA installed a 100-gigabit network connection that's dedicated solely to the FABRIC project, the first FABRIC collaborator to do so.

NCSA already had six 100-gigabit internet connections, says David Wheeler, leader of NCSA's ICI Data Management and Delivery Division and the

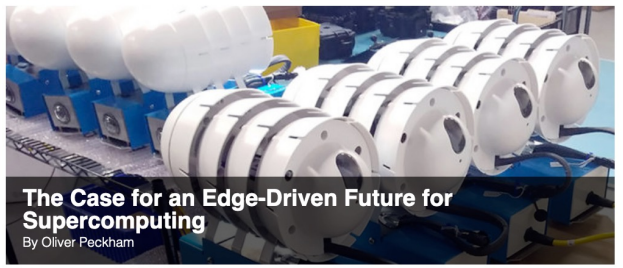

September 24, 2021

"Exascale only becomes valuable when it's creating and using data that we care about," said Pete Beckman, co-director of the Northwestern-Argonne Institute of Science and Engineering (NAISE), at the most recent HPC User Forum. Beckman, head of an Argonne National Laboratory edge computing project called Waggle, was insistent on one thing: edge computing is a crucial part of delivering that value for exascale.

Beckman had opened with a quote from computer architect Ken Batcher: "A supercomputer is a device for turning compute-bound problems into I/O-bound problems." "In many ways, that is still true today," Beckman said. "What we expect from supercomputers is that they're so blindingly fast that really it's bottlenecked on either reading or writing from input or output."

"If we take that concept, though, and flip it over," he added, "then we end up with this idea that edge computing, therefore, is a device for turning an I/O-

Business Wire

GigaIO Awarded Lonestar6 Contract in TACC's First Bid for Composable Disaggregated Infrastructure

di Business Wire

HPC wire

GigaIO Selected to Bring Composability to Bold New National Research Platform

6, 2021

SBAD, Calif., July 16, 2021 — Building on the successes of past
orations with the San Diego Supercomputer Center (SDSC) located at
an Diego, GigaIO, the creators of next-generation data center rack-scale
ecture for artificial intelligence (AI) and high-performance computing
) solutions, is proud to be announcing its low latency universal dynamic
, FabreX, was selected as the technology of choice for the new
type National Research Platform (NRP). This National Science
dation-funded cyberinfrastructure ecosystem is an innovative, all-in-one
m—computing resources, research and education networks, edge
uting devices and other instruments—designed as a testbed to expedite
e and enable transformative discoveries.

plex computational and
workflows underpin many of
ientific research
nges we hope to address
NRP," says Dr. Frank
wein, PI of NRP, and
n director of the San Diego
rcomputer Center. "In areas
erse as public health, high energy physics and wildfire response, this
rch requires that we aggregate disparate computational elements, such
GAs, GPUs, x86 processors and storage systems into highly usable and

GIGAIO

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

The problems we are solving are …

- data and AI-integrated
- heterogeneous
- collaborative

EXISTING SYSTEMS:

- **data** (storage + data/knowledge management)
- **continuum** of distributed computing (edge, cloud, HPC)
- high speed **connectivity** (programmable networks)

GROWING NEEDS:

- Interoperability
- Dynamic scalability
- On-demand interactive access
- Performance measurement
- Composability

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

We need to think and build strategic ecosystems made up of reusable services and solutions toward capability integration!

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

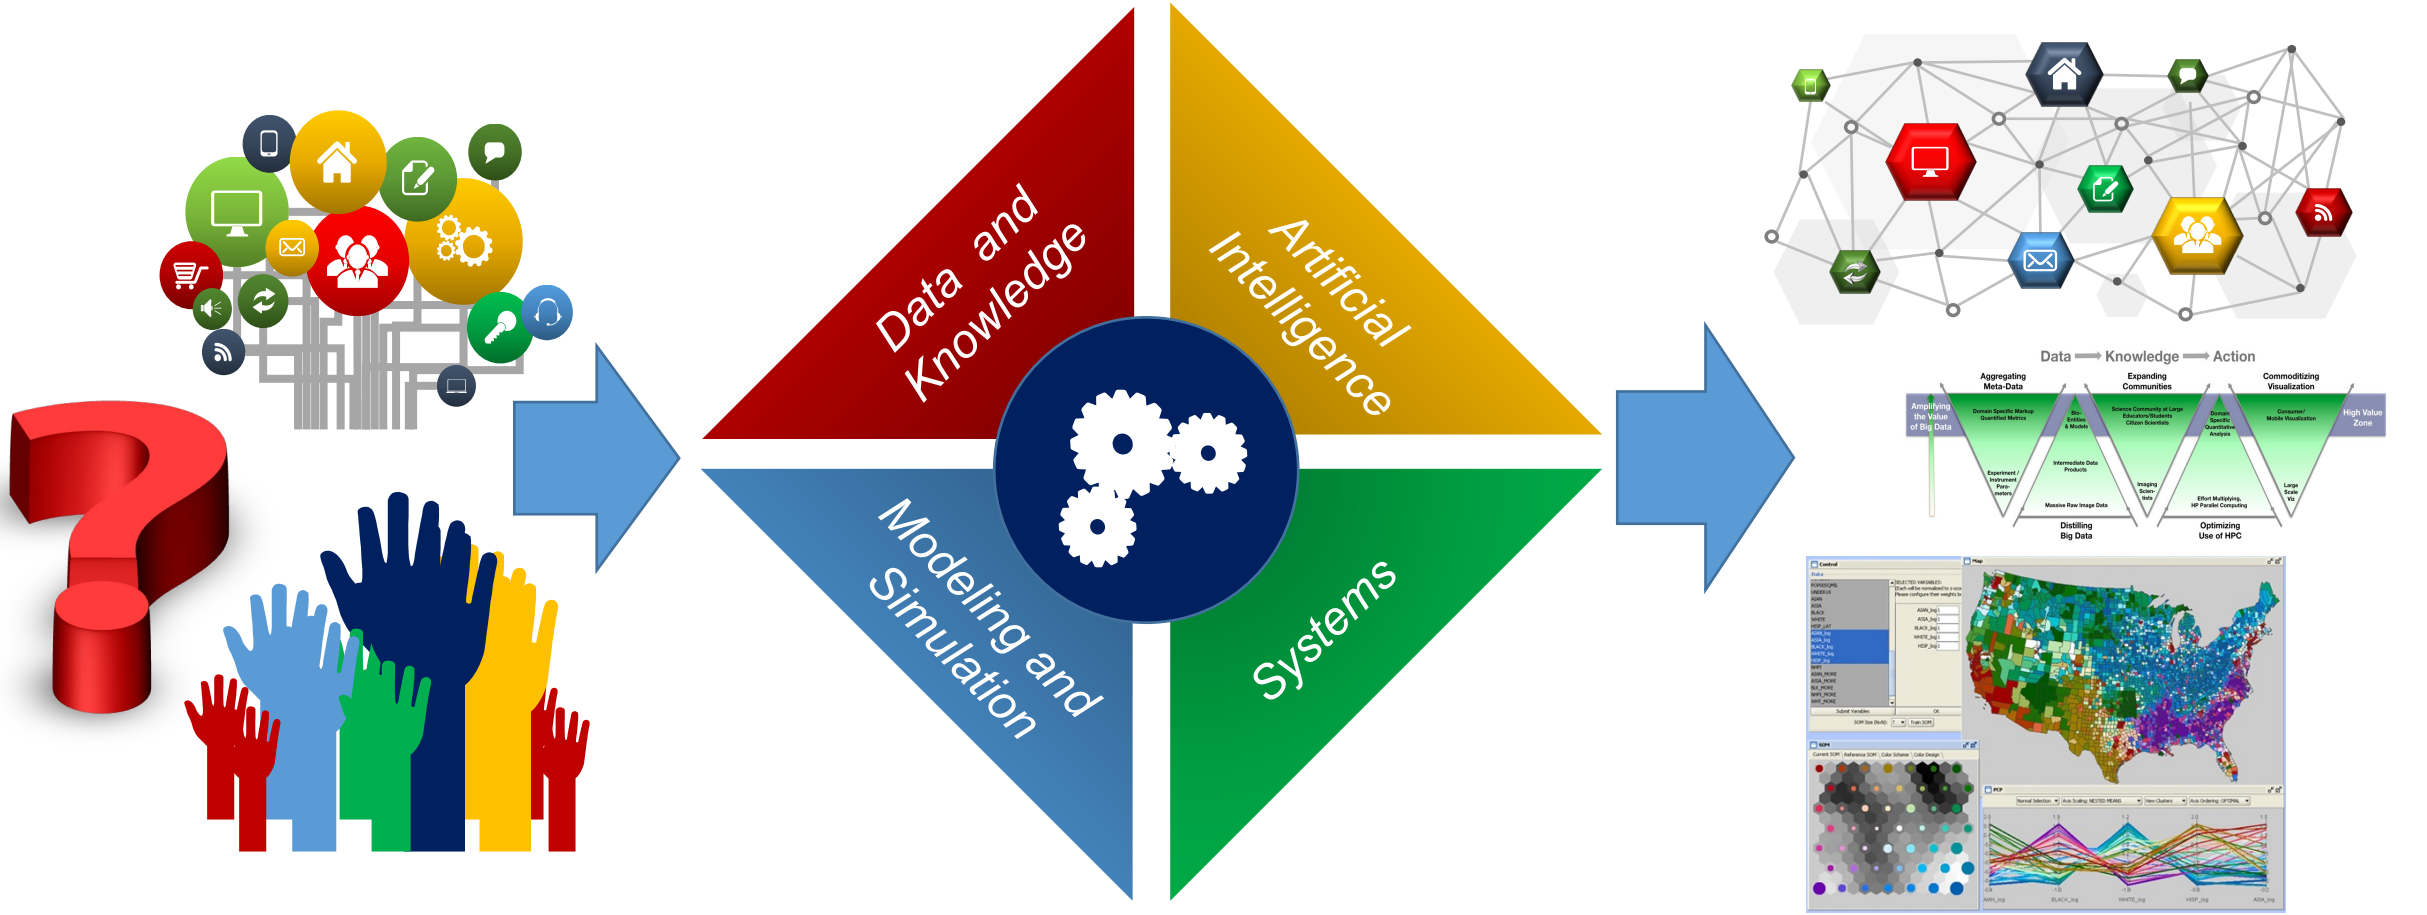

A Typical Application Integration Ecosystem

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Integration requirements…

Dynamic composability matters.

Systems and services are only useful if groups can integrate them into applications.

Tools that enhance teamwork and use need to be coupled with responsible AI systems.

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Dynamic composability matters.

COMPOSABLE SERVICES

e.g., model and data archives, learning and analytics, simulation, training

RESOURCE MANAGEMENT

e.g., container orchestration, optimization

COMPOSABLE SYSTEMS

e.g., GPU, CPU, Big Data, quantum, neuromorphic, SDN, storage

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

DATA LIFECYCLE MANAGEMENT
e.g., active data repositories, long-term archives, knowledge networks, data reuse services

Systems and services are only useful if groups can integrate them into applications.

WORKFLOW MANAGEMENT
e.g., application integration, coordination, optimization, communication, reporting

COMPOSABLE SERVICES

RESOURCE MANAGEMENT

COMPOSABLE SYSTEMS

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

RESPONSIBILITY
e.g., accuracy, privacy, explainability, ethics

REPRODUCIBILITY

TEAM SCIENCE

USE-INSPIRED INTERFACES

e.g., for science, education and scalable practice

Tools that enhance teamwork and use need to be coupled with responsible AI systems.

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

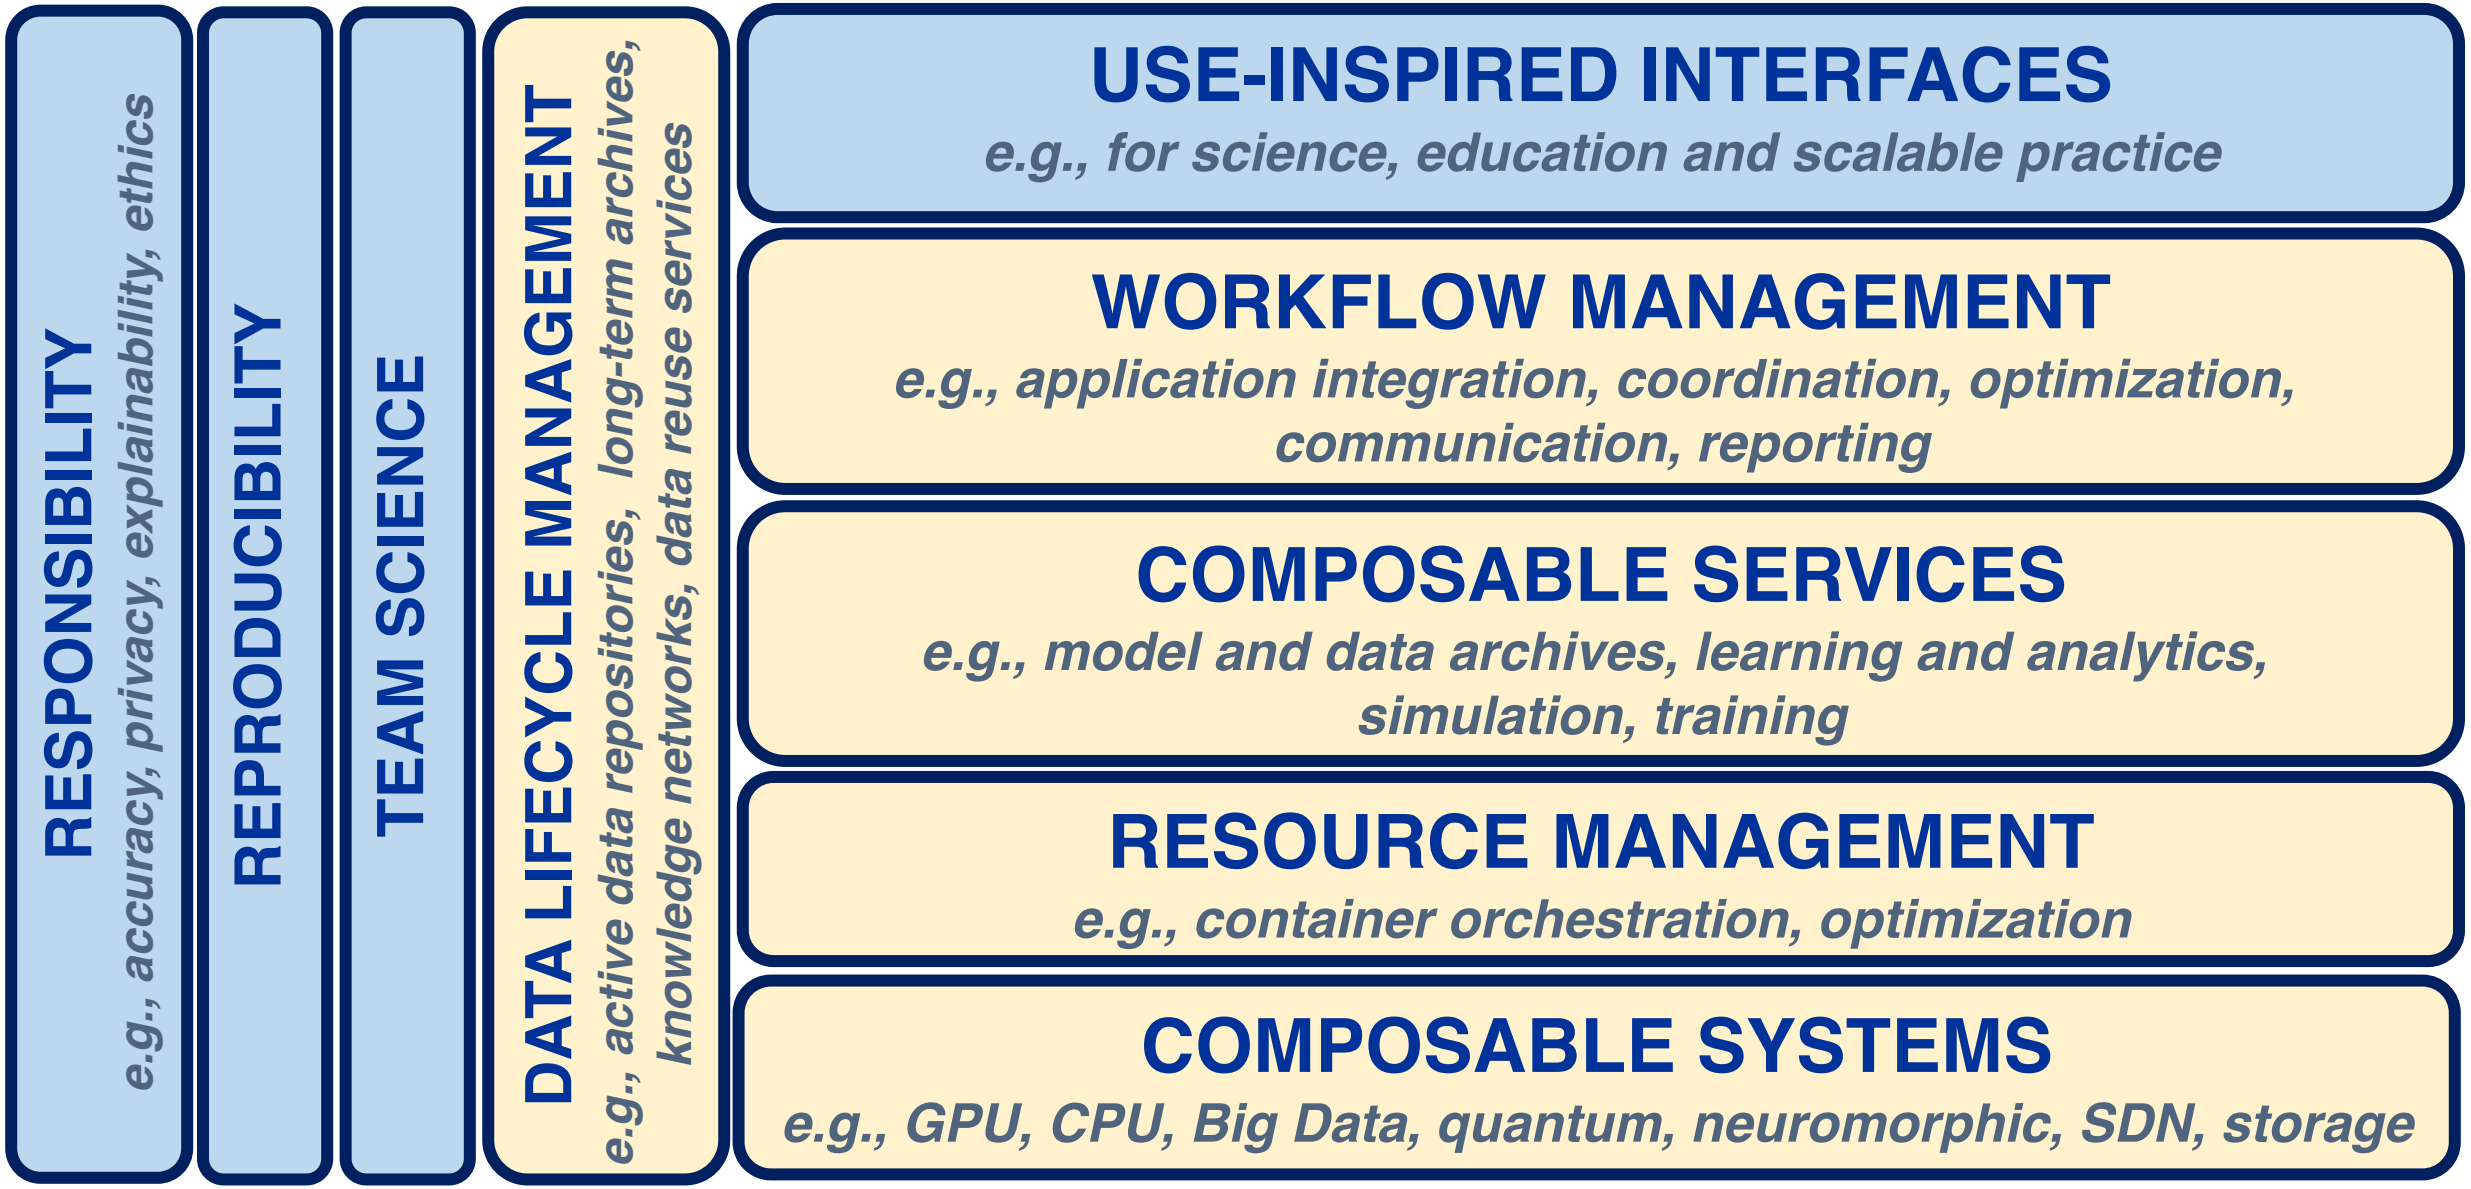

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Smart Composability from Systems to Services

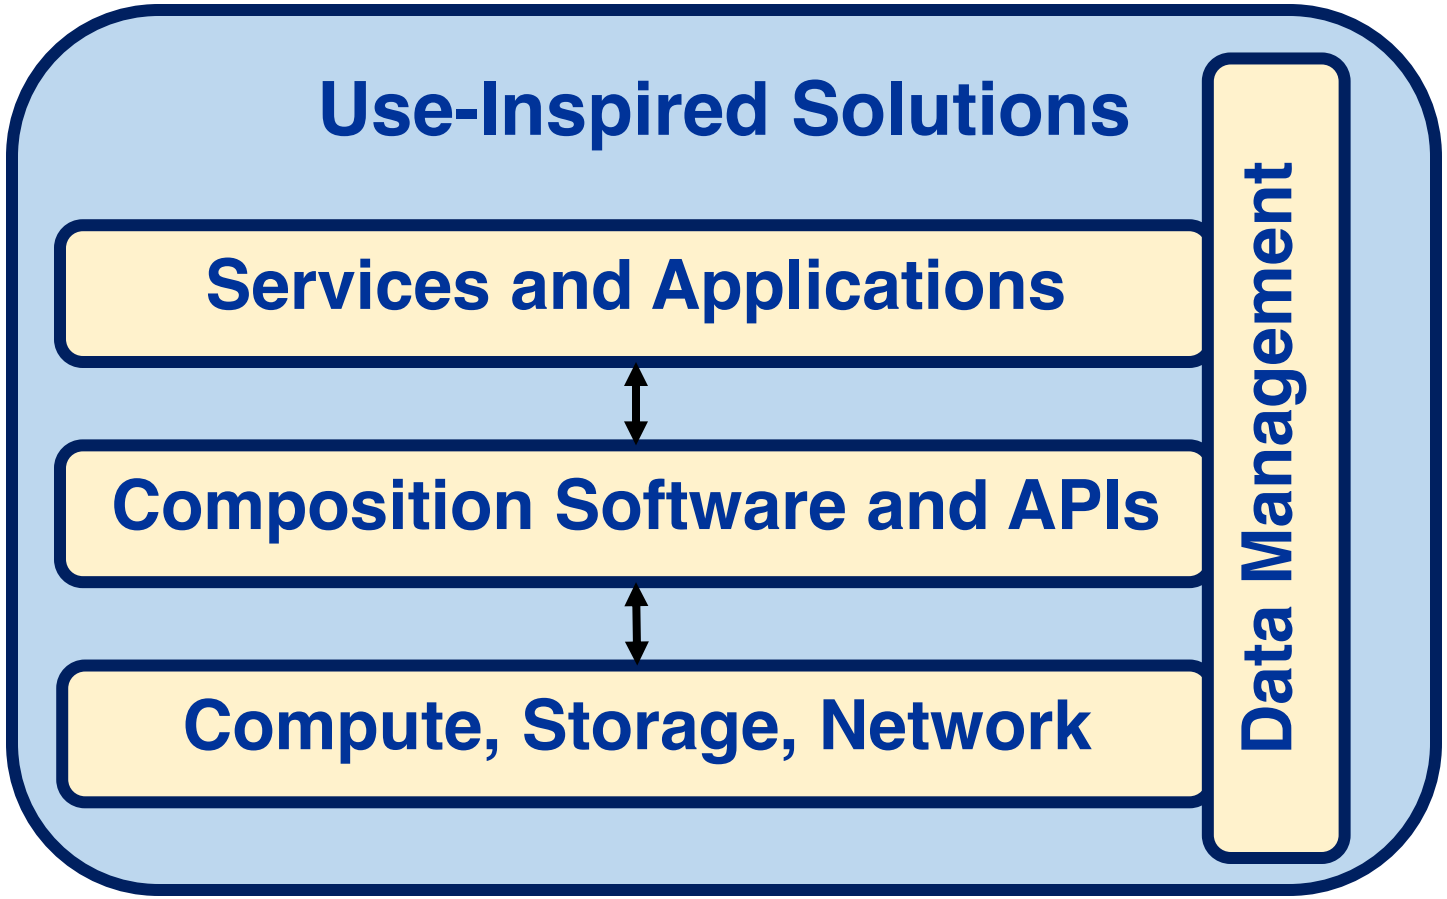

- Dynamically measure, manage and provision resources
- Create plug and play microservices
- Run across many systems
- Capability-based integration
- Improved FAIR data capacity

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Some Composable Systems

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

EXPANSE

COMPUTING WITHOUT BOUNDARIES
5 PETAFLOP/S HPC and DATA RESOURCE

HPC RESOURCE
13 Scalable Compute Units
728 Standard Compute Nodes
52 GPU Nodes: 208 GPUs
4 Large Memory Nodes

DATA CENTRIC ARCHITECTURE
12PB Perf. Storage: 140GB/s, 200k IOPS
Fast I/O Node-Local NVMe Storage
7PB Ceph Object Storage
High-Performance R&E Networking

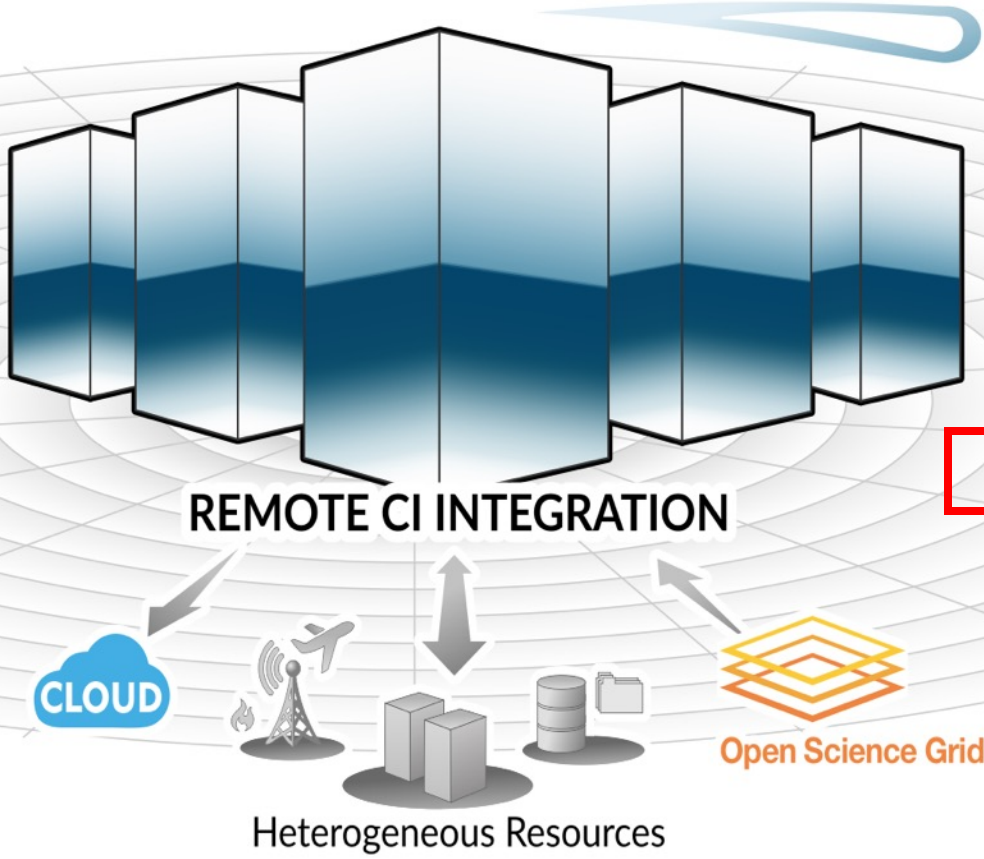

LONG-TAIL SCIENCE
Multi-Messenger Astronomy
Genomics
Earth Science
Social Science

INNOVATIVE OPERATIONS
Composable Systems
High-Throughput Computing
Science Gateways
Interactive Computing
Containerized Computing
Cloud Bursting

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

- Two partitions exist in Expanse
  - Slurm
  - Kubernetes
- Dynamic provision via cm-scale (Bright Computing)
- User Web Portal (ACCESS)

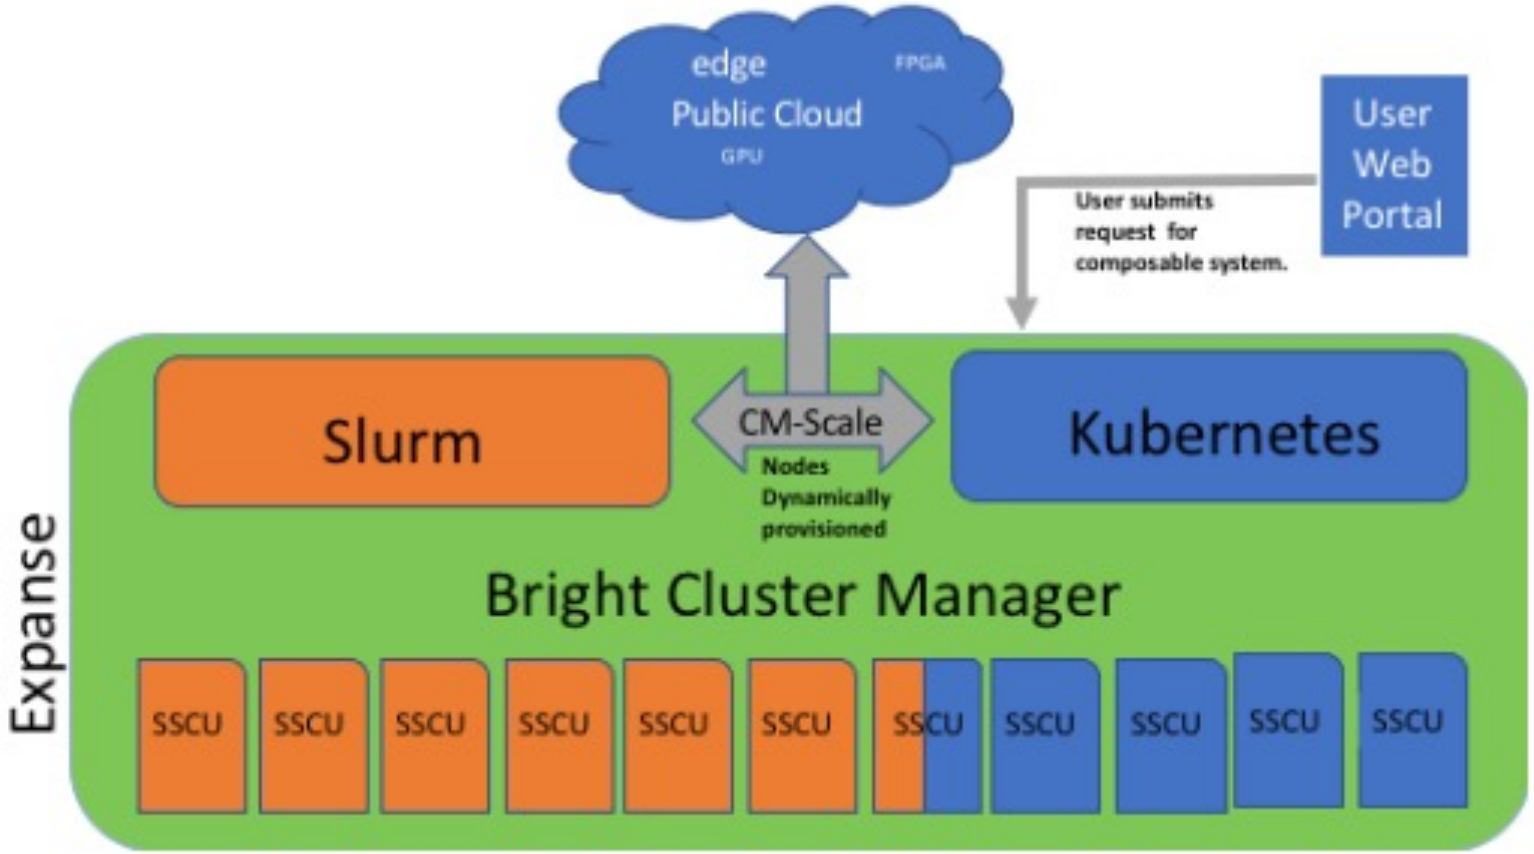

Figure 5.1 Expanse Composable Systems Framework

Nautilus – PRP and NRP

Image Source: https://pacificresearchplatform.org/nautilus

Nautilus is a Multi-Institution Hypercluster Connected by Optical Networks

~200 FIONAs on 25 Partner Campuses Networked Together at 10-100Gbps

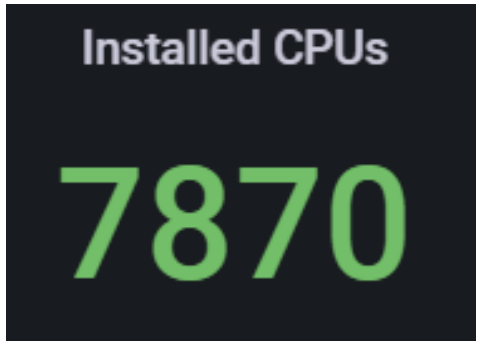

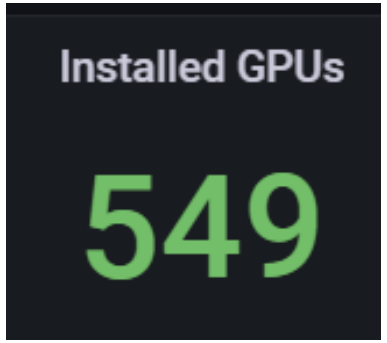

Slide Source: *Larry Smarr, UCSD*

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Composable System using Admiralty for Federation

EXPANSE (Enthalpy) + CHASE-CI (Nautilus)

EXPANSE (Enthalpy)

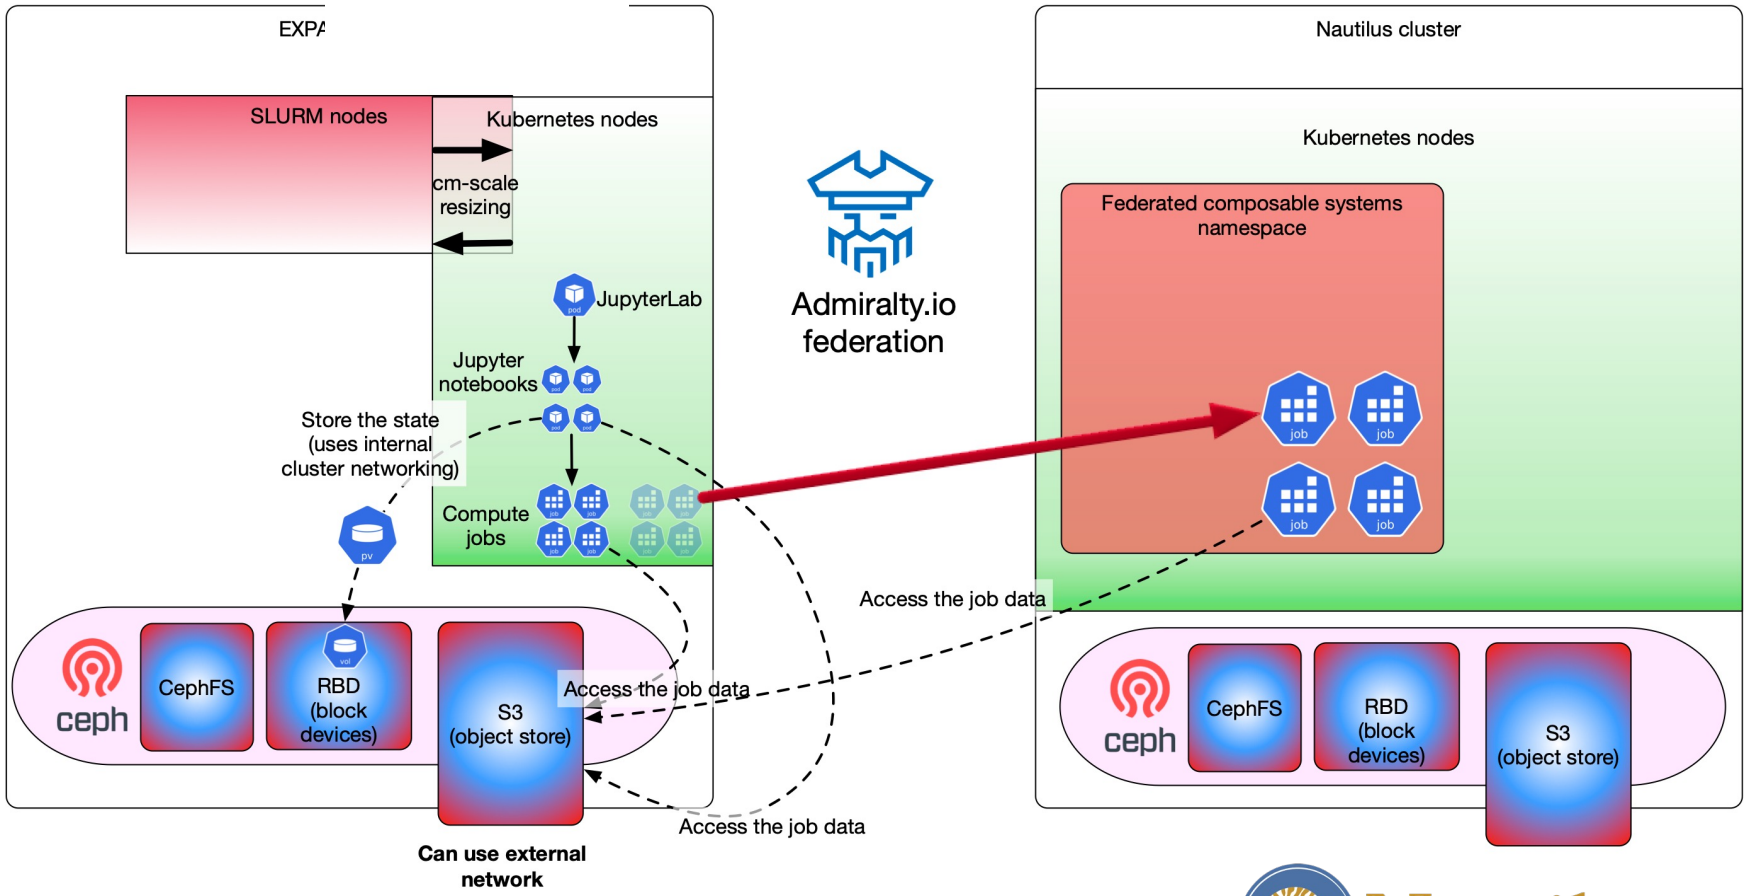

Nautilus

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Admiralty Federation Topologies

- Topologies
  a) Central control plane
  b) Cloud Bursting
  c) Decentralized

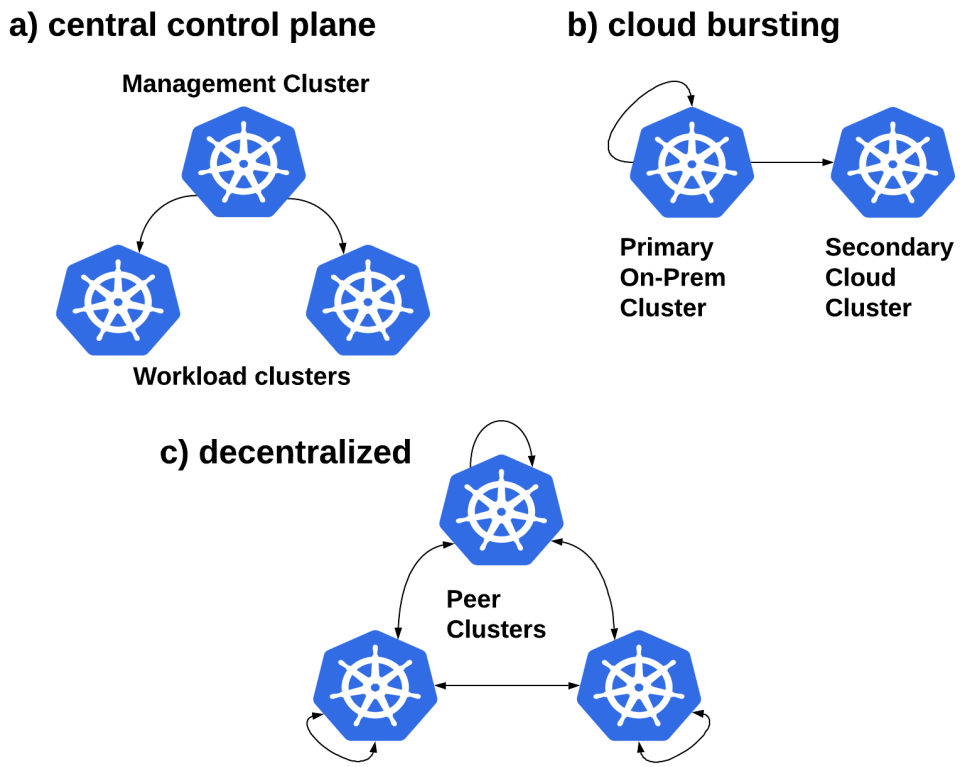

AI@Edge and the Digital Continuum

Slide Source: *Pete Beckman, ANL*

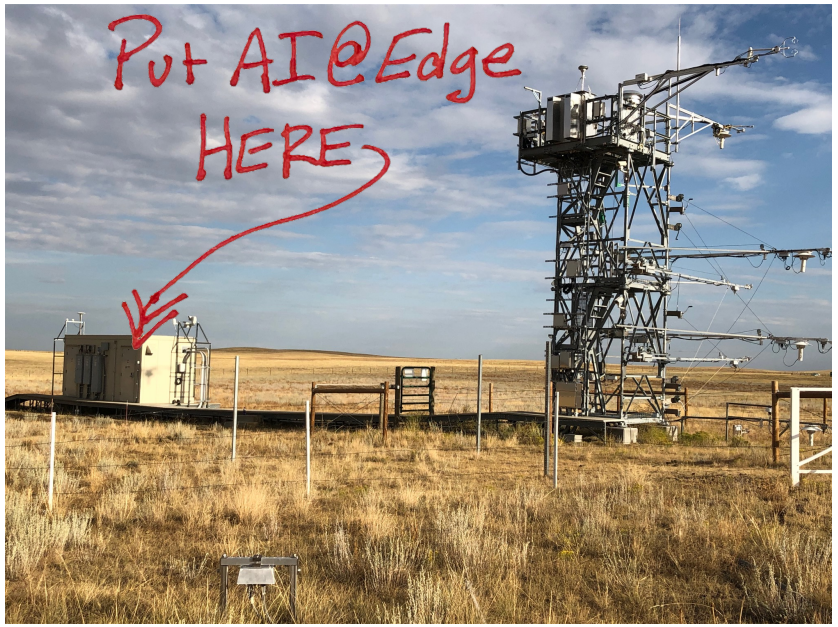

SAGE

Cyberinfrastructure for
AI at the Edge
sagecontinuum.org

NSF

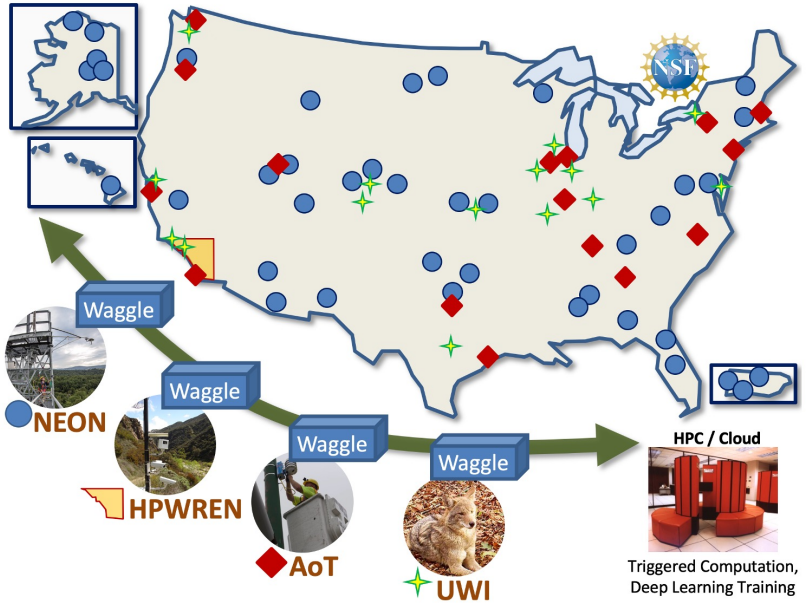

Northwestern University · The University of Chicago · Argonne National Laboratory · Colorado State University · UC San Diego · George Mason University · Array of Things · ARM · neon National Ecological Observatory Network · HPWREN · SDSC San Diego Supercomputer Center · Lincoln Park Zoo For Wildlife. For All.

Education & Training




Leadership Team


Pete Beckman (NU; Director)


Nicola Ferrier (NU; Deputy Dir.


Ilkay Altintas (SDSC; Data)


Charlie Catlett (UIllinois; AoT)


Scott Collis (NU; ARM)


Valerie Taylor UChicago; Broader Impacts)


Jim Olds (GMU; Life Sci, Risk)


Dan Reed (Utah; Architecture)


Eugene Kelly (CSU; NEON)


Irene Qualters (LANL; Advisory Committee Chair)

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Scientific Digital Continuum in Sage

Sensors

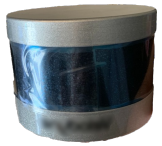

LIDAR

Software Defined Radios

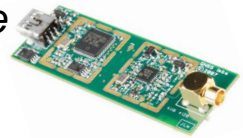

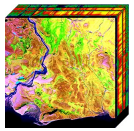

Hyperspectral Imaging

Facilities

Actuators

Servos

Dynamic adaptation

Edge Computing

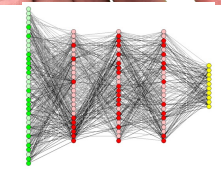

Scientific Data Analysis & Control

Artificial Intelligence
Deep Learning Inference
Lightweight Training

Slide Source: *Pete Beckman, ANL*

admin.sagecontinuum.org

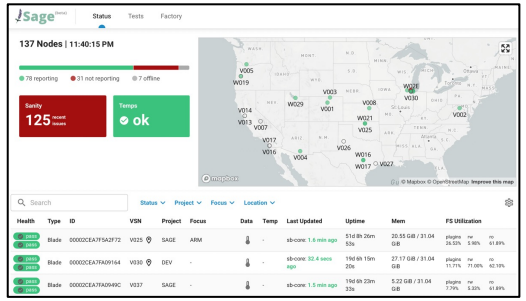

Advanced Wireless & Networking

New inference (model)
Adaptive steering

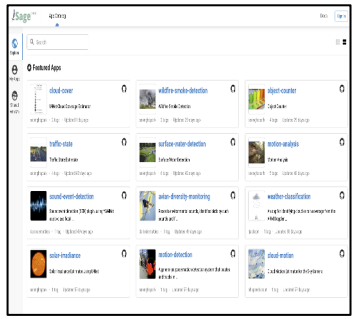

portal.sagecontinuum.org

Computation

Digital Twin
Data Analysis
Machine Learning

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego HALICIOĞLU DATA SCIENCE INSTITUTE

Using Expanse, NRP and Sage for Edge AI Development

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

Integration of NSF EXPANSE, PRP and Sage
A Composable System Deployment of JupyterHub

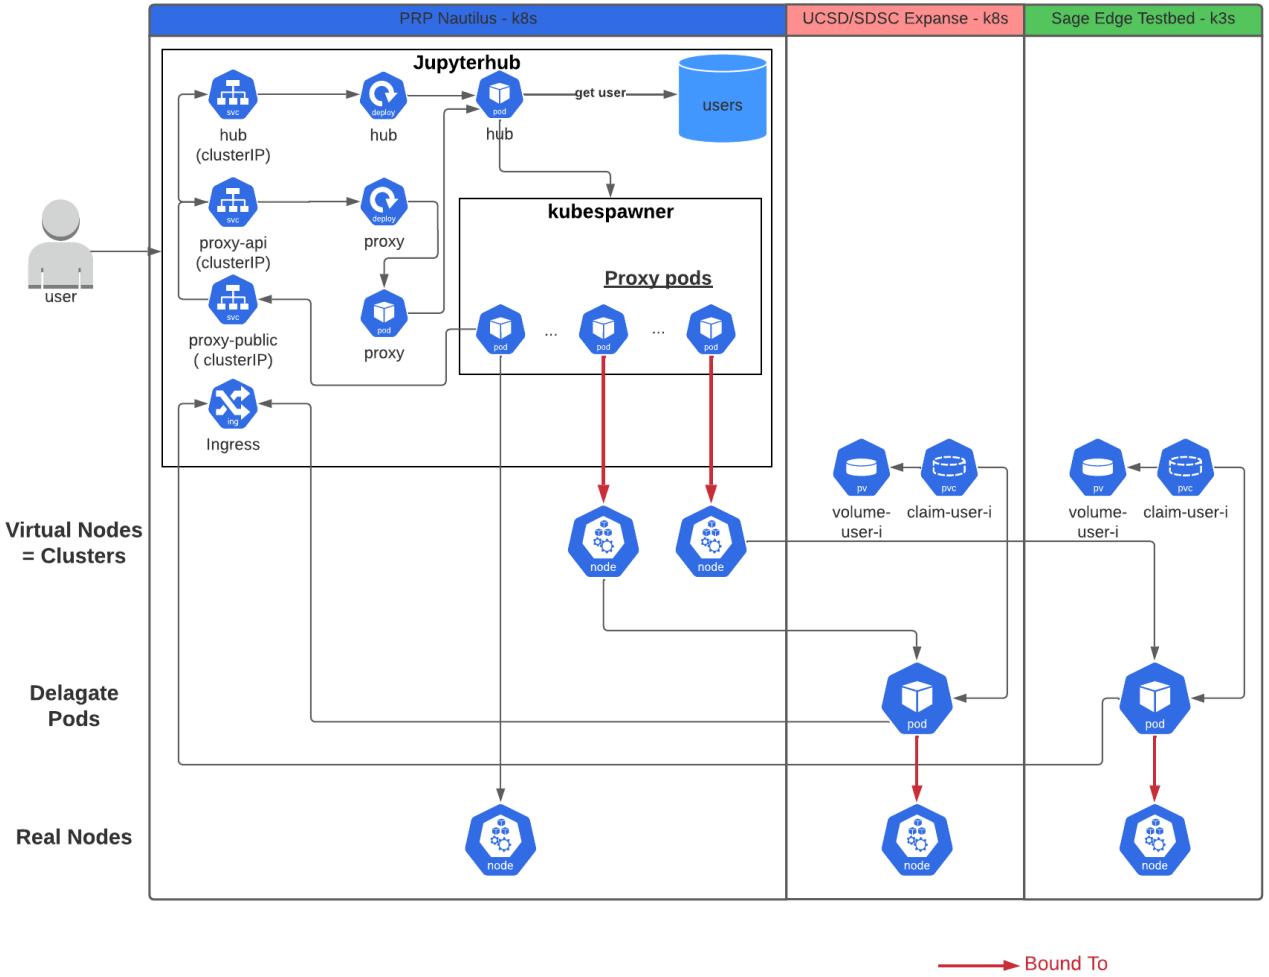

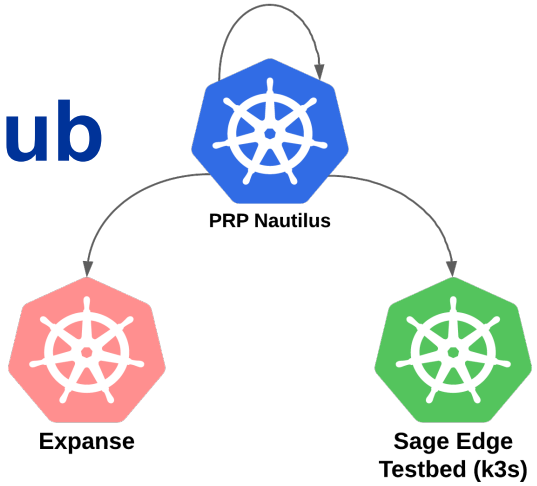

- Edge-Cloud Unified Environment for prototyping AI models to deploy on the Edge
- A user can easily be provided the right environment for developing their AI Edge Application

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

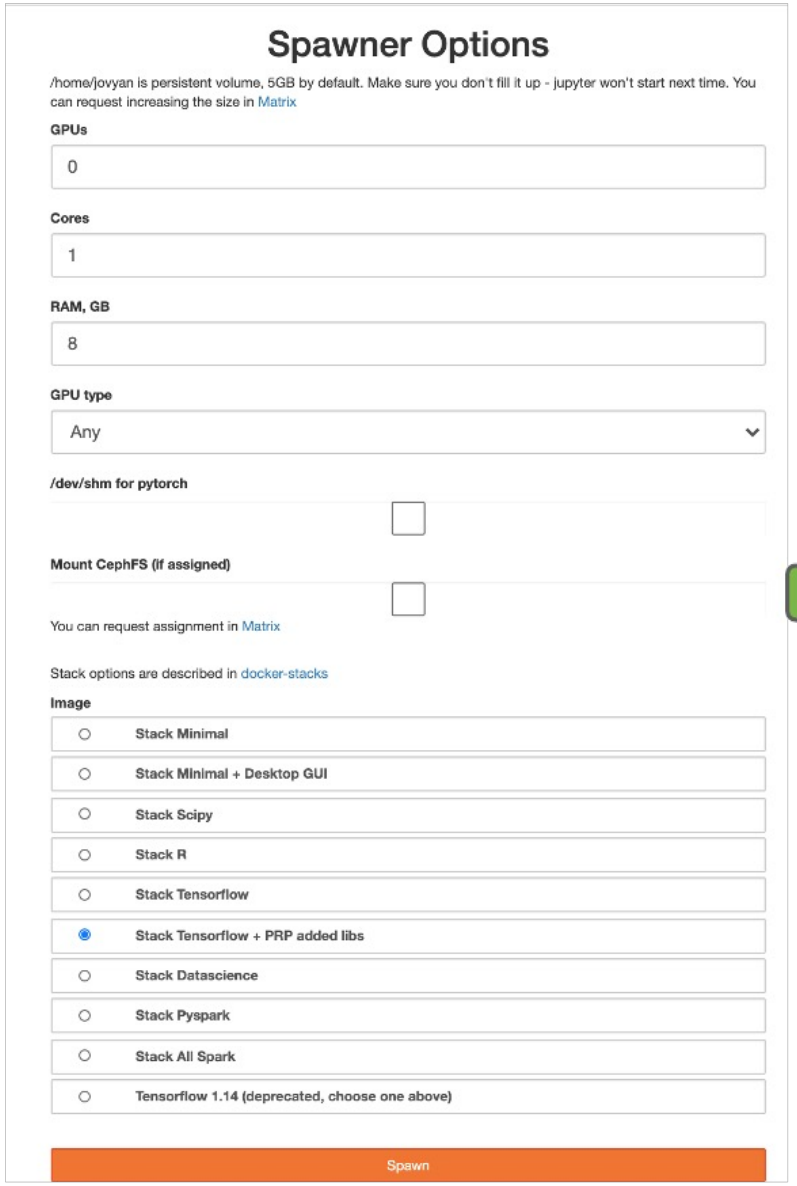

Kubernetes Pod Spawned for Exploration

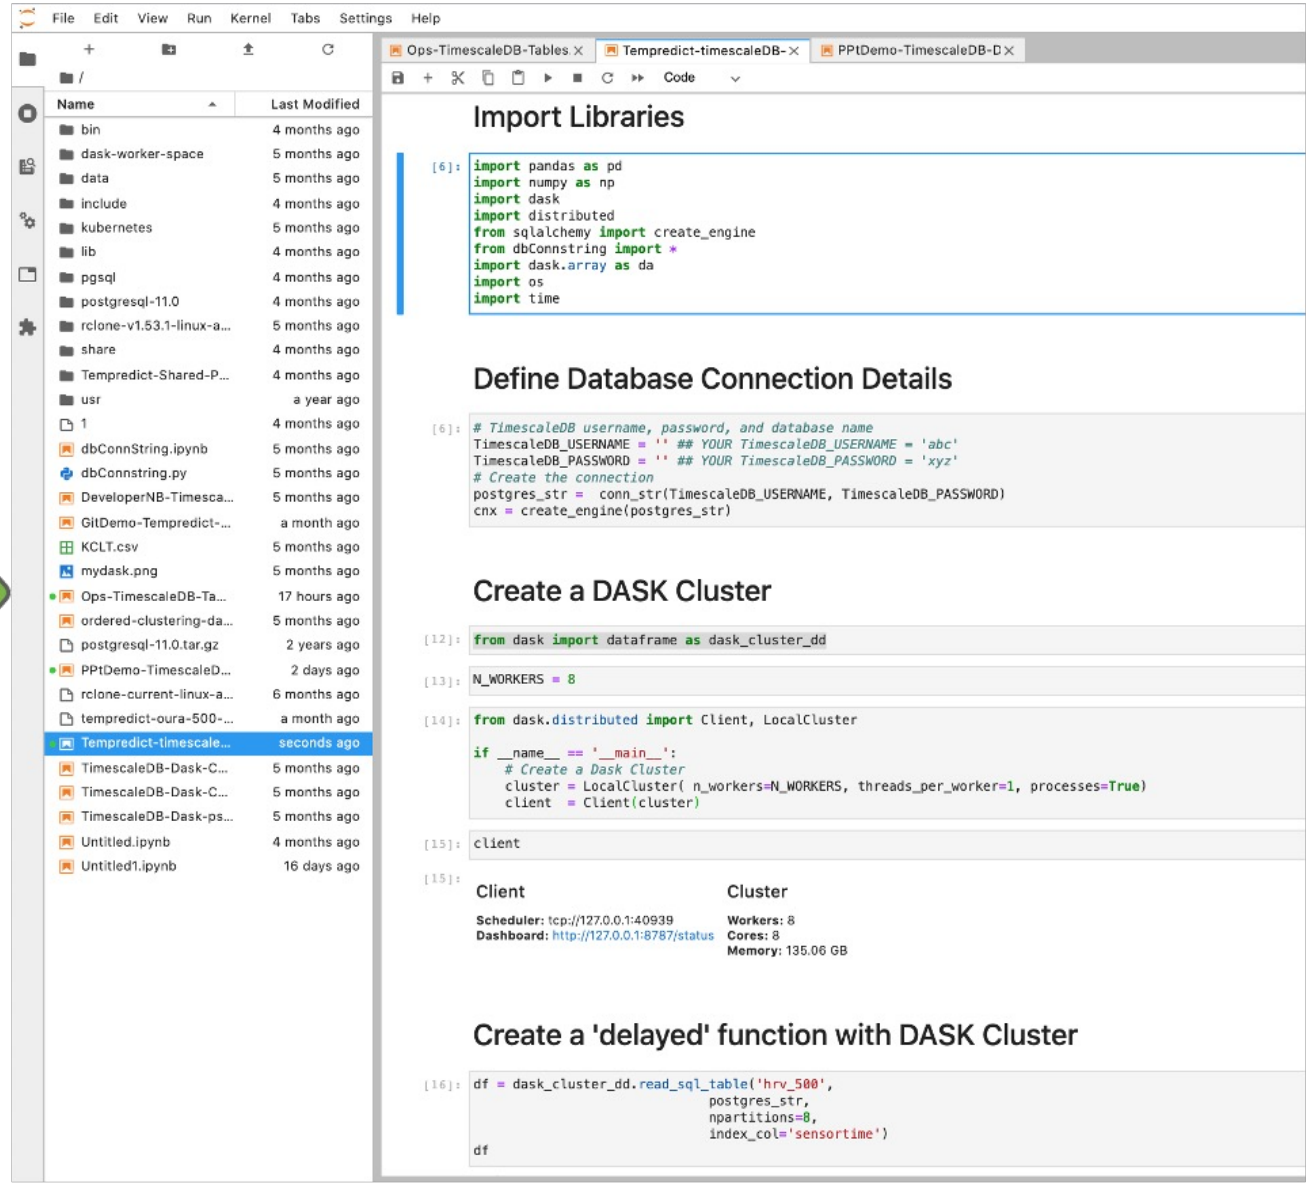

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Data federation is still a big deal!

Needs to be coupled with data systems, knowledge networks, and commons frameworks.

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Application of Composable Systems in Fire Science and Management

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

Fire is an inevitable part of nature, but megafires are not.

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Next generation fire prediction combines emerging fire science with data from many sources.

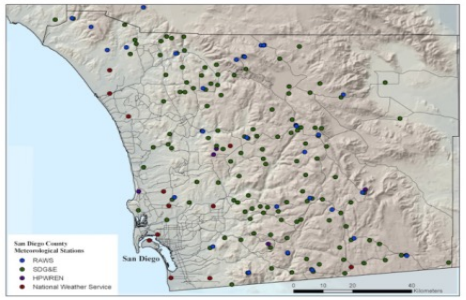

Ground-based real-time weather and camera imagery

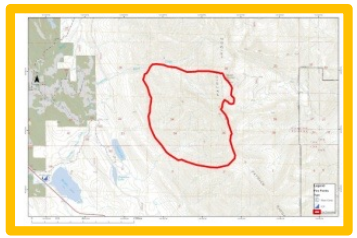

Fire perimeter

Remote sensing

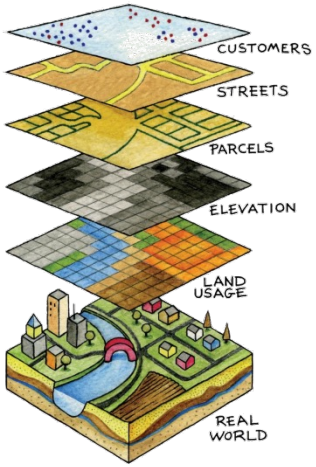

Landscape data

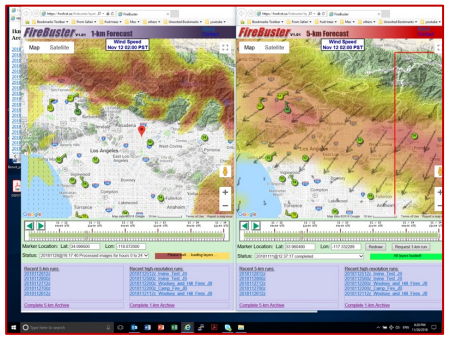

Weather forecast and modeling

Data-Driven Fire Preparedness, Mitigation and Response

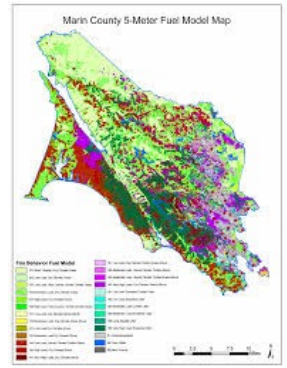

Land cover and fuel

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego HALICIOĞLU DATA SCIENCE INSTITUTE

WIFIRE Lab at-a-Glance

wifire.ucsd.edu

FIREMAP

Initial Attack Fire Response

WIFIRE Edge

Edge AI services and data plugins

BurnPro3D

Prescribed Fire Optimization

Cal OES GOVERNOR'S OFFICE OF EMERGENCY SERVICES

NSF SDGE

Other funders:

NOAA NASA SILICON VALLEY community foundation ZONEHAVEN

Powered By

WIFIRE Commons

FAIR Data and Model Catalog and AI Gateway

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Teamwork and AI

WIFIRE Commons

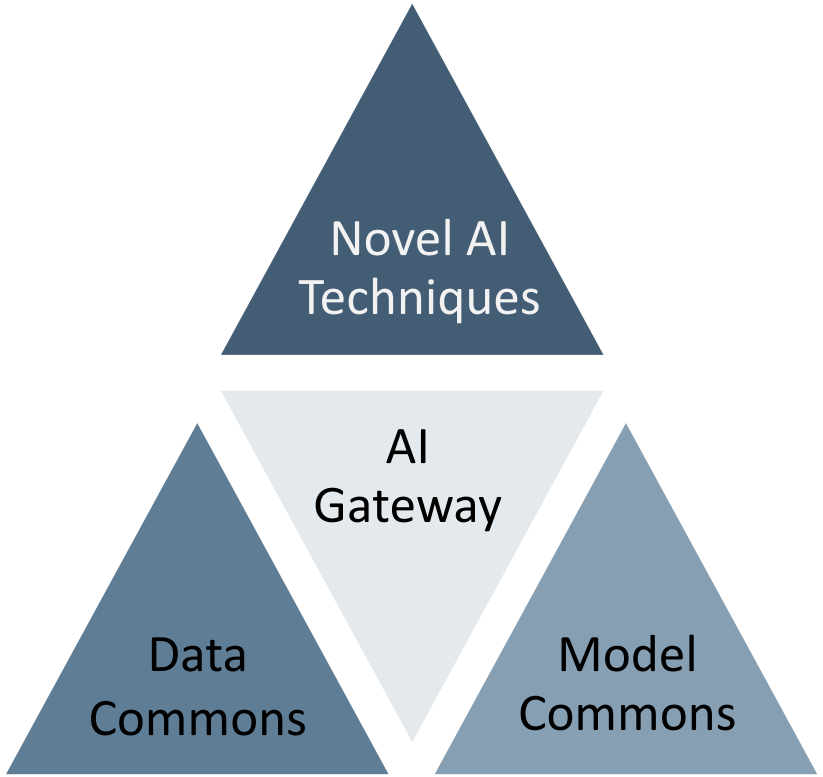

https://wifire.ucsd.edu/commons

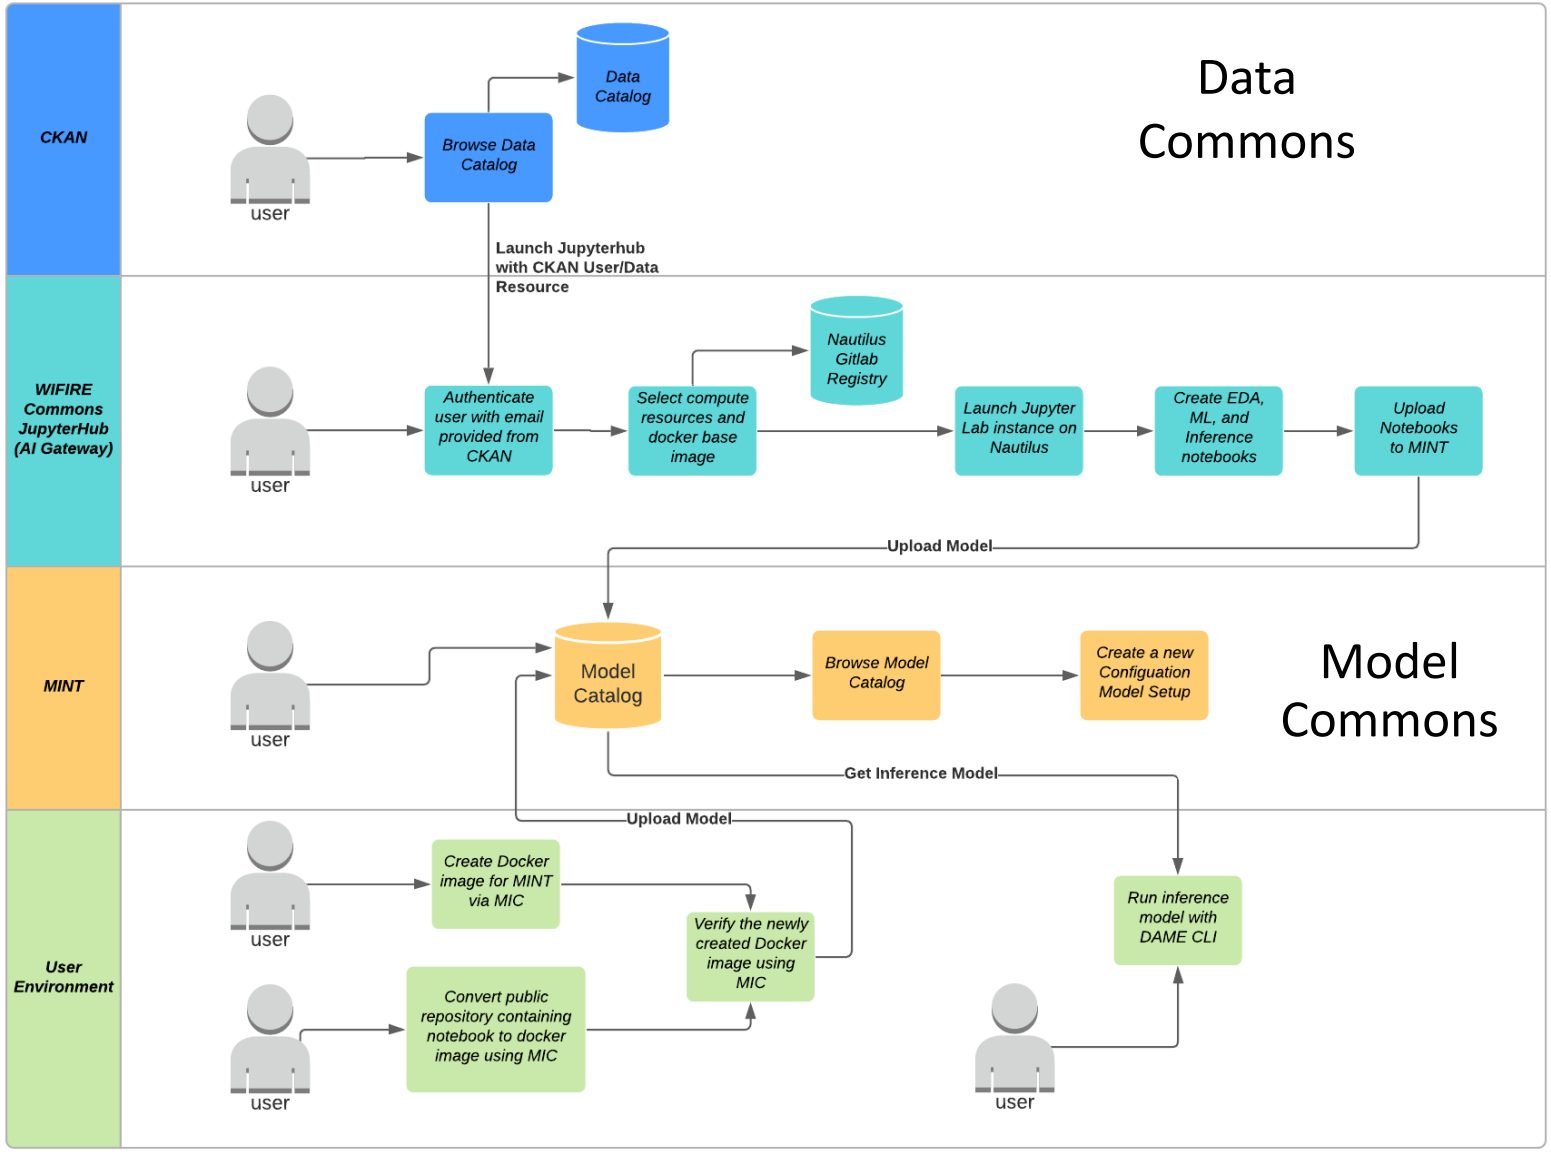

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

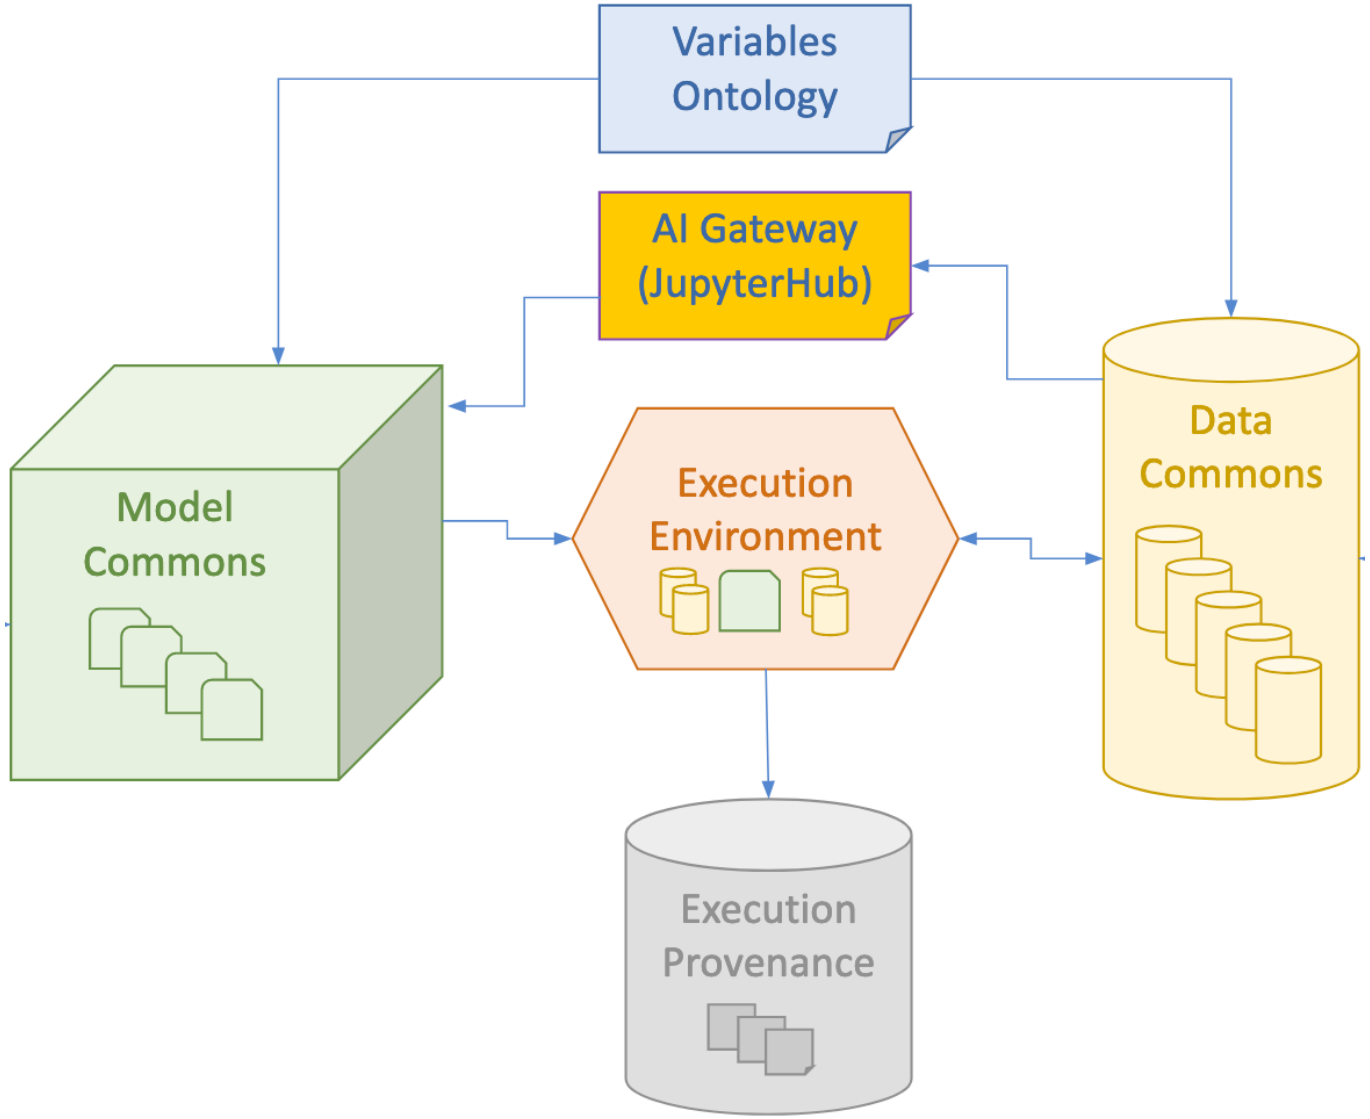

Linking Data and Models to Open Knowledge Networks

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

DOIs for Models and Datasets

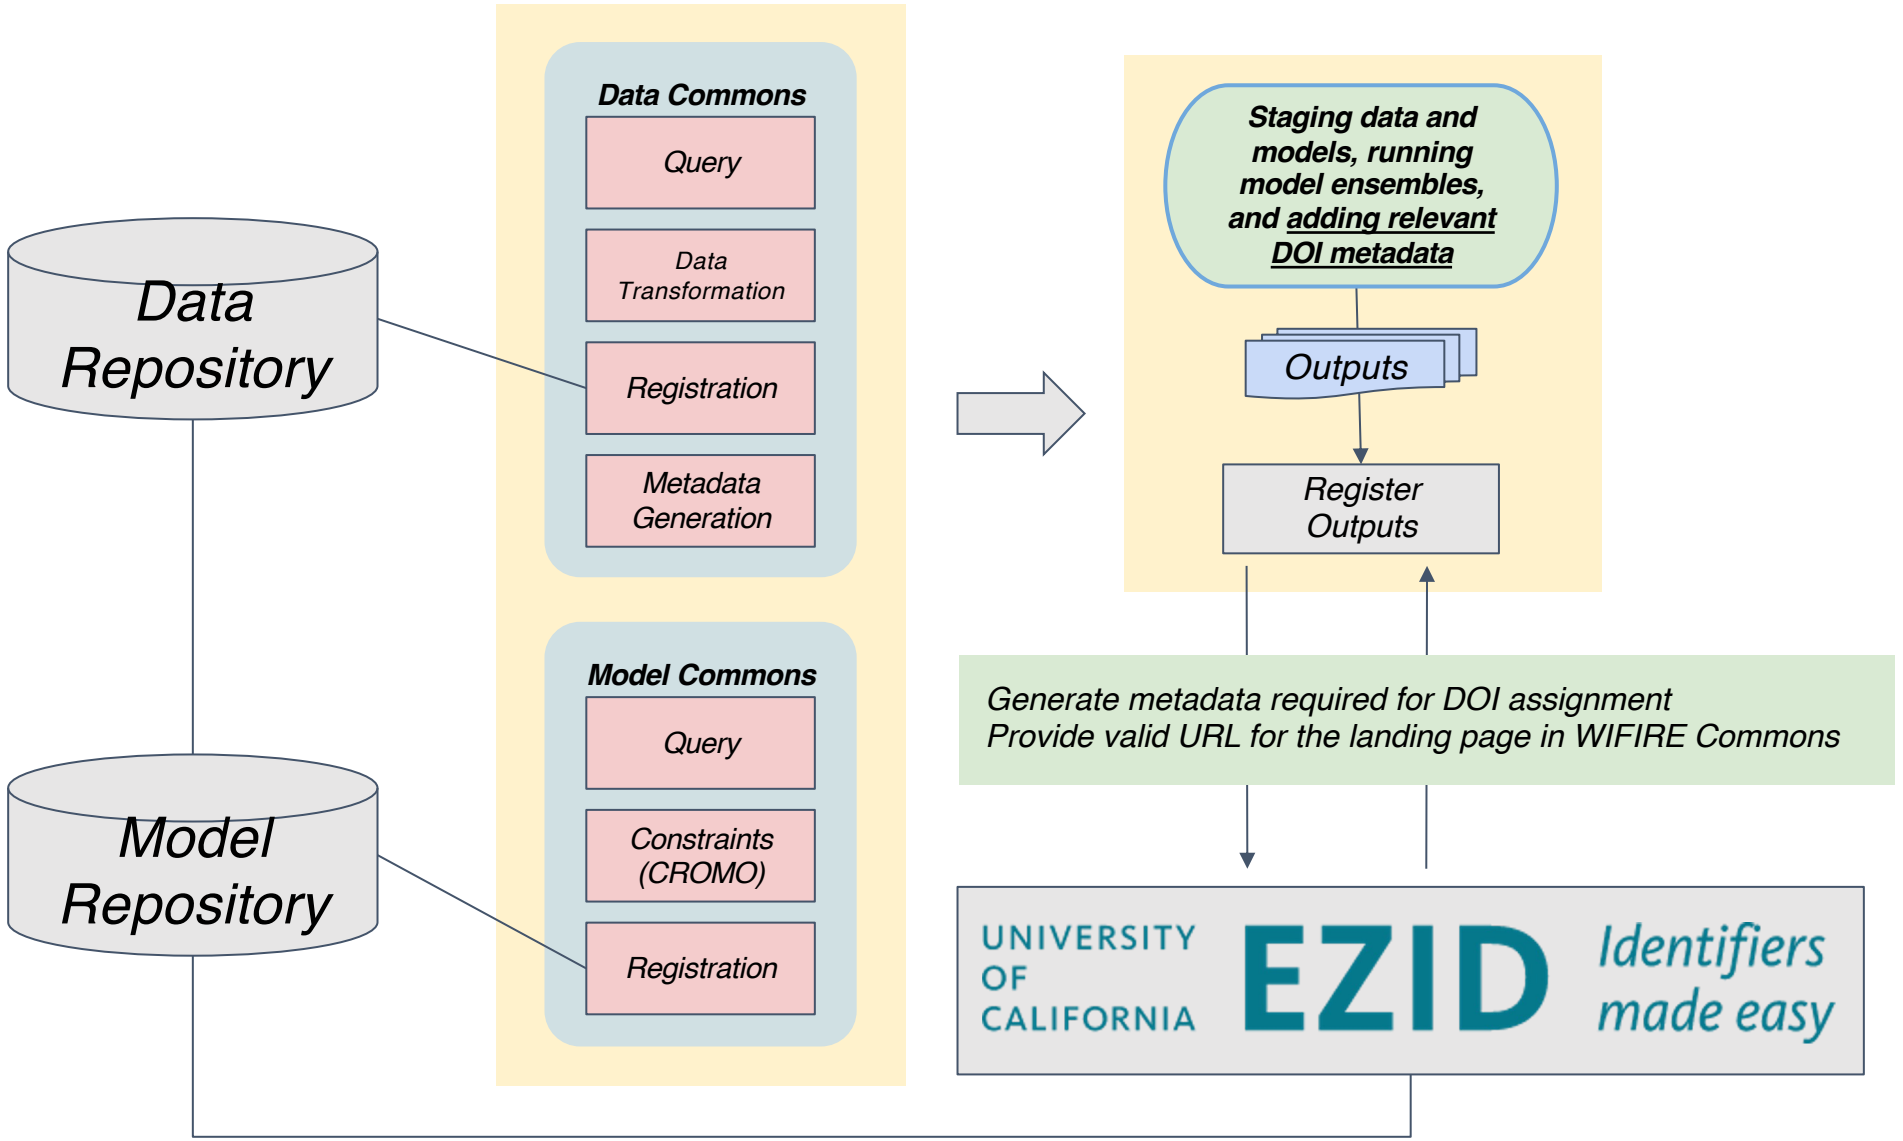

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Data to Value in WIFIRE Commons

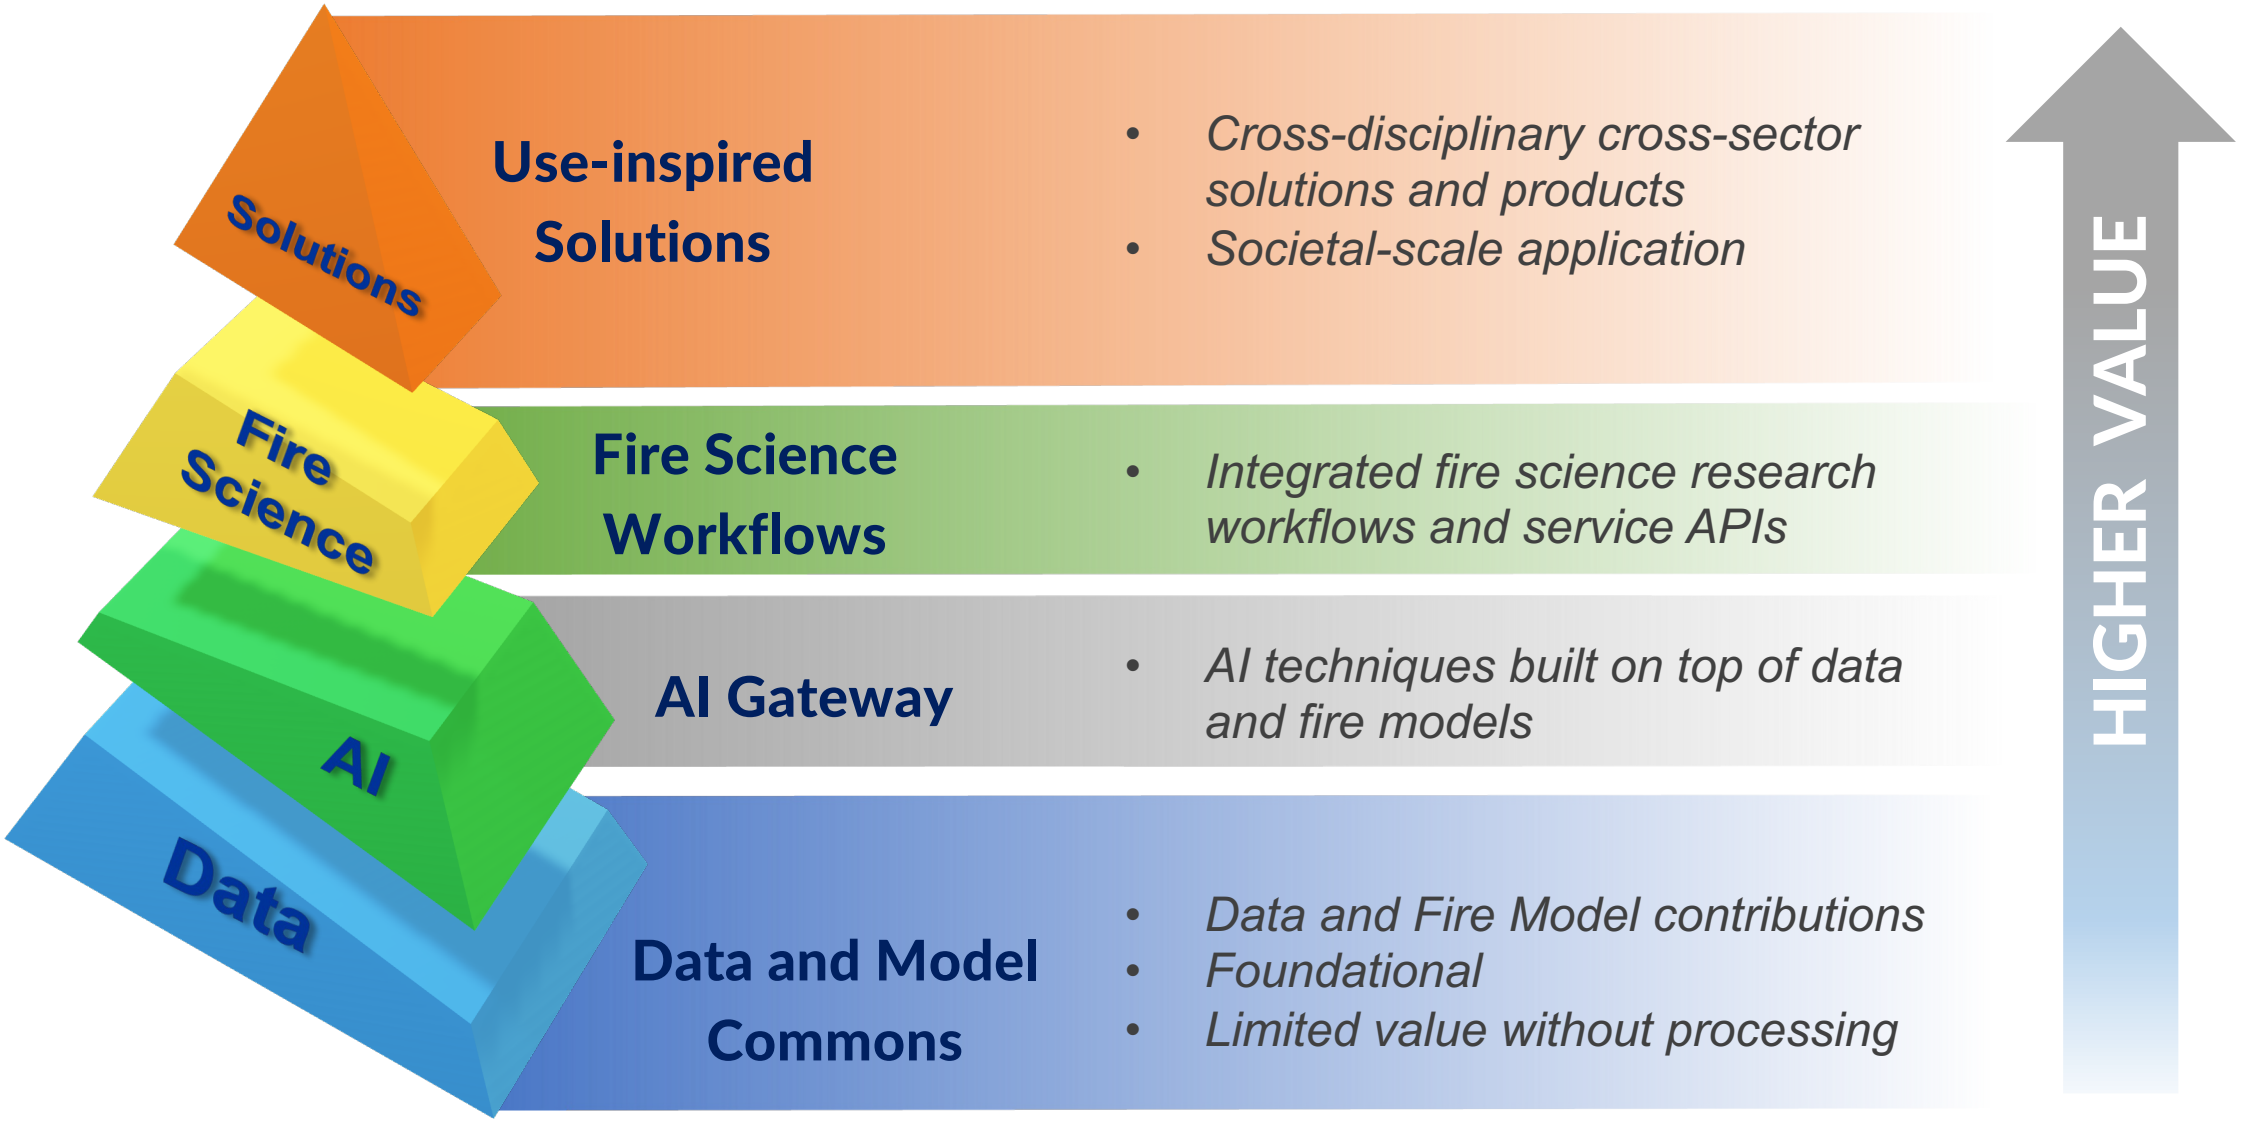

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

AI Gateway → Pangeo Customized for AI and Fire Science

https://wifire-commons-pangeo-jupyter.nrp-nautilus.io/

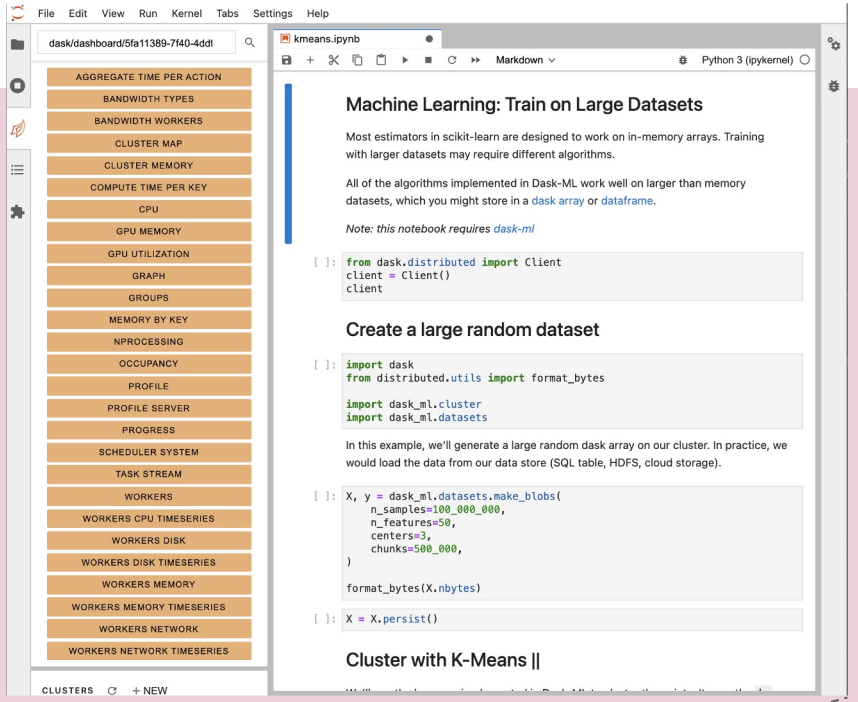

AI-Ready Ensemble Datasets *stored in cloud-optimized Zarr format*

FastFuel

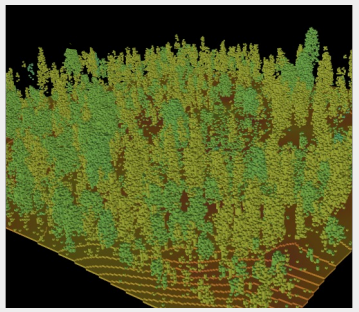

Zarr

distributed Ceph storage

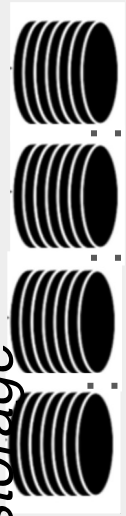

- *format for cloud-optimized storage of chunked, compressed, N-dimensional arrays*

Tech stack supporting ***scalable analysis of large wildfire datasets***

Nautilus PRP

xarray

- *Label-based indexing and arithmetic*
- *Wide range of input/output options: netCDF, HDF, geoTIFF, Zarr*

INTAKE

- *Data cataloging utility*

DASK

- *Parallel processing*

JupyterHub facilitates interactive data analysis on remote systems

jupyter

web browser

end user

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

WIFIRE

Focus on Two Ways to Manage Wildland Fires

REACTIVE

to

PROACTIVE

Fighting Severe Wildfires

Conducting Prescribed Burns

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

AI-Integrated Real-time Fire Response In WIFIRE

WIFIRE Firemap

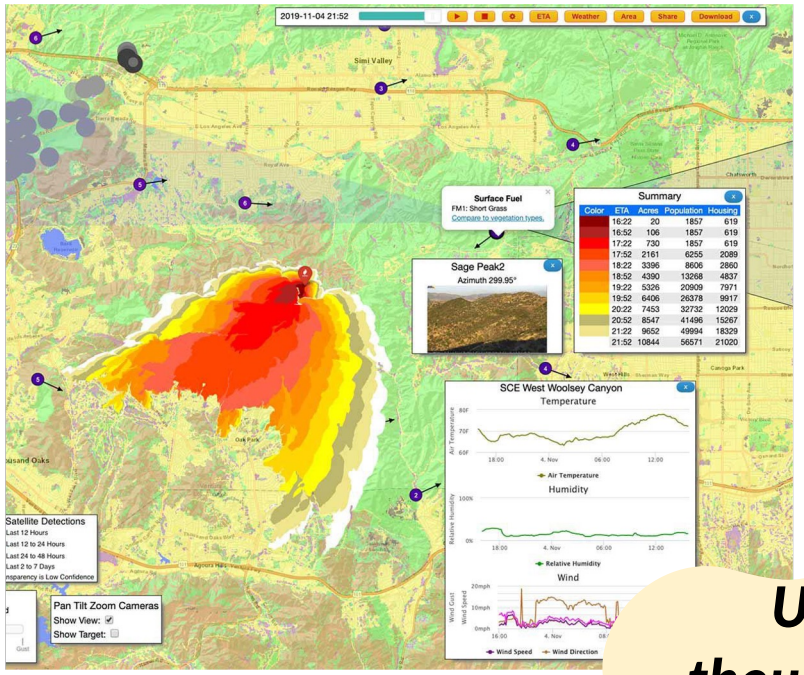

Operational in 100+ fires

Used by thousands of firefighters

FIRIS

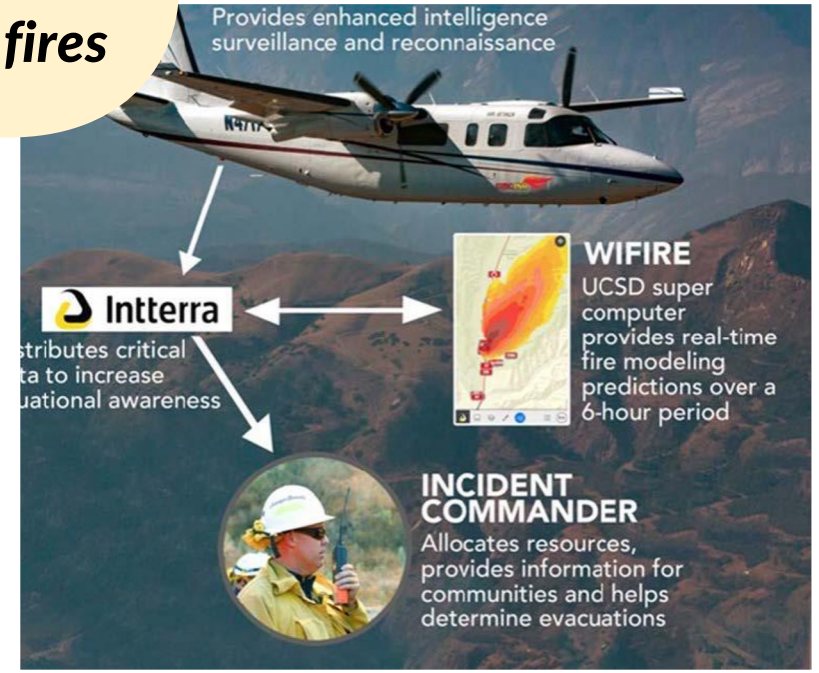

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

Dynamic Data-Driven Fire Modeling

Data and AI Needs

-- Characterizing the dynamic fire environment : Variation of wind, smoke, moisture, fuels, fire perimeter, …
-- Detection of fire ignitions
-- Decision support for fire management
-- Prediction of potential fire ignitions

Heterogenous Computing Needs

-- Edge
-- AI processors
-- Cloud
-- HPC

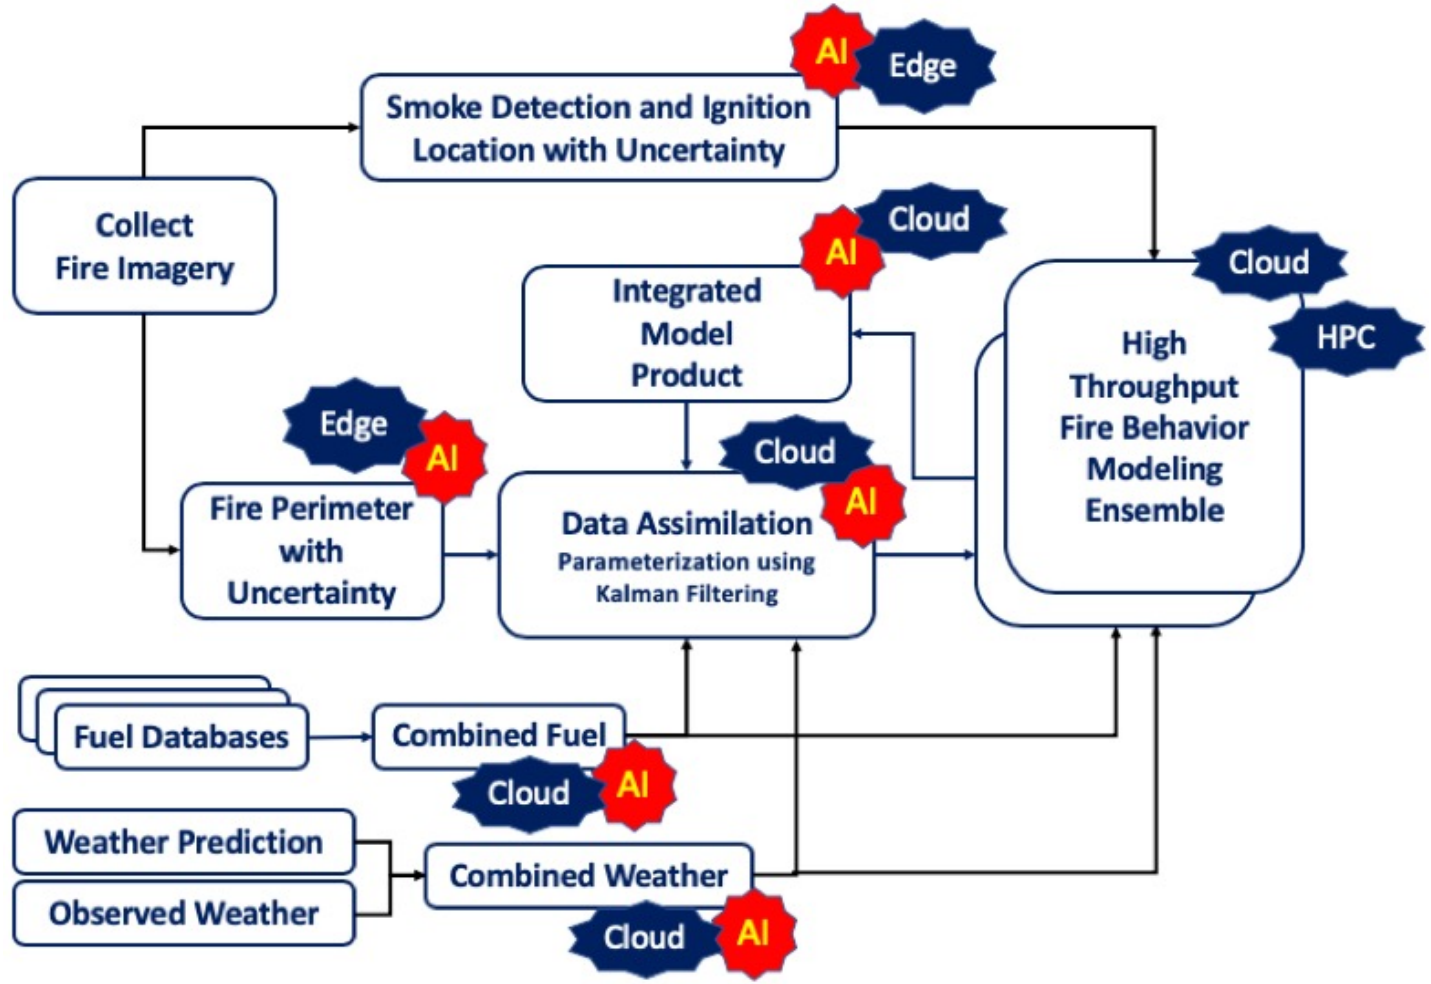

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Deep Learning-Based Smoke Detection

CITE: https://www.mdpi.com/2072-4292/14/4/1007

Data: *Camera images of wildland fires*

- **System Architecture:** Several deep learning models to extract spatio-temporal information from camera imagery to detect smoke plumes.

Training Data: *Labeled camera images*

Computing: *Edge and composable systems*

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Fire Simulations using Composable Systems and Edge Smoke Detection

- Three workflows
  - Smoke – Sage Edge App
  - Fire simulator
  - AI Training

- Both the fire simulator and training workflows are can be run on Expanse or Nautilus through the federation layer

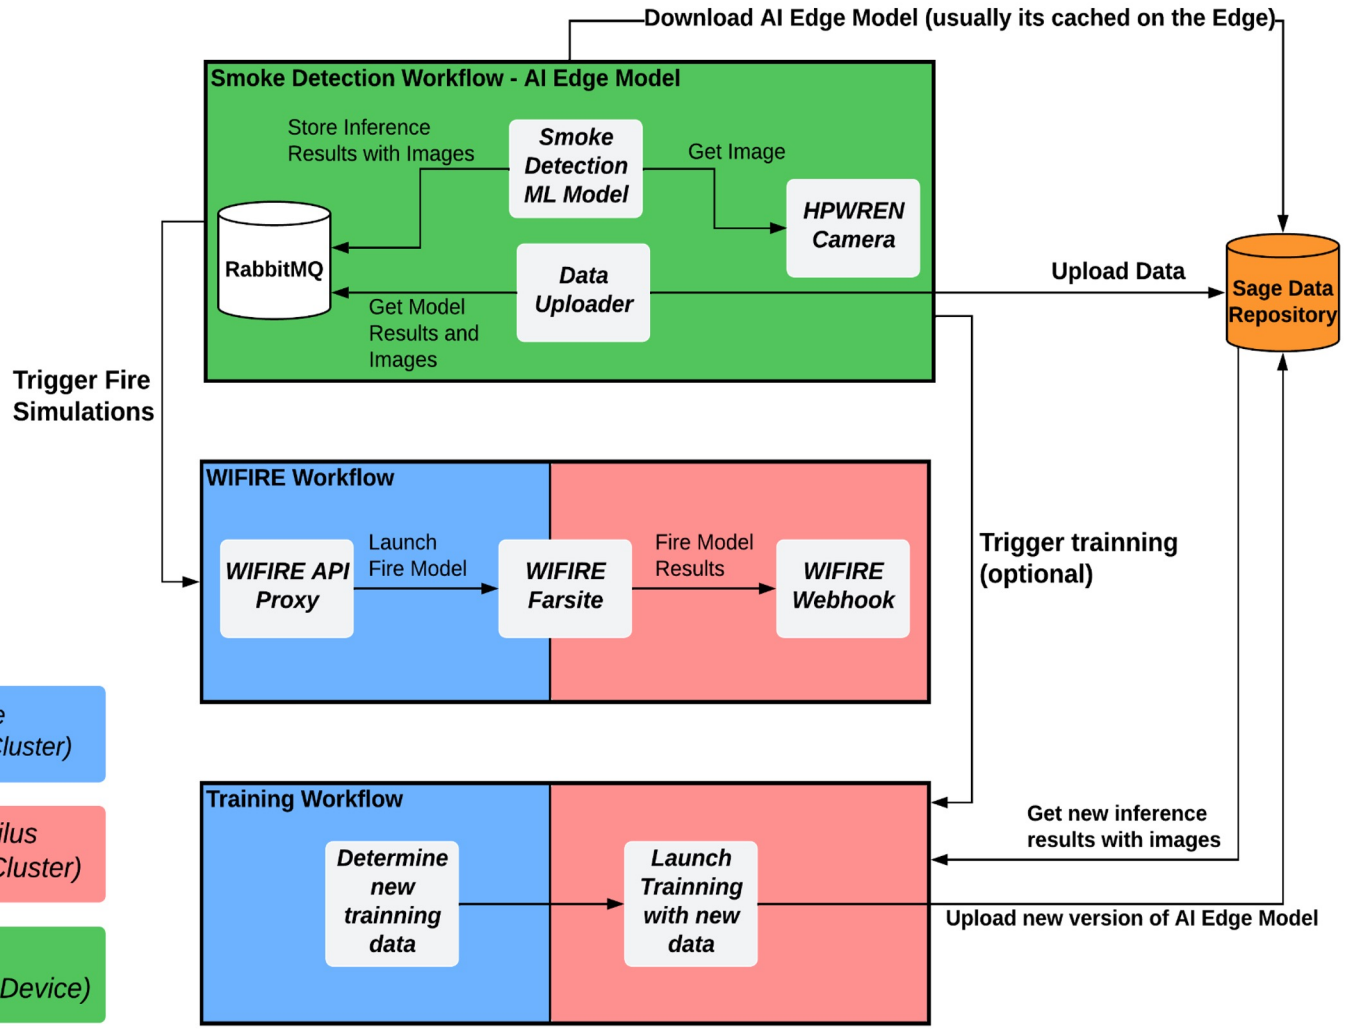

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Prescribed Burns

Transformational Science

Fast 3D is better than 2D!

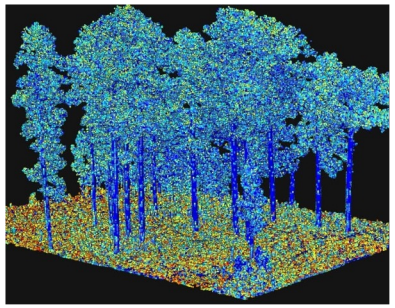

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Next-generation fire models provide the right science basis for proactive mitigation and management decisions.

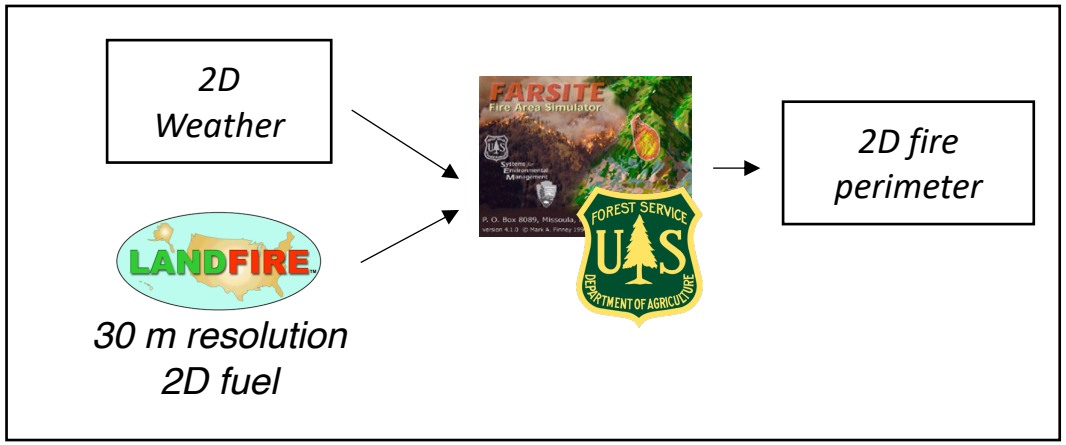

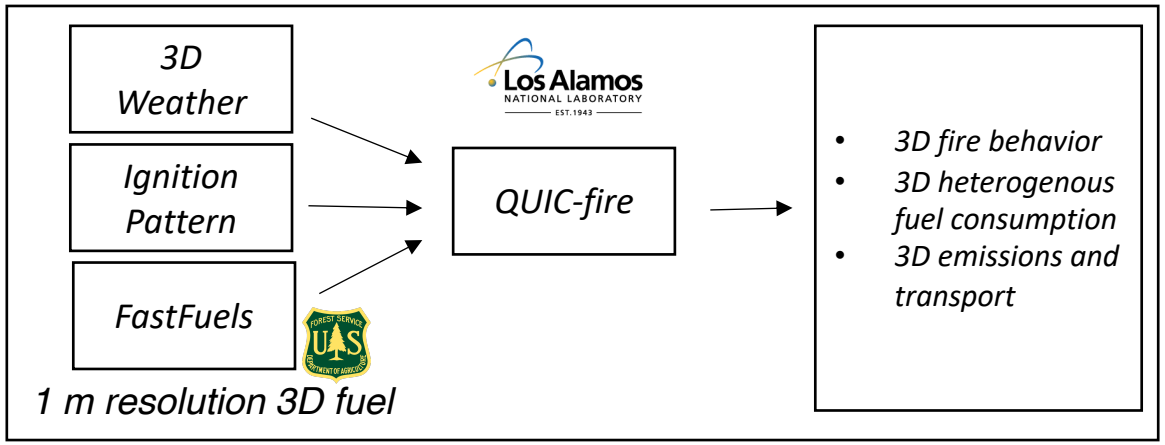

Collaboration with Rod Linn (LANL), Kevin Hiers (TTRS) and Russ Parsons (USFS)

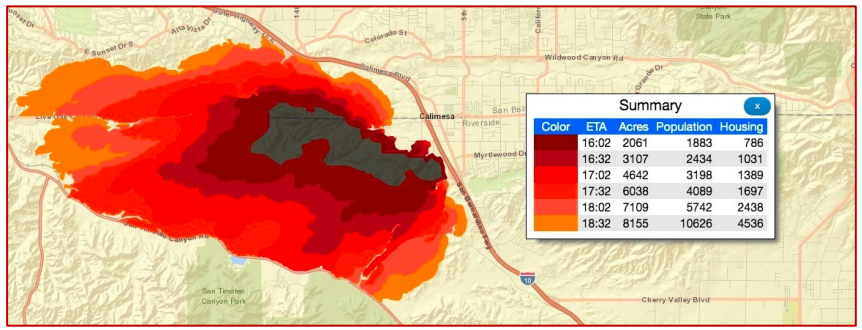

***From** 30m resolution in 2D ... in 3D*

***To** 1m resolution*

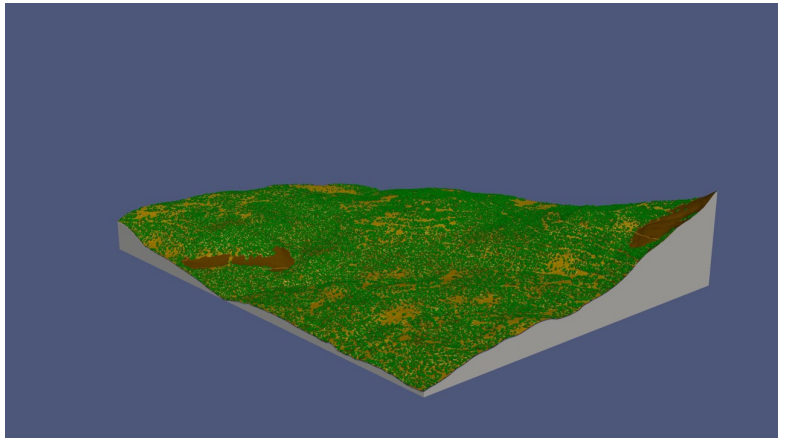

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

Ensemble Run Analysis

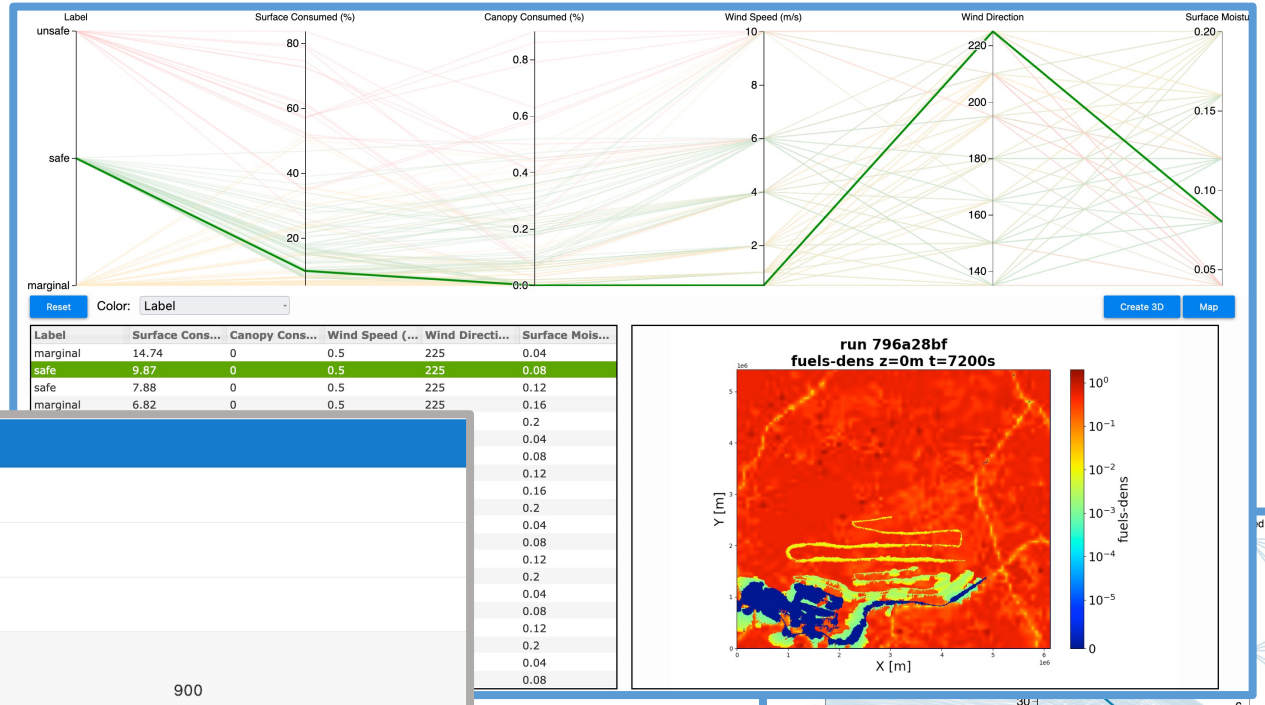

Ensembles

- Idealized Grass
- Piedmont-Hitchiti
- Yosemite Studhorse BU 1-4

Runs	900
Fuel	FastFuels, 1.8 km x 1.1 km
Sim Time	53 min
Ignitions	3: Contour, Ring, Strip
Surface Moisture	5%, 10%, 20%, 35%
Wind Direction	180, 195, 225, 240, 255
Wind Speed	2, 3, 5, 8, 12
Outputs	Fuel Density, Surface Energy, Wind

ANALYZE

- Yosemite Studhorse BU 5

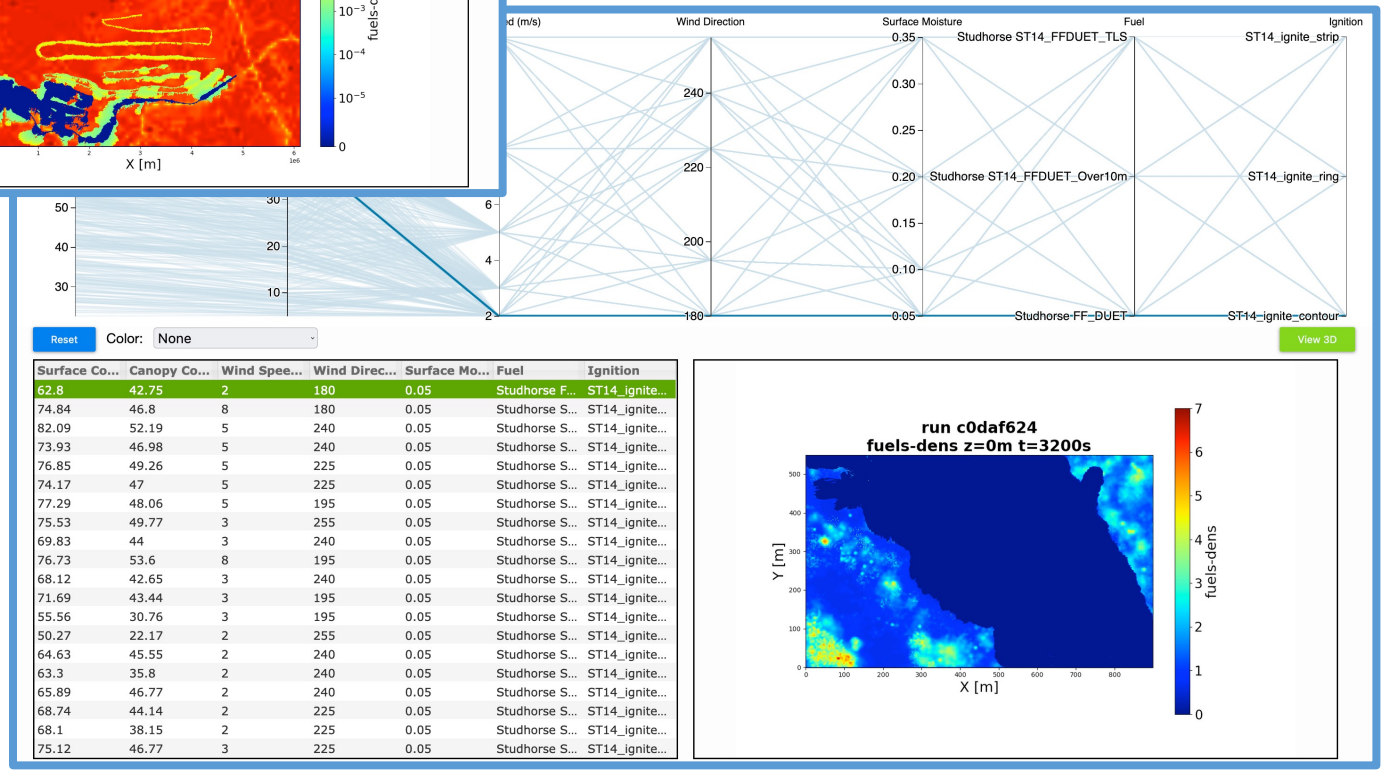

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

Emerging AI Needs for Next Generation Fire Modeling

- High-res, 3D and dynamic fire environment data
- Speed, flexibility, and interpretability of fire models

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

The BurnPro^{3D} platform gives our public sector partners next-generation fire science using data and AI to optimize prescribed burns at an unprecedented scale.

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

3D Fire Simulation Ensembles using Composable Systems in BurnPro3D

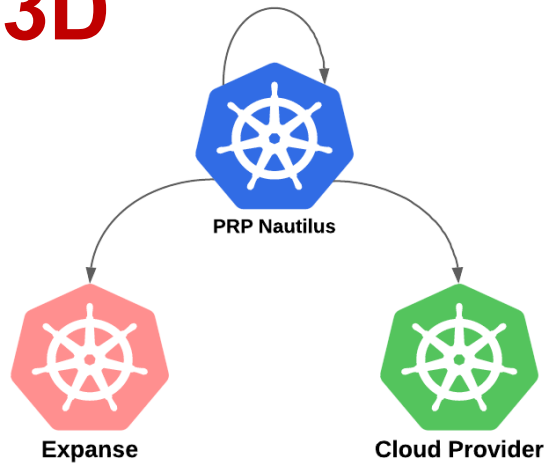

- Instead of deploying all the services on each cluster we only federate over the pods that are performing the fire simulation

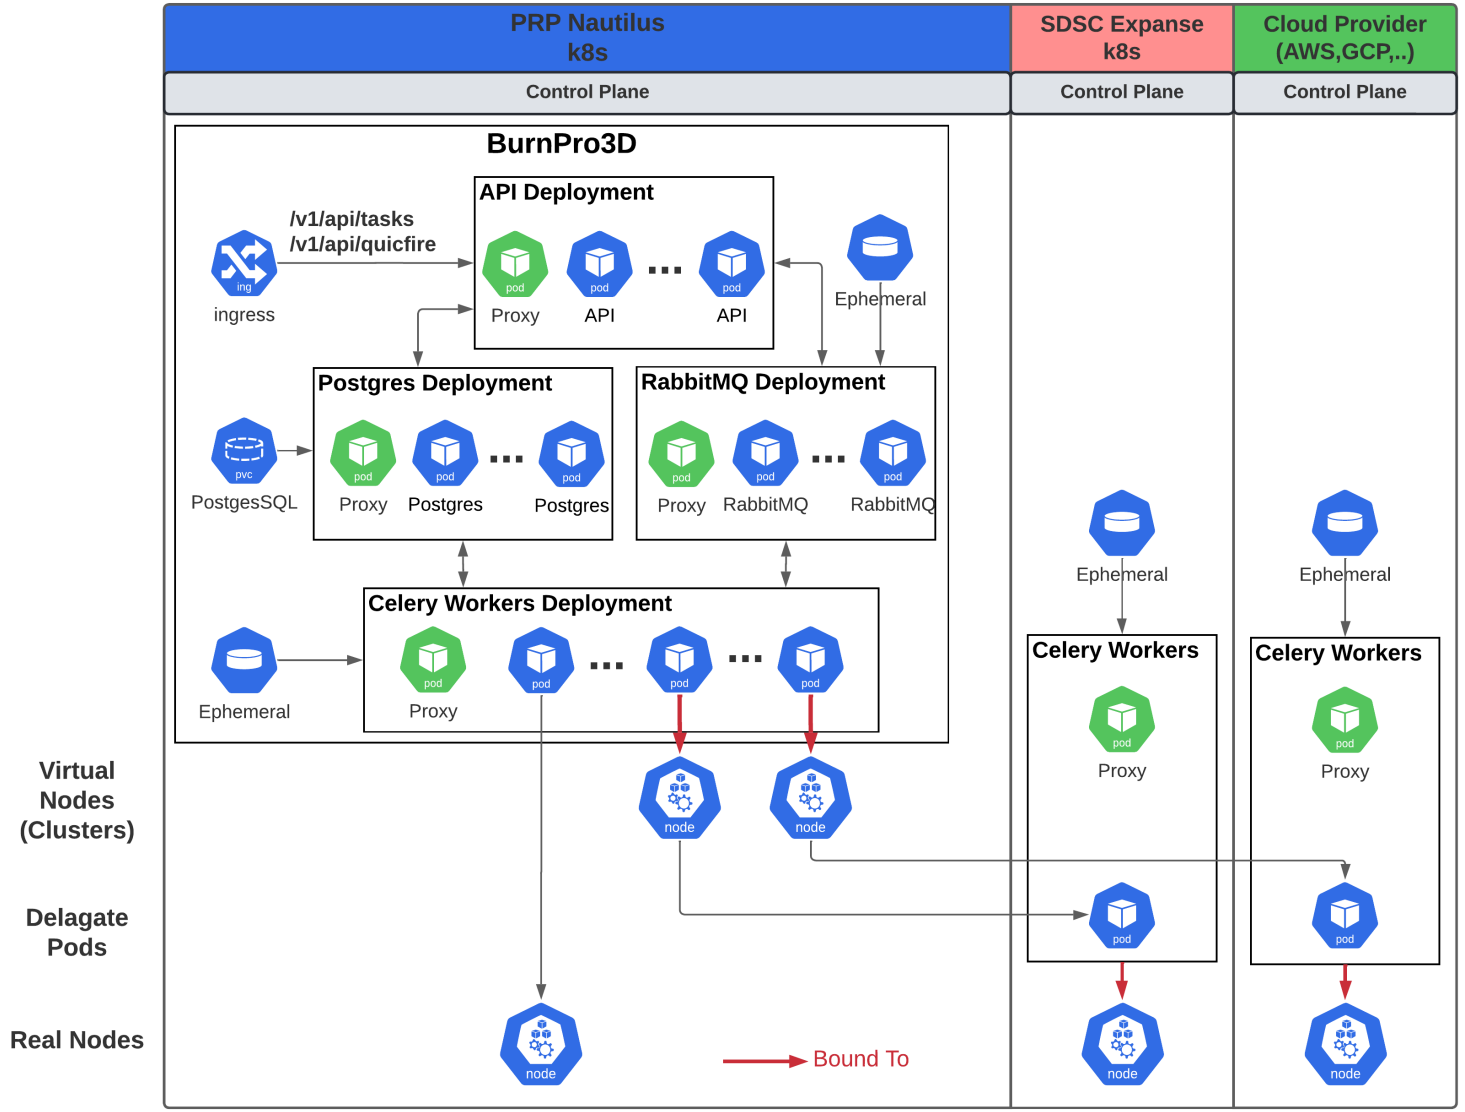

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

AI Techniques to Condition Data and Improve Model Accuracy

3D vs 2D

900x more detailed

Collaboration with Rod Linn (LANL), Kevin Hiers (TTRS) and Russ Parsons (USFS)

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

BEFORE AFTER

Yosemite National Park

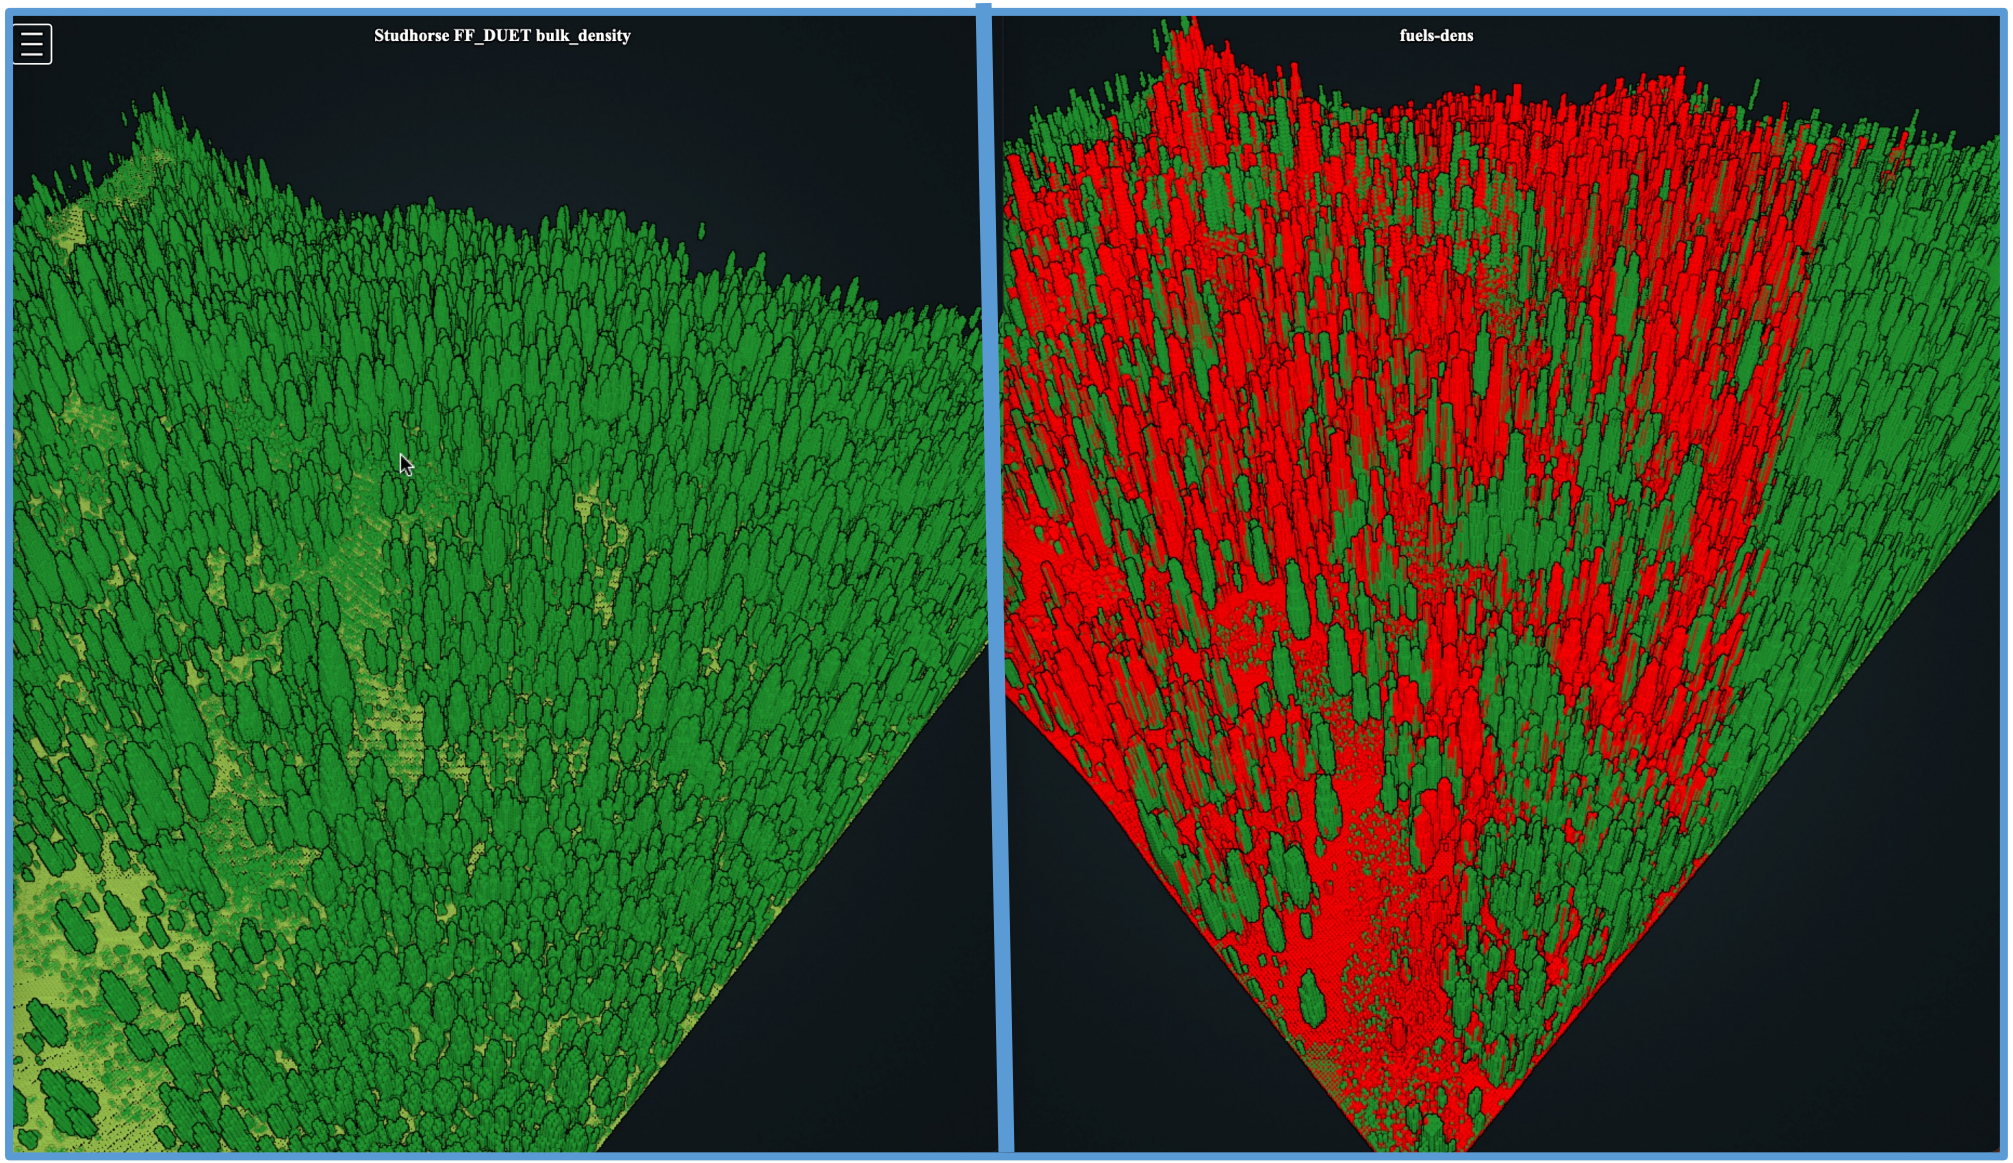

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

AI Techniques to Improve Decision Making

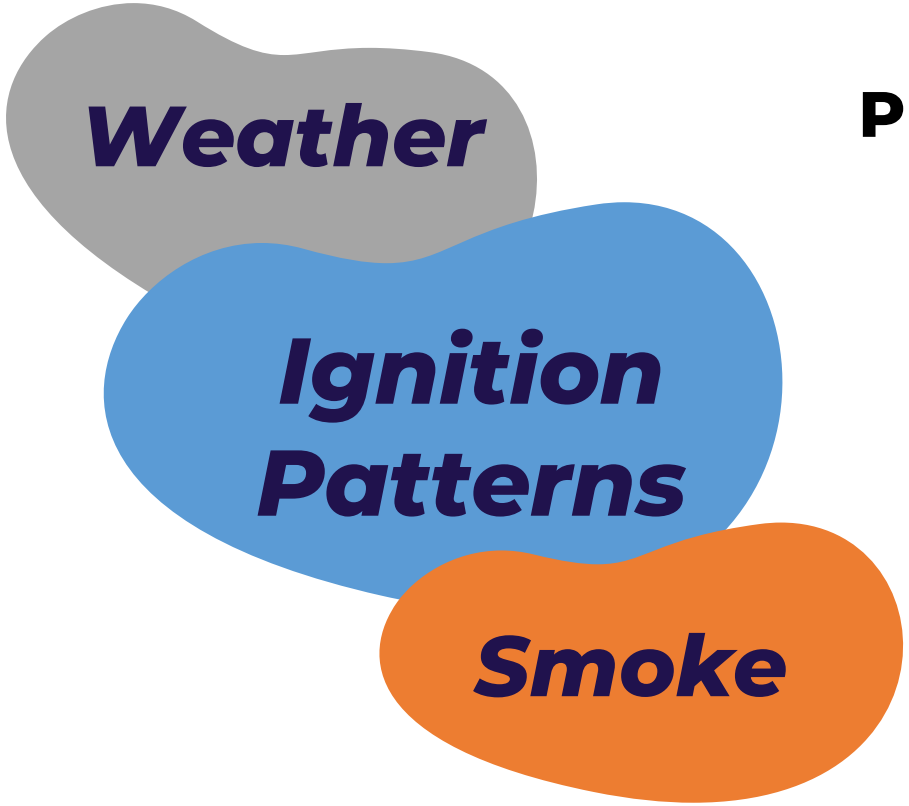

PHYSICS-GUIDED MACHINE LEARNING

To improve predictive fire behavior models

OPTIMIZATION

To address complex tradeoffs and prioritization

EXPLAINABLE AI

To increase scientific understanding and interpretability all along the decision-making chain

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

High resolution data from Sage is needed for understanding of the prescribed fire environment at the resolutions required.

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

NEON Mobile Deployment Platform (MPD) with SAGE Konza Prairie for controlled burn: April 2022.

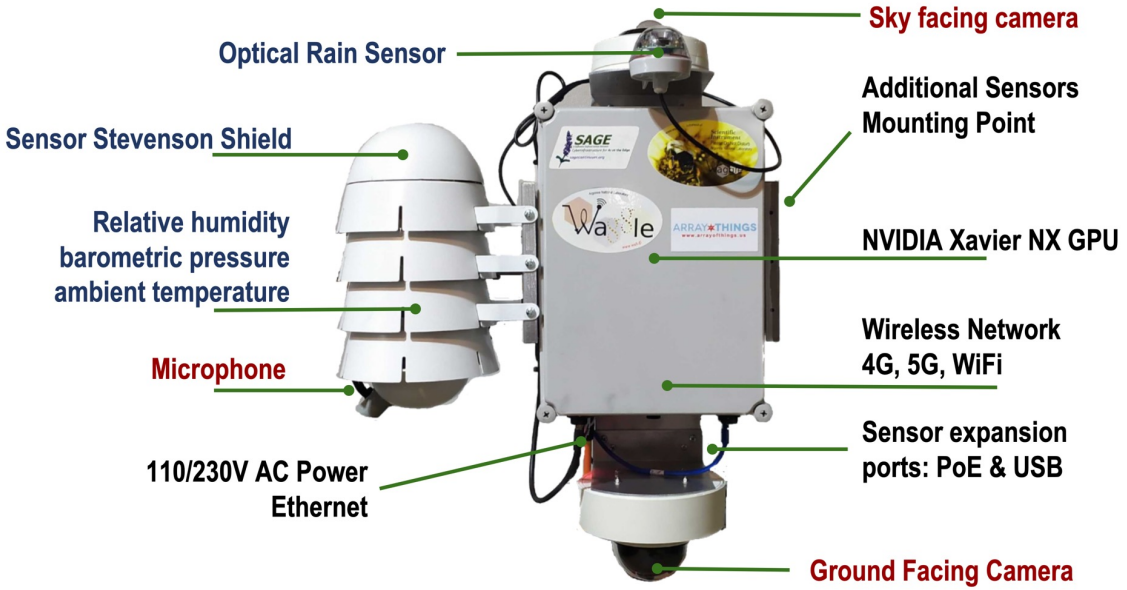

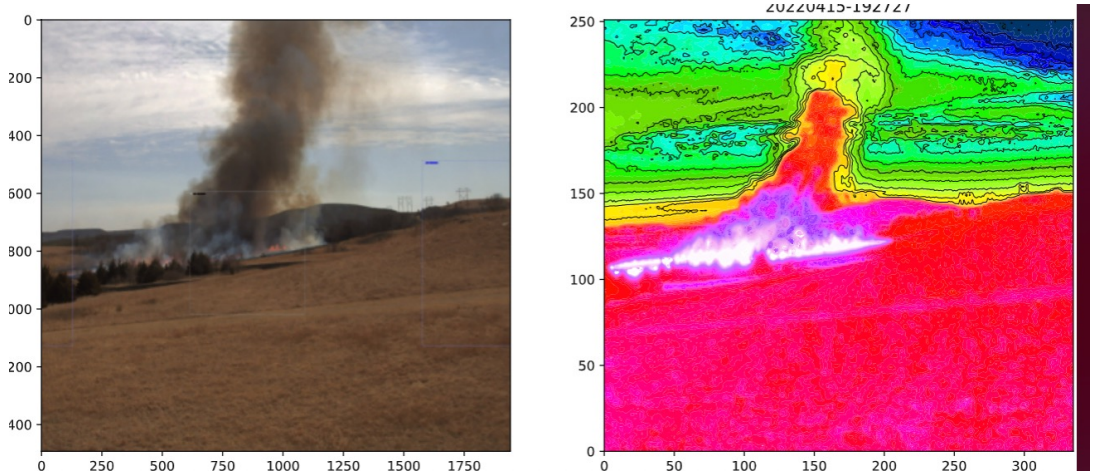

Konza Site Prescribed Fire Scenario

BEFORE

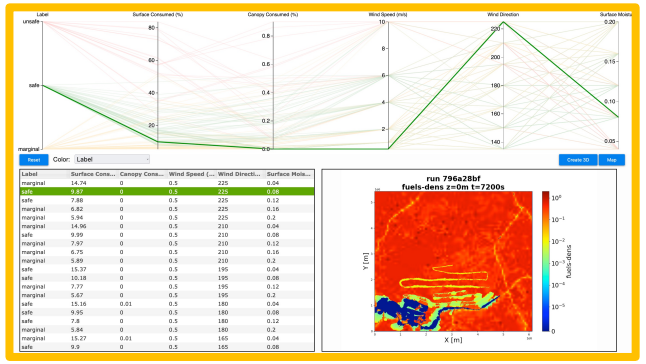

- ***Run ensemble models of the burn***
- ***Decide on burn prescription***
- ***Make Sage NEON mobile unit request***

DAY OF

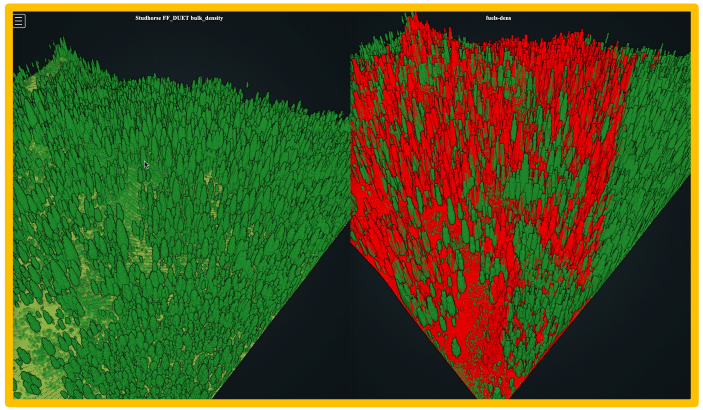

- ***Make Go/No Go decision***
- ***Deploy edge AI for smoke detection and environment monitoring***
- ***Set up Sage dashboard***

AFTER

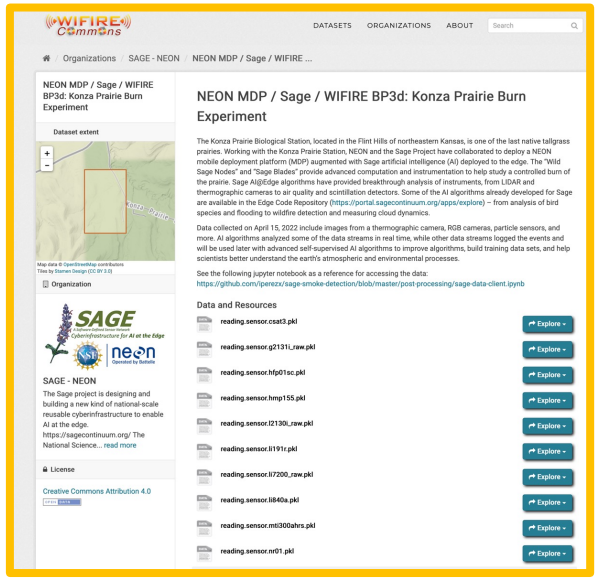

- ***Publish data in WIFIRE Commons***
- ***Conduct post-fire model evaluation***
- ***Create adaptive prescribed burn strategies***

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego HALICIOĞLU DATA SCIENCE INSTITUTE

Use Case: Generalized Application Development

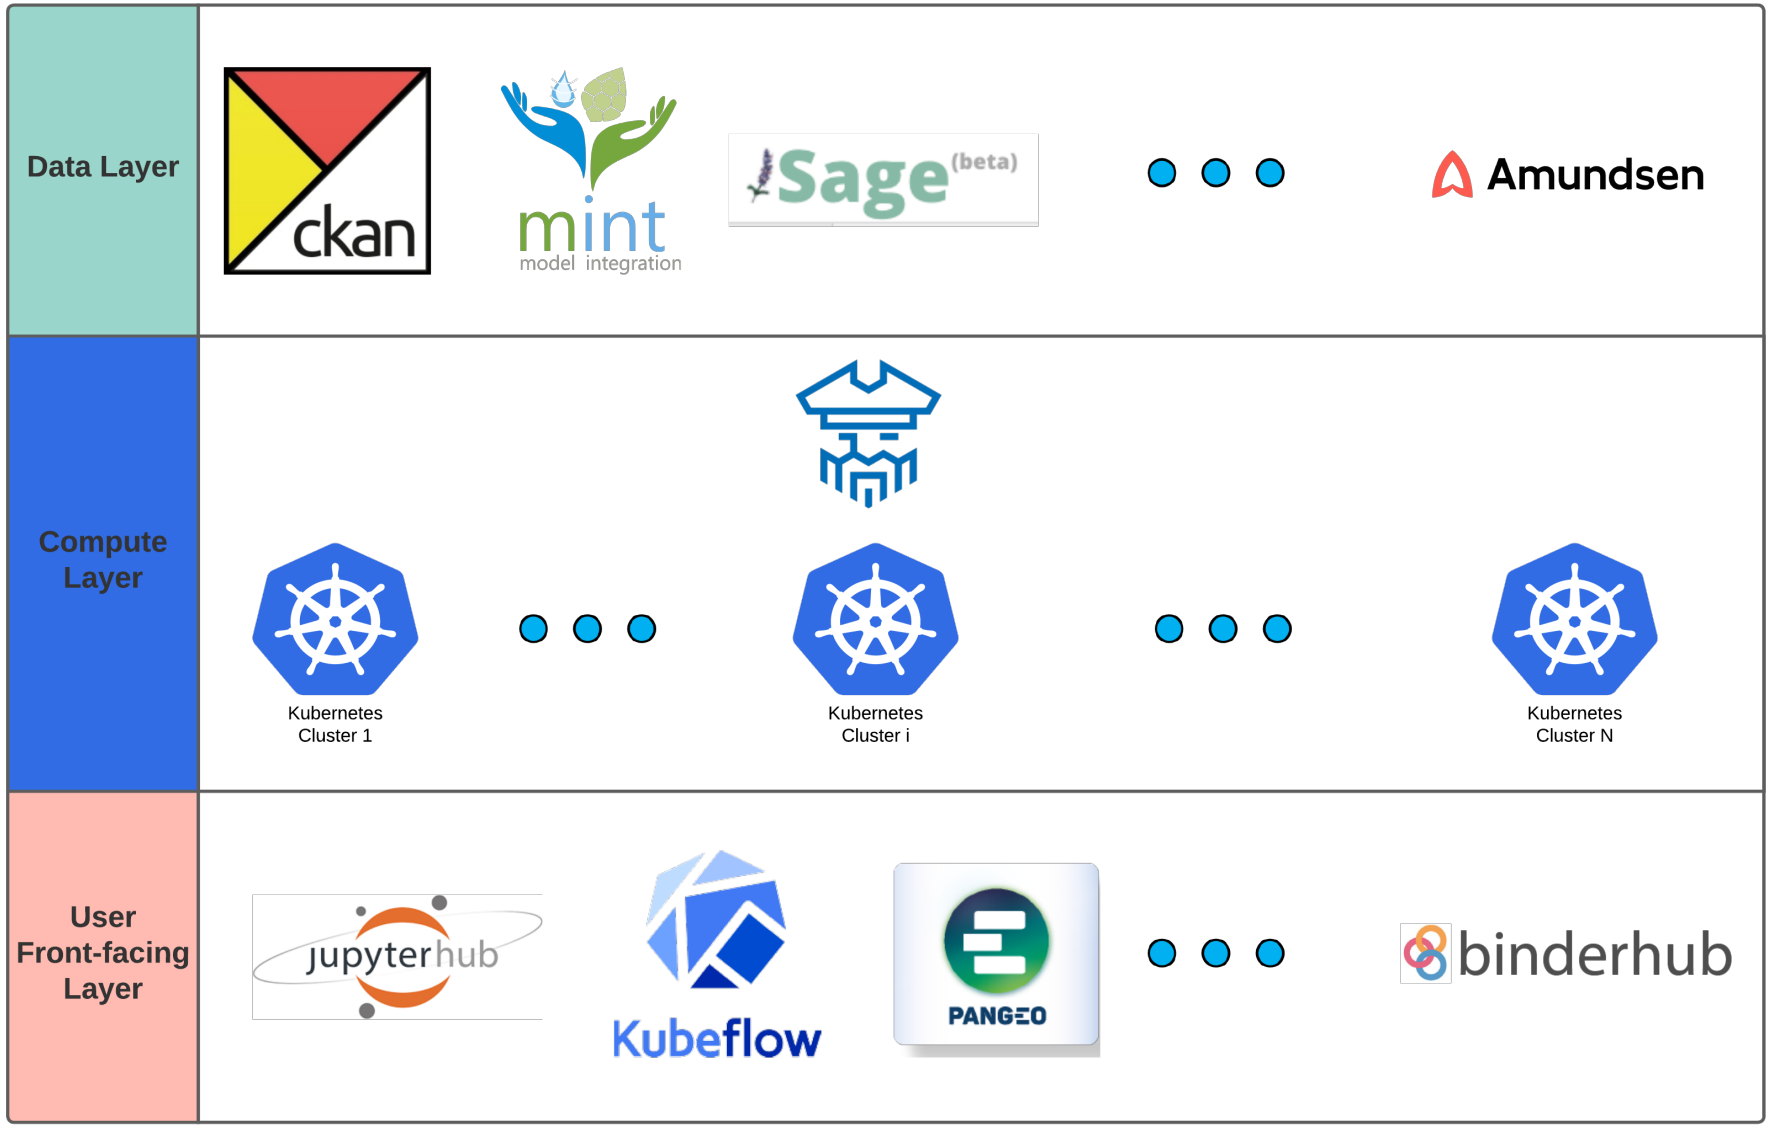

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Use-Case: Many Fire Simulators Workflow

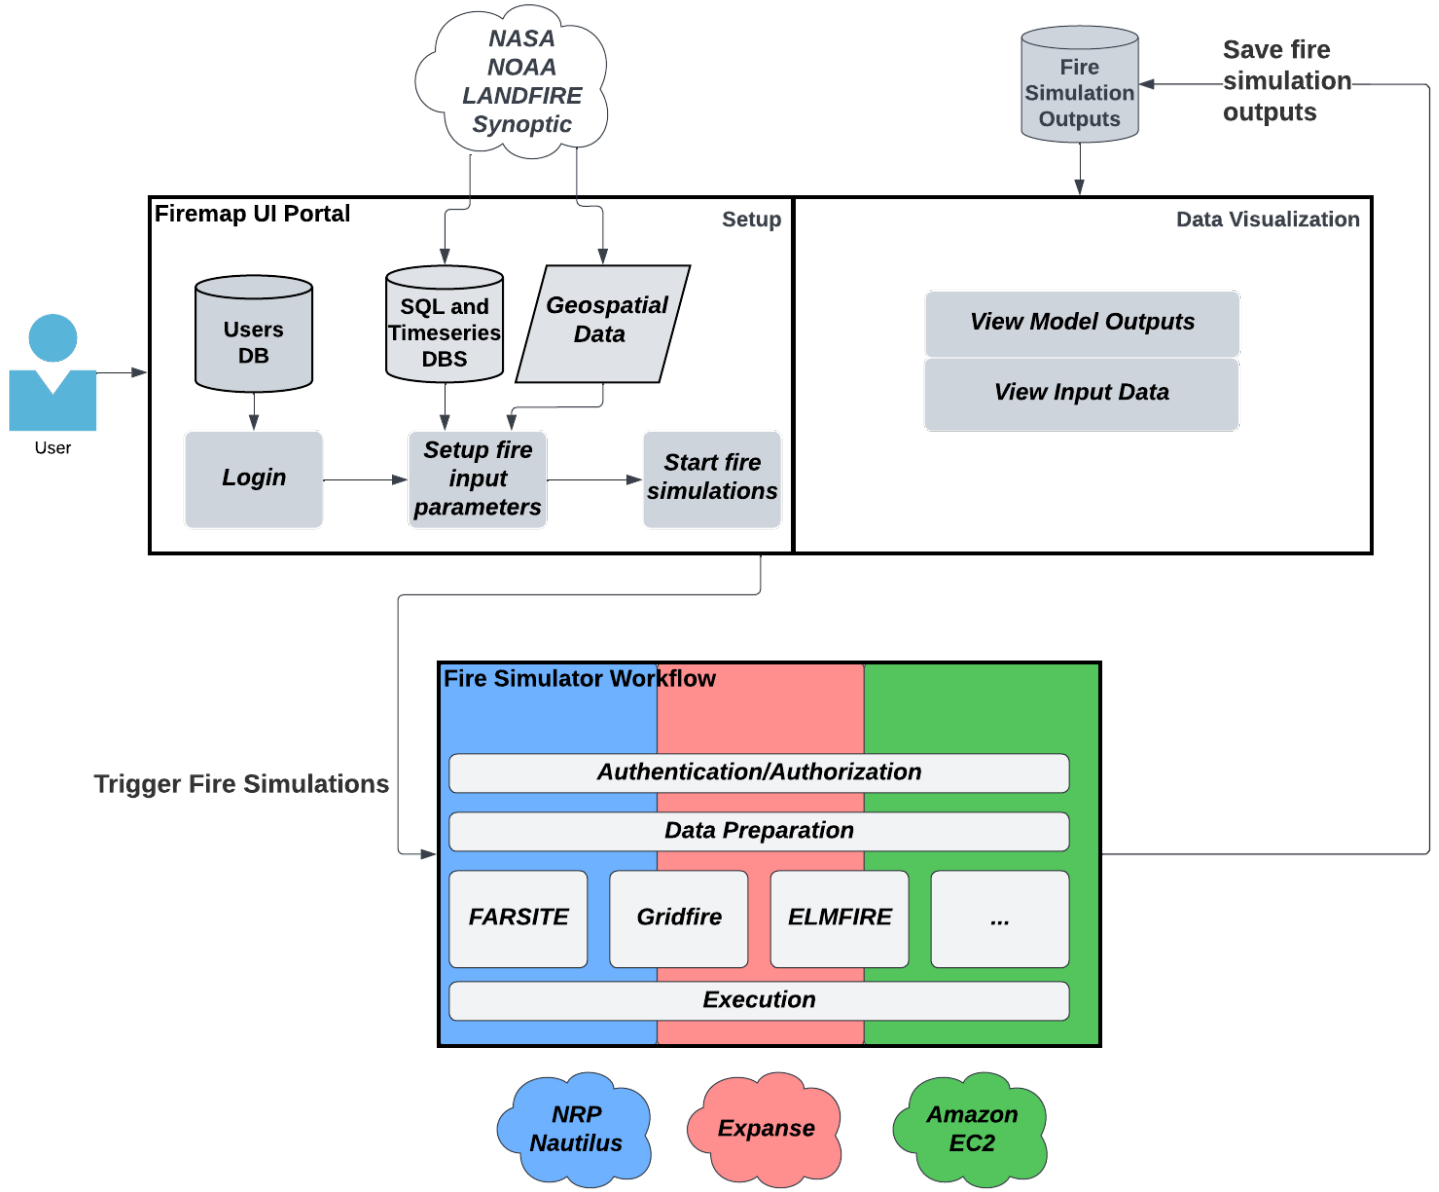

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

powered by WIFIRE Commons

A Data and AI Framework for Convergence in Wildland Fires

Data Sharing and AI as a bridge.

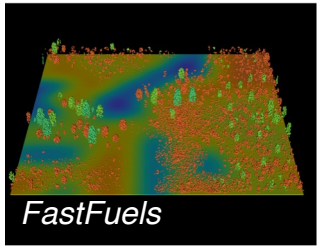

Fire Science

- *Knowledge Management*
- *Physics-Guided ML*
- *Optimization*
- *Interpretability*

Application to Preparedness, Mitigation, Response, Recovery

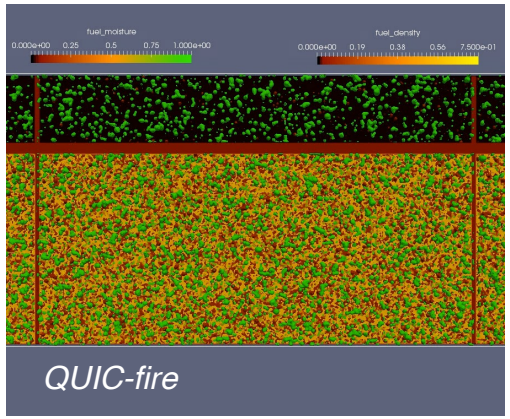

3D Next Generation Fuel and Fire Models

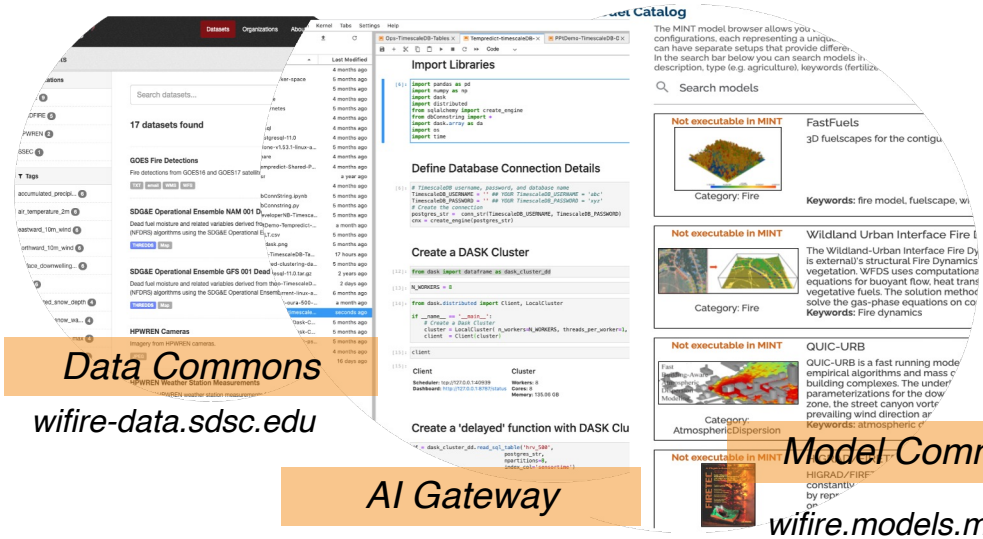

Data Commons

wifire-data.sdsc.edu

AI Gateway

Model Commons

wifire.models.mint.isi.edu

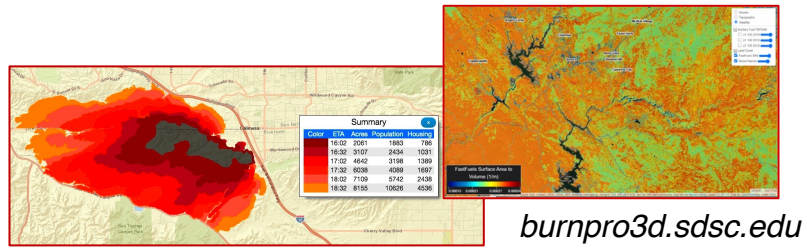

firemap.sdsc.edu

burnpro3d.sdsc.edu

Use-Inspired Interfaces

SDSC SAN DIEGO SUPERCOMPUTER CENTER

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

UC San Diego HALICIOĞLU DATA SCIENCE INSTITUTE

Mapping between Heterogenous Workflows to Composable Systems

- Burnpro3D: Prescibed Fire Planning Platform
- JupyterHub: Sage Edge App Development and WIFIRE AI Gateway
- Other projects
  - Neurokube: AI Neuroscience imaging
  - TemPredict for IoT use in COVID detection

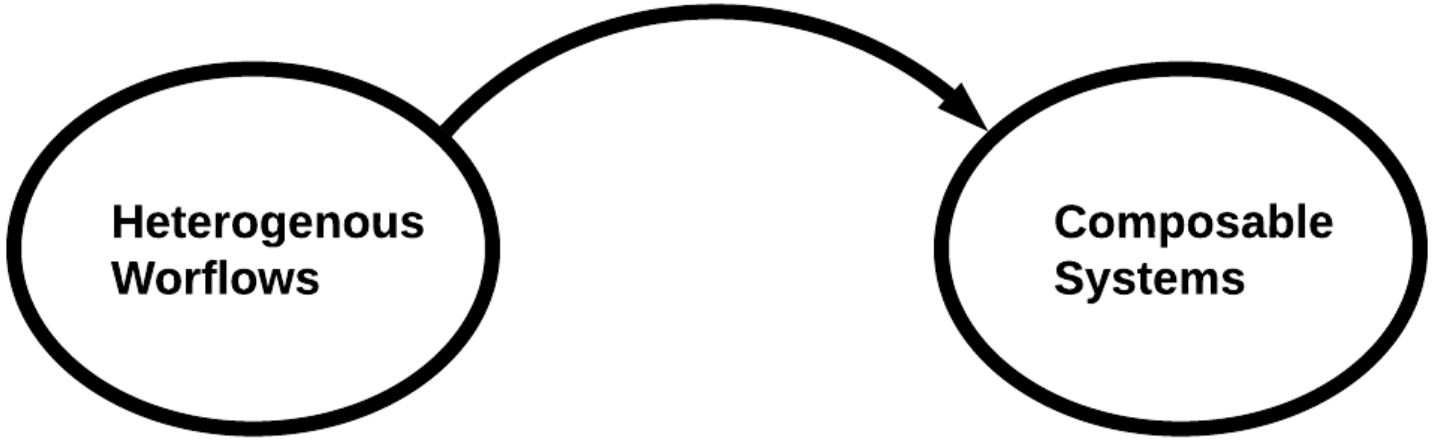

M. Madany, K. Marcus, S. Peltier, M. H. Ellisman and I. Altintas, **"NeuroKube: An Automated and Autoscaling Neuroimaging Reconstruction Framework using Cloud Native Computing and A.I.,"** *2020 IEEE International Conference on Big Data (Big Data)*, Atlanta, GA, USA, 2020, pp. 320-330, doi: 10.1109/BigData50022.2020.9378053.

S. Purawat et al., "TemPredict: A Big Data Analytical Platform for Scalable Exploration and Monitoring of Personalized Multimodal Data for COVID-19," 2021 IEEE International Conference on Big Data (Big Data), 2021, pp. 4411-4420, doi: 10.1109/BigData52589.2021.9671441.

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

To sum up…

Emerging new applications require integrated AI in dynamically composed workflows.

Artwork: ***Jen Stark, Cosmographic, 2014****, acid-free paper, holographic paper, glue, wood, acrylic paint, 34 x 37 x 4 in.*

Embrace Complexity!

Complexity comes at a cost!

- Composable systems is not a turnkey functionality
- Requires collaboration with and between infrastructure providers
- End-to-end data pipelines need to be defined for each application along with microservice execution

İlkay Altıntaş, PhD (ialtintas@ucsd.edu)

Contact: Ilkay Altintas, Ph.D. *Email: ialtintas@ucsd.edu*

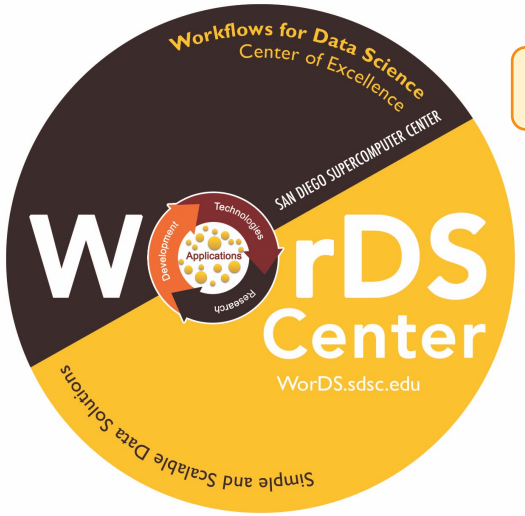

https://words.sdsc.edu/

https://wifire.ucsd.edu/

WIFIRE

We are hiring! -- https://words.sdsc.edu/careers

Questions?

NIH National Institutes of Health

NSF

U.S. DEPARTMENT OF ENERGY | Office of Science

The presented work is collaborative work with many wonderful individuals, and parts of it are funded by NSF, DOE, NIH, UC San Diego and various industry, government and foundation partners.