Extreme Data Management Analysis and Visualization for Exascale Supercomputers and Experimental Facilities

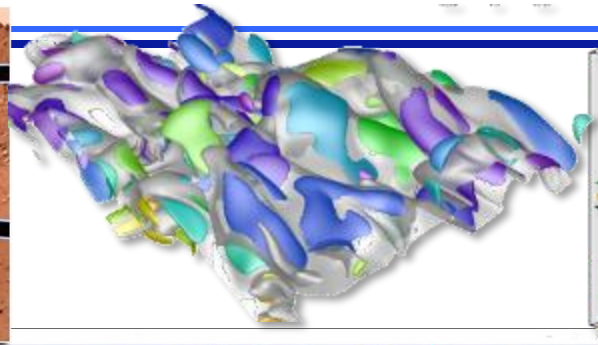

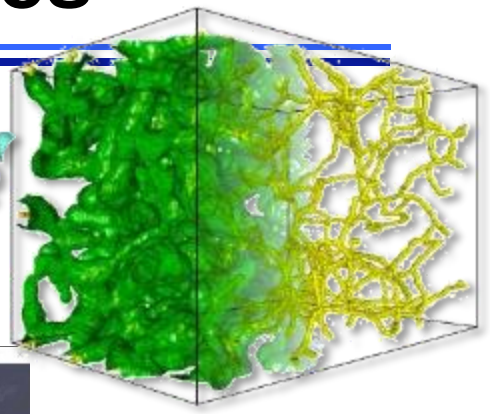

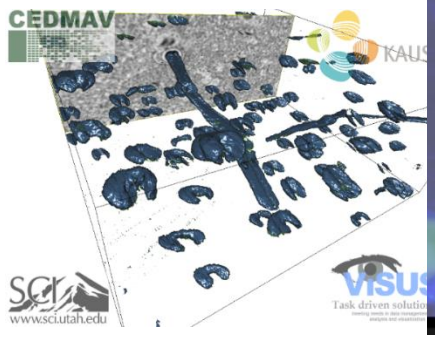

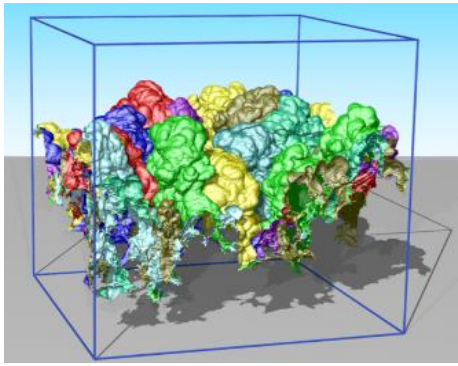

Valerio Pascucci

Inaugural John R. Parks Endowed Chair of the University of Utah
Director, Center for Extreme Data Management Analysis and Visualization
Professor, SCI institute and School of Computing, University of Utah
Laboratory Fellow, PNNL, CEO, Visus LLC

CEDMAV

SCI THE UNIVERSITY OF UTAH Pacific Northwest NATIONAL LABORATORY

Center for Extreme Data Management, Analysis, and Visualization

- 10 Faculty + scientists, developers, students, …
- Primary partners: UU & PNNL
- Other partnerships: NSA, INL, LLNL, ANL, Battelle, ….
- Involvement in national Initiatives

SDAV Scalable Data Management, Analysis, and Visualization

EXACT

SUPER

THE UNIVERSITY OF UTAH Pacific Northwest NATIONAL LABORATORY

Lawrence Livermore National Laboratory

Sandia National Laboratories

INL Idaho National Laboratory

$1.6B NSA data center
(1.5 million-square-foot facility)

Massive Simulation and Sensing Devices Generate Great Challenges and Opportunities

BlueGene/Q

Satellite

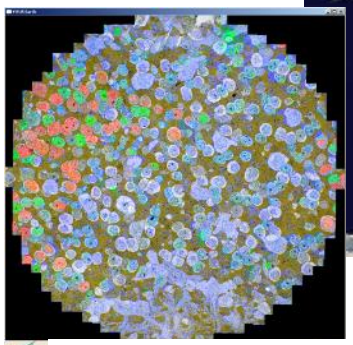

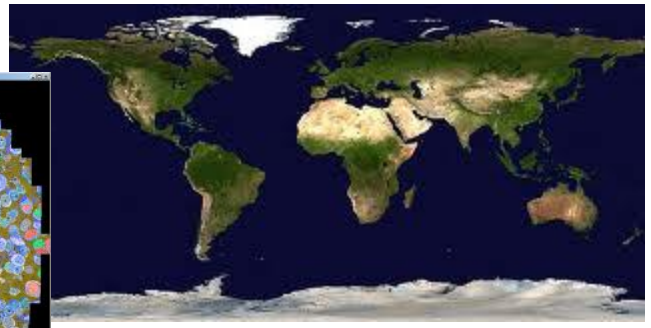

Earth Images

Retinal Connectome

Cameras

EM

Jaguar

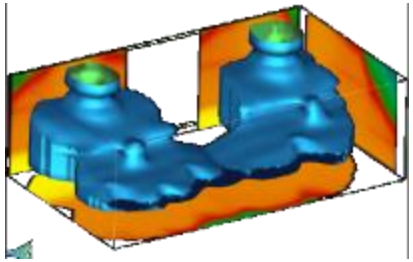

Carbon Seq. (Subsurface)

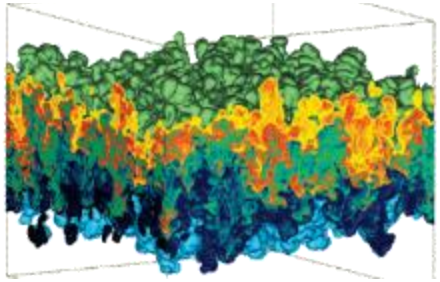

Hydrodynamic Inst.

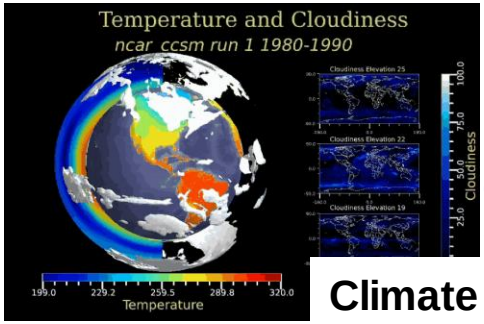

Climate

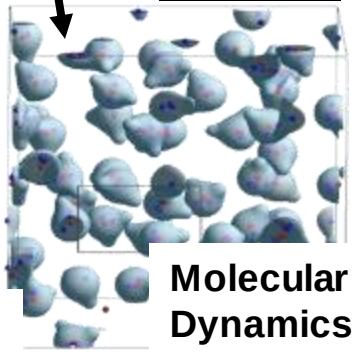

Molecular Dynamics

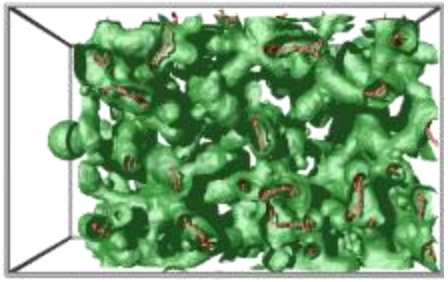

Porous Materials

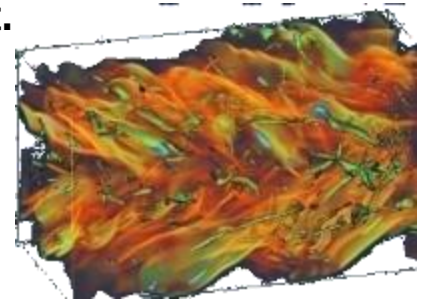

Turbulent Combustion

Photography

Understanding the Dynamics of Rayleigh-Taylor instabilities

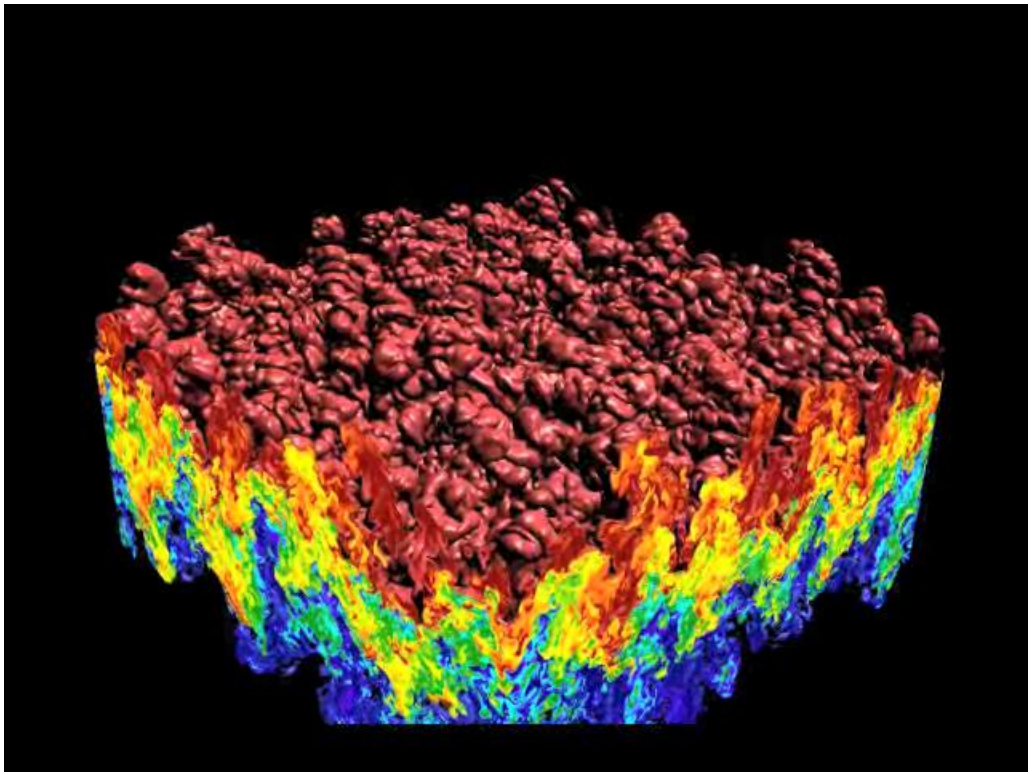

Rayleigh-Taylor instabilities arise in fusion, super-novae, and other fundamental phenomena:

- start: heavy fluid above, light fluid below
- gravity drives the mixing process
- the mixing region lies between the upper envelope surface (red) and the lower envelope surface (blue)
- 25 to 40 TB of data from simulations

Simulation of a Satellite Deflagration for Space Awareness

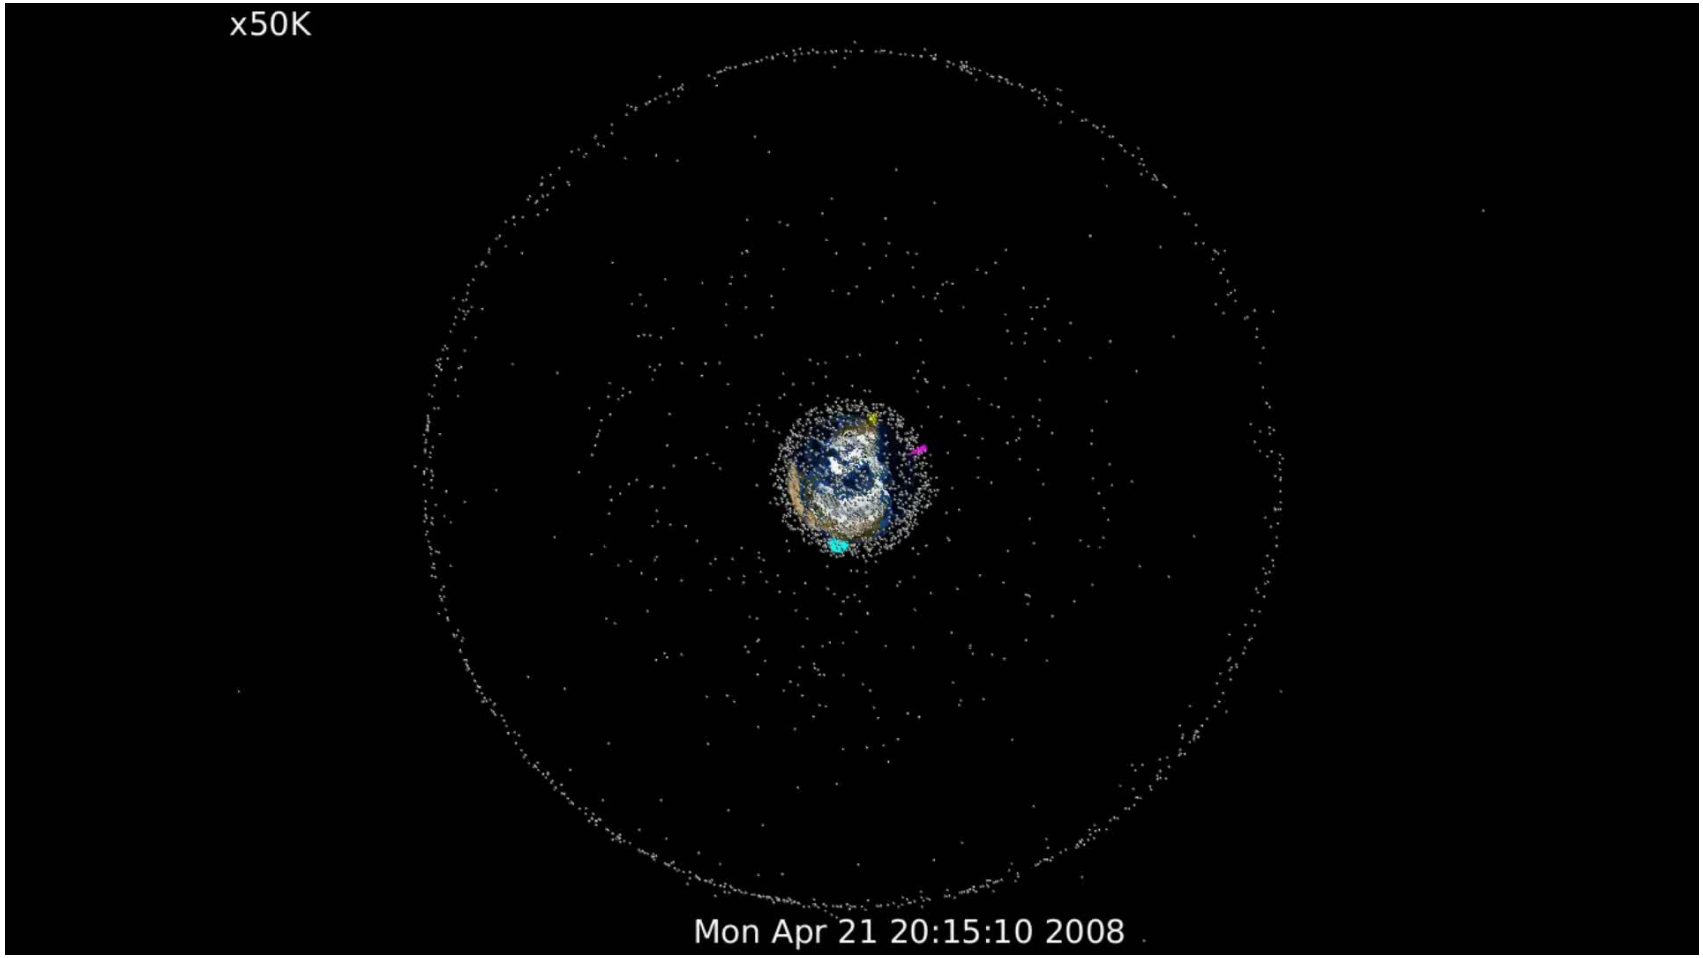

Computational Infrastructure for Information Discovery is Highly Interdisciplinary

- Performance
  - remote data access
  - Scaling for HPC
  - Progressive techniques
  - In-situ analytics
  - Compression
  - Asynchronous computing
- User Access
  - Usability
  - Platform portability
  - Collaboration
  - Data abstractions
  - Visual metaphors
  - User interactions
- Analytics
  - Statistic
  - Topology
  - Geometry
  - Data mining
  - Machine learning
  - Feature extraction/tracking
- Applications
  - Smart Cities: services, population, healthcare, …
  - Simulations: climate, combustion, astrophysics, …
  - Experiments: microscopy, light sources, tomography, …
  - Data Collection: agriculture,…

A Information Cyberinfrastructure Requires Efficient Big Data Management and Analytics

- **Advanced data storage techniques:**
  - **Data re-organization.**
  - **Compression.**

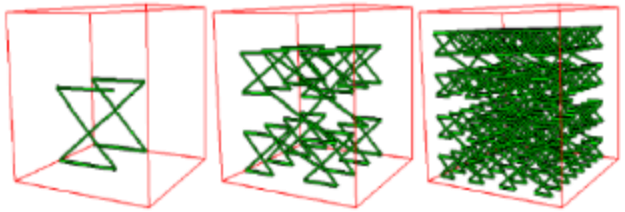

- **Advanced algorithmic techniques:**
  - **Streaming.**
  - **Progressive multi-resolution.**
  - **Out of core computations.**

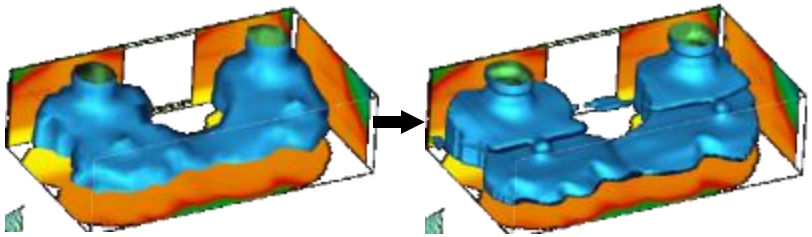

- **Scalability across a wide range of running conditions:**
  - **From laptop, to office desktop, to cluster of PC, to BG/L.**
  - **Memory, to disk, to remote data access.**

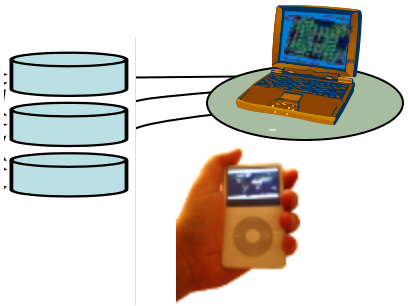

Topology Has Been Successful for Analysis and Visualization of Massive Scientific Data

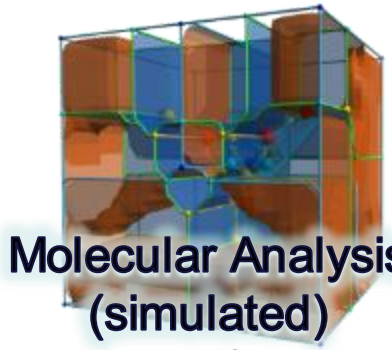

Molecular Analysis (simulated)

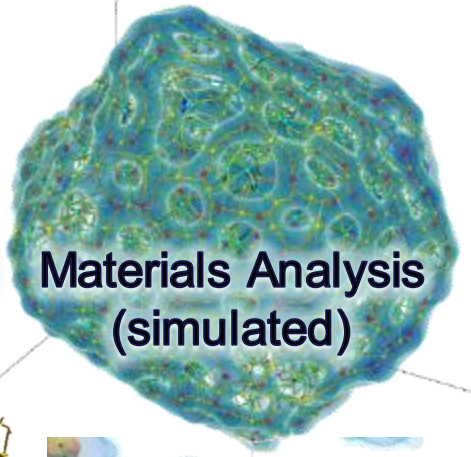

Materials Analysis (simulated)

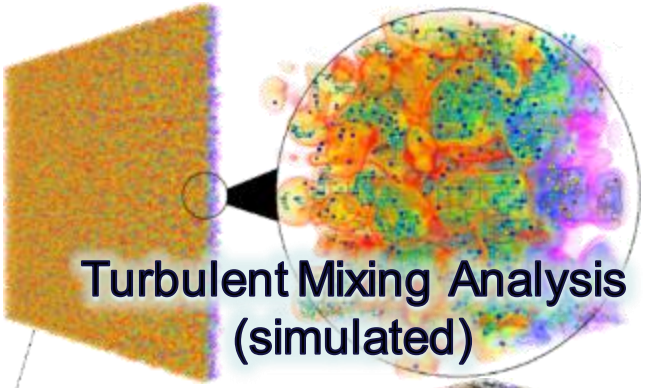

Turbulent Mixing Analysis (simulated)

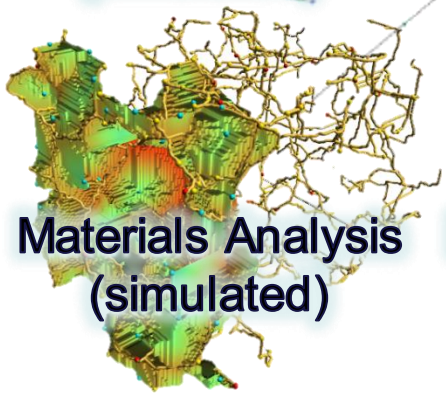

Materials Analysis (simulated)

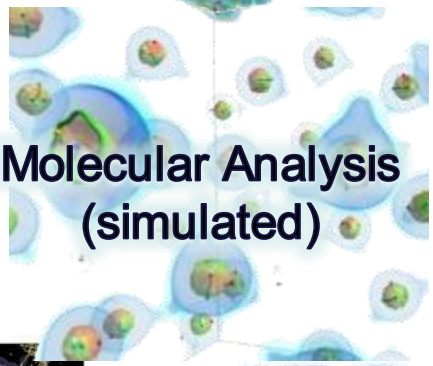

Molecular Analysis (simulated)

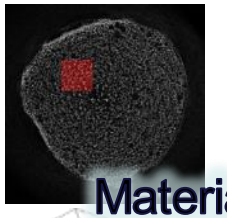

Materials Analysis (imaged)

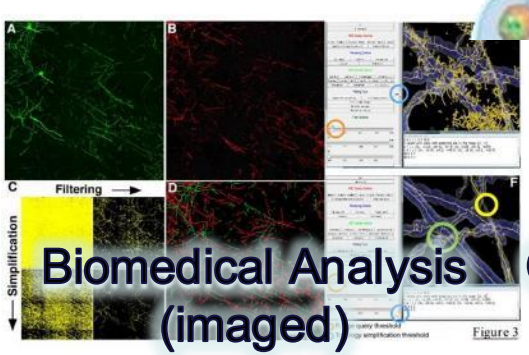

Biomedical Analysis (imaged)

Cosmology Analysis (simulated)

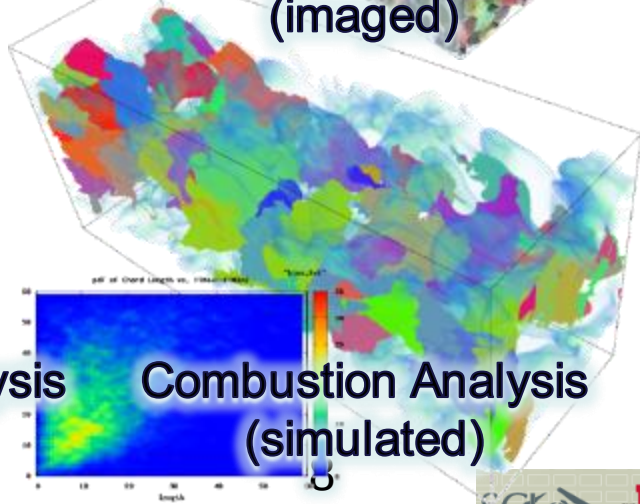

Combustion Analysis (simulated)

Quantitative Analysis of the Impact of a Micrometeoroid in a Porous Medium

- **Many possible applications:**
  - **NASA's Stardust Spacecraft**
  - **National Ignition Facility Targets**
  - **Light and Robust Materials**
  - **many more...**

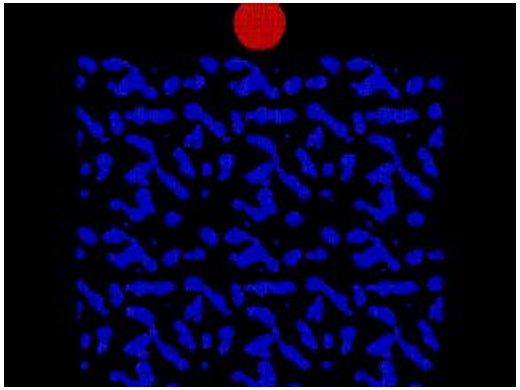

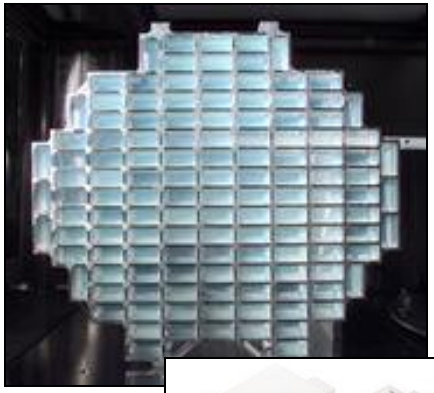

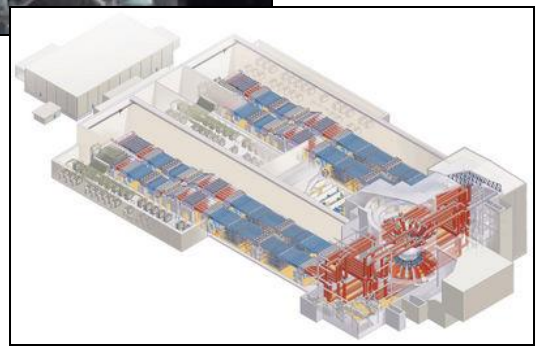

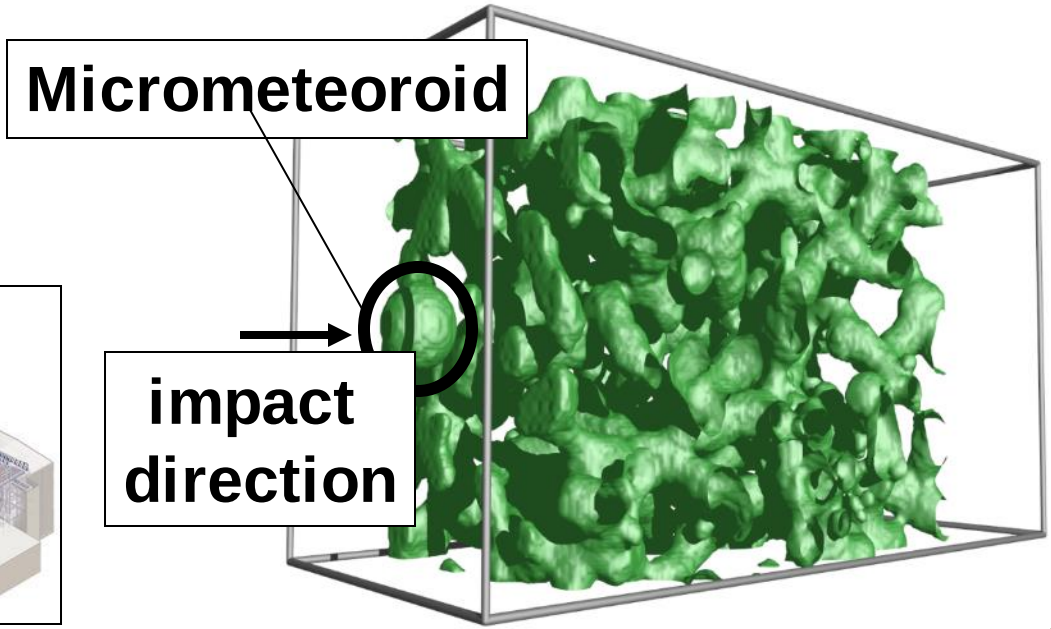

Porous Medium

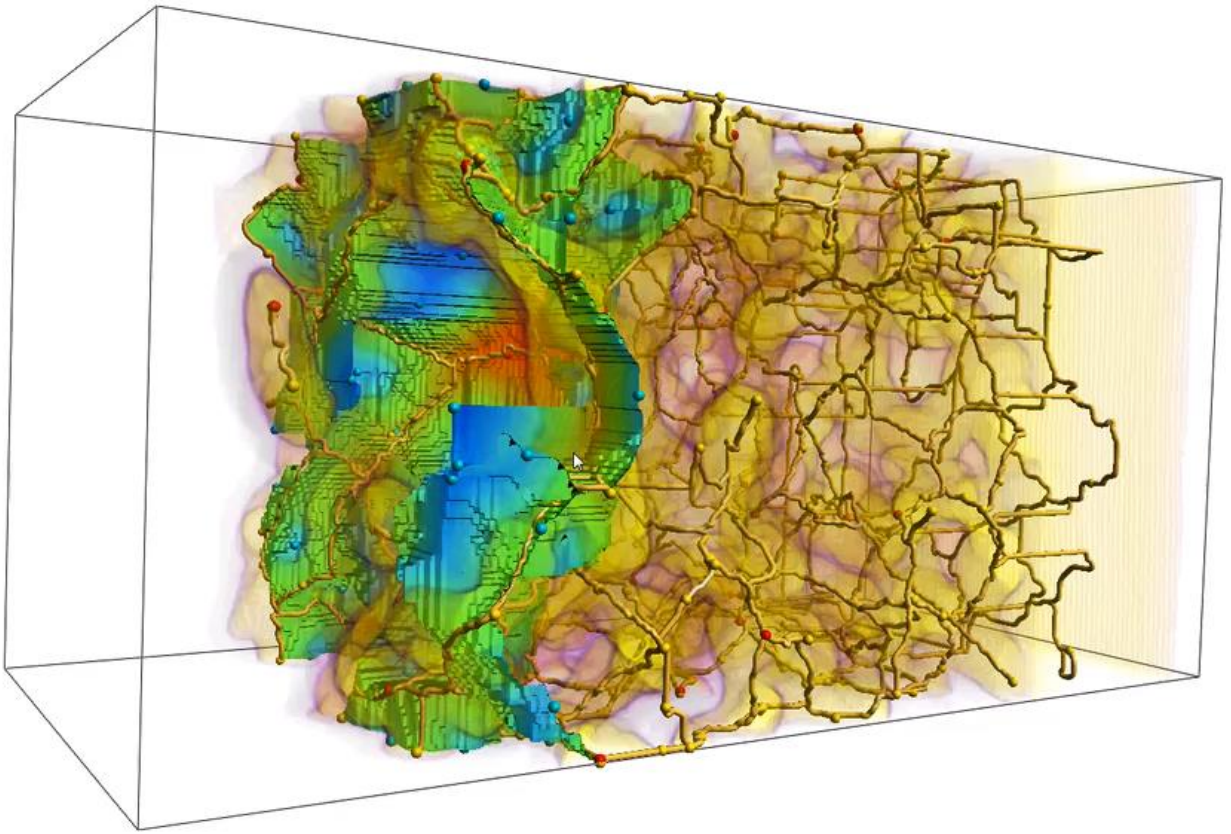

Lithium-Ion Battery

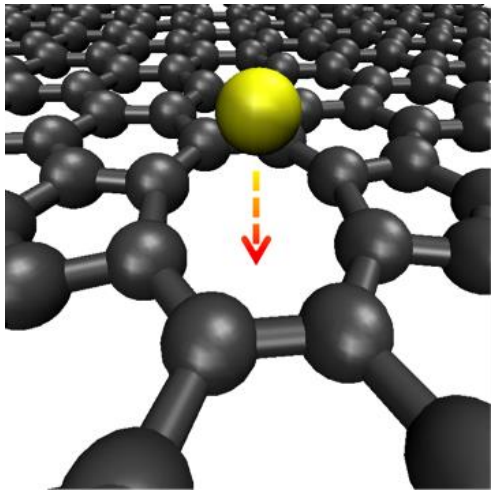

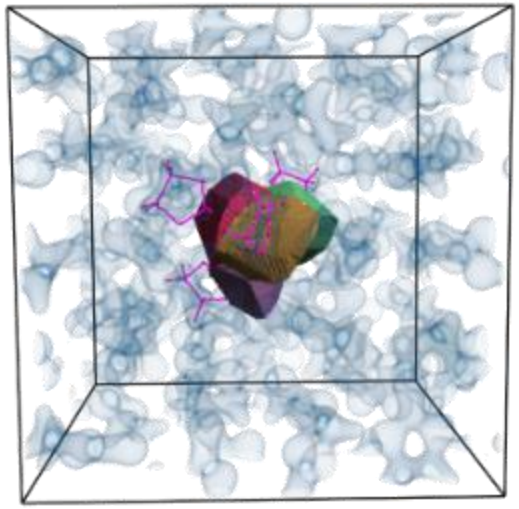

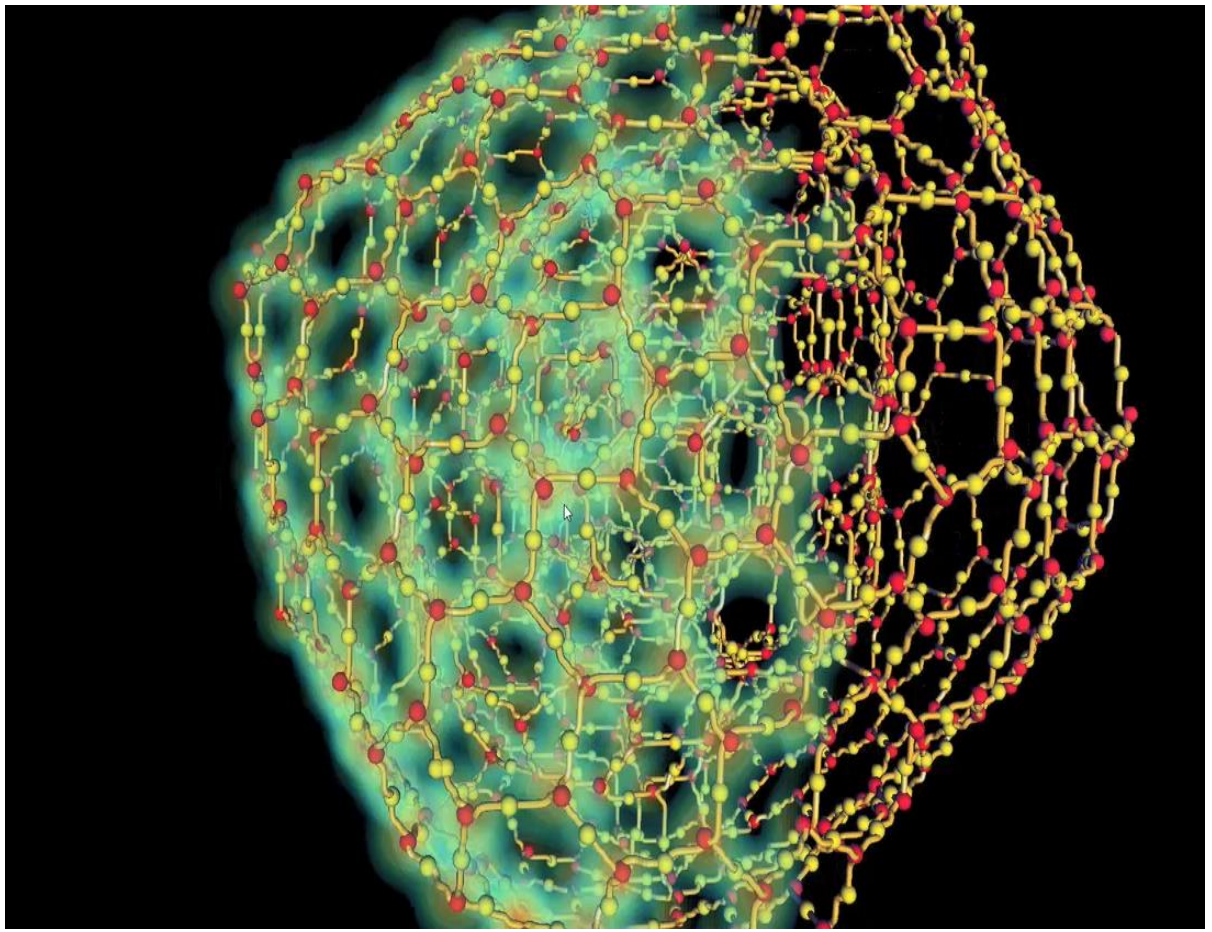

Analysis of Extinction and Reignition Regions for Hydrogen Combustion

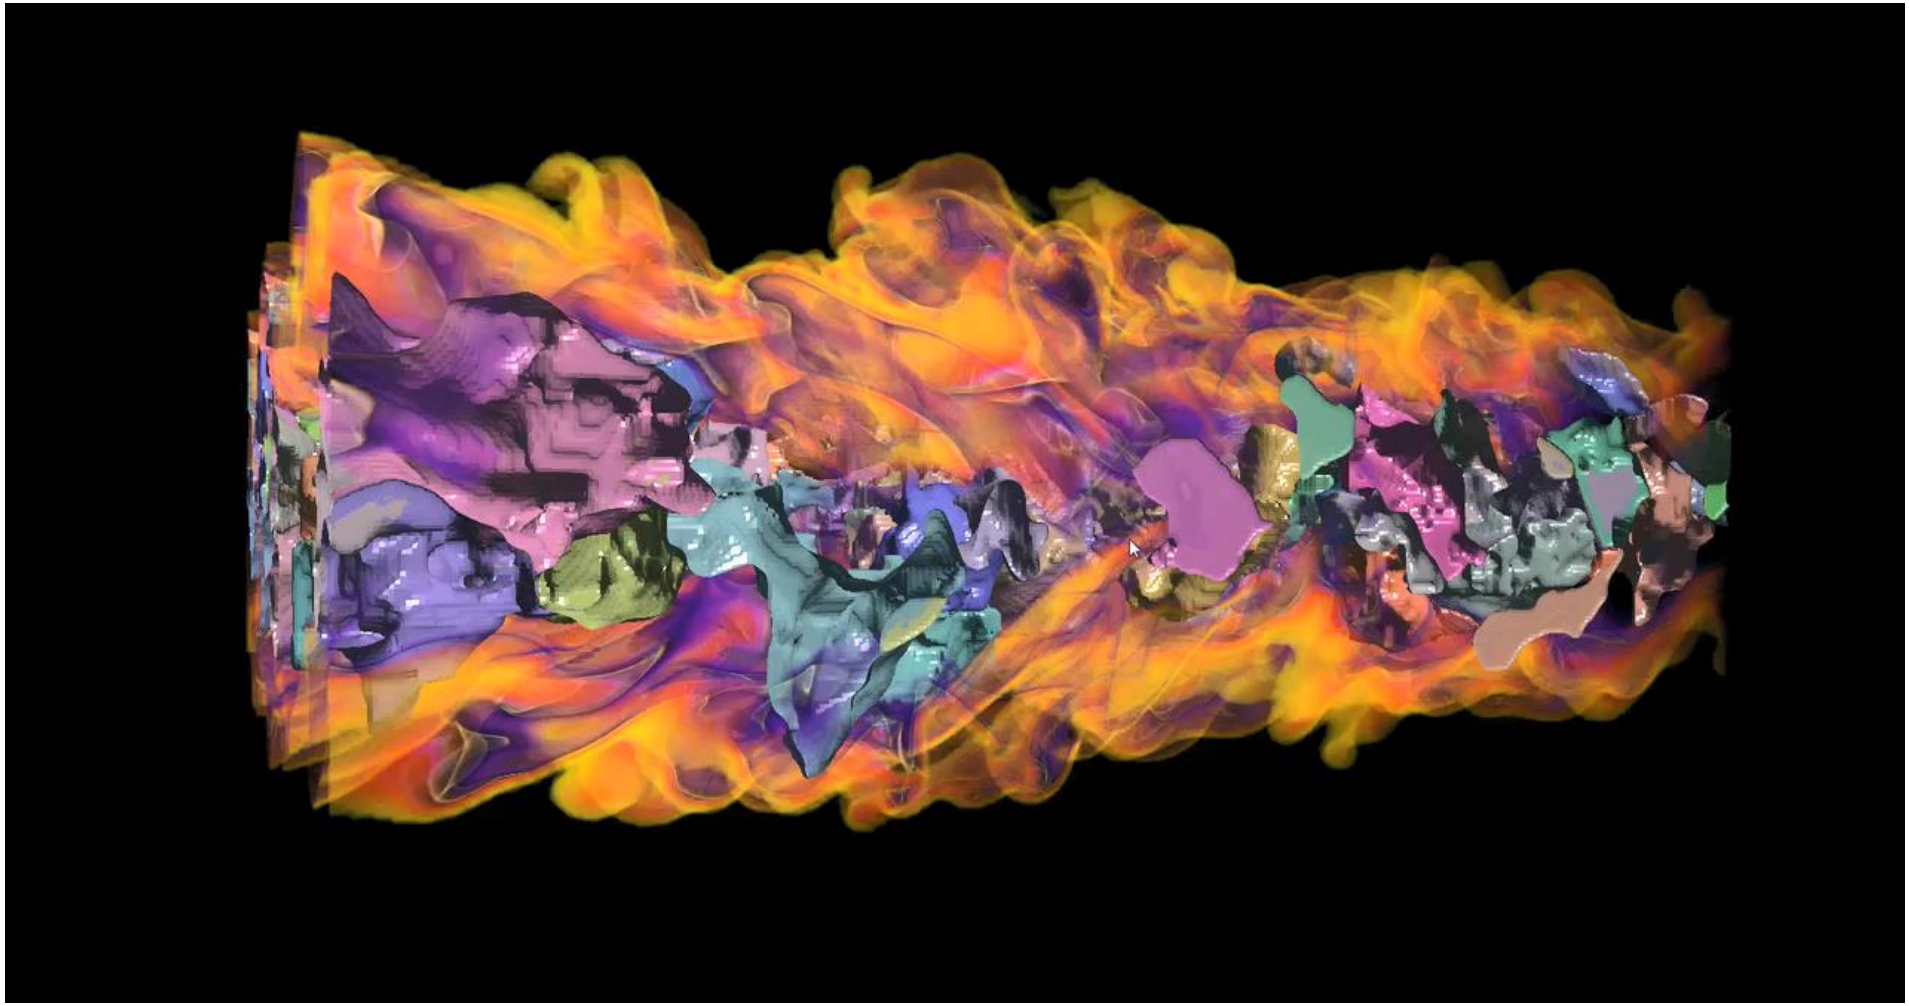

Integrated Topological Analysis and Visualization

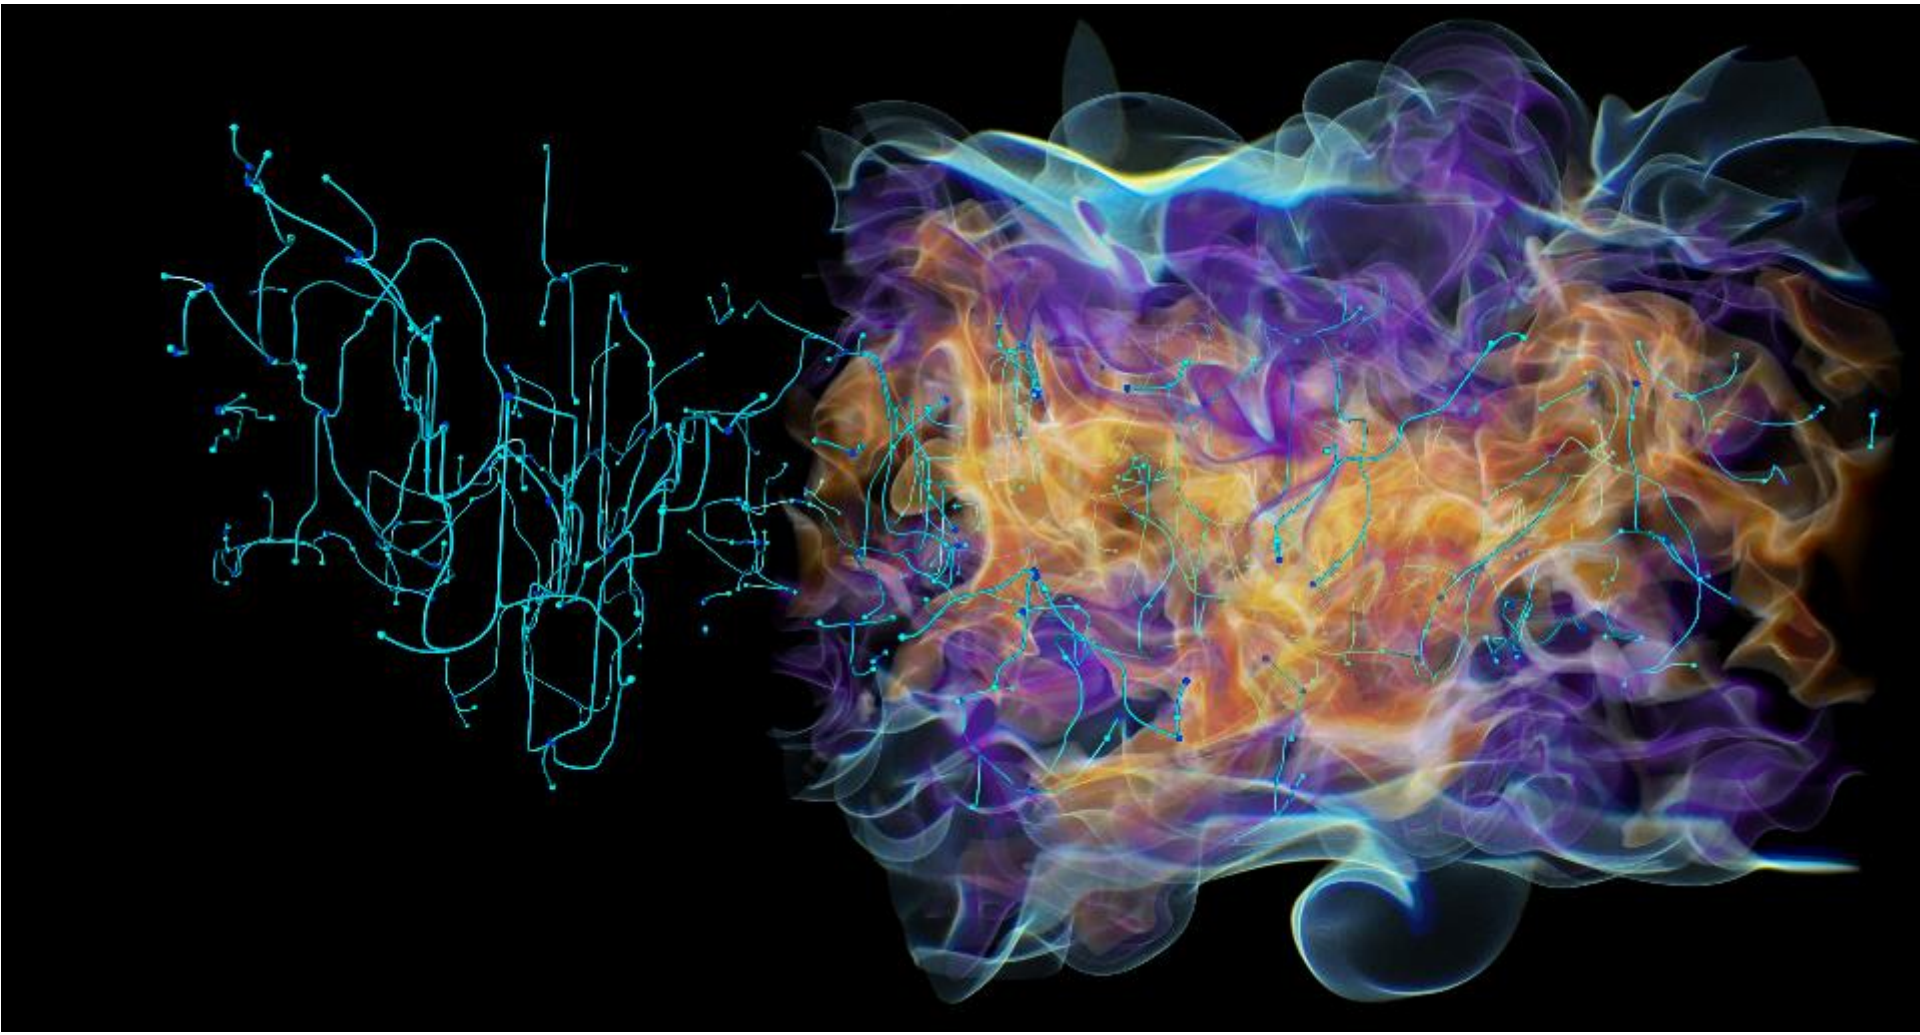

Now on Display at Cité des Sciences's Exhibit on Fire in Paris

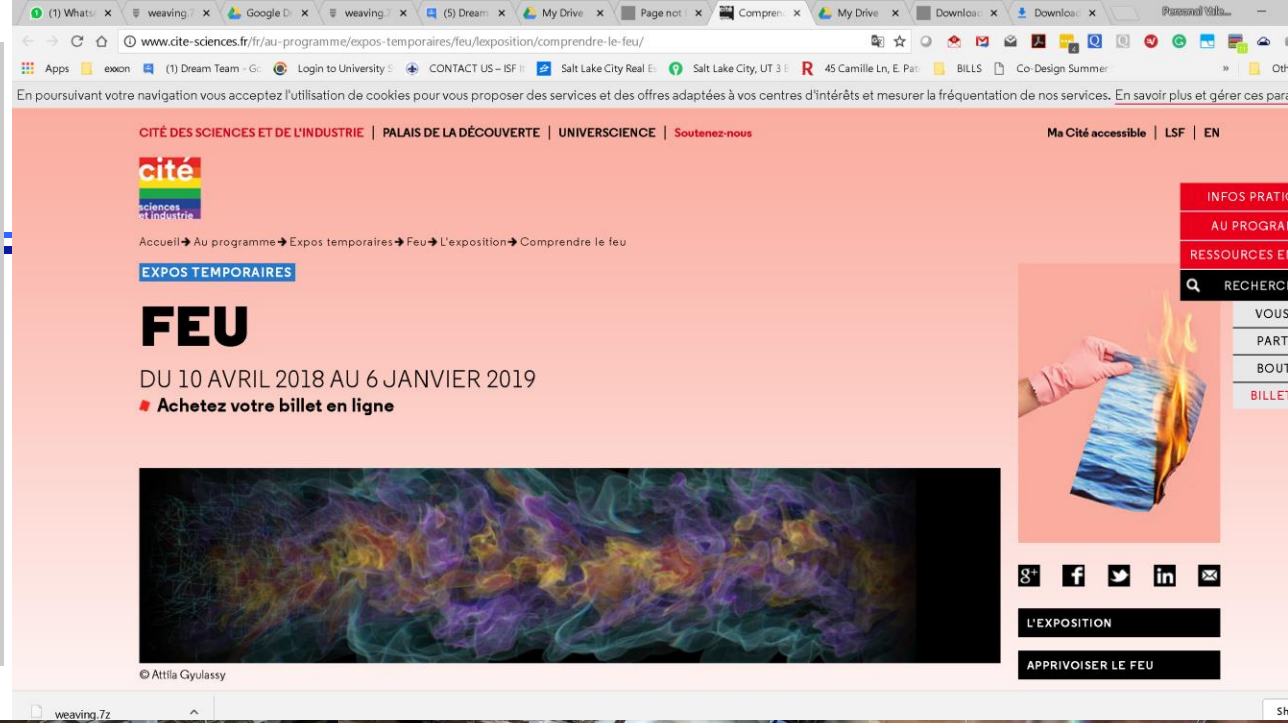

Streaming Analytics and Visualization

Live demonstration from ANL to SLC

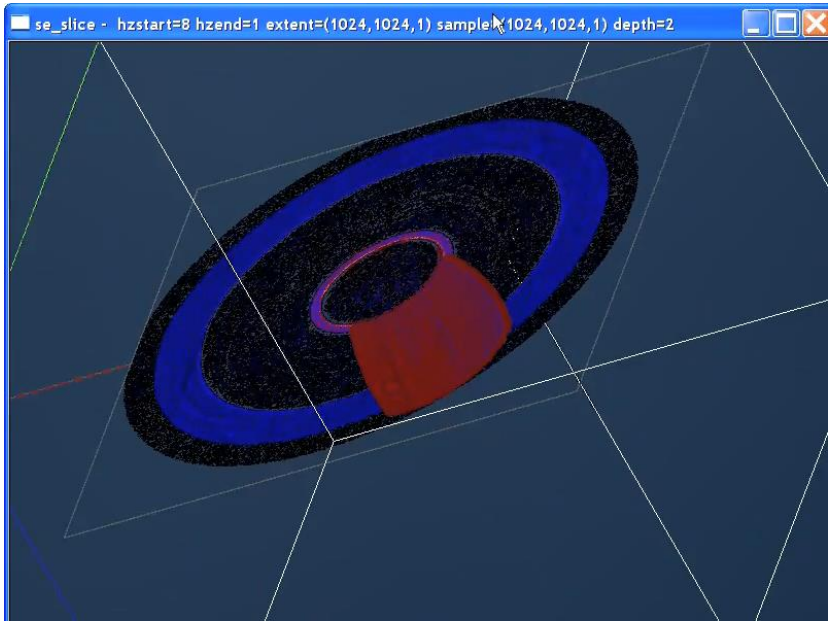

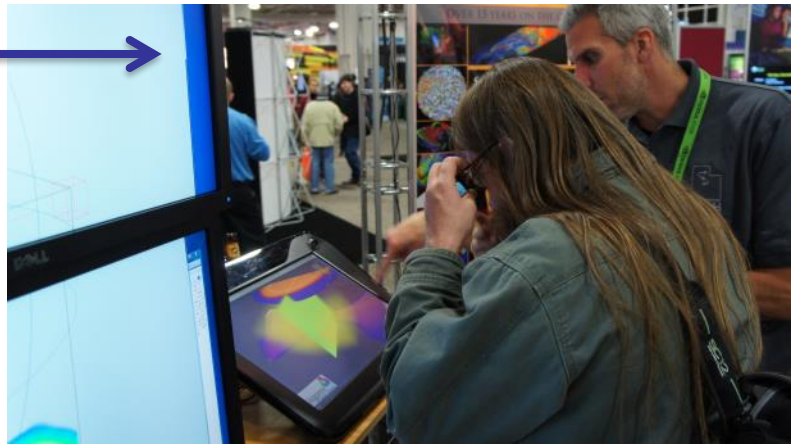

Infrastructure that scales gracefully with available hardware resources

Cores available

Scalable Deployment: Real Time Exploration of 3.5 Petabytes of Weather/Climate Data

Workflow

- ***Data creation***
- ***Data Management***
- Processing
- Analysis
- Visualization

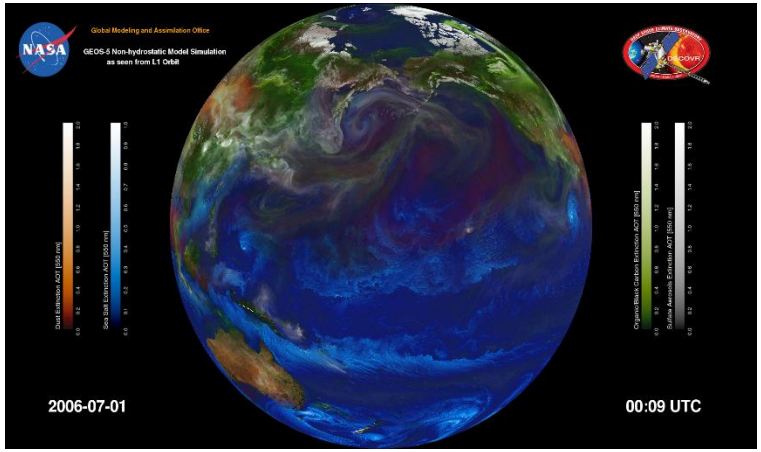

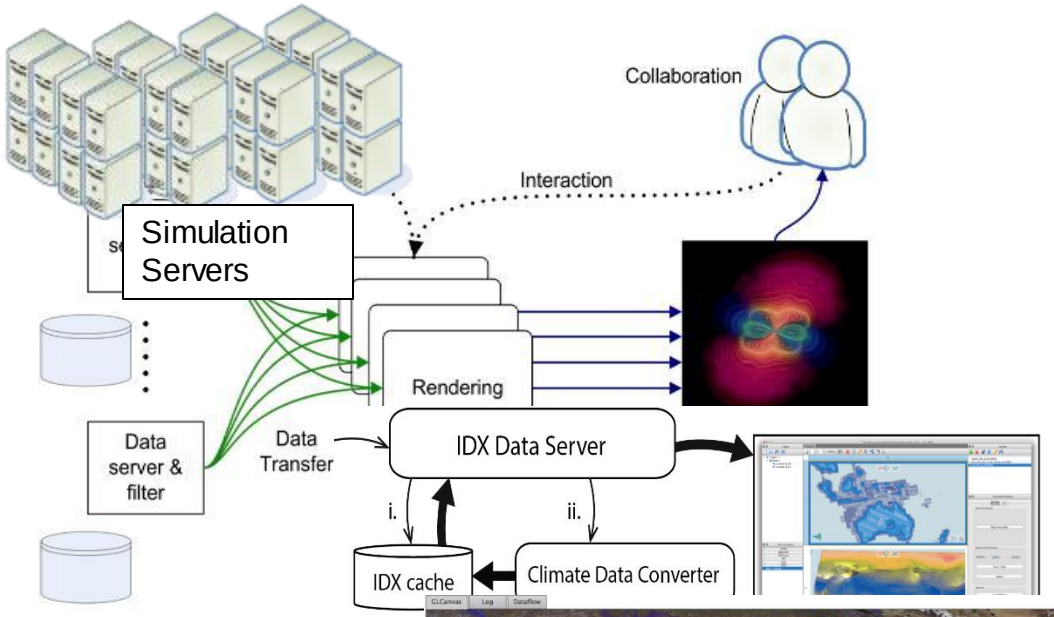

- 7km GEOS-5 "Nature Run" -> 1 dataset, 3.5 PB
- theoretically: openly accessible -> practically: precomputed pics

Distributed Resources

- 3.5 PB of *data store in NASA*
- *Primary ViSUS server in LLNL*
- *Secondary ViSUS server in Utah*
- *Clients connect remotely*
- *Work without additional HPC resources*

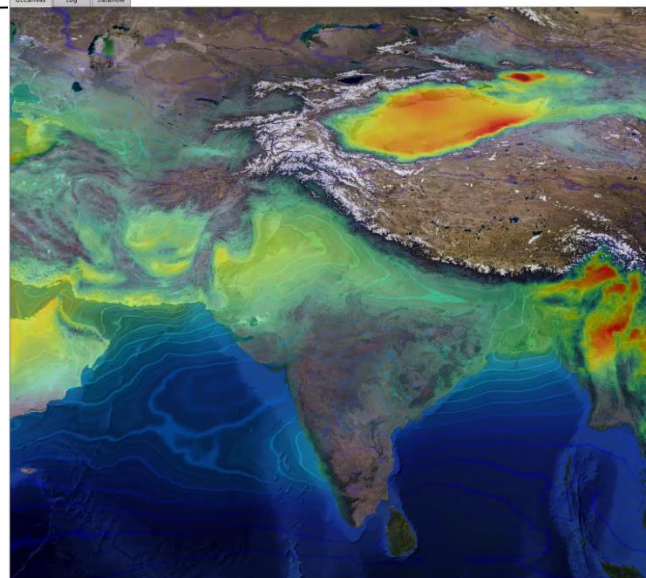

http://atlantis.sci.utah.edu/visus/webviewer/
nature_2007_aer1_hourly

High Performance Data Movements for Real-Time Access to Large Scale Experimental Data

- Experiment run at Advance Photon Source at ANL
- Materials Scientists at University of Utah

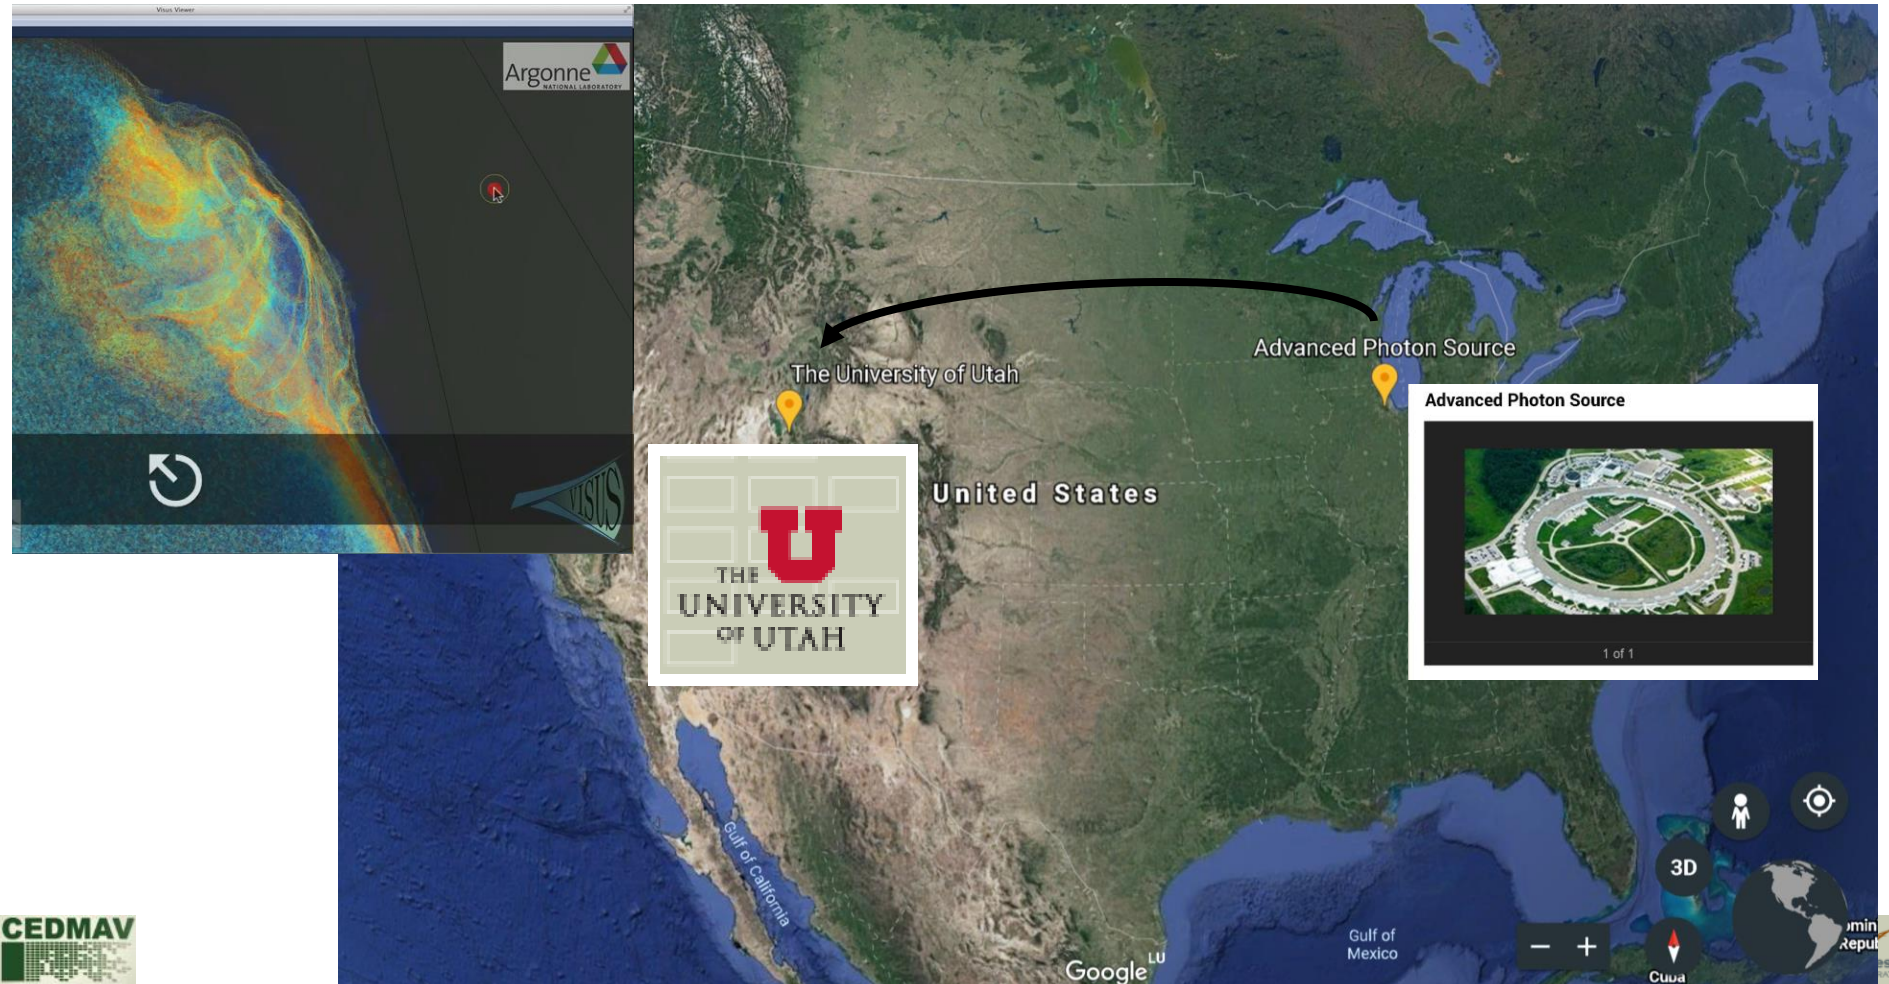

Data Acquisitions at Advanced Photon Source facilities at ANL

- April 2-6th 2018 Prof. Spears's students collected X-ray and CT images of size 1560x1024x1024 (per experiment)
- Traditional workflow (typical for long tail of science):
  - store data in a hard drive -> go back to the lab (Utah)
  - reconstruct the volume
  - limited data exploration (no visualization clusters …)
  - Data sharing? FedEx!
  - What if data are not as desired or partially corrupted?
    - Go back to the facility and repeat the acquisition…
- The lead could not join the team during acquisition
- How to determine if the acquisition was good before leaving the facilities?
- How to get any feedbacks from the lead or any collaborators during the acquisition?

Deployment of data processing and streaming capabilities at APS

- **Easily deployed as a Docker container (on the beam line) including:**
  - **Server for data streaming**
  - **Data processing utilities (stich images, data format conversion)**
  - **Web viewer for fast exploration of data acquired (from your browser)**
- **Timeline**
  - **Local test of data stitching, conv. and stream: < 1 week**
  - **Deployment of the same setup on the beam line: < 1 day**
  - **Image data pre-processing and ready to stream : < 1 hour**
- **Real time visualization vis VPN!!!!**
  - **Alternative though same VPN: >> 1 month**
    Someone is still be waiting :-)

A new practice for acquiring, collecting and sharing big data

- **Using a desktop client or a webviewer Prof. Spears was able to see the data being acquired at APS from her office in UoU**
- **Webviewer Demo: Aluminum Foam of similar size**

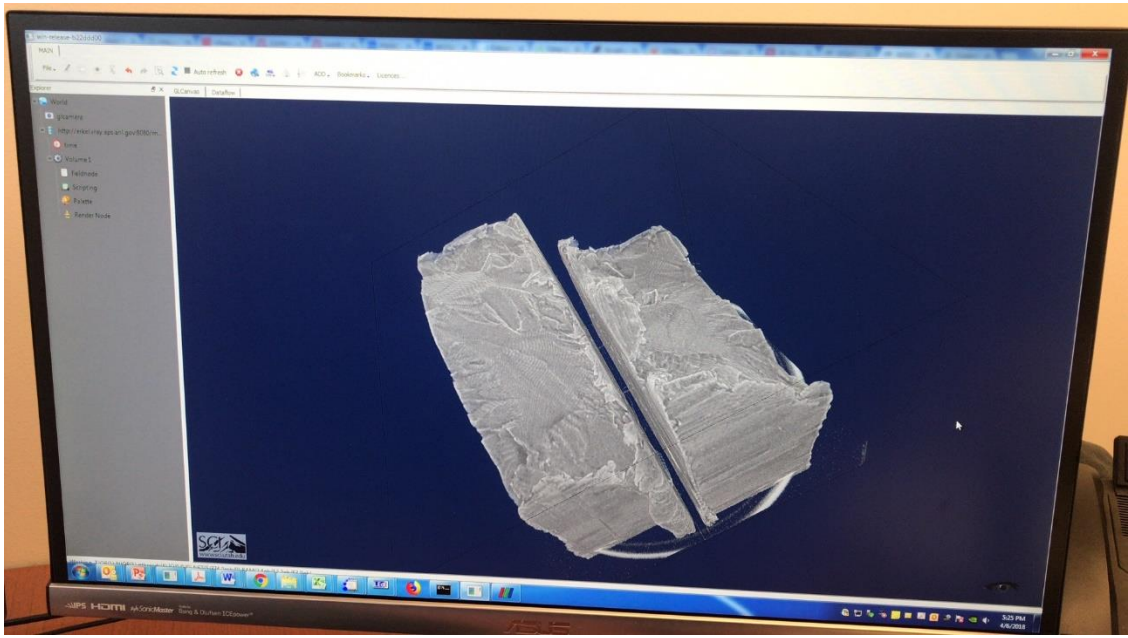

Online Acquisition and Interactive Visualization of Terascale Microscopy

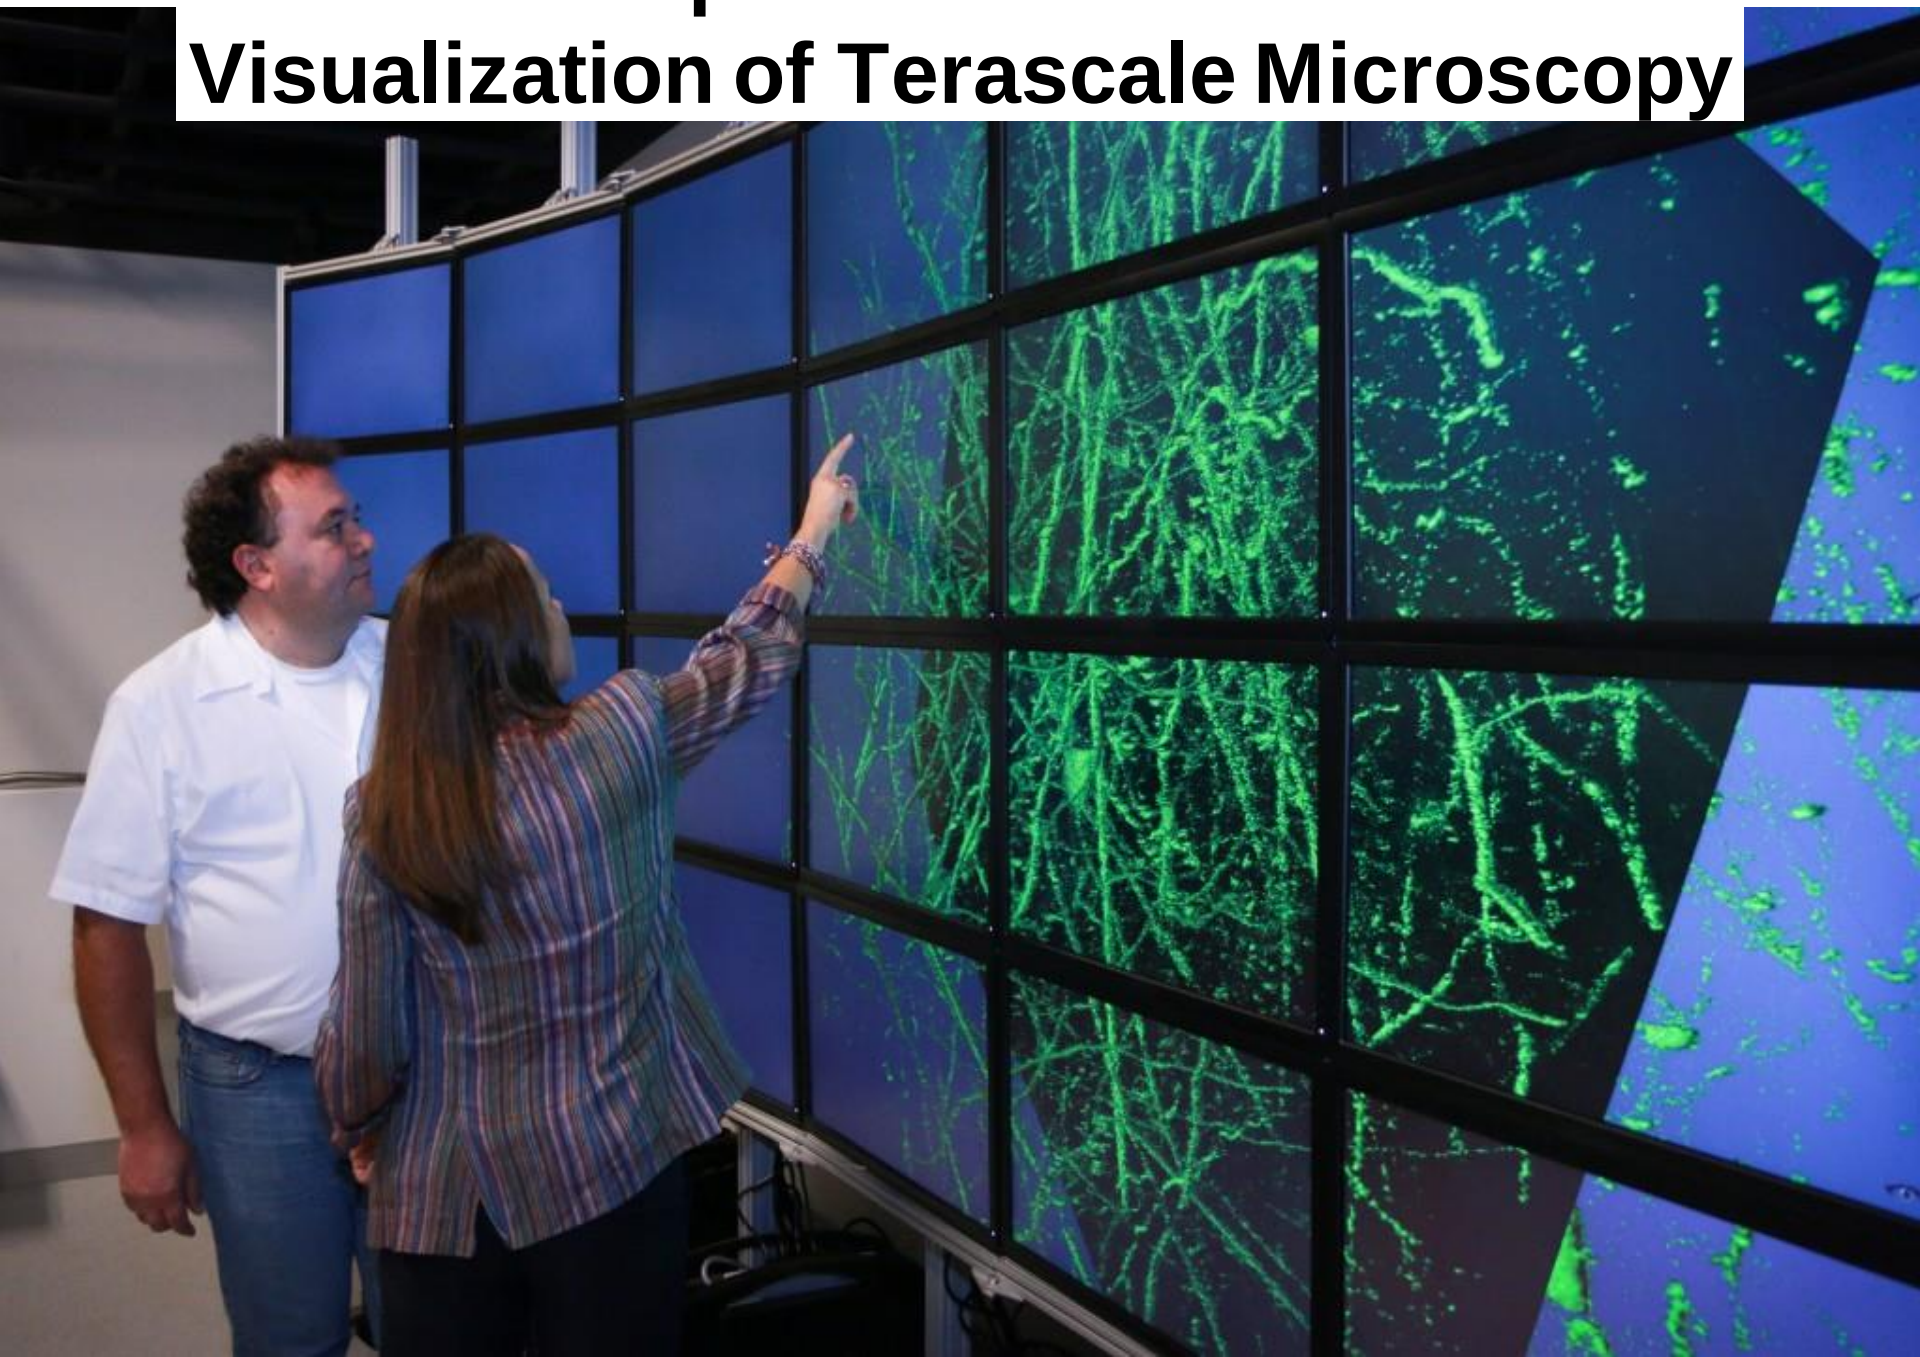

Remote Monitoring of Data Quality During Acquisition

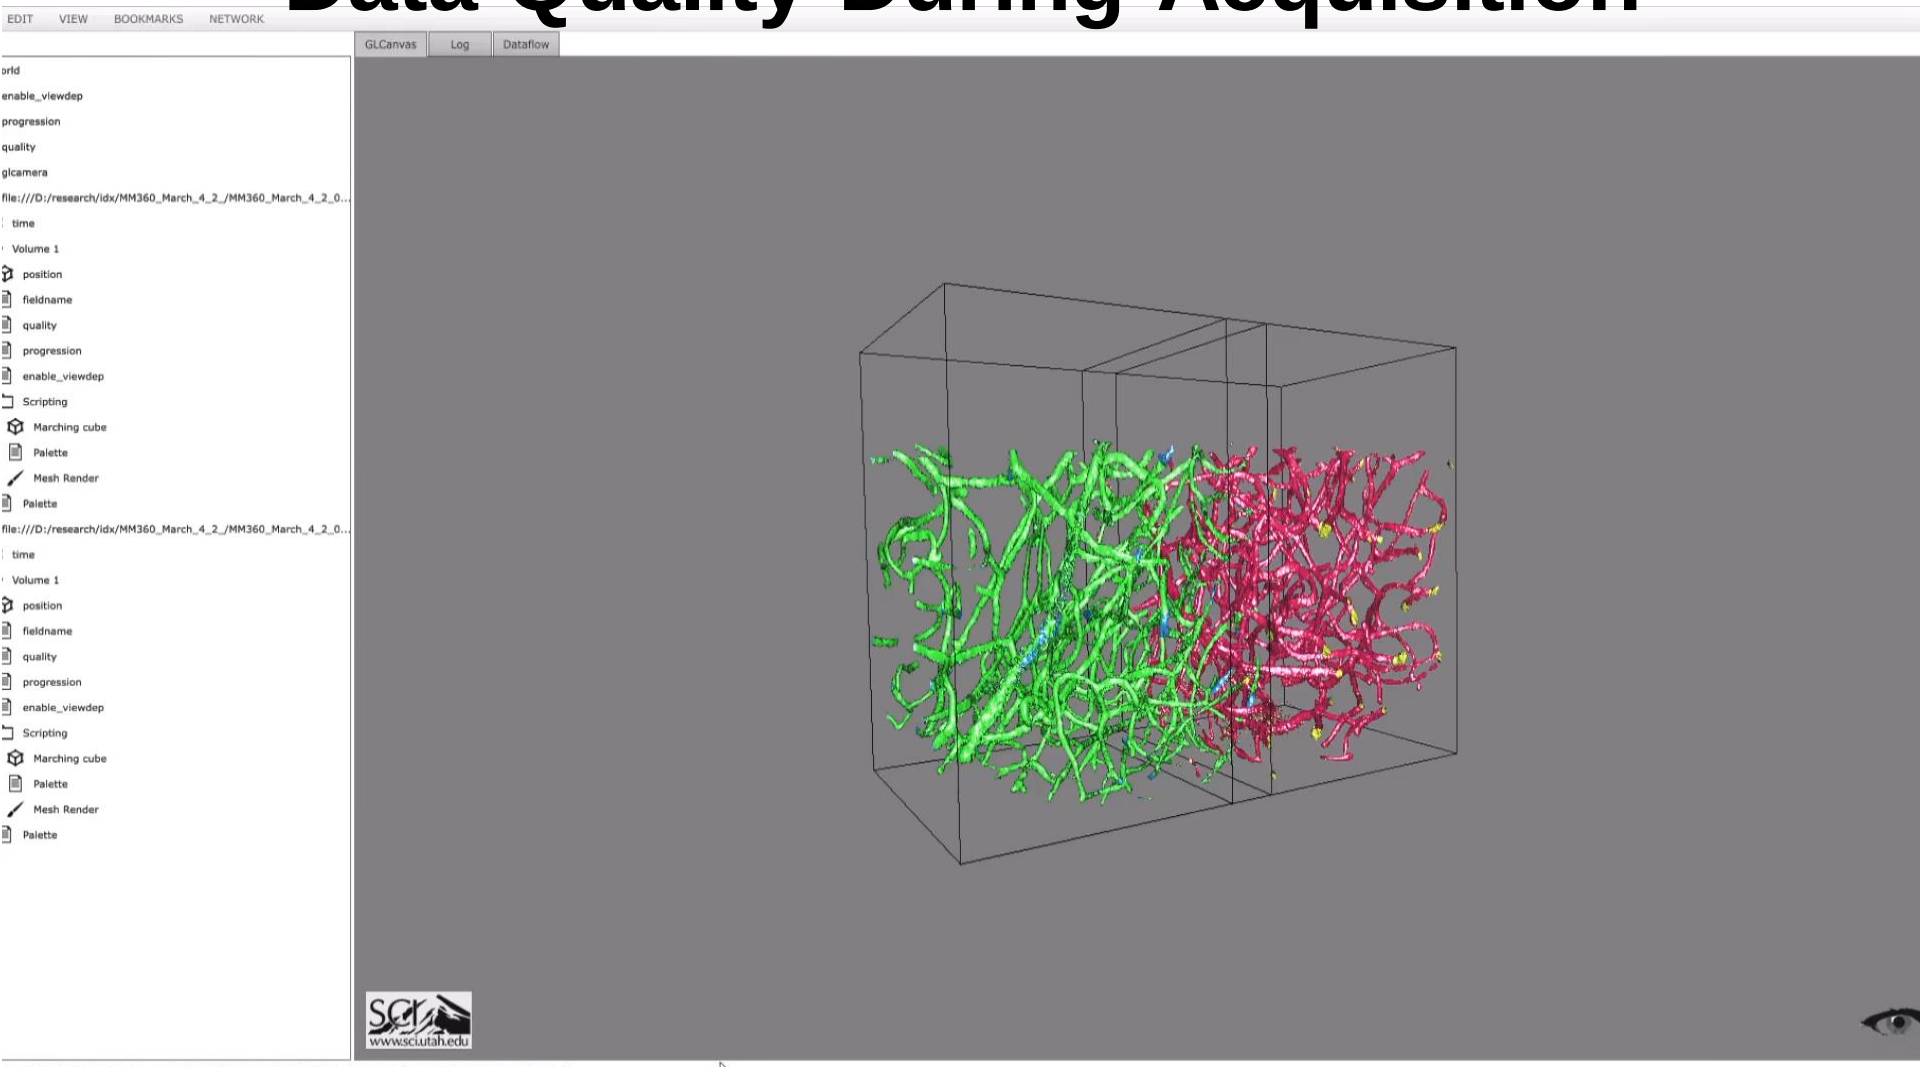

Demo: large Scale Geology Data

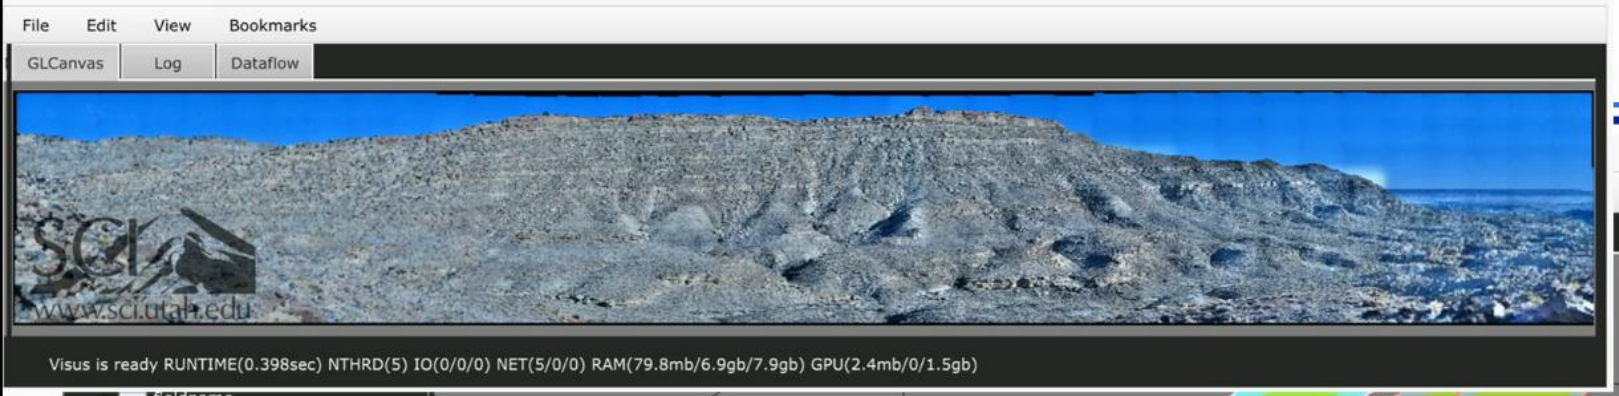

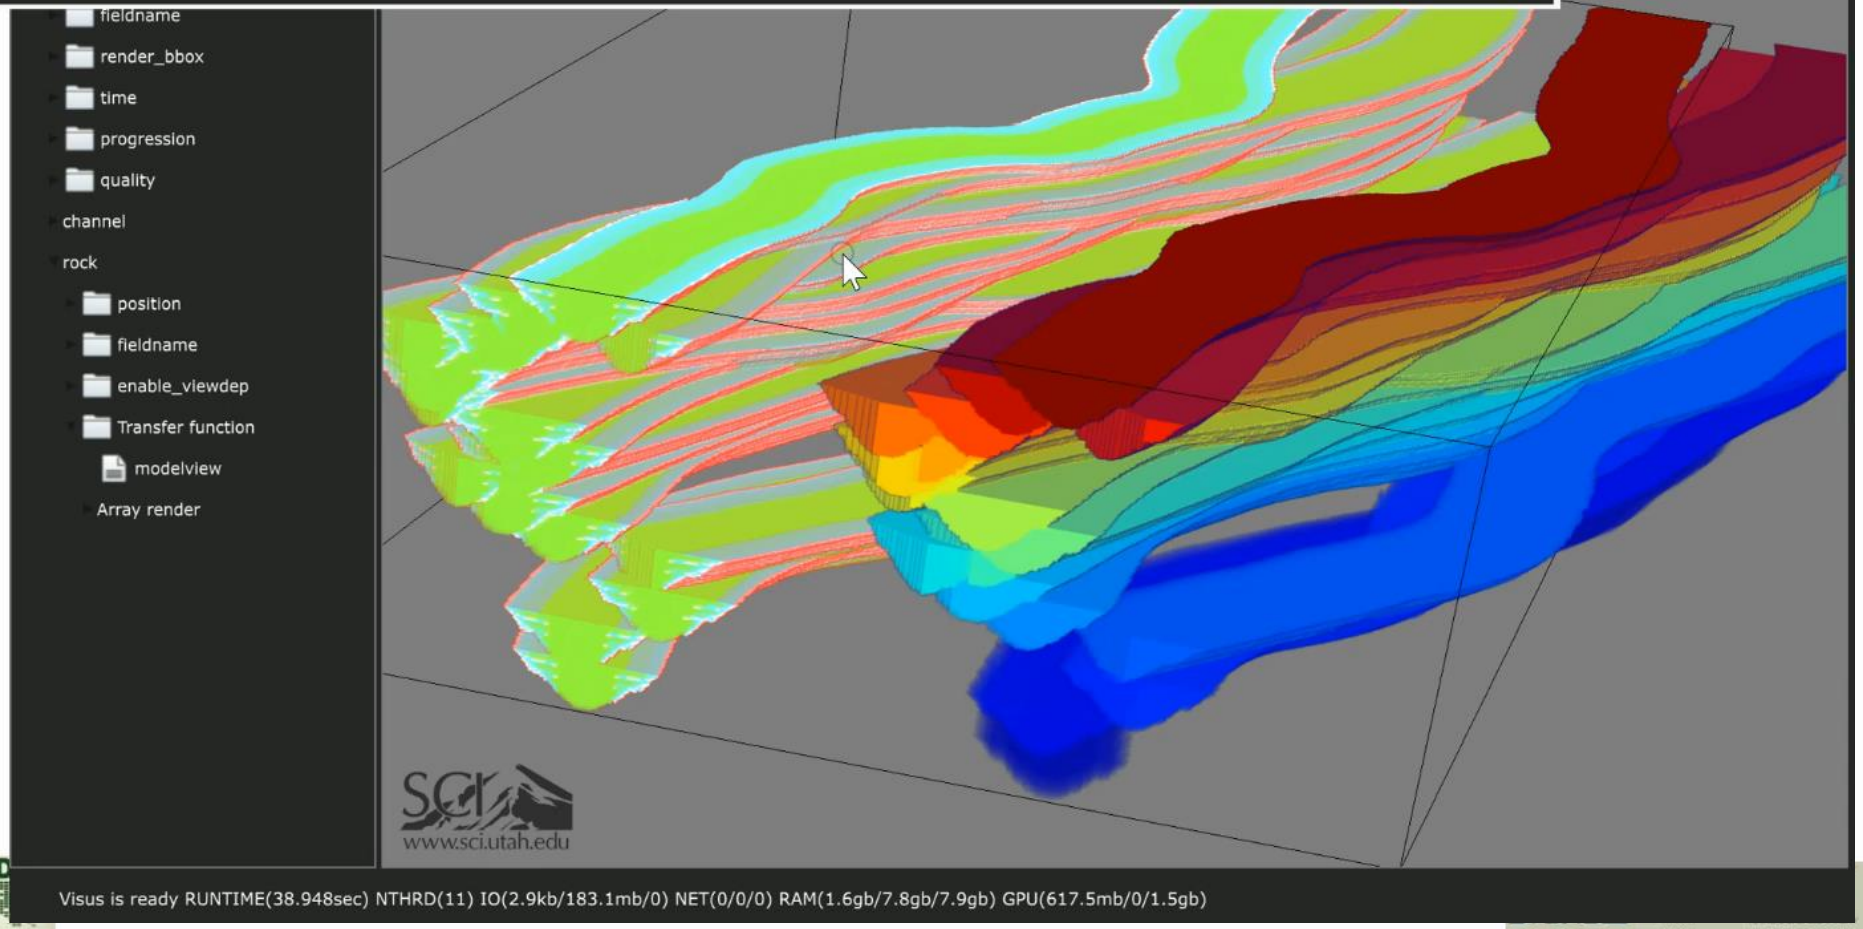

High Resolution Display Platforms for High Resolution Outcrop and Seismic Data

KAUST PowerWall: Installed and Fully Operational in a Few Hours

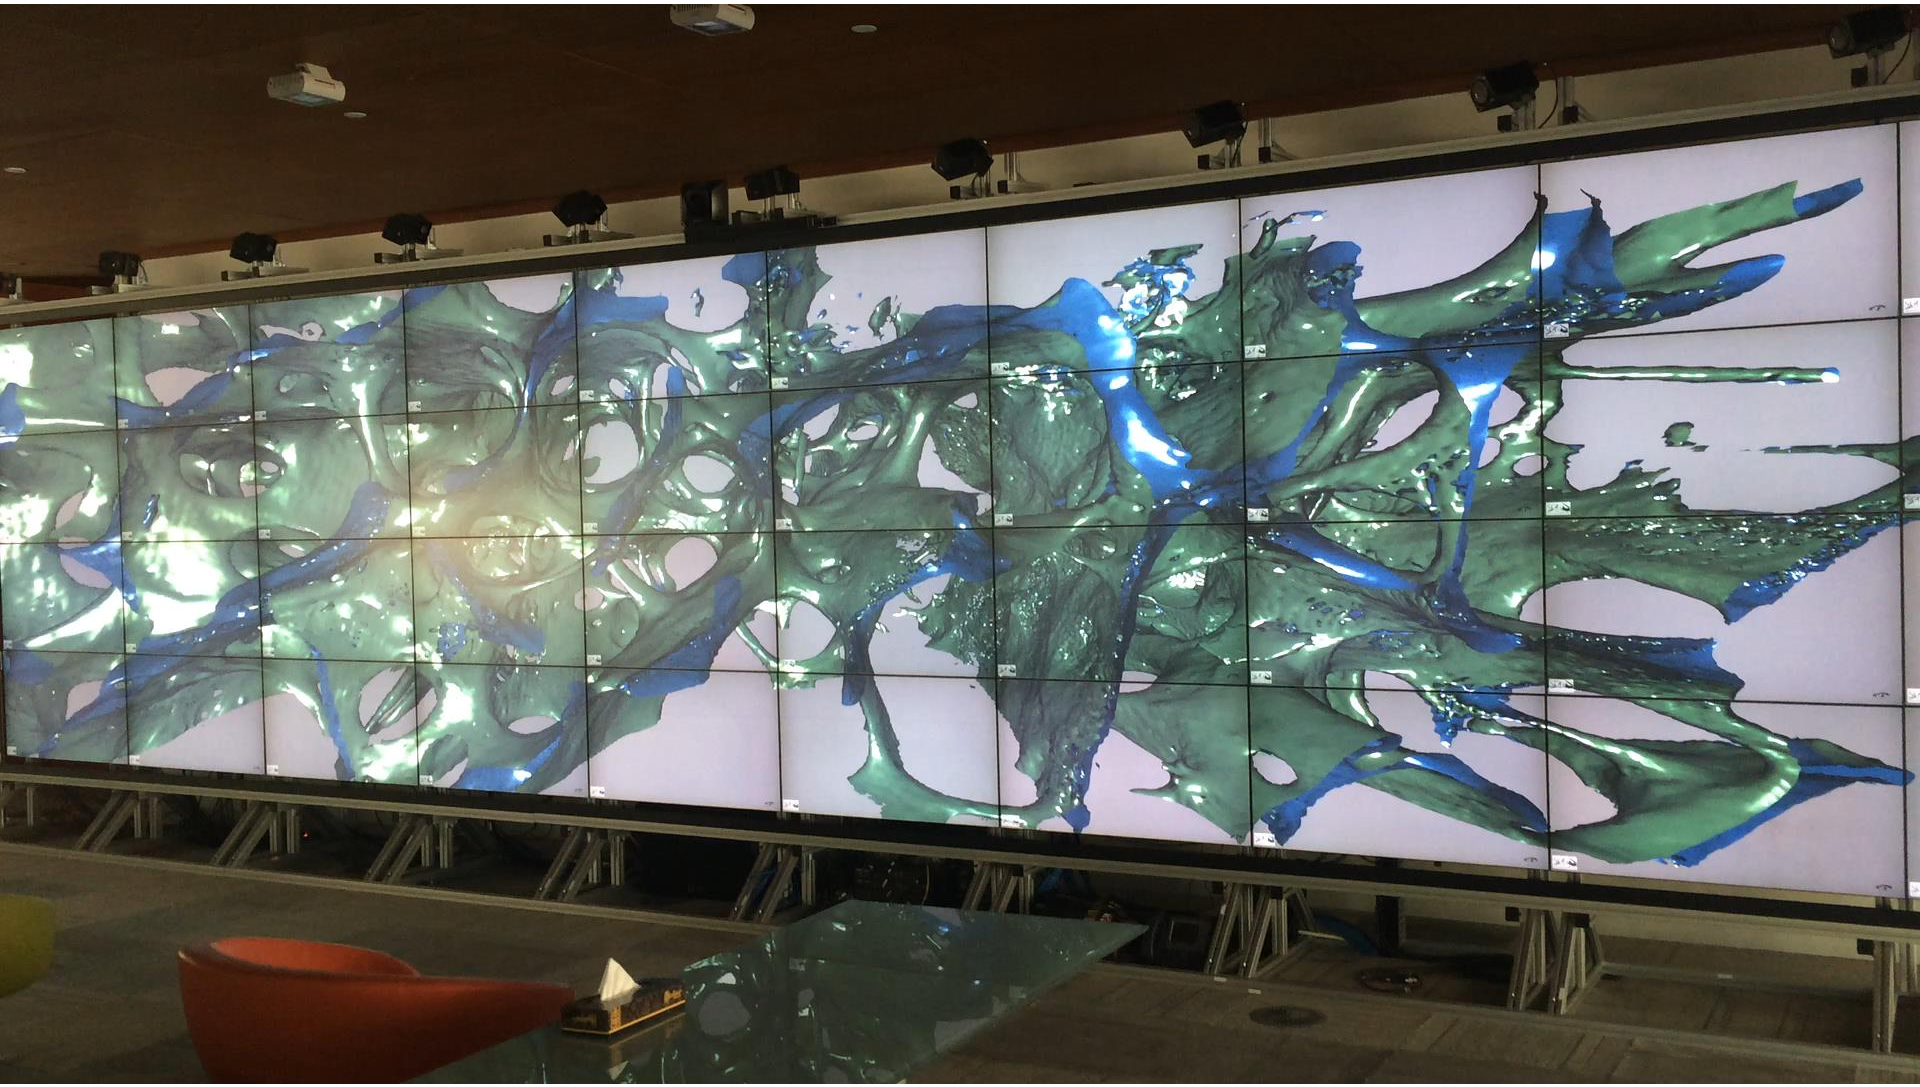

Main Components of Current PIDX Library for Fast, Scalable Simulation I/O

The different stages (some optional) of the PIDX data management pipeline allow for flexible deployment and optimization

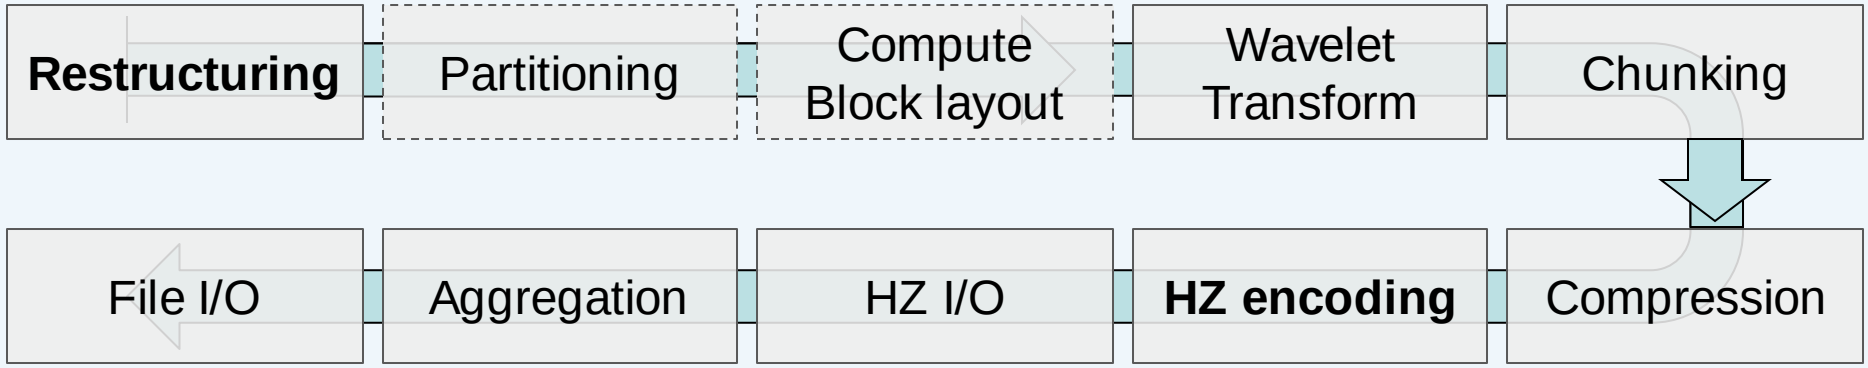

High Performance Data Movements for Real-Time Monitoring of Large Scale Simulations

Hopper (fat tree)

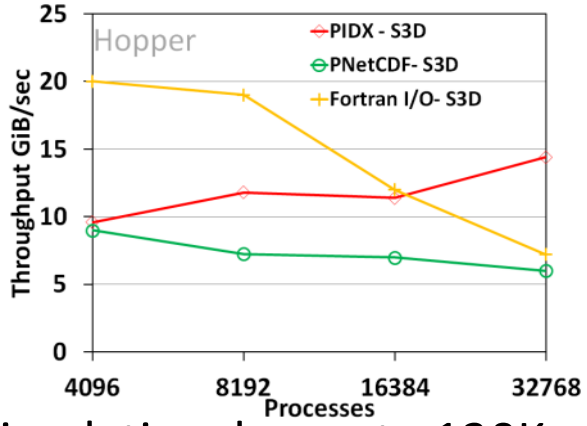

Intrepid (torus)

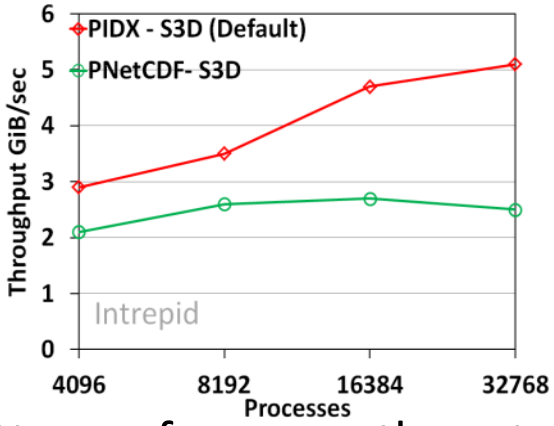

Scale simulation dumps to 130K cores with better performance than state of the art libraries while enabling real-time, remote visualization

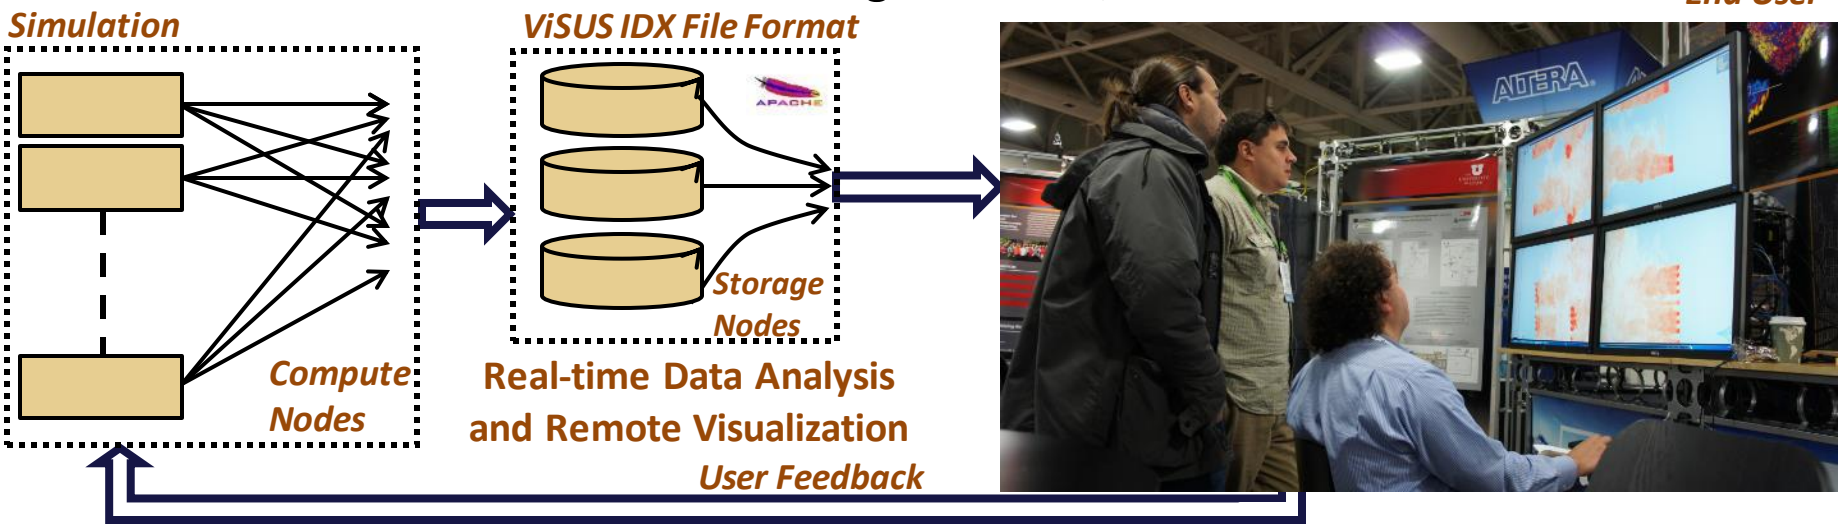

High Performance Data Movements for Real-Time Monitoring of Large Scale Simulations

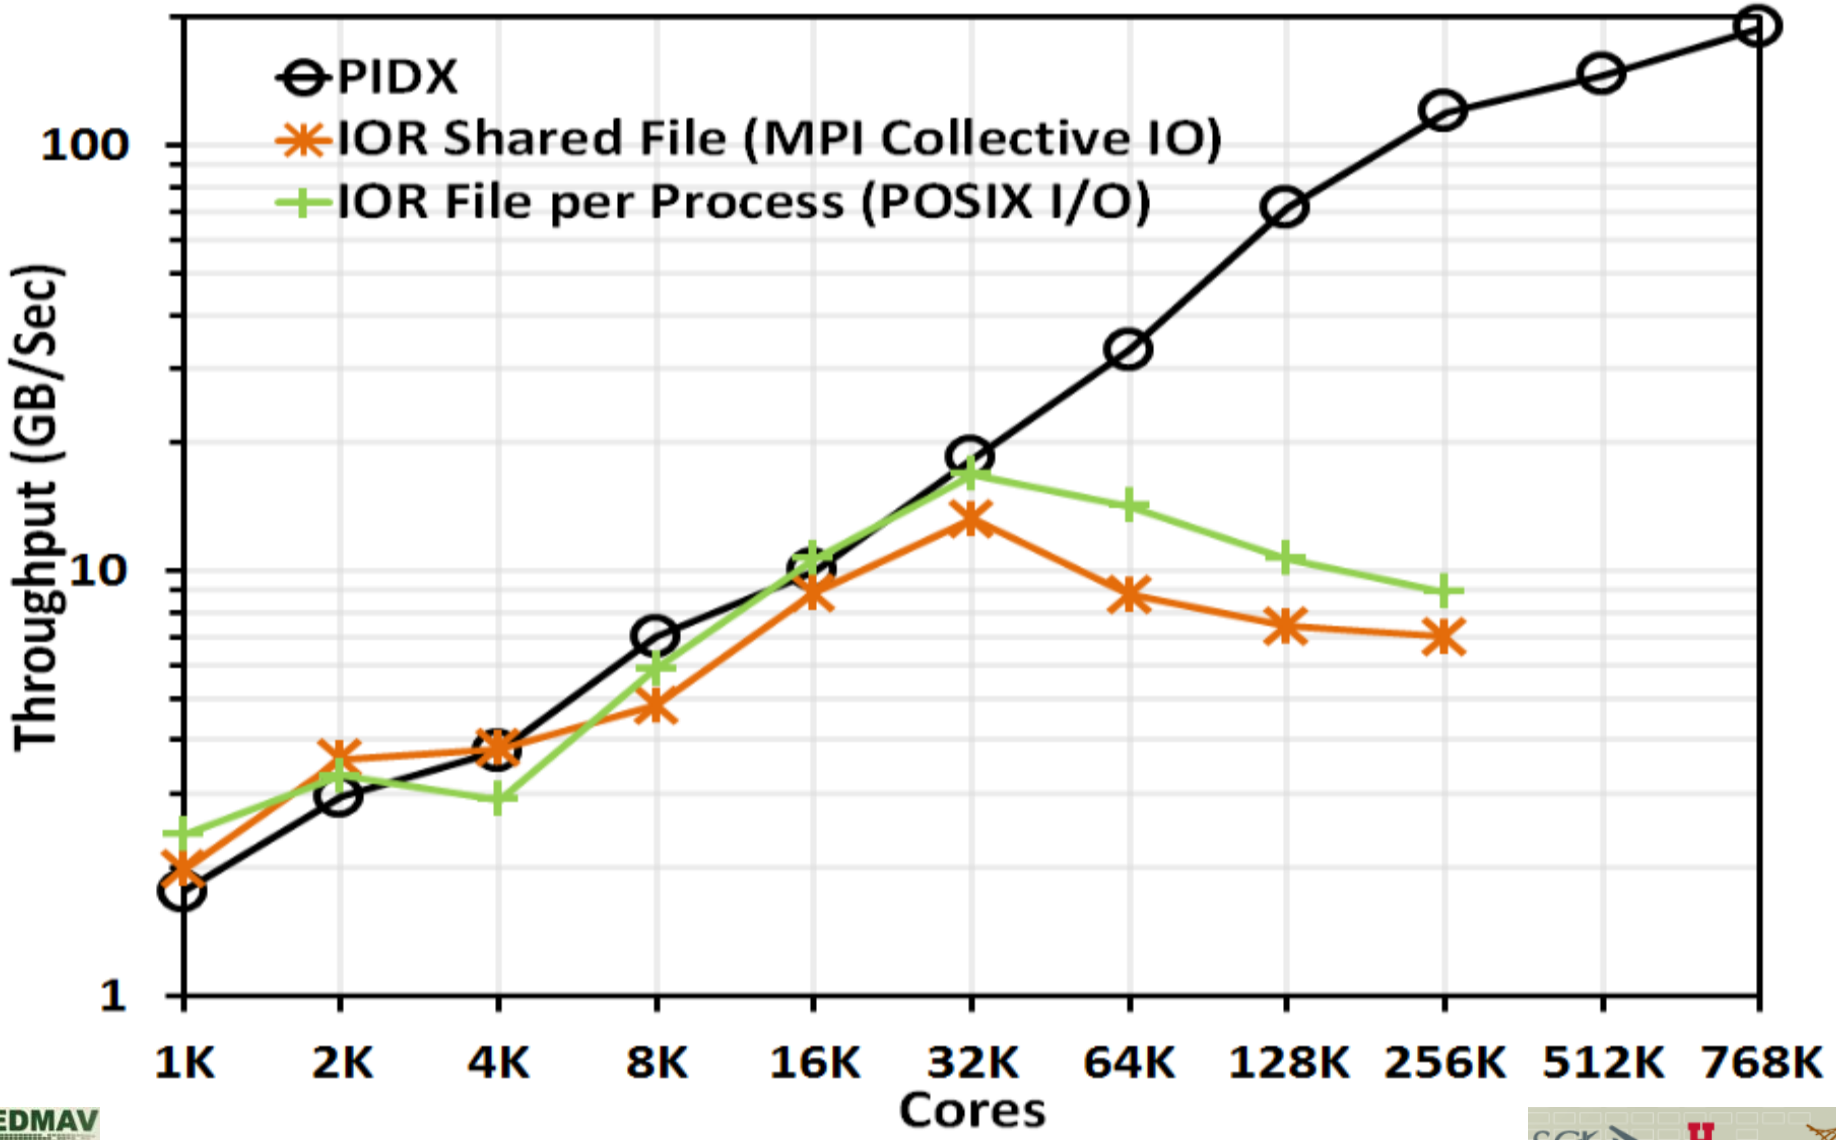

Analytics Domain Specific Language

- Analytic workflows often contain very well known communication patterns
- DSL-like representation based on dataflow abstraction
- Many runtime systems
  - Different API
  - Different data model
  - Different execution model
  - Different learning curves
  - Different performance
  - Very limited portability across them

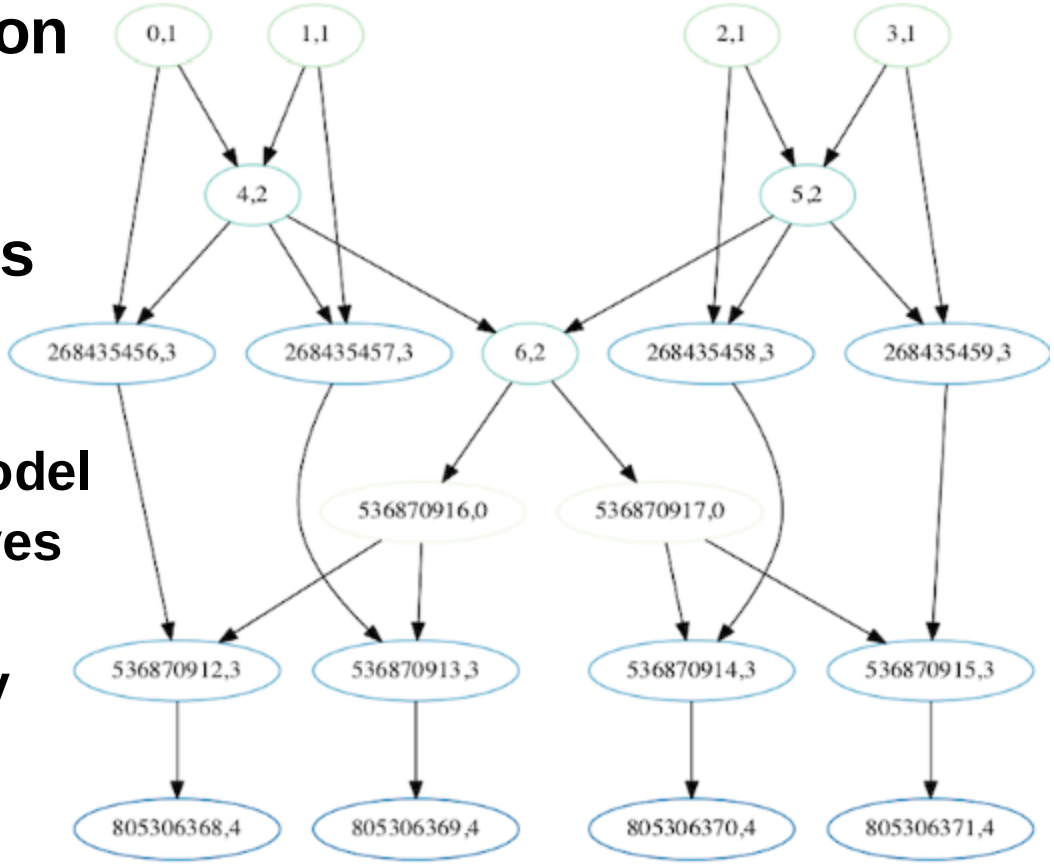

Scalable and Portable Data Analytics

- **Decouple the definition and implementation of the workflow**
- **Simple Dataflow graph that represents the communication patterns between tasks**

- Idempotent tasks
- Flexible mapping of tasks and dataflow nodes
- Ability to run the same exact workflow on different runtime
- Each controller is an actual runtime-based application (easy to integrate in a native runtime environment)
- Distributed graph representation allows scalability

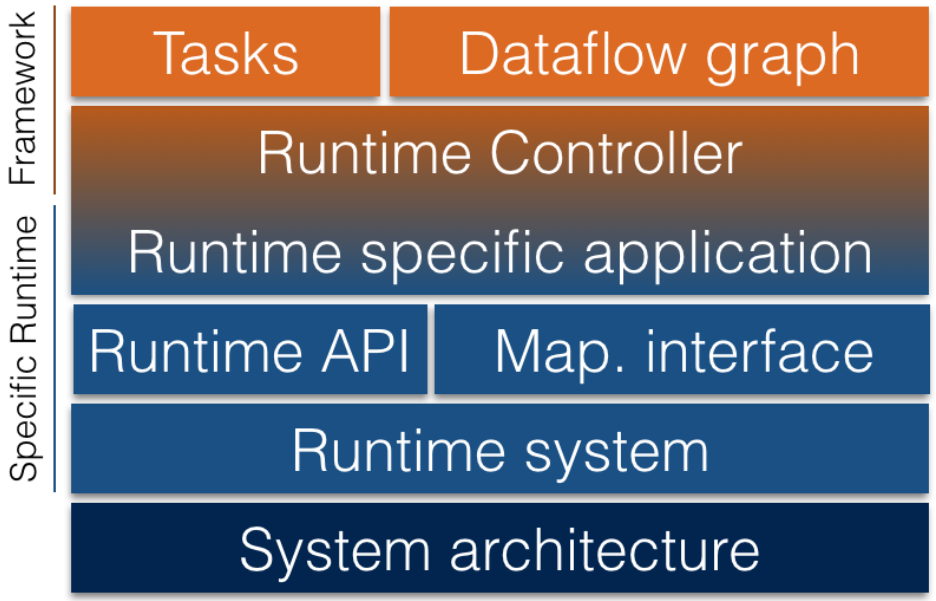

Flexible Targeting of Specific Runtimes

- Legion
  - Dynamic task allocation
  - Preprocessing of graph to extract rounds of tasks and ordering
  - Communication exposed through data dependencies
  - SPMD implementation
- MPI
  - Static task allocation
  - Each rank manages a subset of the tasks
  - Communication using global Ids
- Charm++
  - Dynamic task allocation using chare arrays and asynchronous remote calls
  - Independent tasks procedurally derive inputs and outputs and communicate using global Ids

Representative Use Cases

- **In-Situ Topological analysis: Parallel merge tree**
- **In-Situ Visualization: Image compositing**
- **Parallel Processing of Brain data: volume alignment**

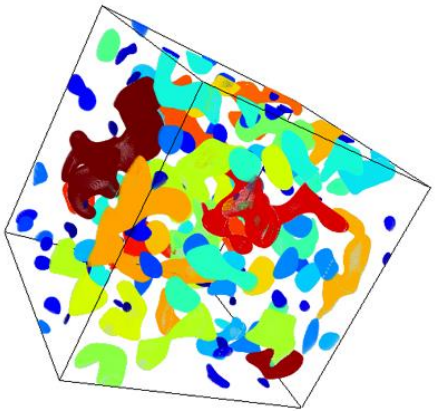

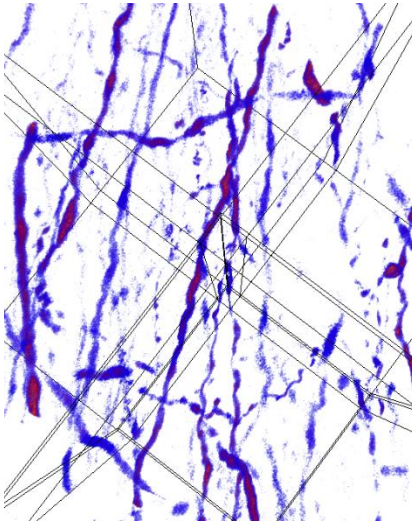

Use Case 1 – Image Compositing
Binary swap dataflow

- **Binary swap composite**
  - **High utilization**
  - **Results are tiles: Extra collection step**

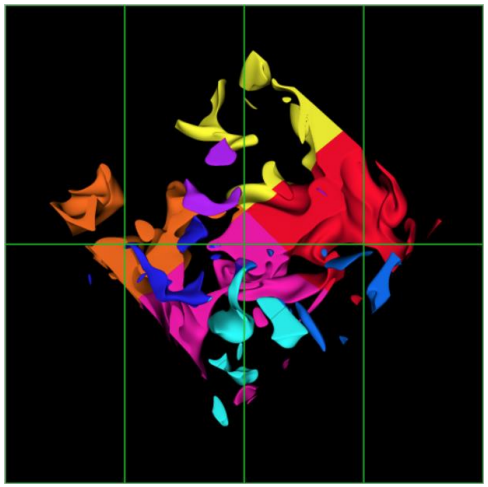

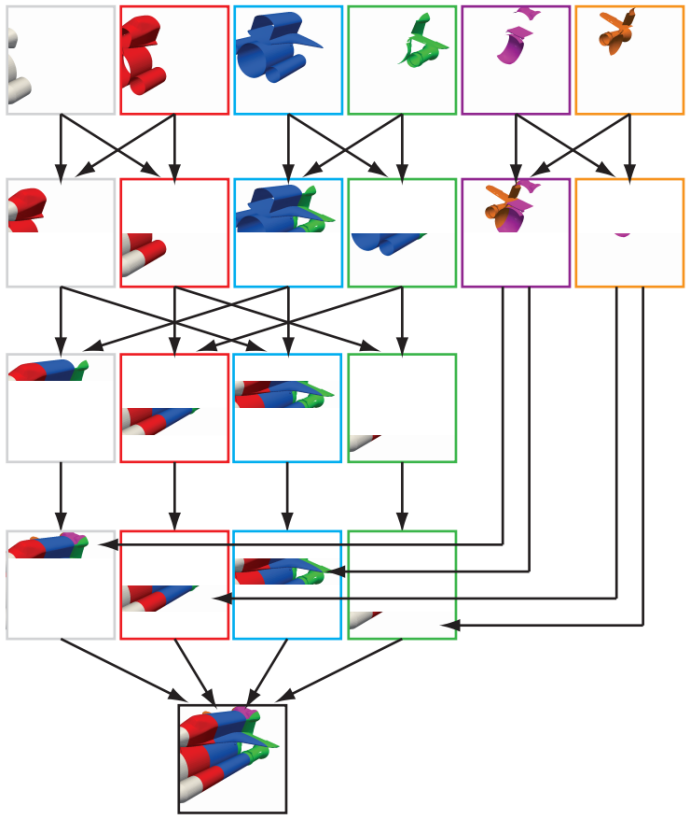

Use Case 2 – Features extraction using topological analysis

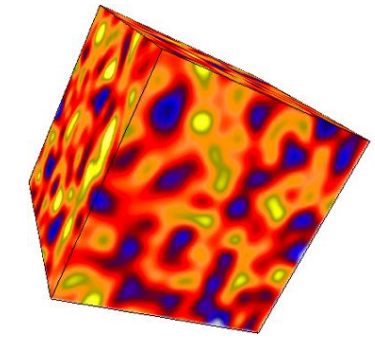

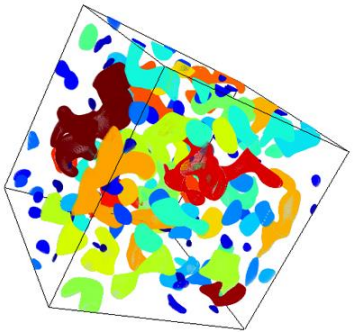

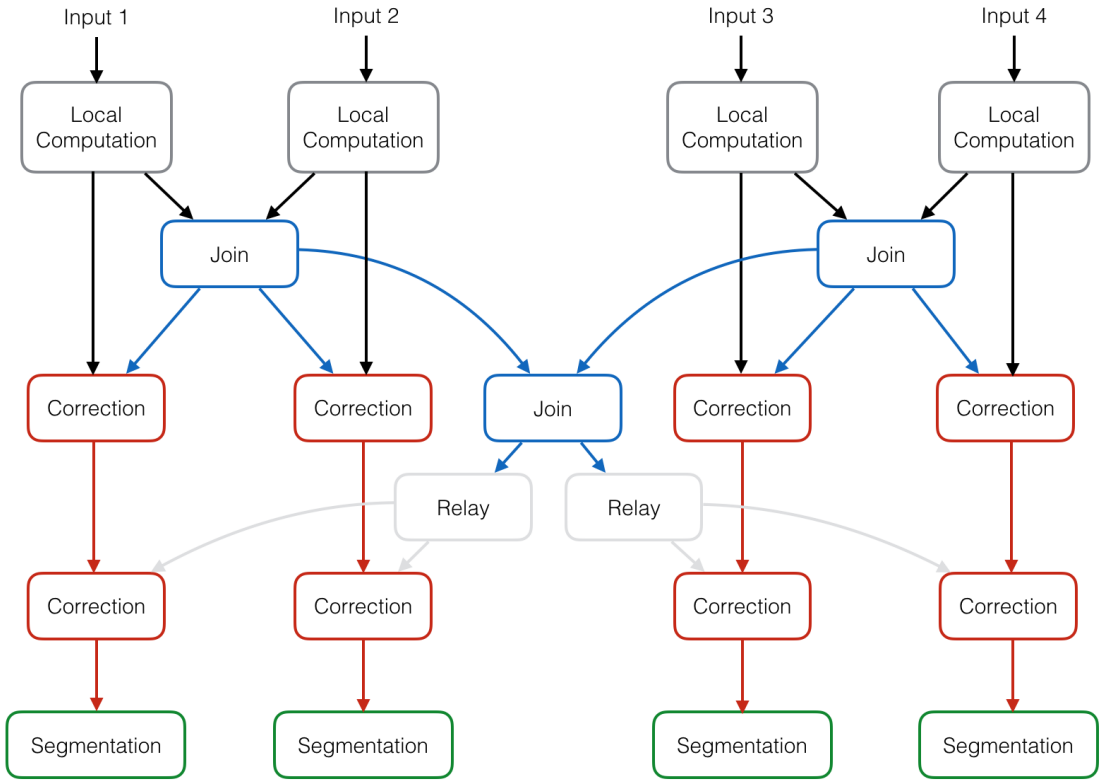

Use Case 3 – Brain Data Analysis Workflow

- Neighbors communication pattern
- No correlation repetition
- Results collection over volume
- Correlation graph
- Global alignment using minimum spanning tree

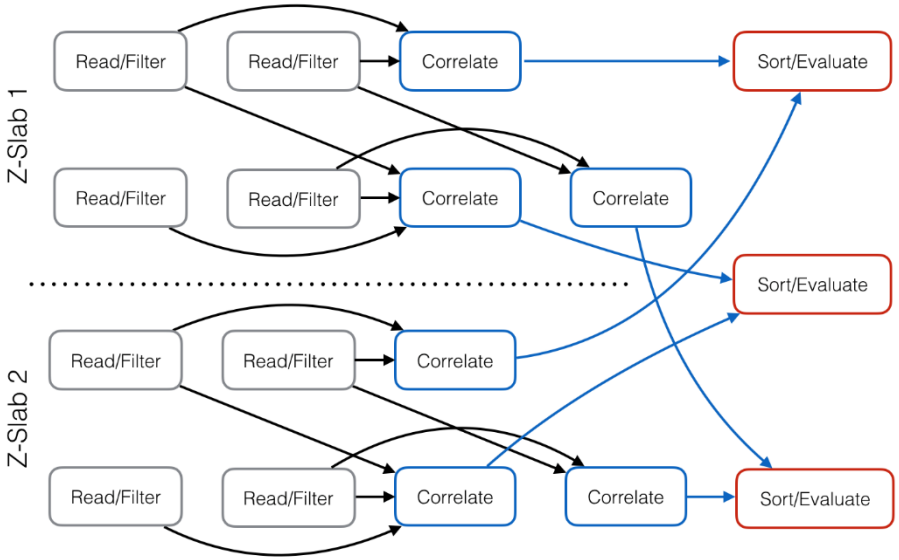

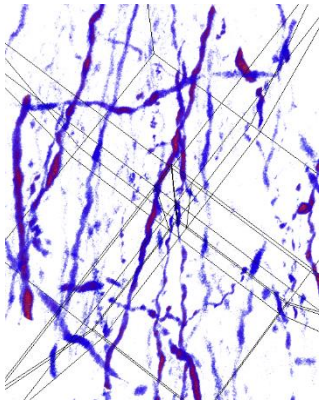

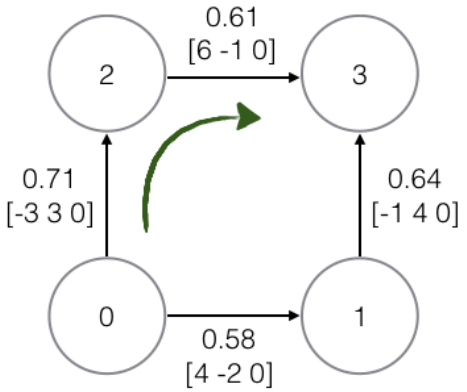

Scaling In-Situ Analytics to Full Titan with Computational Overhead < 1%

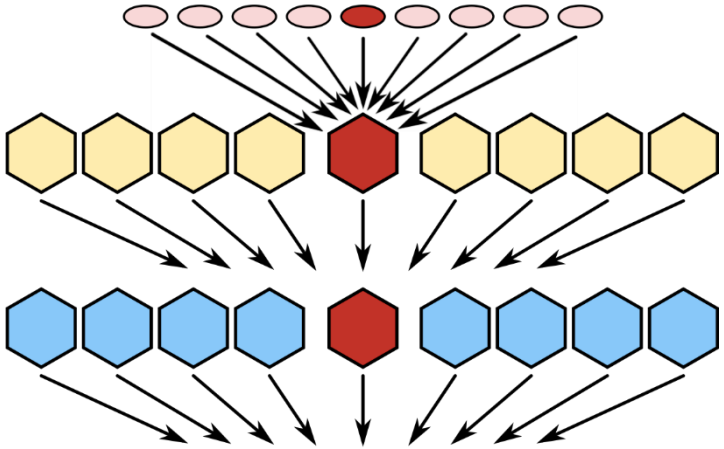

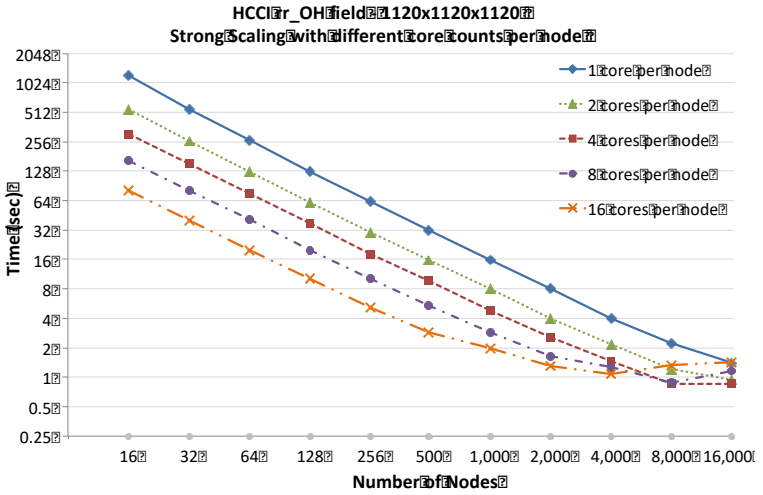

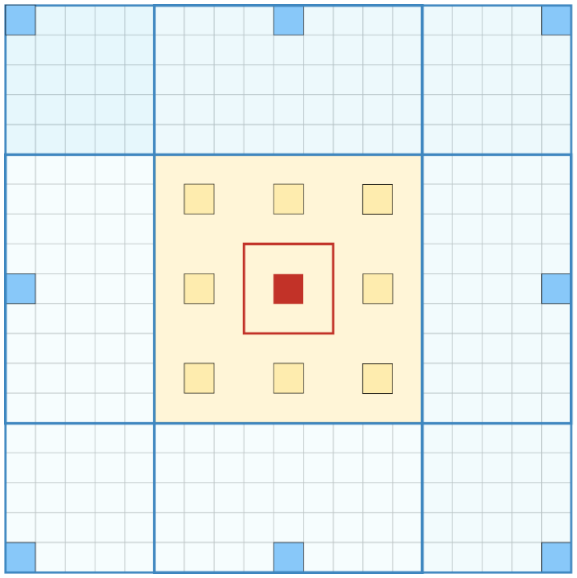

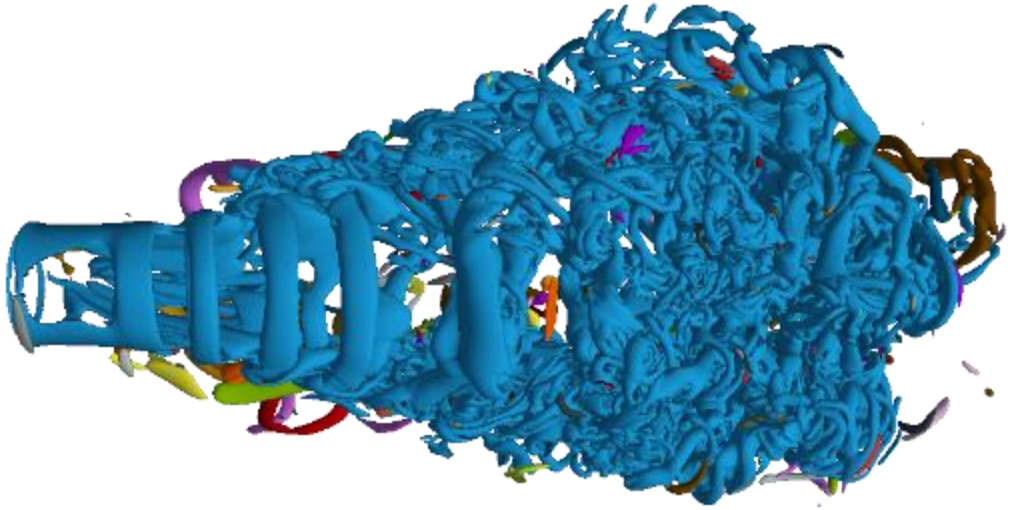

Parallel registration

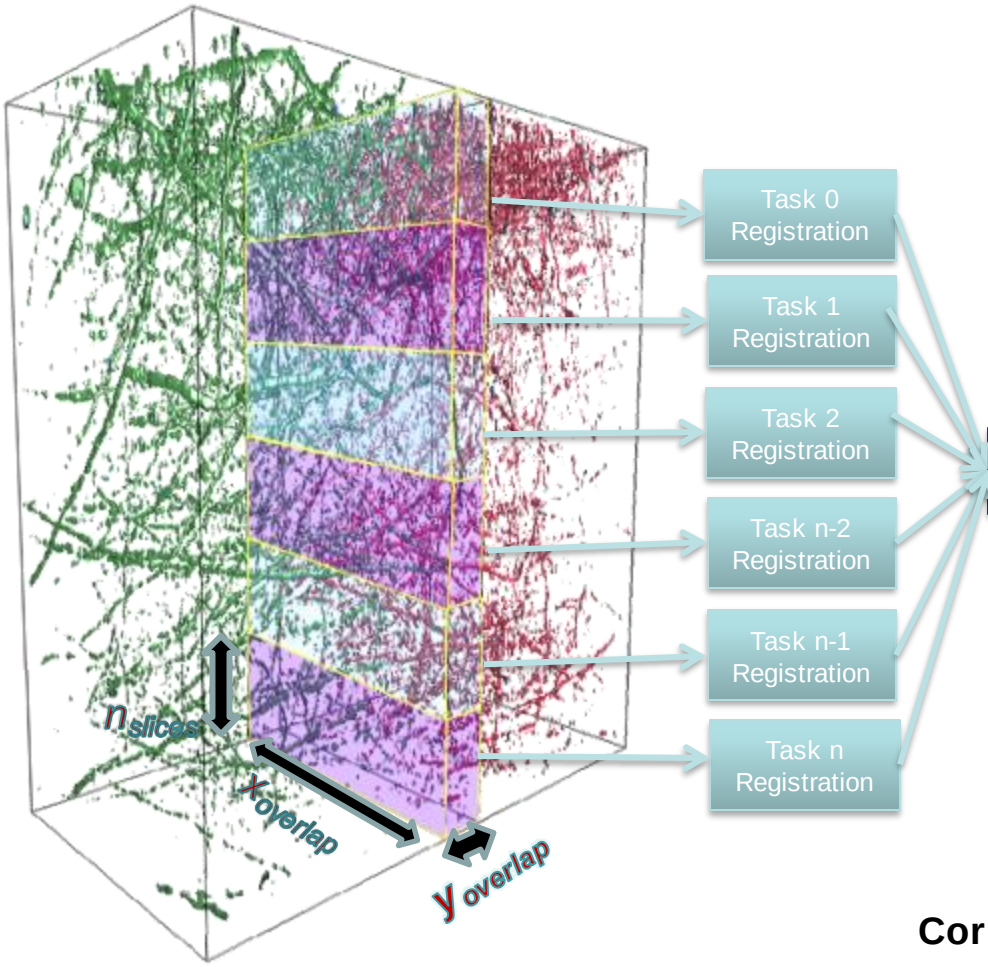

Step 1: *Decompose the overlapping regions into sub-blocks of size* $x_{overlap} \times y_{overlap} \times n_{slices}$.

Step 2: *Perform 3D registration using normalized crosscrrelation in the frequency domain.*

Step3: *Aggregate correlation results and find global alignment.*

Correlation graph

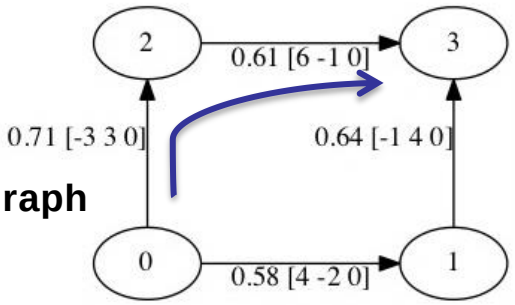

Performance scaling

Strong scaling on Shaheen II, a supercomputer in KAUST (King Abdallah University of Science and Technology), for varying number of cores (256 to 2048) using two 2K cube volumes.

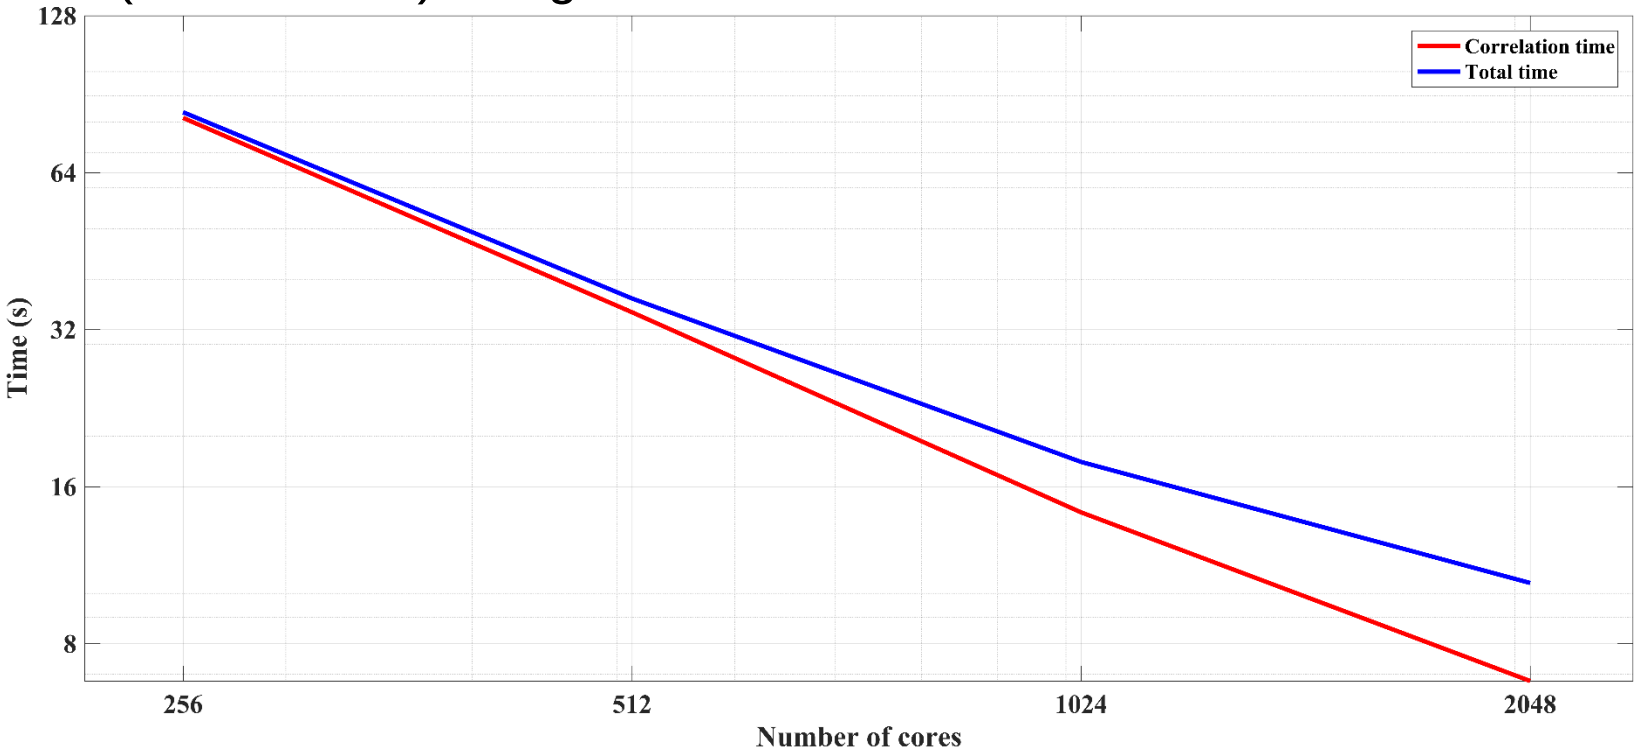

Correlation time: time spent to compute the optimal tile offsets. Total time: includes time required for data read, transfer, correlation result collection and final graph production.

We Develop an Integrated Data Acquisition, Management and Computation for Neuroscience

(1) Data Source

(2) Preliminary Interactive Analytics

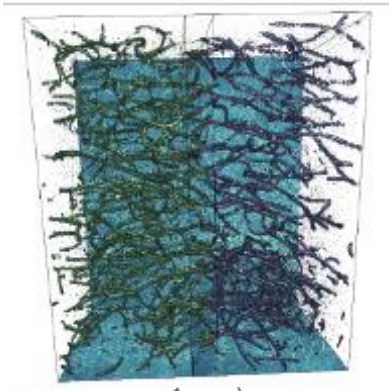

(4) Interactive, Exploratory Assessment and Feedback

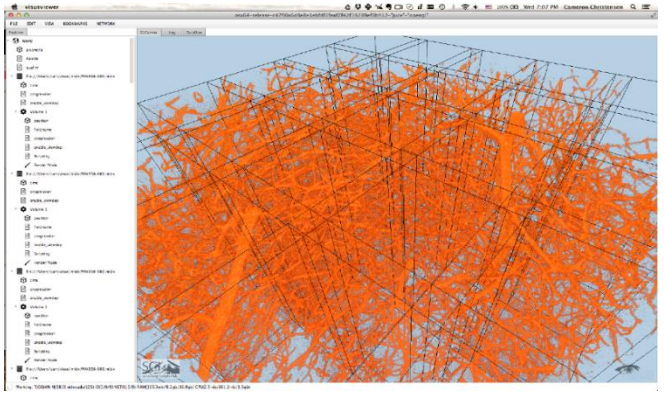

(3) Asynchronous Parallel Processing

Intuitive Understanding of Relationships in High-Dimensional Multivariate Data

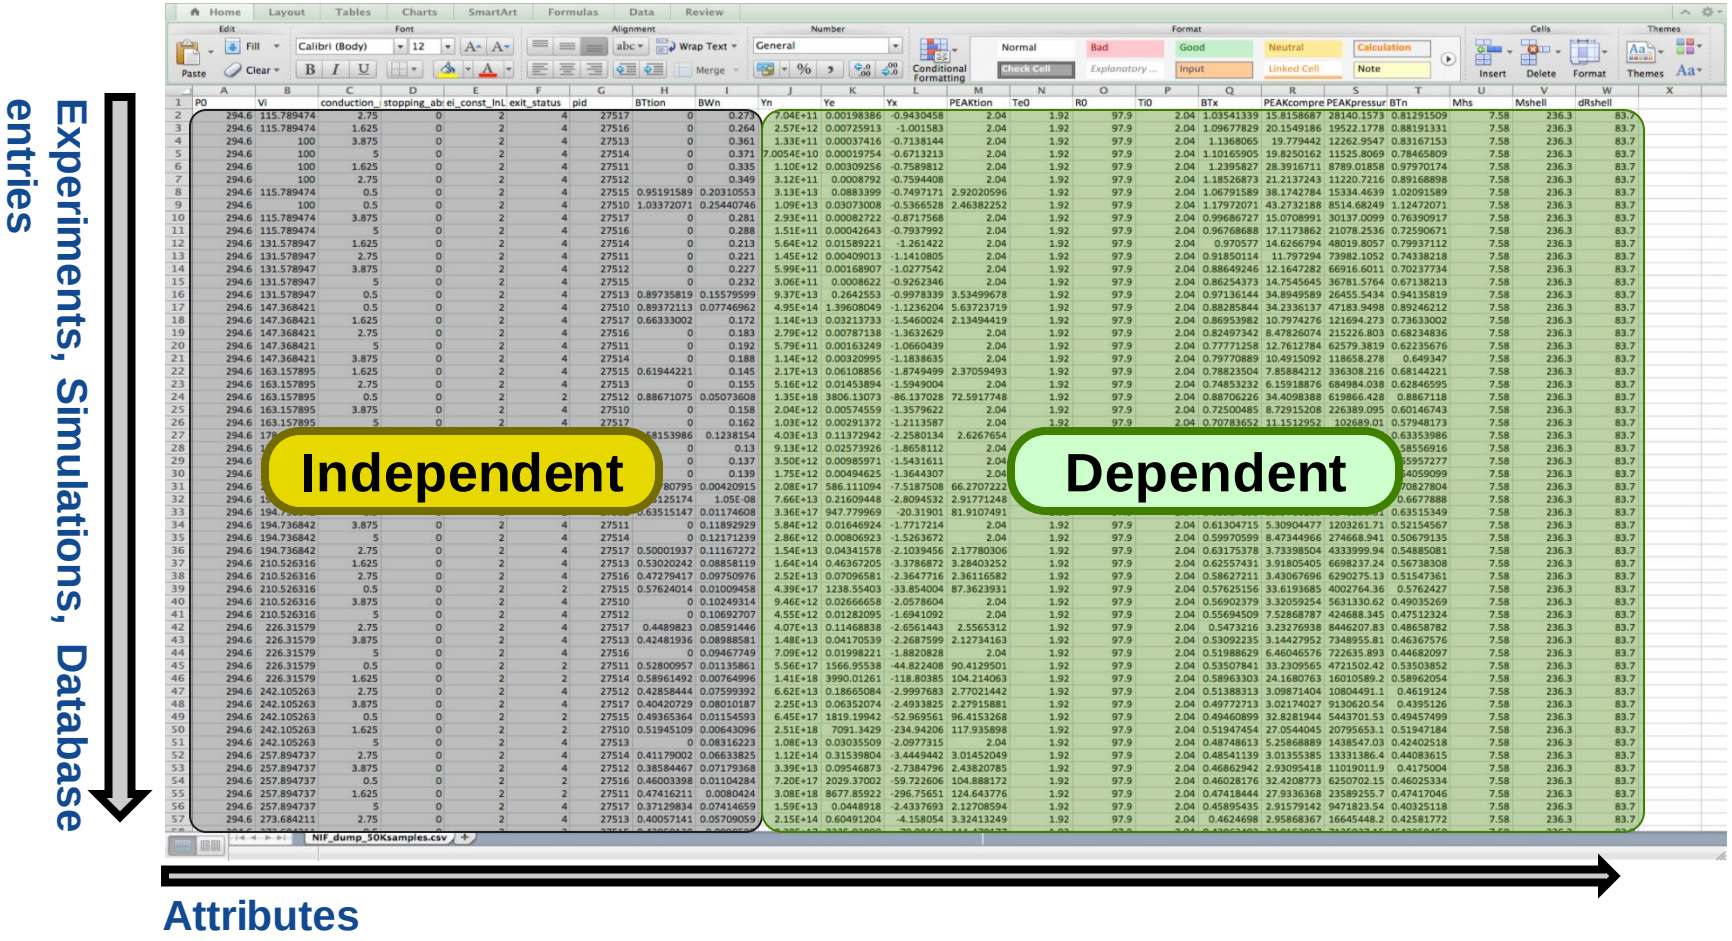

Repeated analysis by selecting a "Domain of Interest" X and a "Figure of Merit" Y

$$F : X \rightarrow Y$$

Given and Domain X and Figure of Merit Y, X is Partitioned in a Topological Complex

Repeated analysis by selecting a "Domain of Interest" X and a "Figure of Merit" Y

$$F : X \rightarrow Y$$

Simple, intuitive data Abstraction valid in any dimension based on experience in navigating terrains

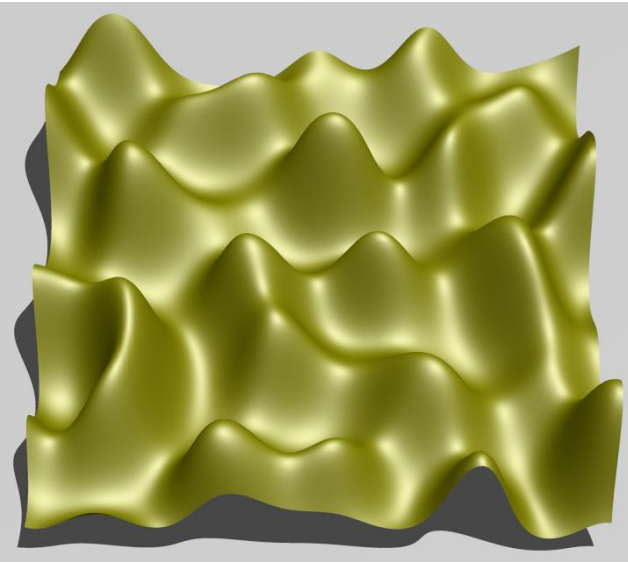

$F : X \rightarrow Y$

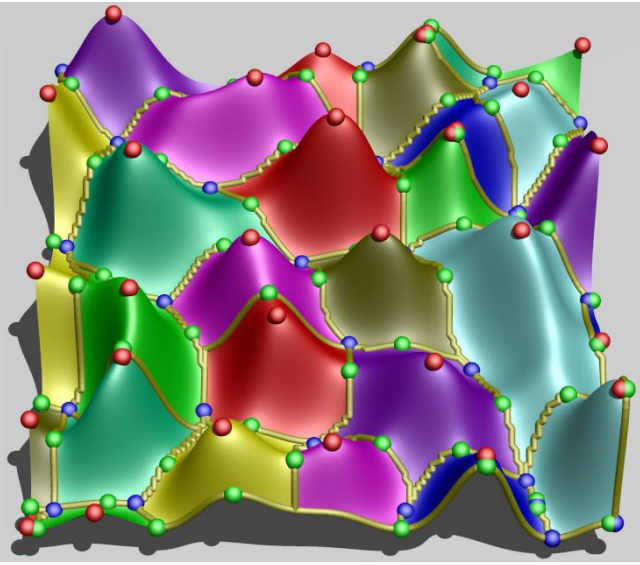

Morse Complex

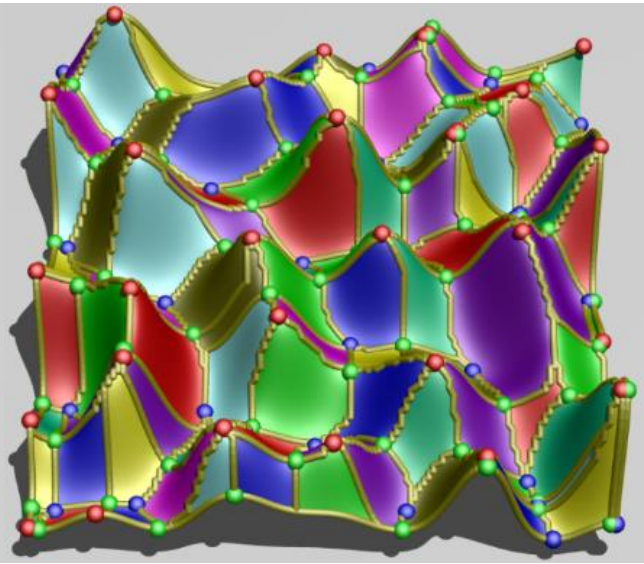

Morse-Smale Complex

Simplification Allows to Create a Multi-Scale Representations of the Data

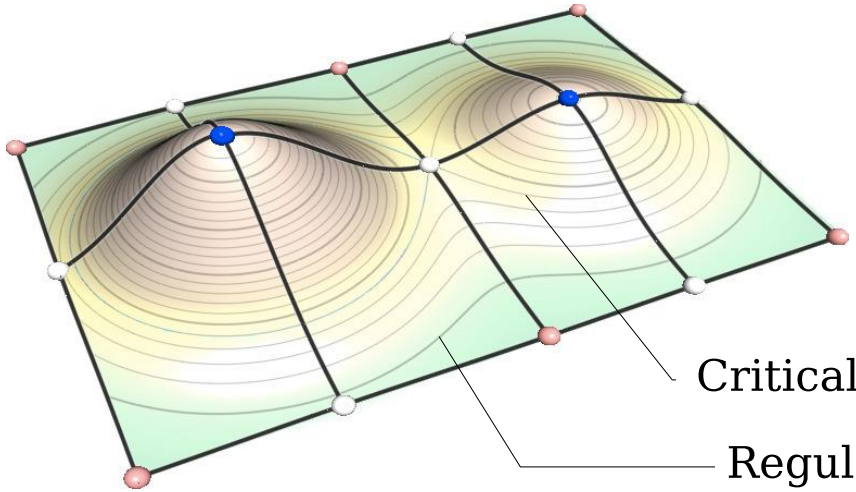

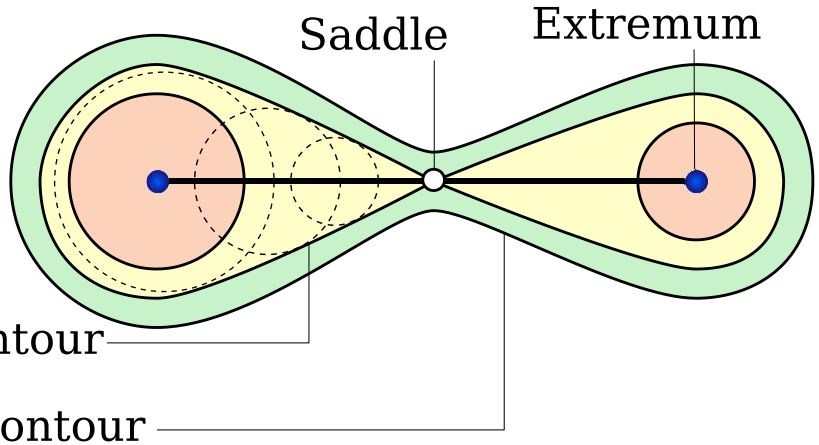

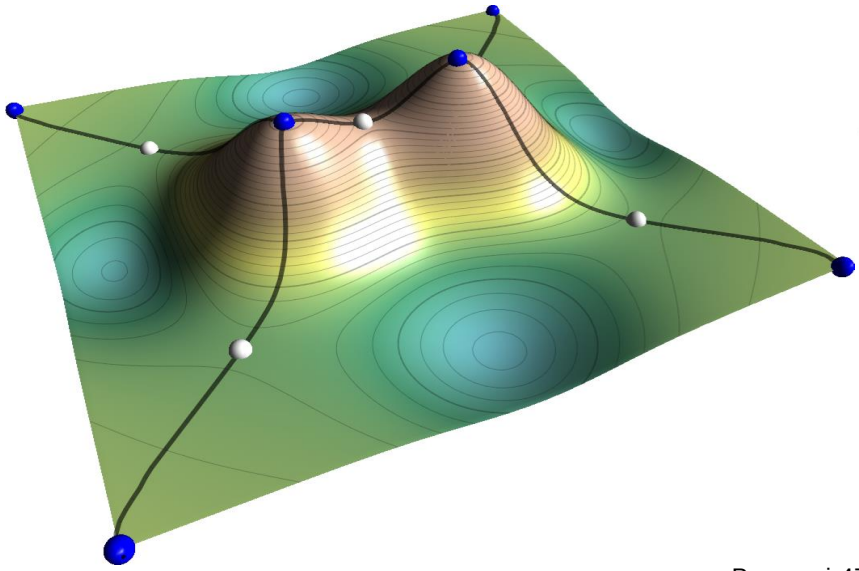

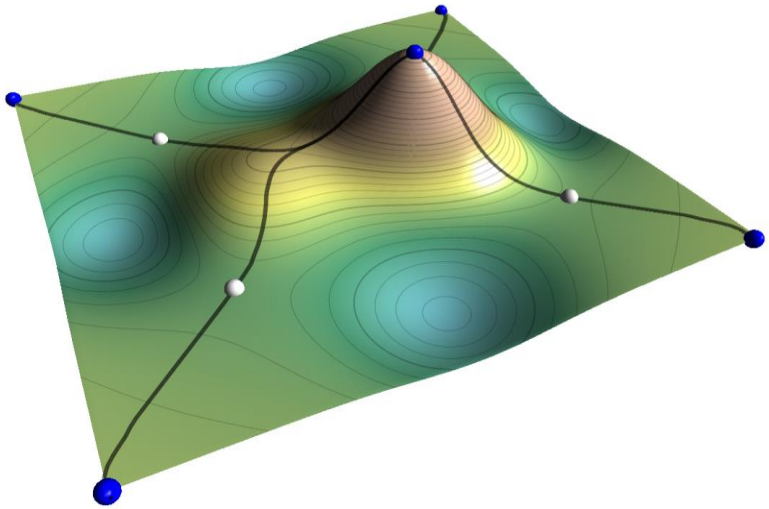

Cancellation Trees Identify Prominent Ridges as Part of the Simplification Process

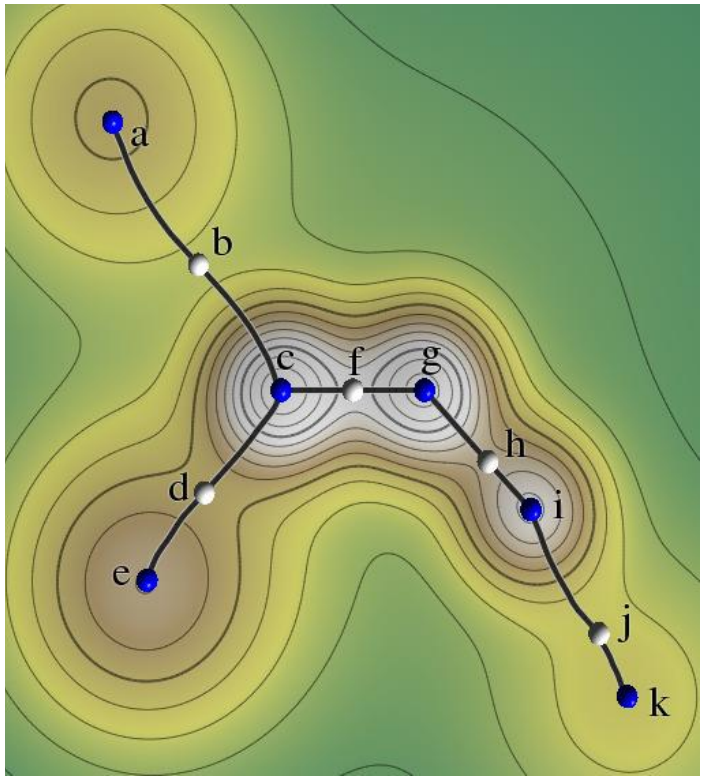

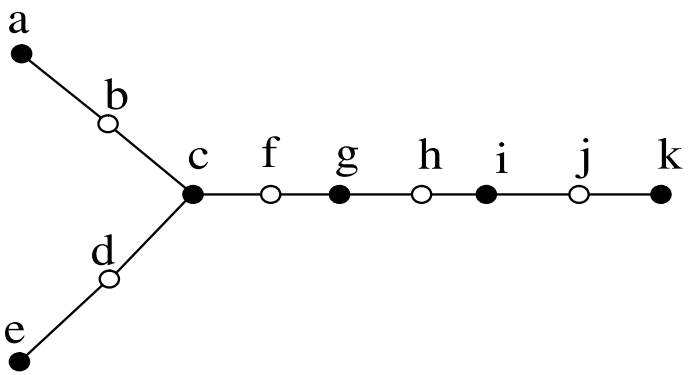

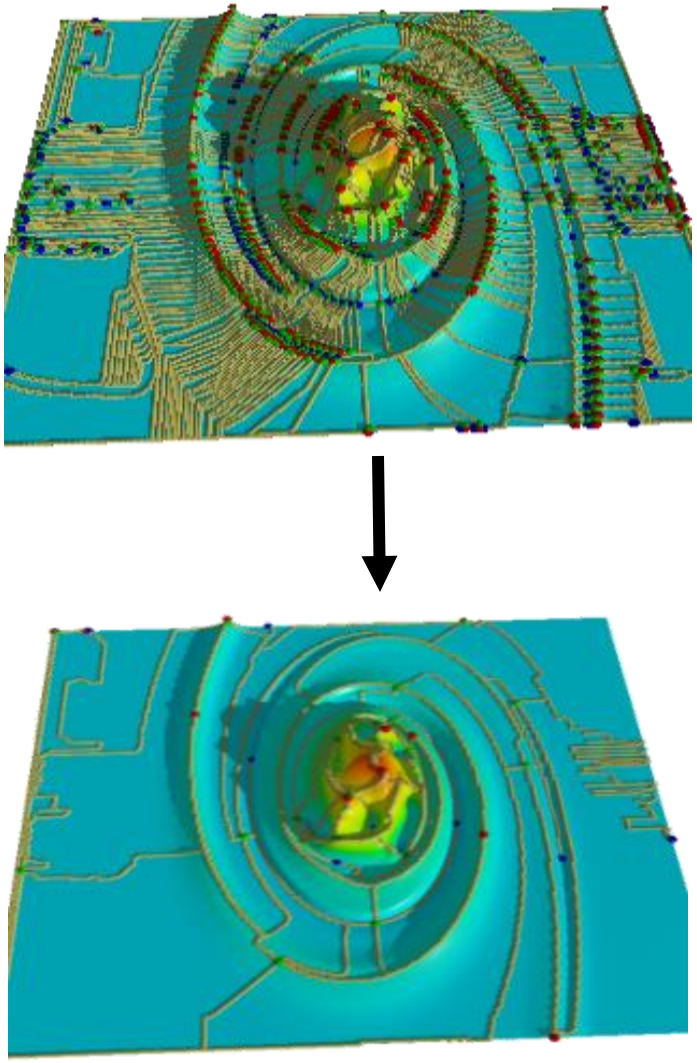

Exploration of High Dimensional Functions Data

Integrated presentation of statistics and topology

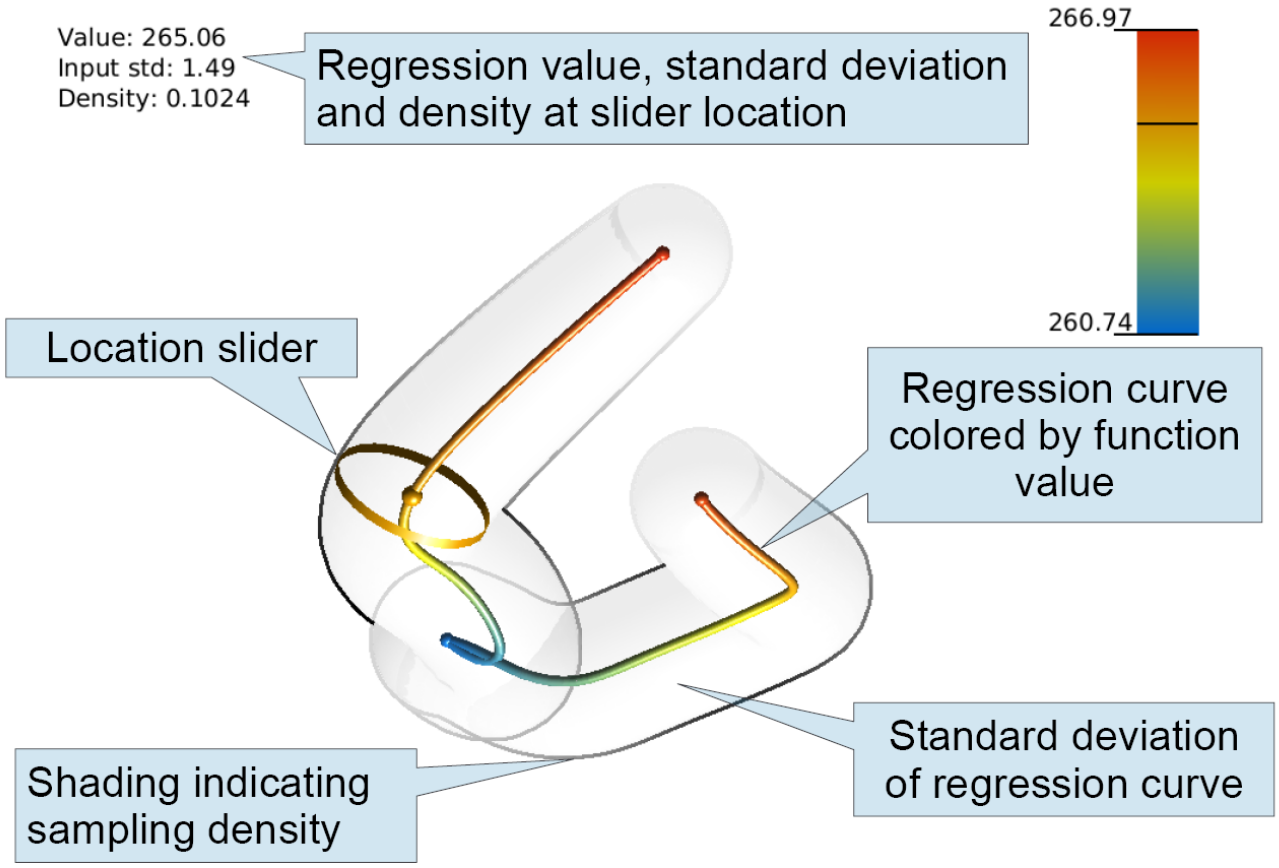

Improve Techniques for Visualizing and Exploring High Dimensional Data

Reduce dimensionality
and then extract structure

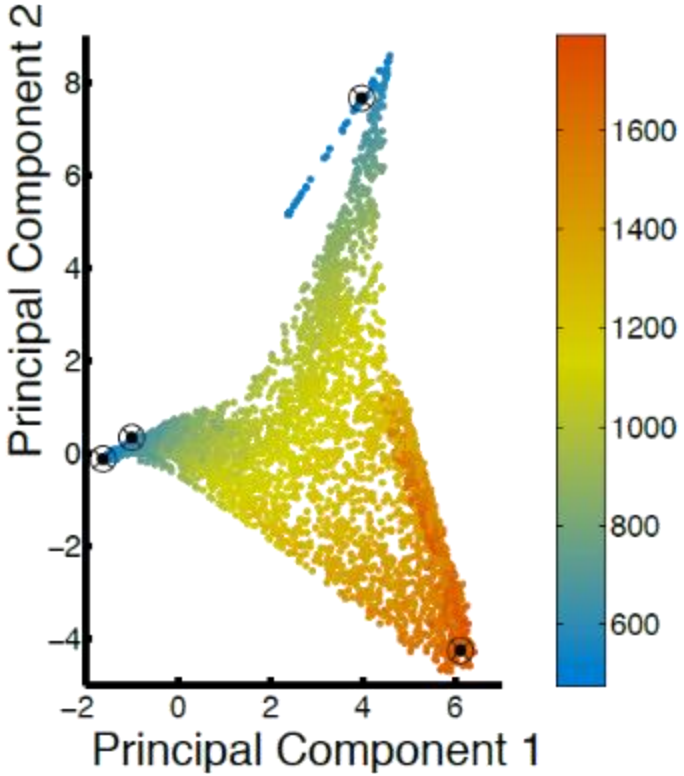

Gerber et al, 2010

Improve Techniques for Visualizing and Exploring High Dimensional Data

Reduce dimensionality
and then extract structure

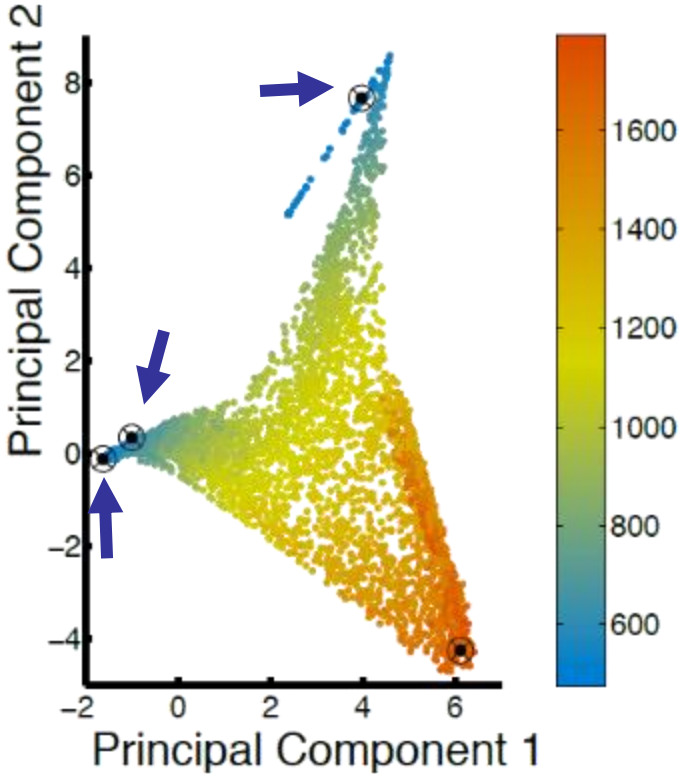

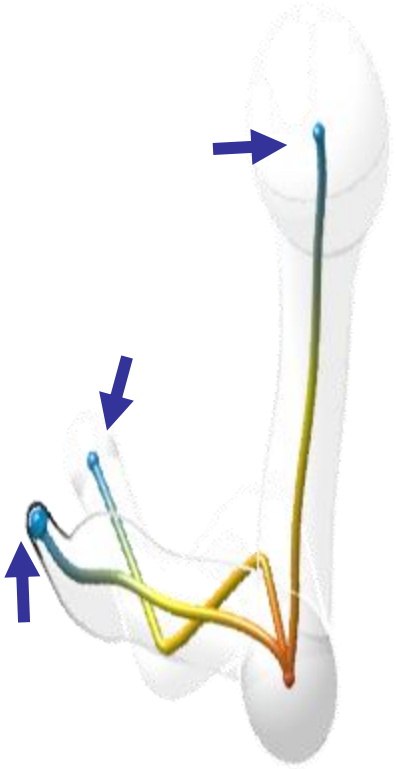

Gerber et al, 2010

Analysis of Combustion Simulations

Combustion Simulation of Jet CO/H2-Air Flames

Input: Composition of 10 chemical species

Output: Temperature

The Framework Allows Detailed Visualization and Analysis of High Dimensional Functions

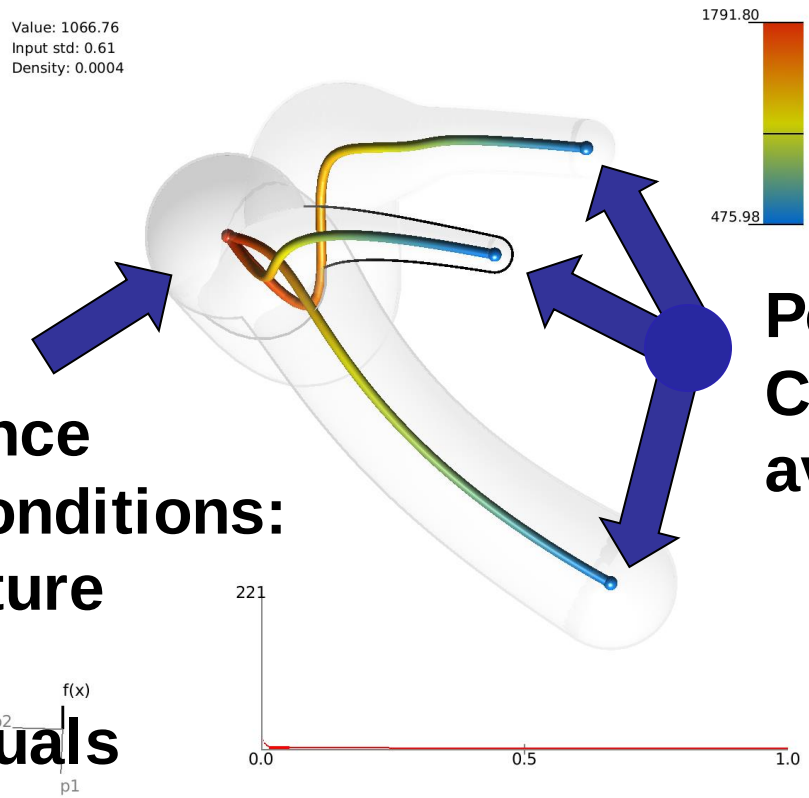

Poor Combustion Conditions to be avoided

High Performance Combustion Conditions:

- **high temperature**
- **low emission**
- **low fuel residuals**

10 dimensional data set describing the heat release wrt. to various chemical species in a combustion simulation.

Combustion Simulation of Jet CO/H2-Air Flames

Input: Composition of 10 chemical species

Output: Temperature

The SCRAM Event: Analysis of Nuclear Reactor Safety

- 10,000 individual simulations of a SCRAM event.
- Input = 6 dimensions:
  - PumpTripPre - pressure in heat exchange pump causing SCRAM
  - PumpStopTime - the relaxation time (sec) of pump's phase-out
  - PumpPow - end power of the pump
  - SCRAMtemp - the maximum temperature (K) in the system
  - CRinject - control rod position at the end of SCRAM
  - CRtime - the relaxation time of the control rod system.
- Output = Peak Coolant Temperature (PCT)

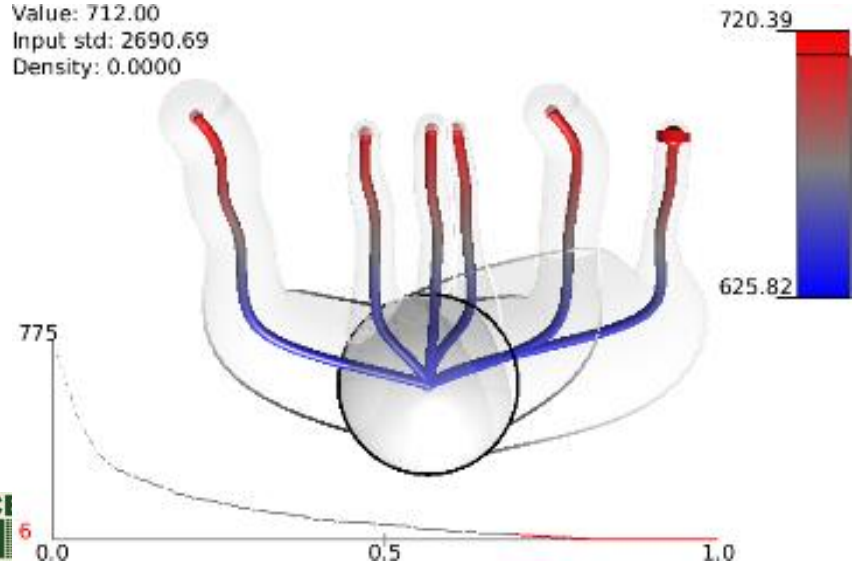

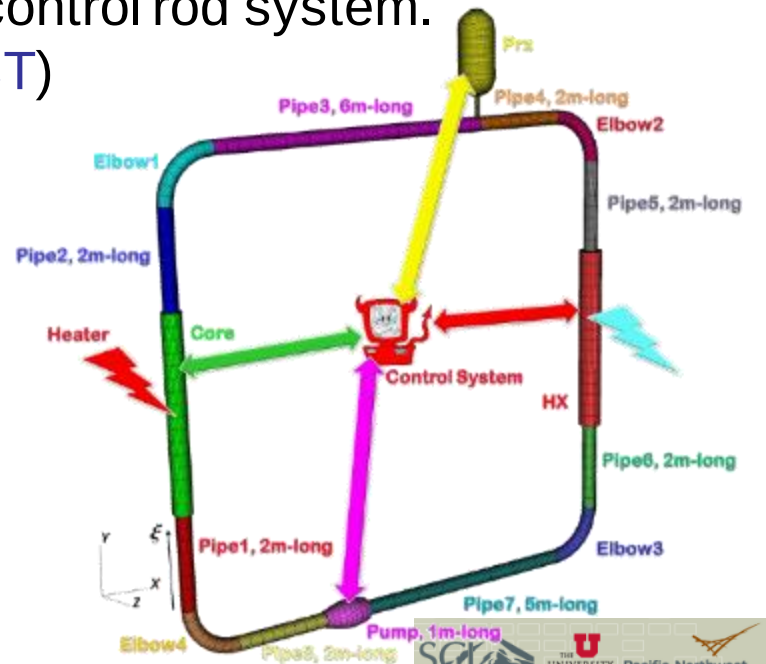

The SCRAM Event: Analysis of Nuclear Reactor Safety

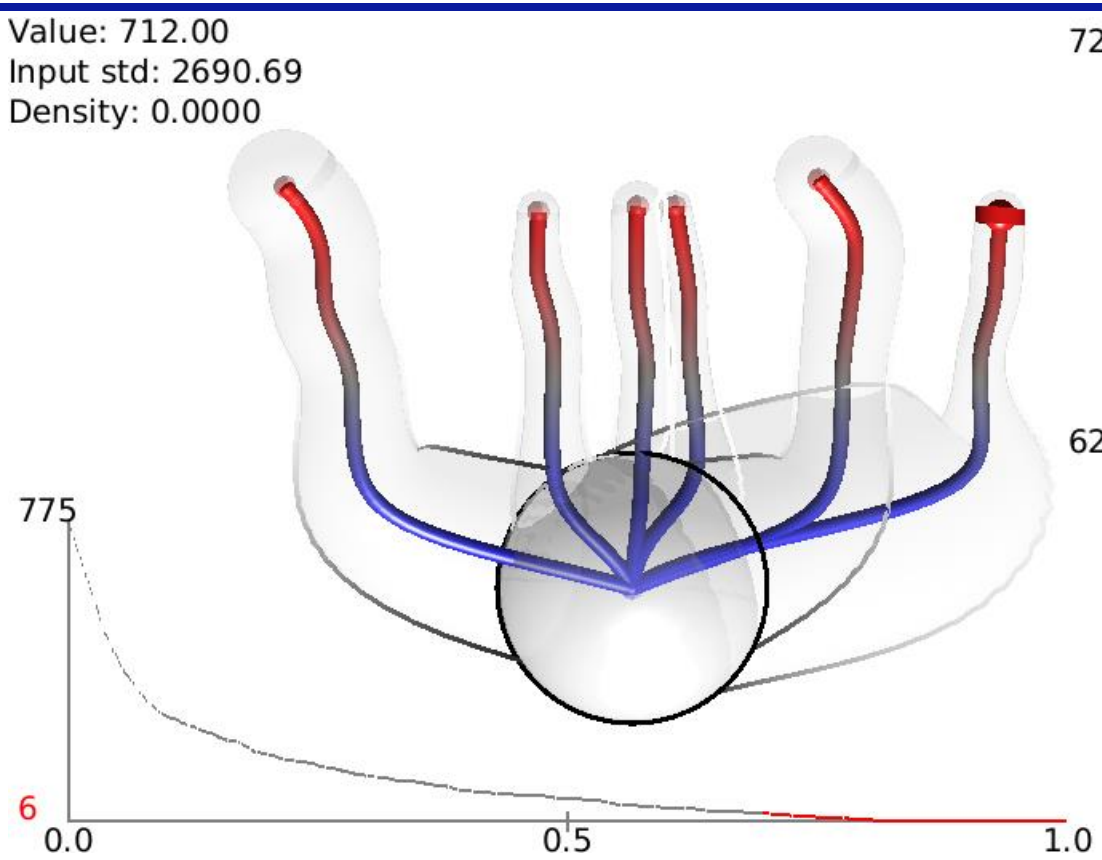

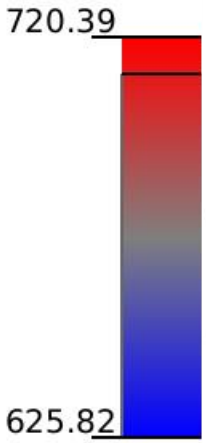

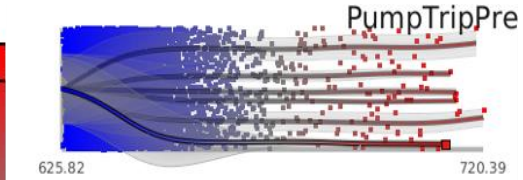

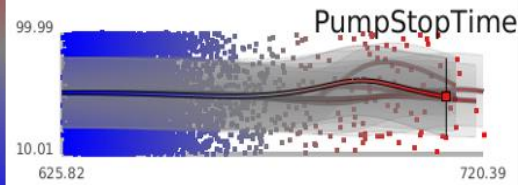

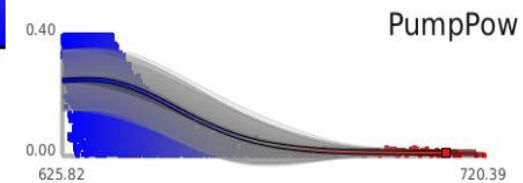

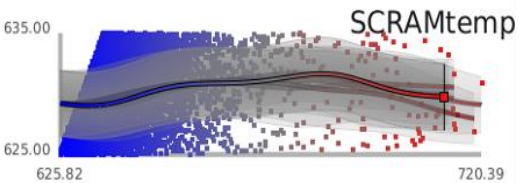

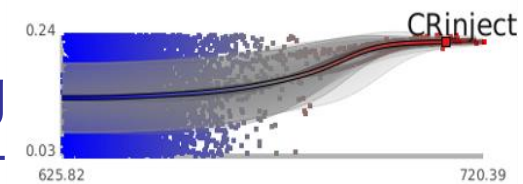

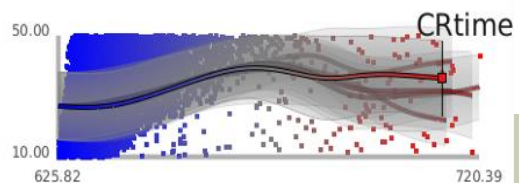

Pump trip pressure is the key distinguishing factor differentiating the high Peak Coolant Temperature (PCT) model

Topological Analysis of the Space of Composite Materials of a Given Class

- Features in experimental data show unexpected structures and are used to plan future experiments.

Stakeholder: A. Karim, PNNL.

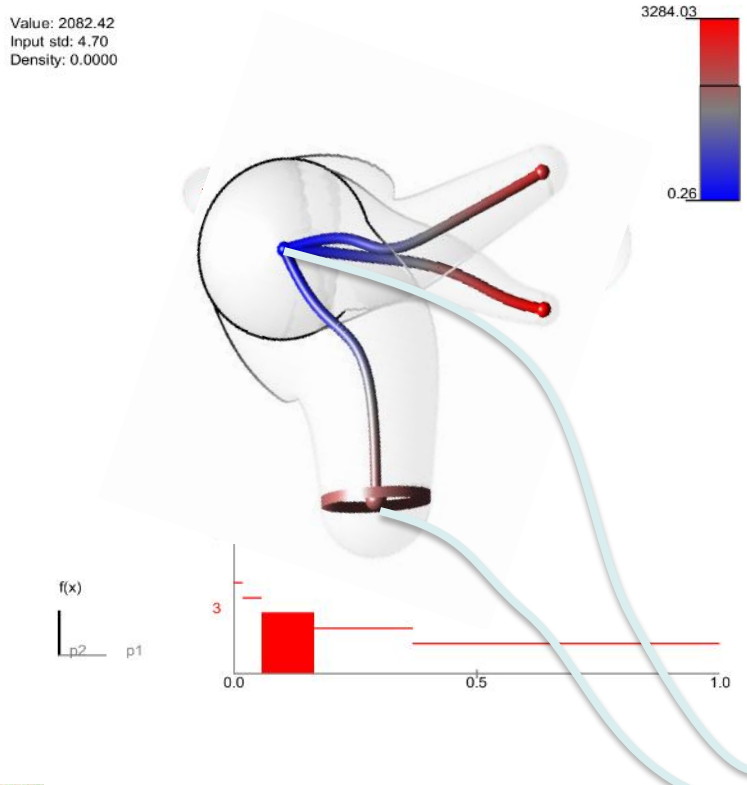

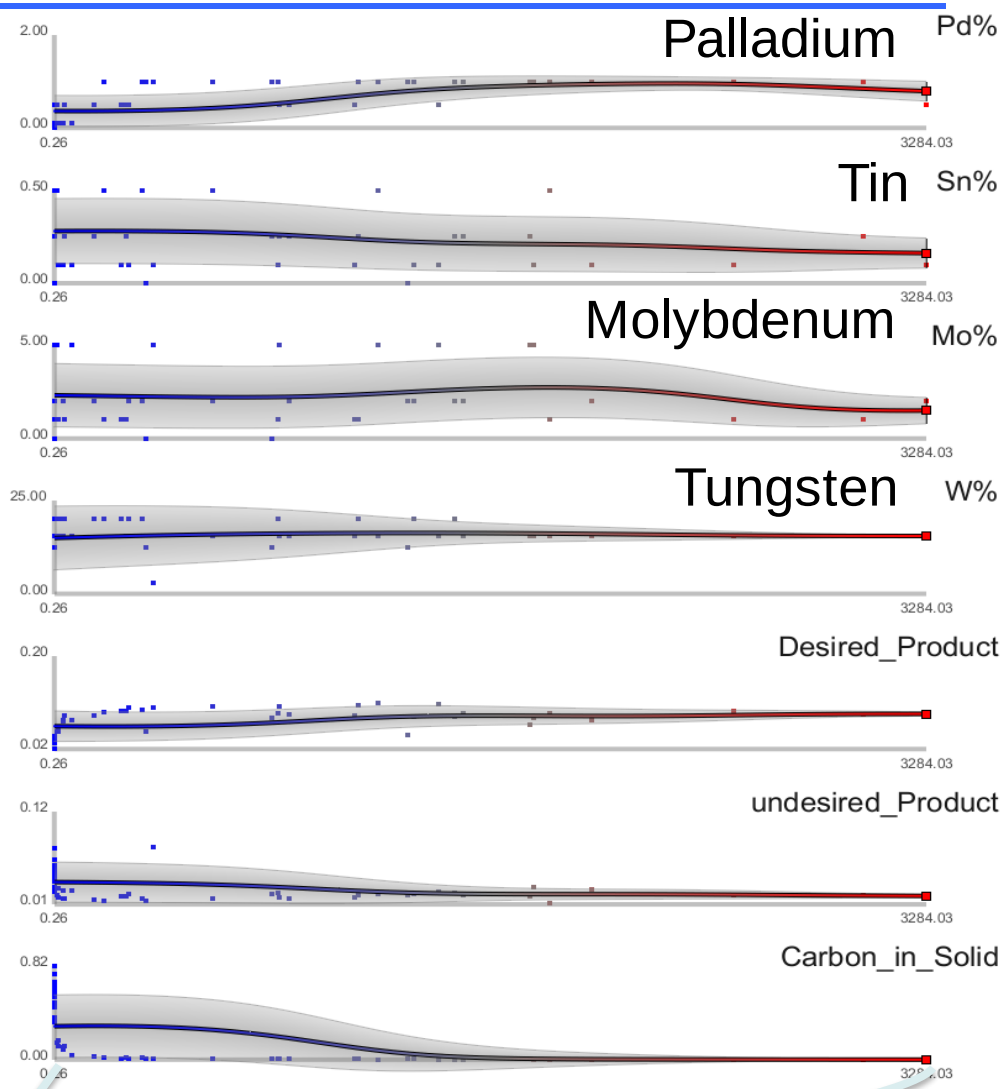

Design of a Composite Materials Topological Analysis Used to Plan Experiments

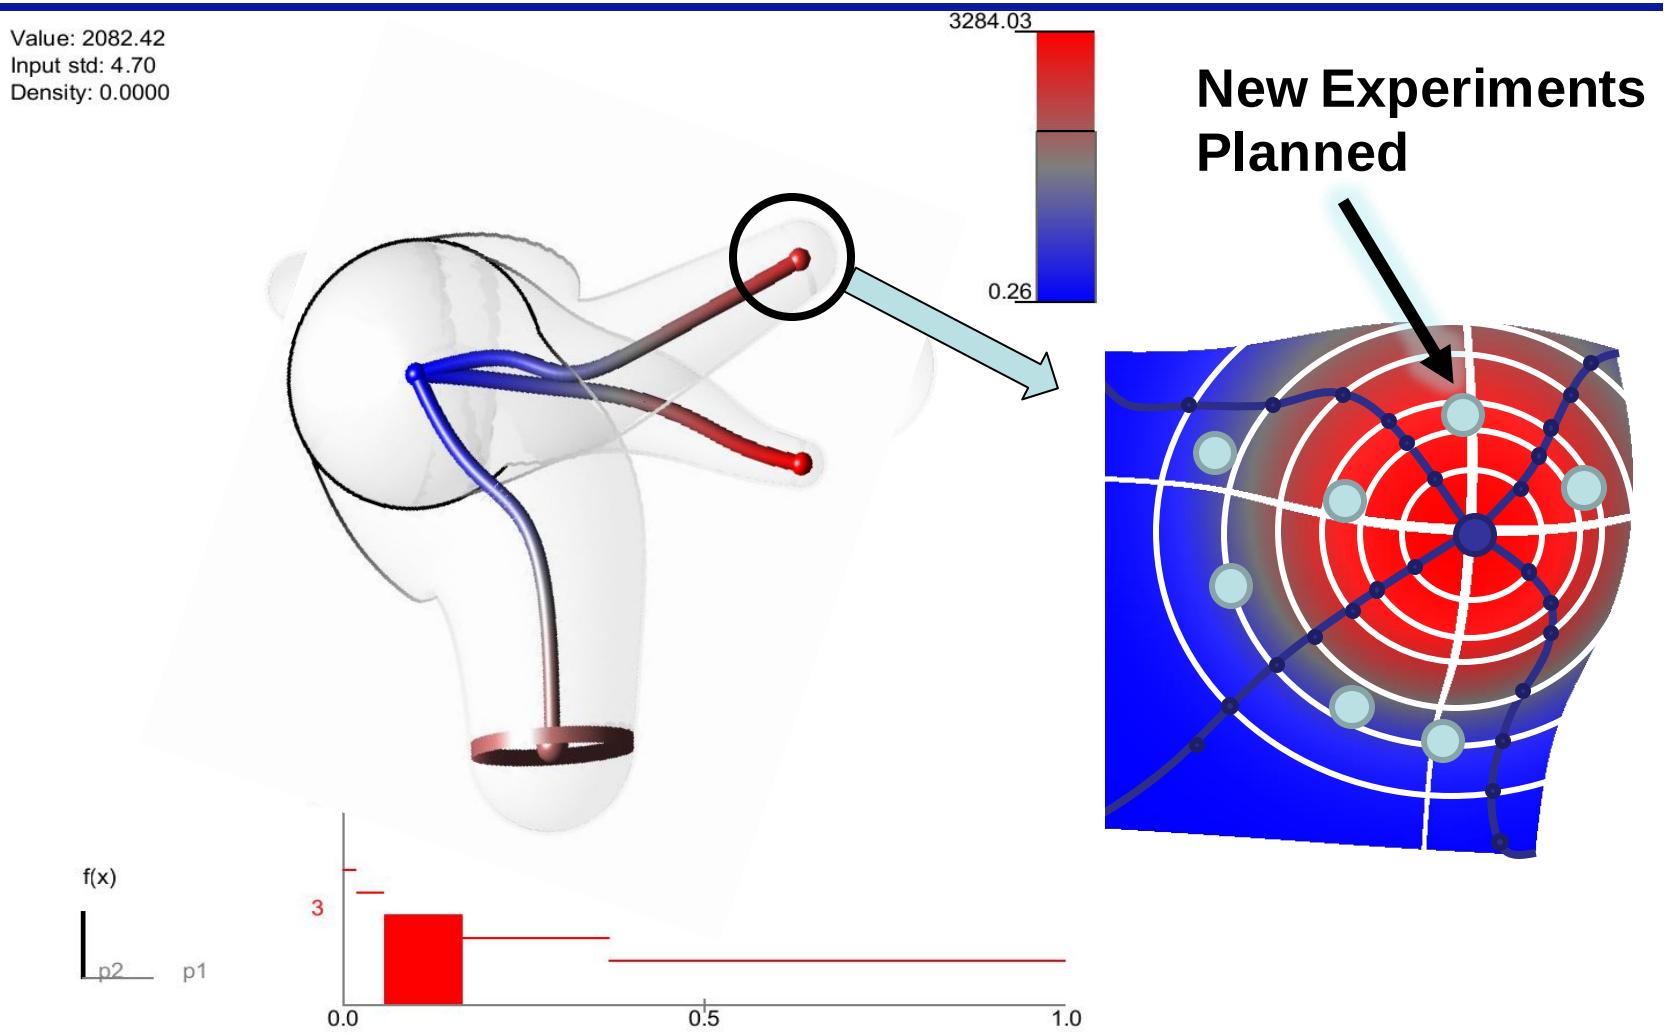

What is the Best Way to “Look” at The Data?

Viewpoint Selection is Equivalent to Exploration of Space of All Linear Projections

2D->1D projections
Parametrized on a circle

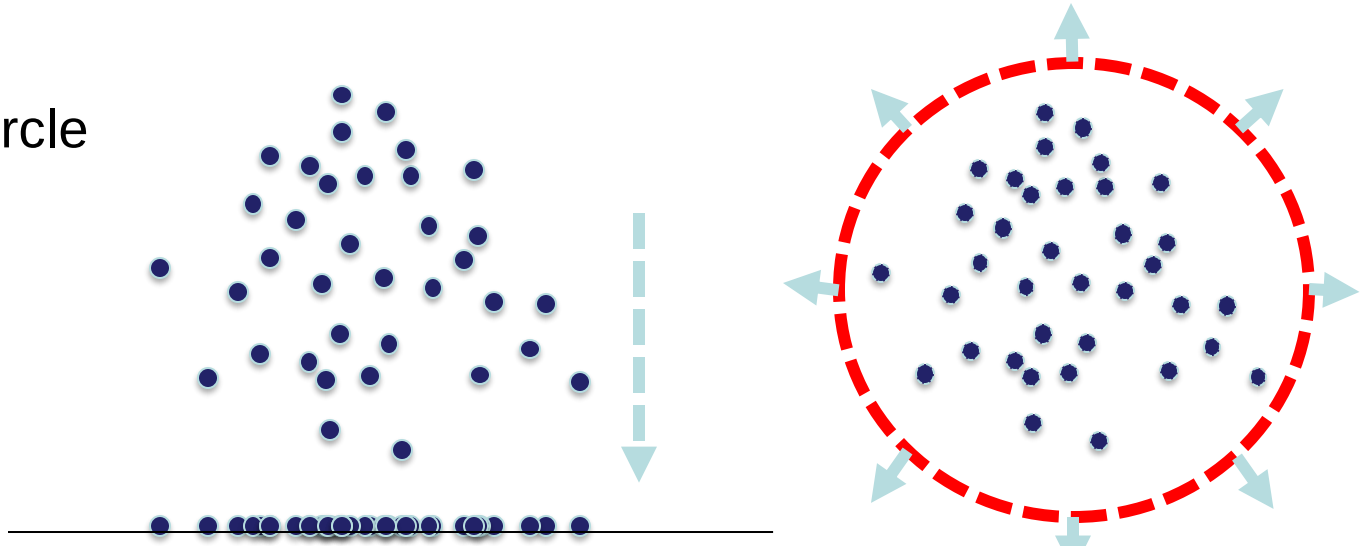

3D->2D projections
Parametrized on a sphere

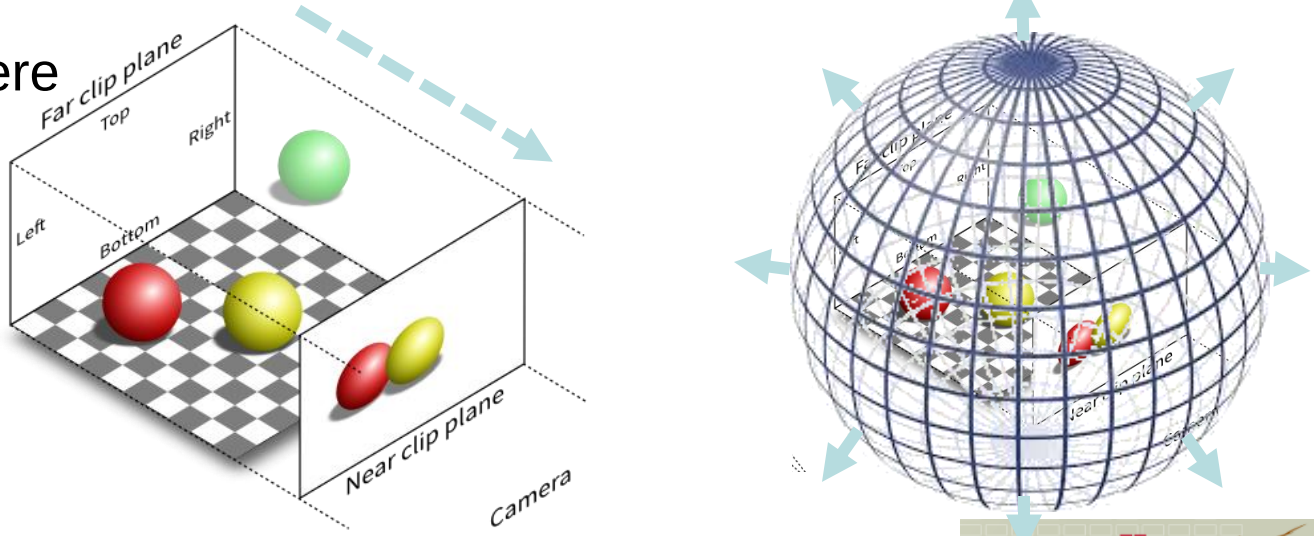

High Dimensional Data Exploration via Grassmannian Sampling

Sample the Grassman

Build neighborhood graph

Calculate quality measures on sampled views

Construct topological model

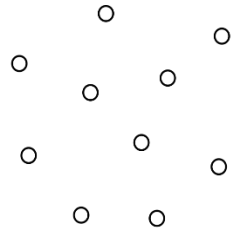

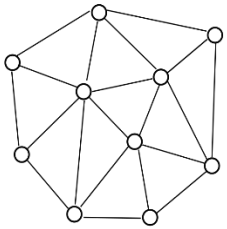

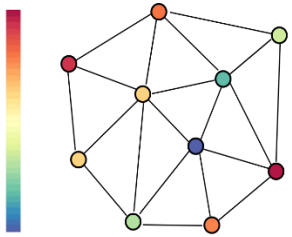

. . . .

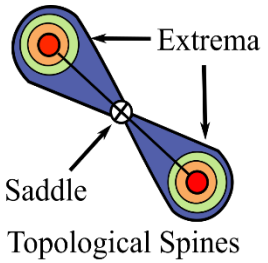

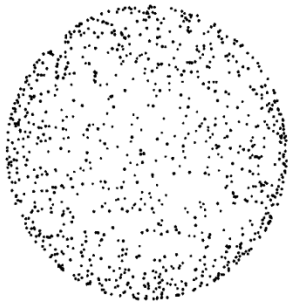

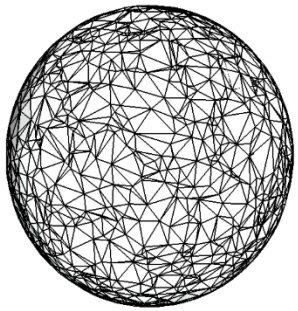

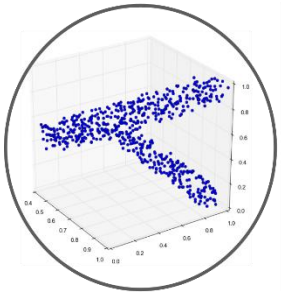

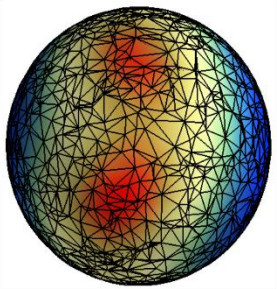

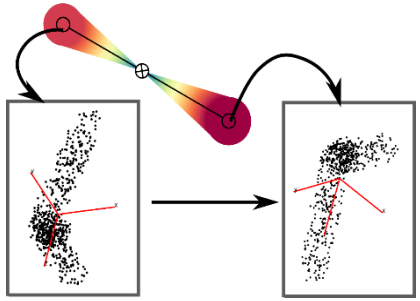

The Grassmannian: The Space of Linear Subspaces

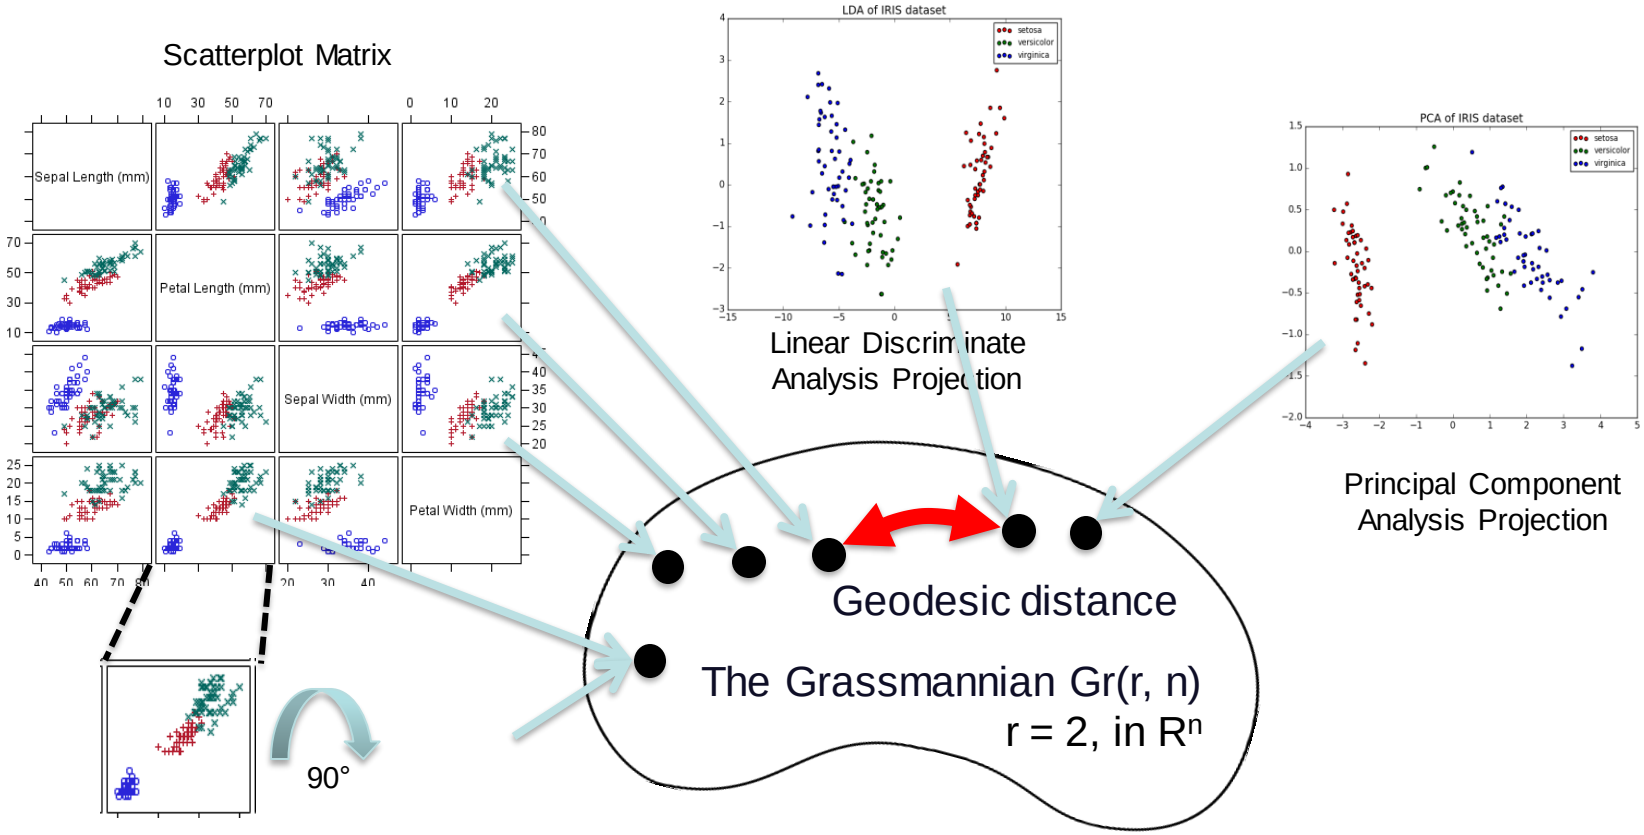

The Dimension of the Grassmannian Worsens the Curse of Dimensionality

The Space of All 2D Linear Subspaces is a **(2n-4)-Dimensional** Space

Data Dimension:	Dimension of the Grassmannian:
n	**2n-4**
2	**0**
2D dataset	A single 2D projection
3	**2**
3D dataset	From the 2D surface of a sphere to the origin

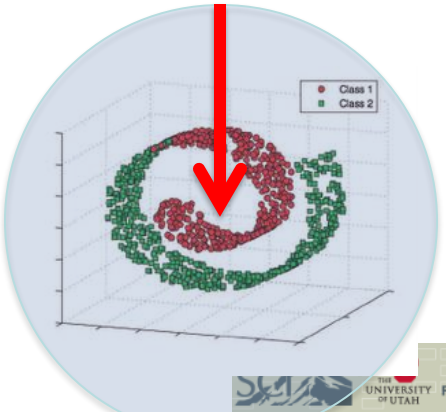

The Dimension of 2D Linear Projections Space Grows Twice as Fast as the Dimension of the Data

High-Dimensional data

n-dimensional

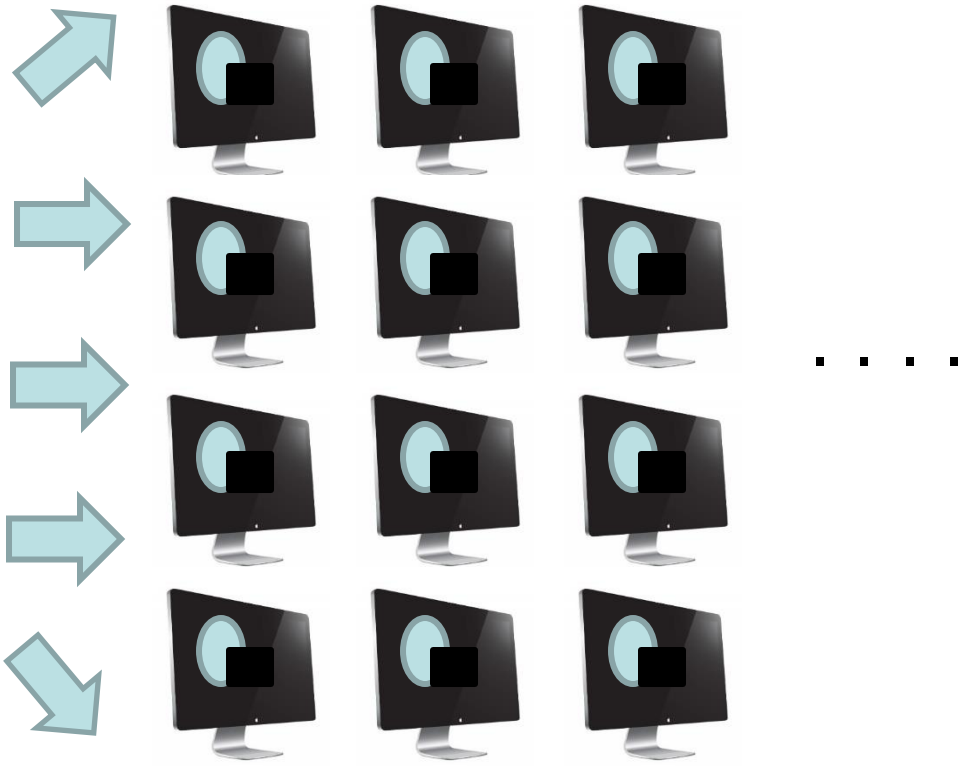

The dimension of the space of all 2D linear subspaces :

(2n-4)-dimensional

Interpreting the 2D Representations Is Difficult and Prone to Error

Information loss is usually unavoidable in a 2D representation

High-Dimensional Data

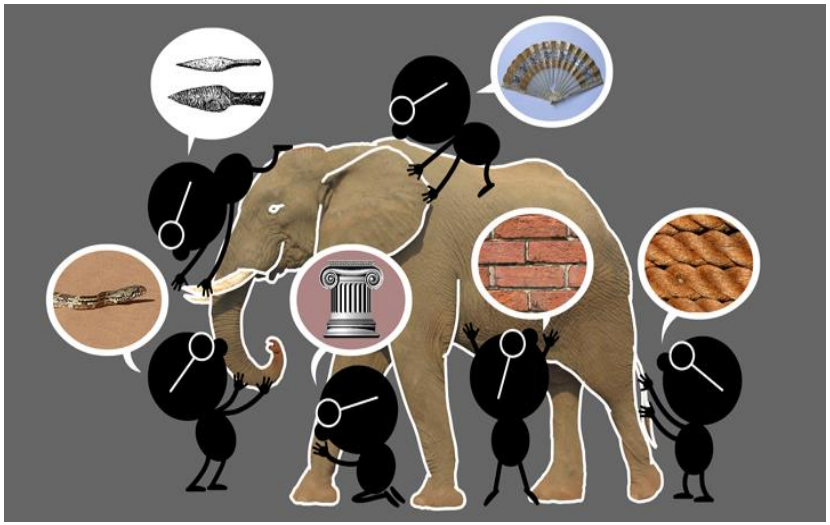

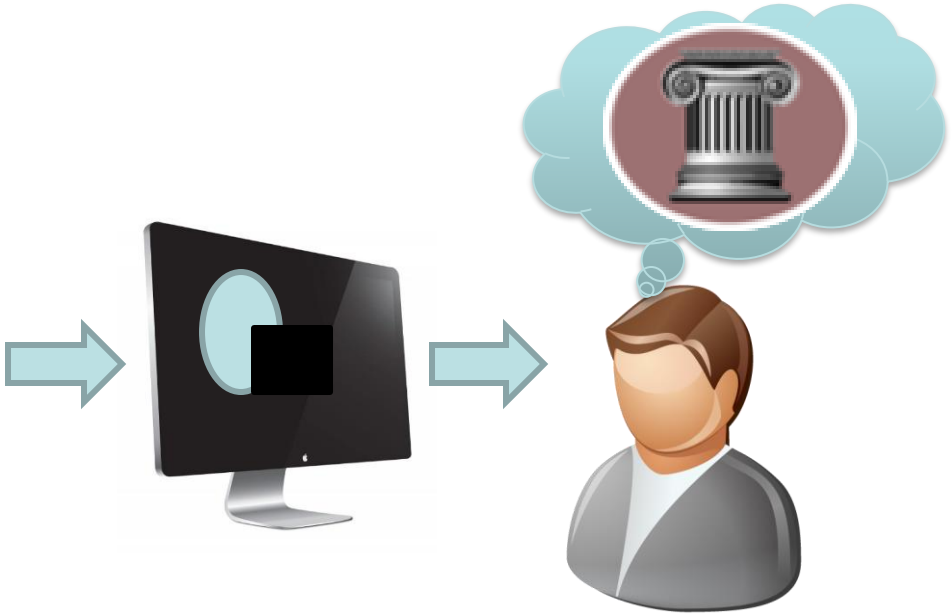

What information is lost?
How to avoid misleading observations?

Subspace Analysis Helps Identify Informative Projections

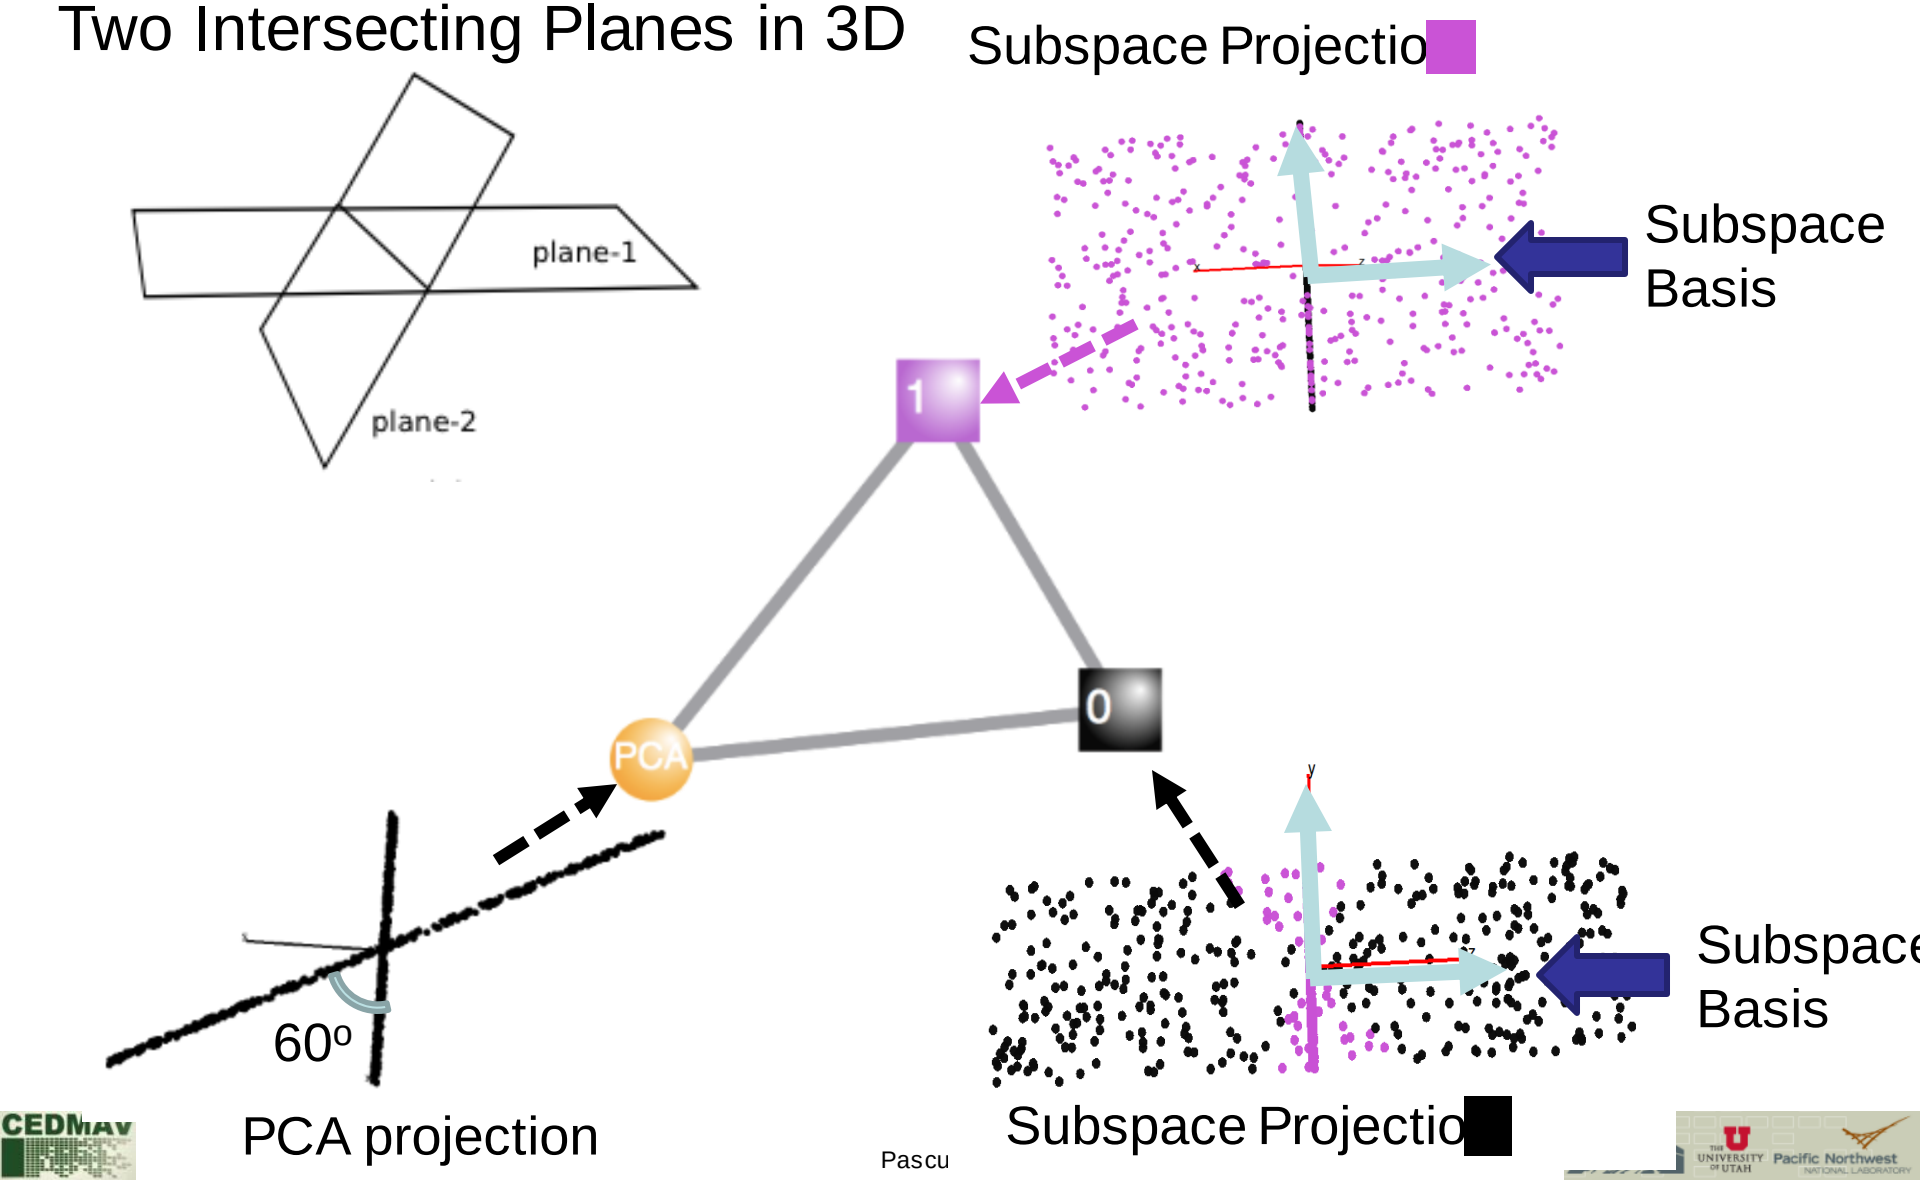

Views Navigation Graph Provides A Mental Map for Exploration

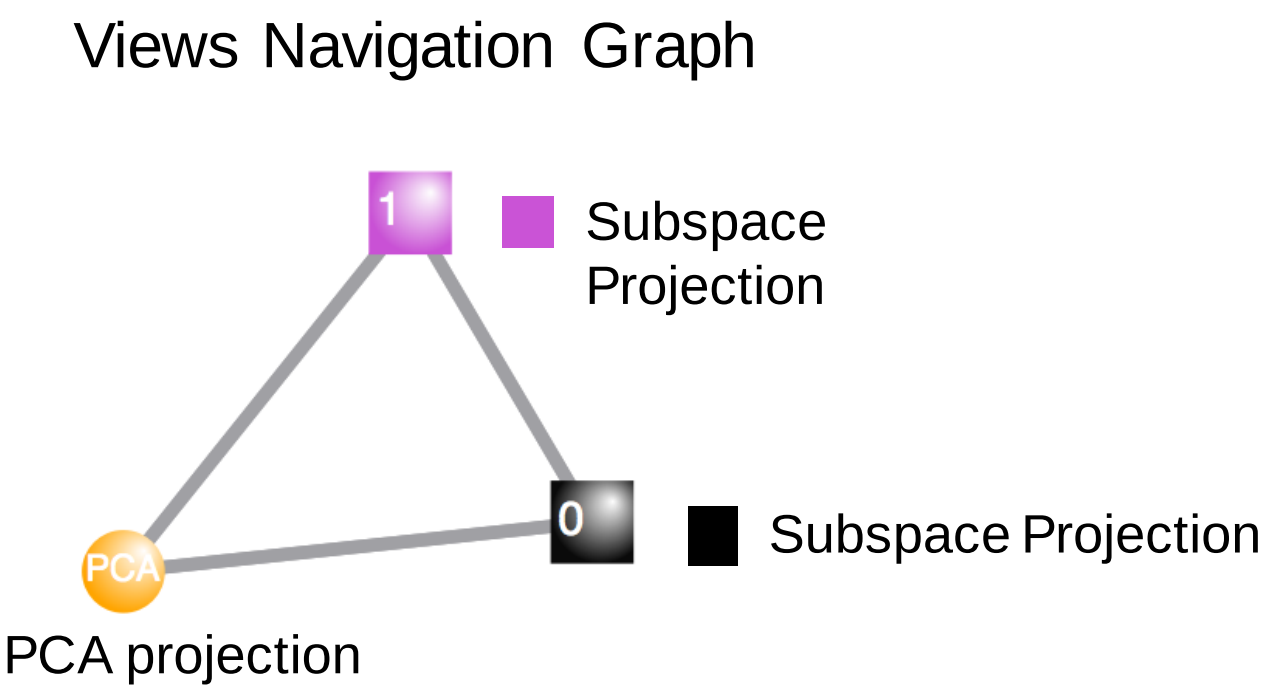

The projections (views) are connected by a k-Nearest Neighbor Graph based on their Grassmann distance

Lacking Frame of Reference During Transitions Hinders Exploration

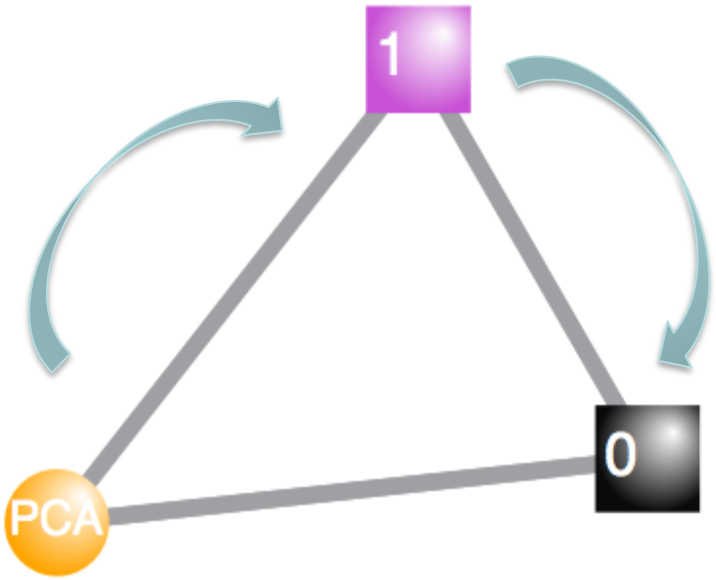

Views Navigation Graph

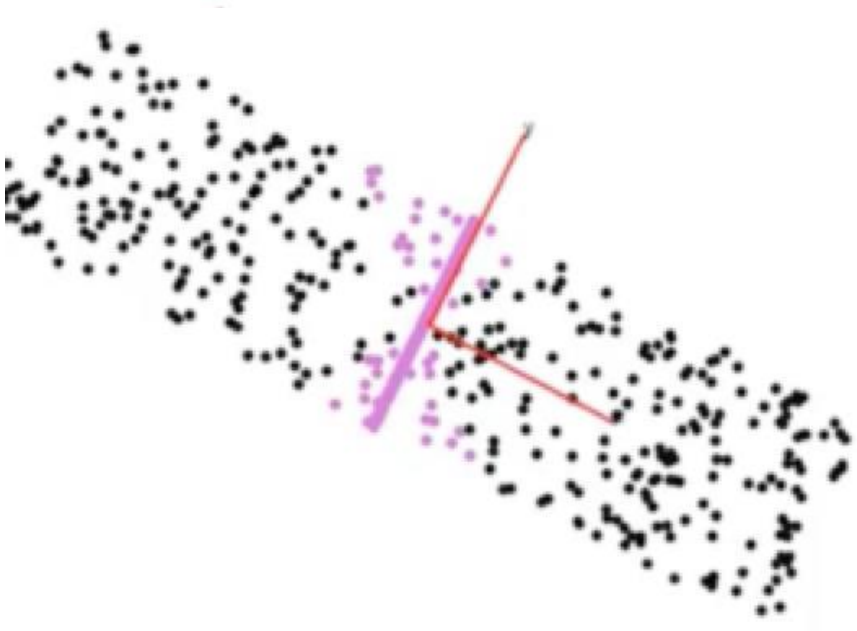

Projections (Views)

Dynamic Projection Generates Smooth Transitions Among Linear Projections

Each frame is a linear projection along the shortest path on the Grassmannian

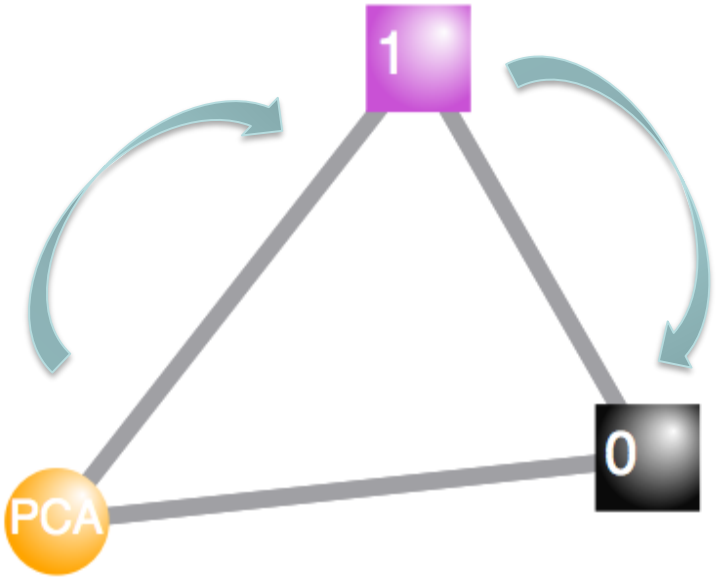

Views Navigation Graph

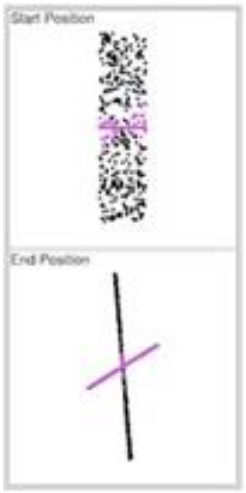

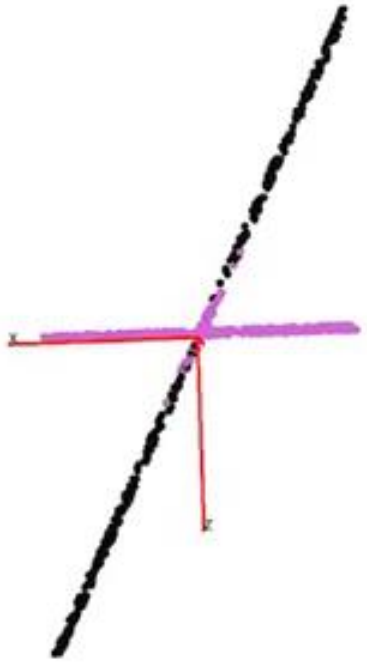

Dynamic Projection

Application Example: Yale Face Dataset

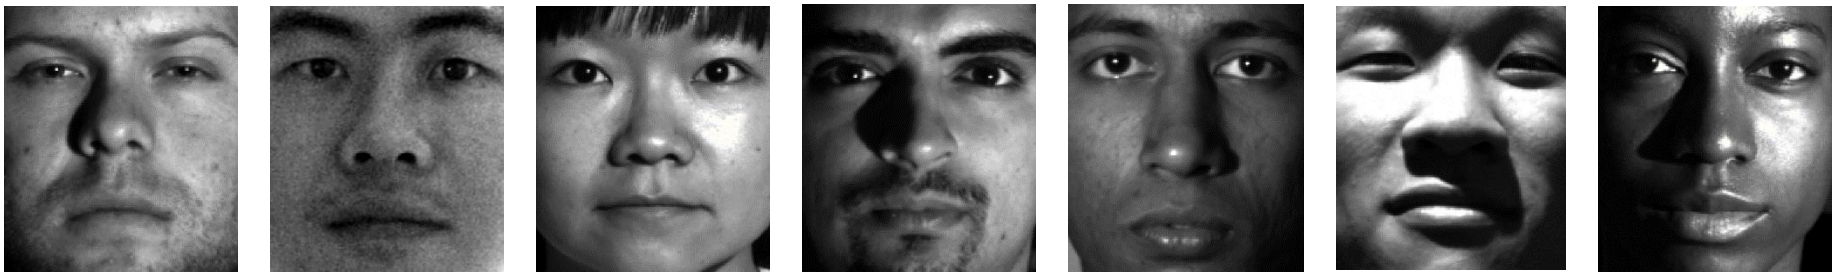

- Subsample from the Yale Face database
- Consists of 439 face images from seven people
- Use random projection to reduce their resolution to 10 x 10
- Explore the image feature space of a computer vision task

Views Navigation Graph for Exploring Yale Face Dataset

Views Navigation Graph

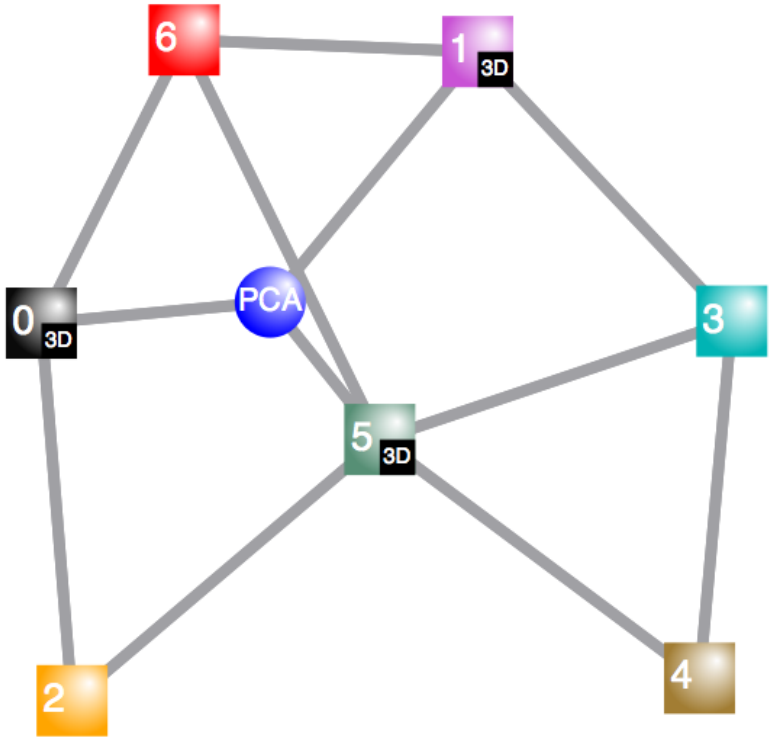

- Each color corresponds to a cluster
- The edges in the layout encode the Grassmannian distances among the subspace projections

PCA Projection Gives Poor Cluster Separation

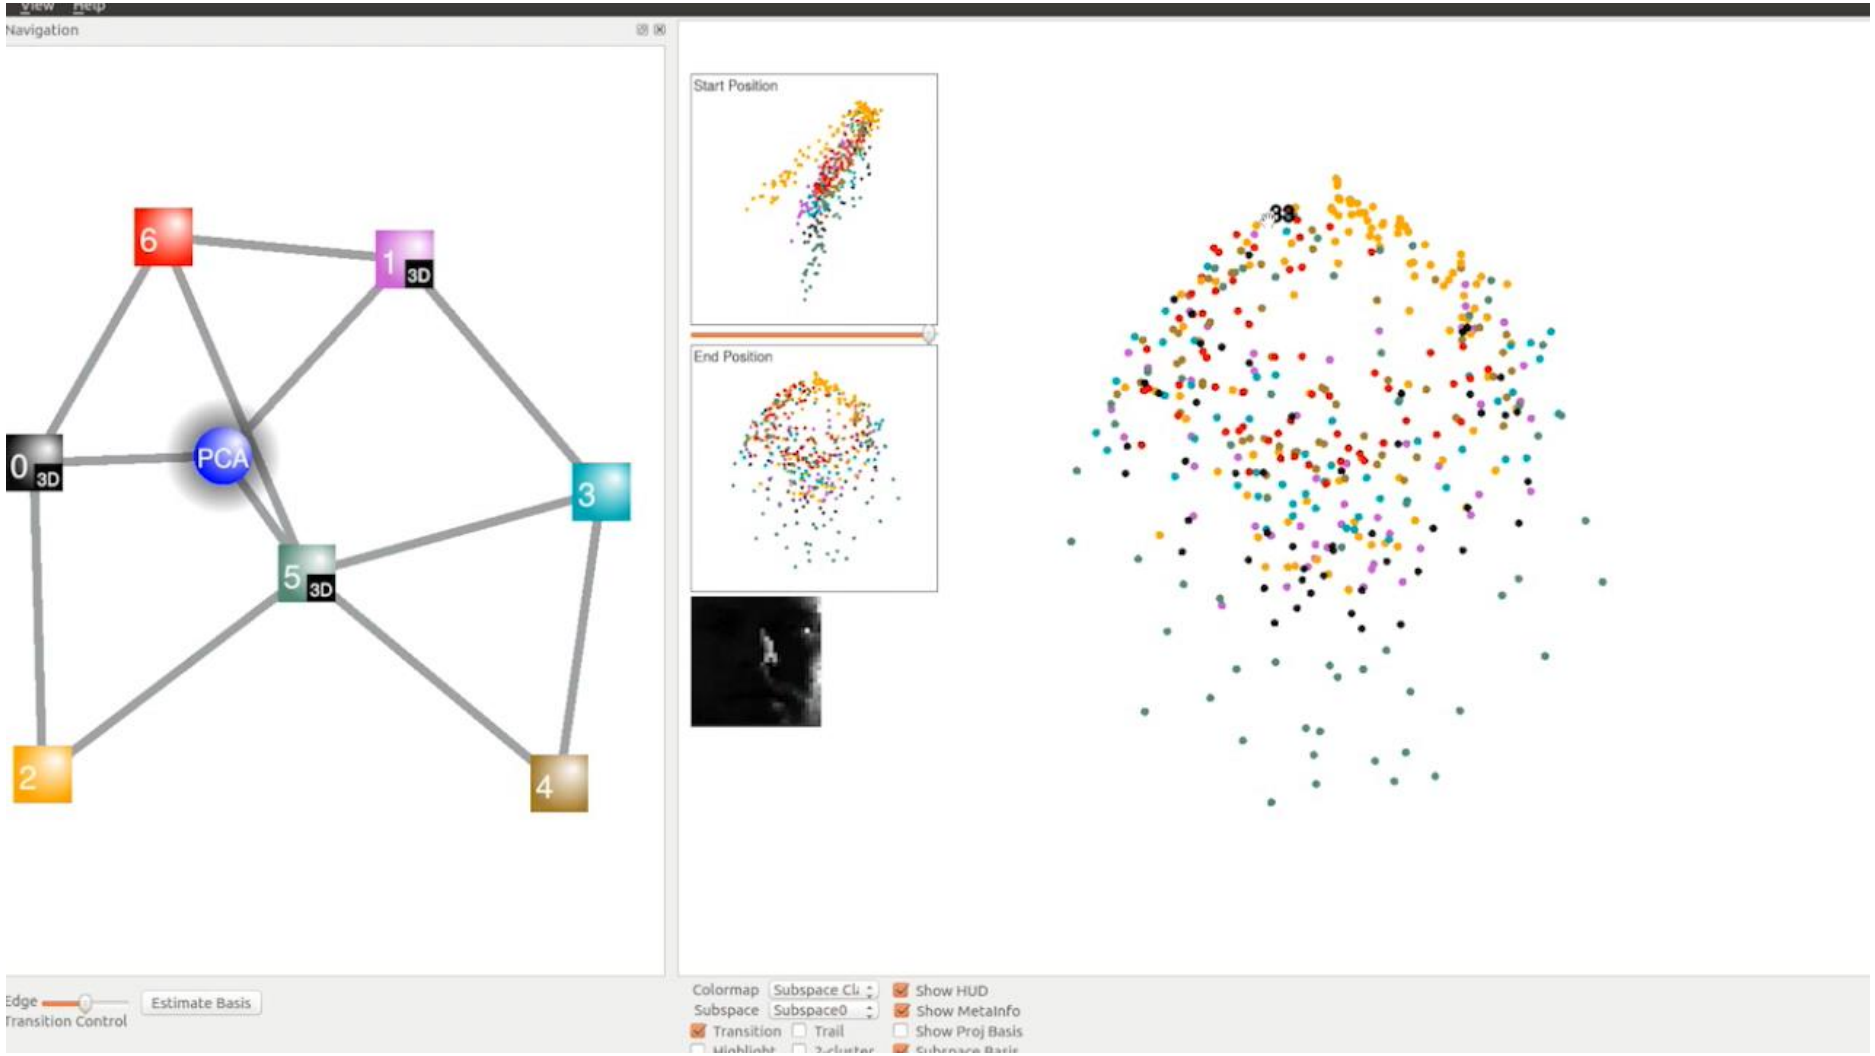

PCA Shows Faces Laid Out in Circular Pattern According to the Varying Lighting Directions

“Rotational” Transitions From PCA to Brown, Cyan, and Orange Subspace Projections

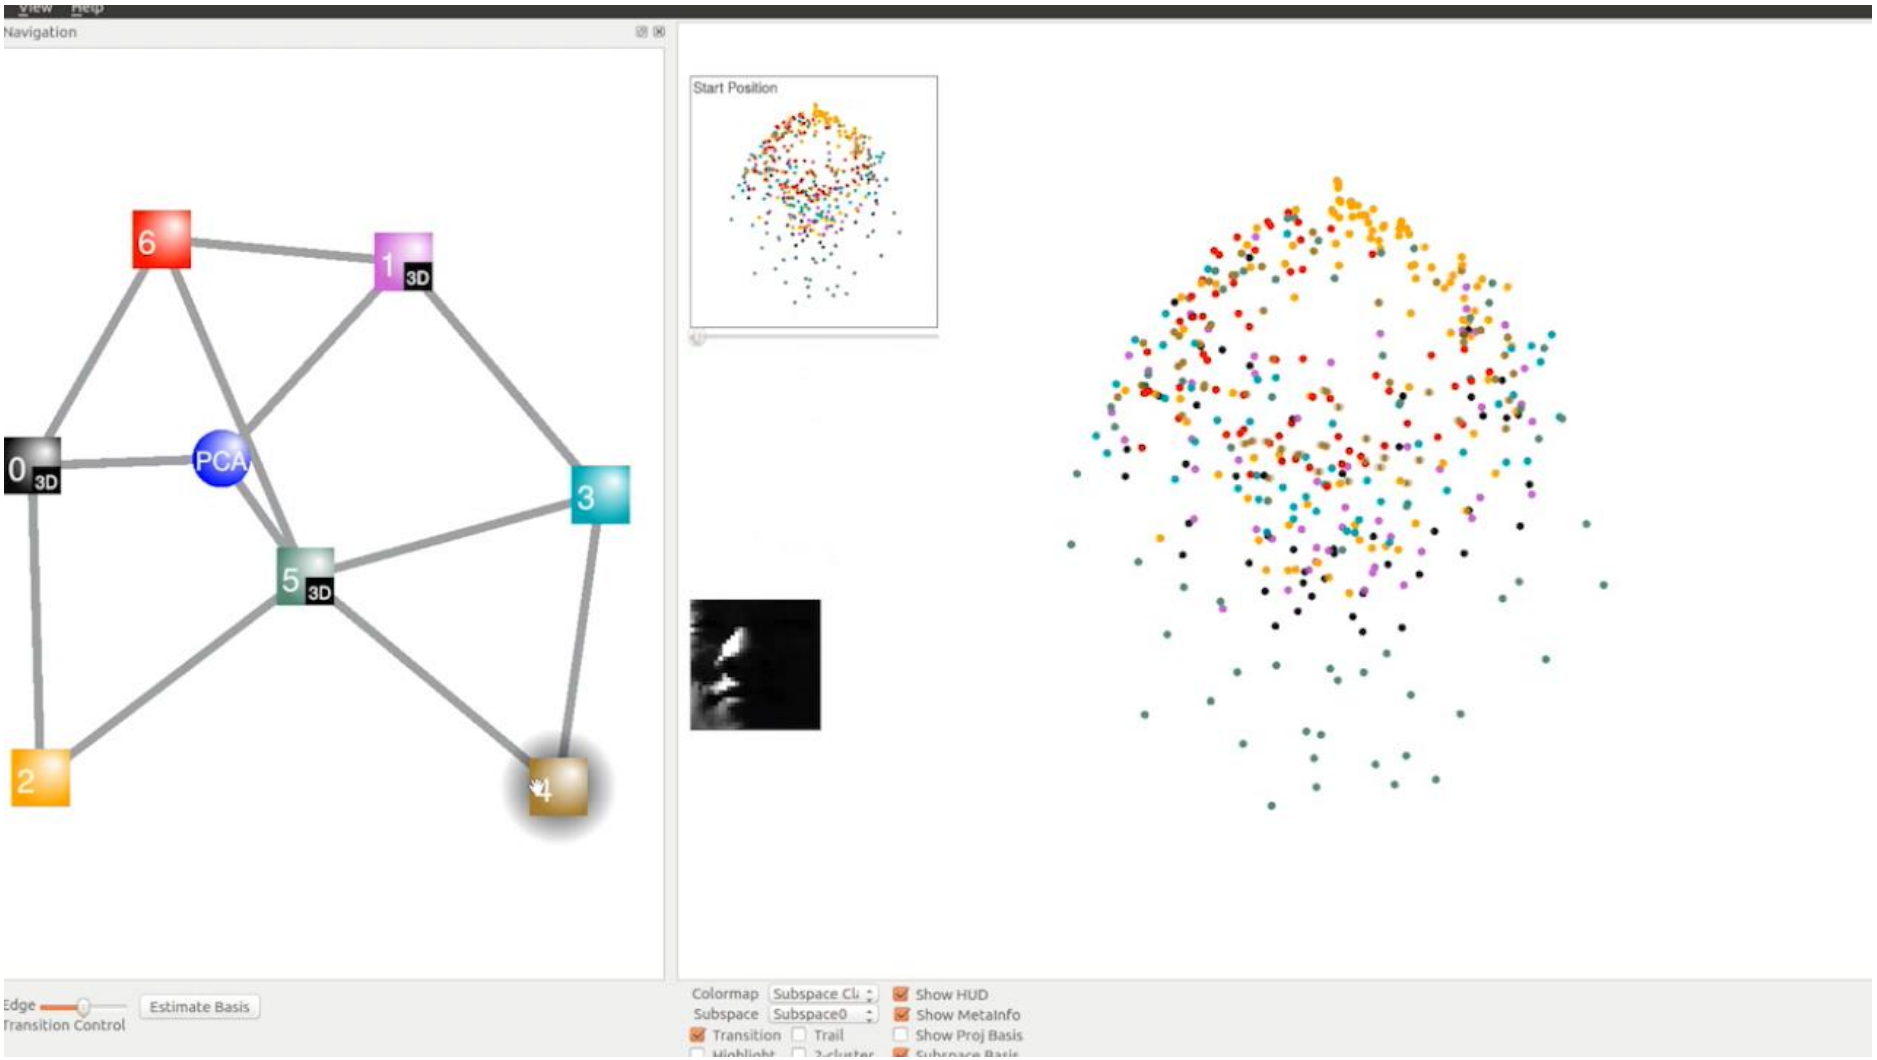

Orange Cluster (Asian Female) Is Well Separated from the Remainig Points

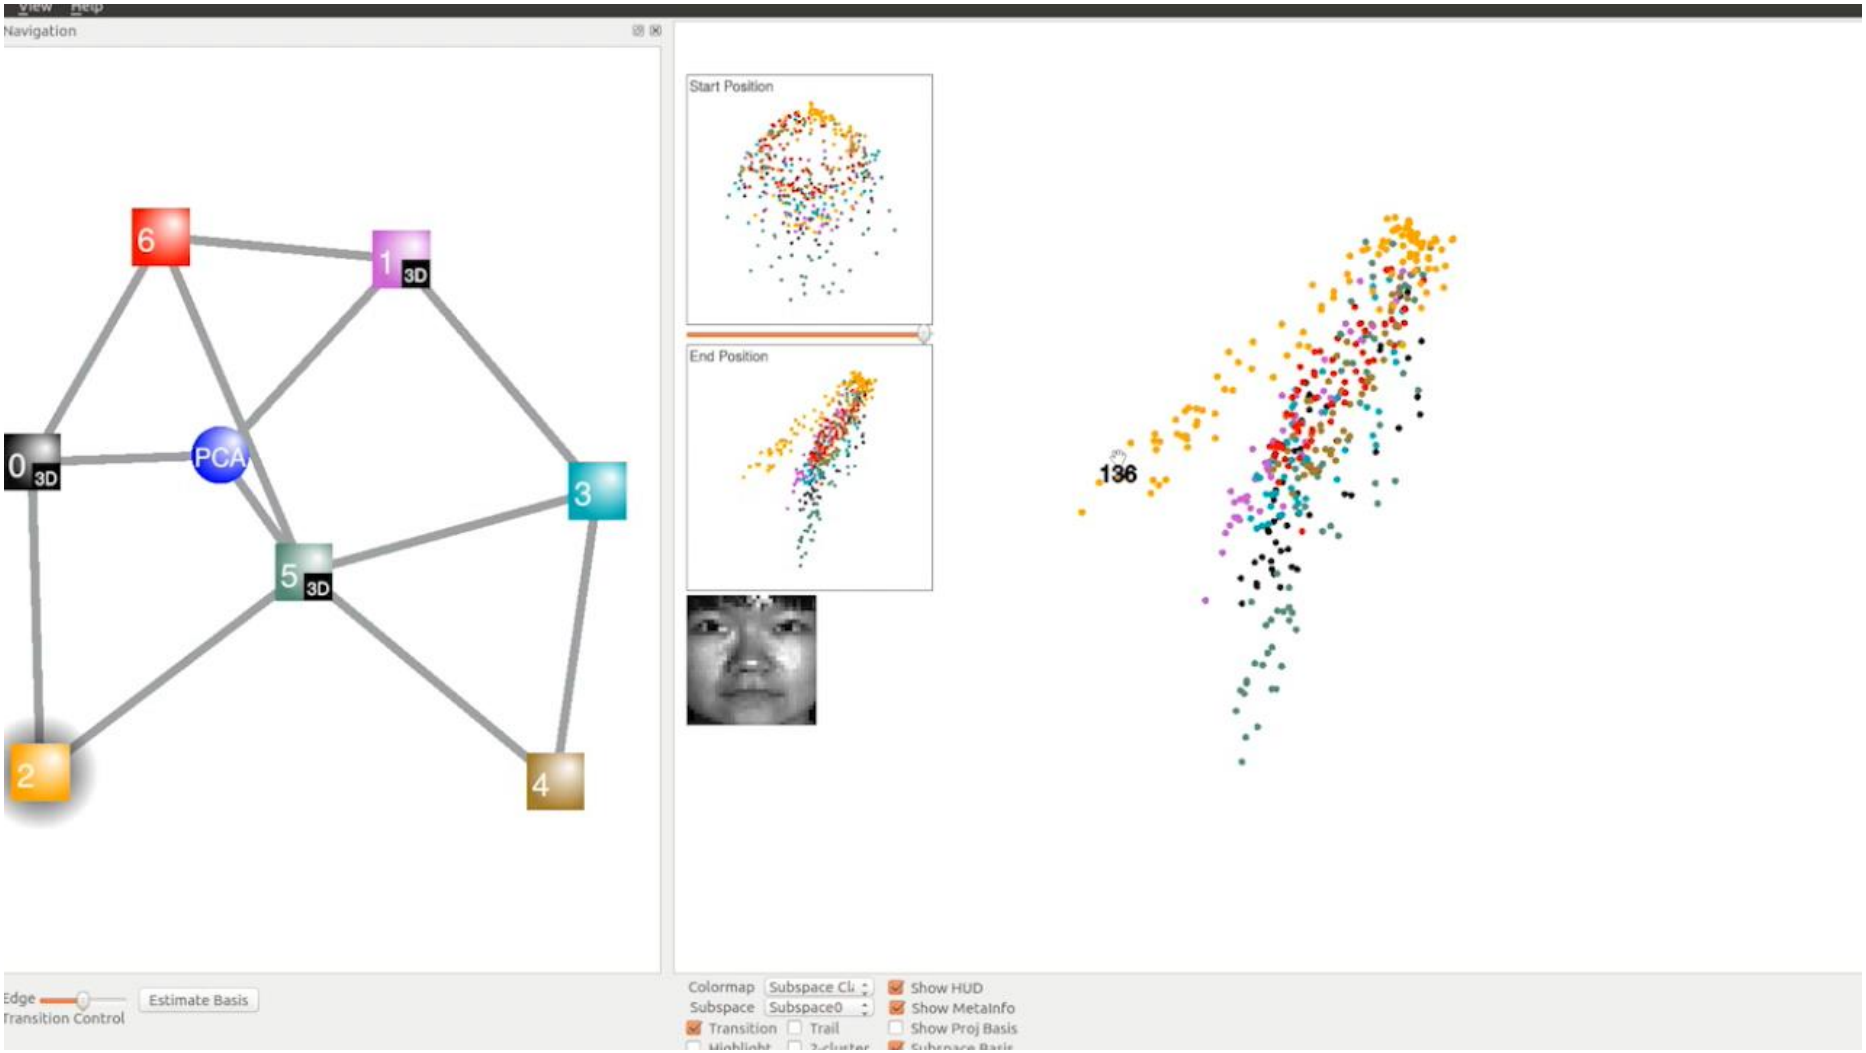

The Subspace Projection Captures the Dominant Local Structure

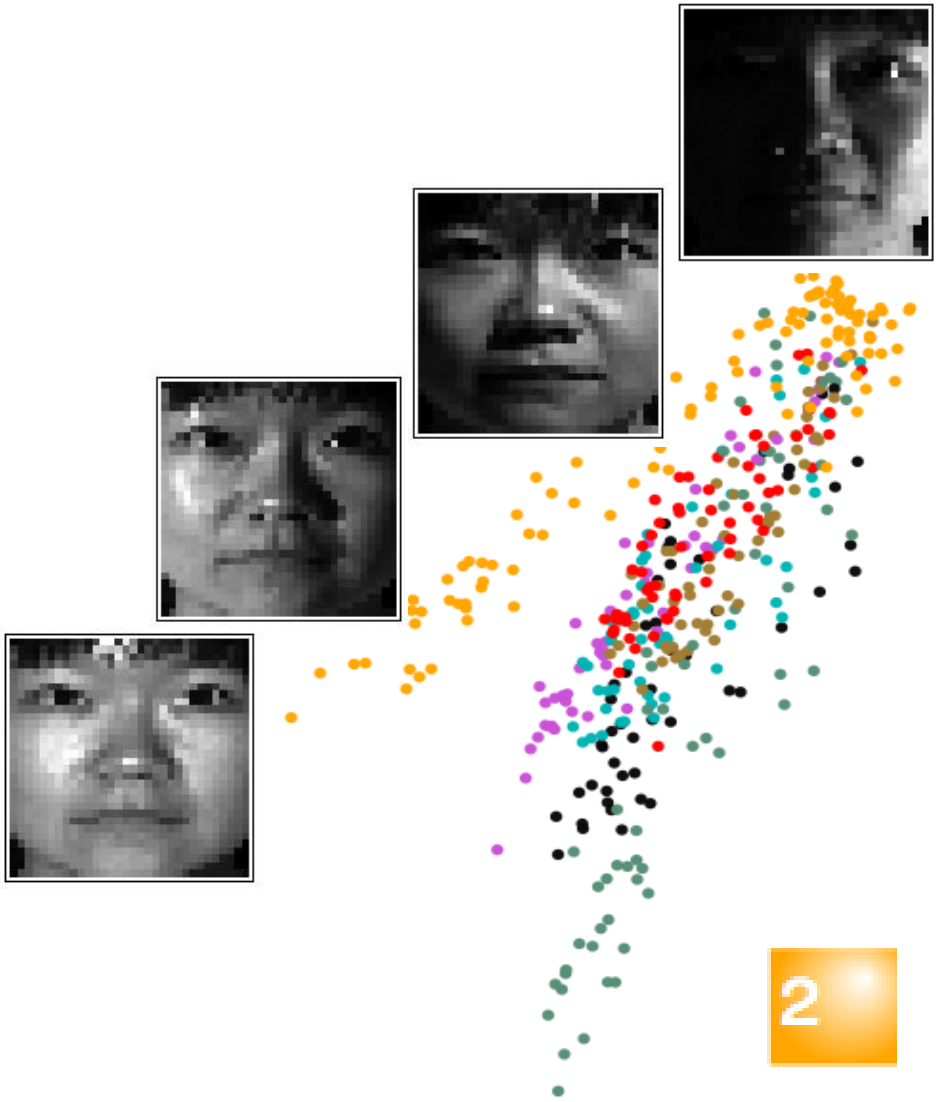

Subspace Clusters Label vs. Ground Truth Label

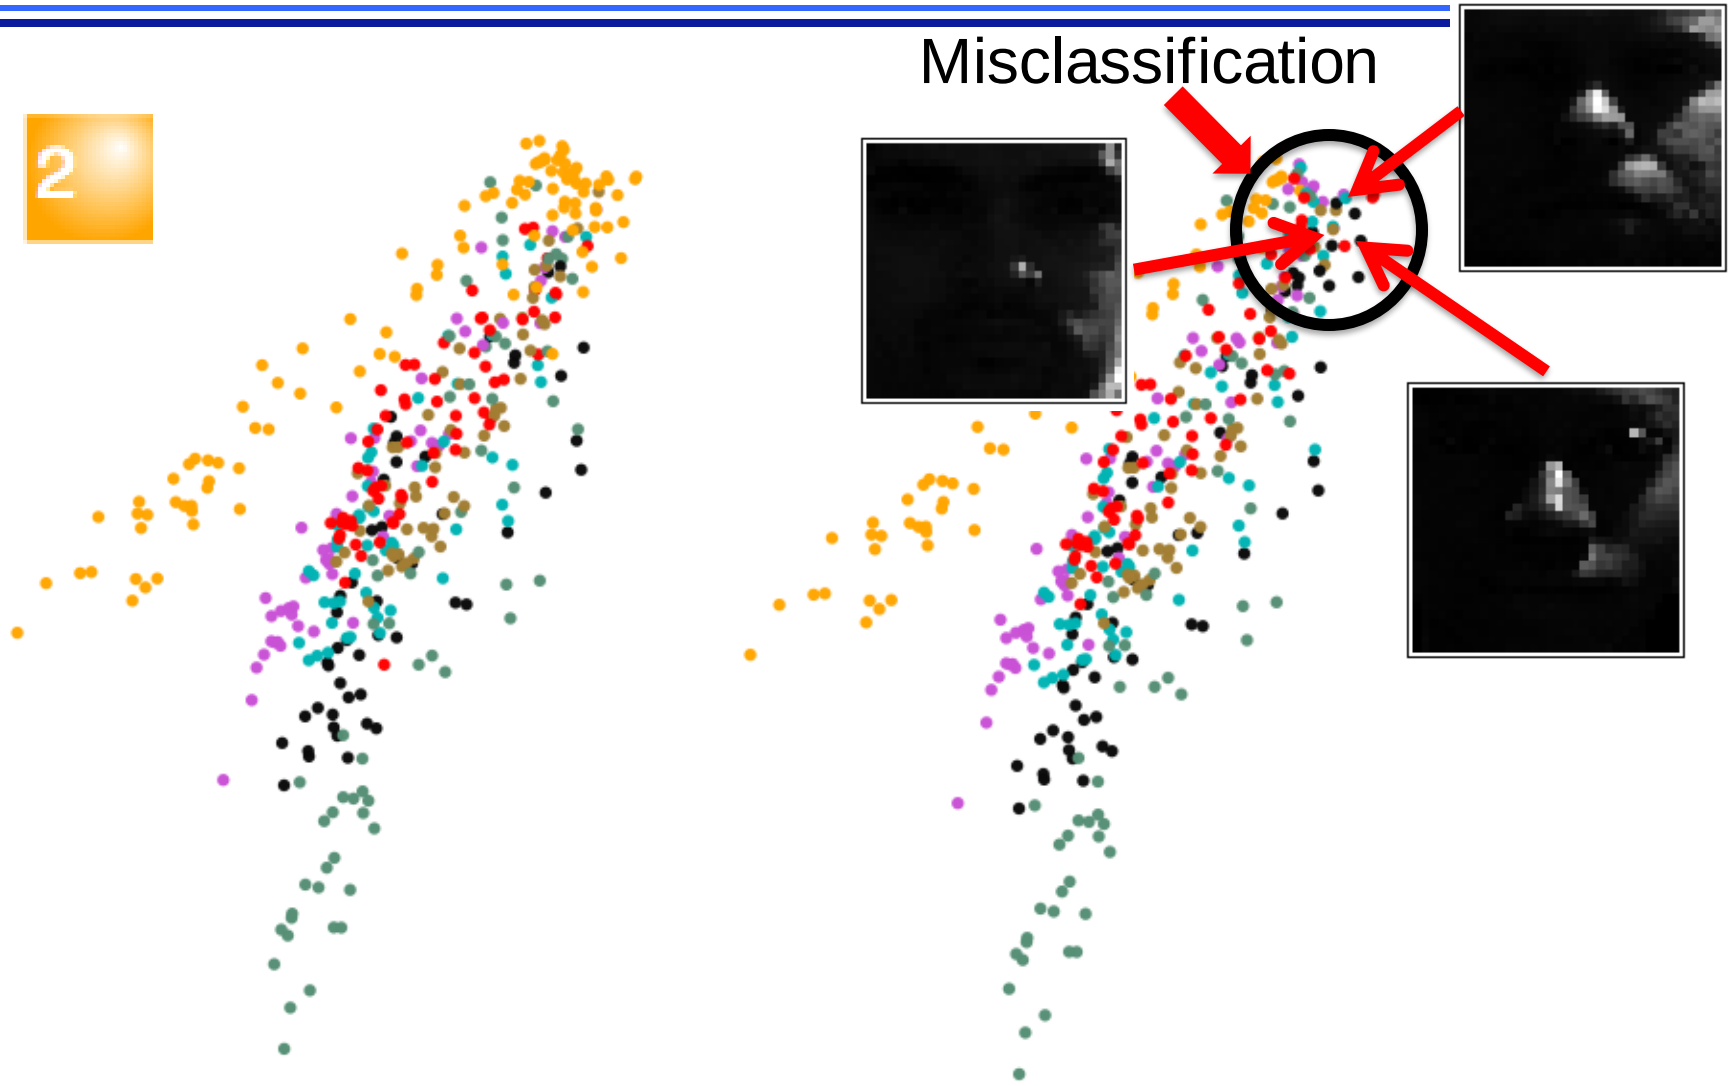

Subspace Clusters

Ground Truth Label

Transition to the Red Subspace Projection Exhibits a Different Rotational Pattern

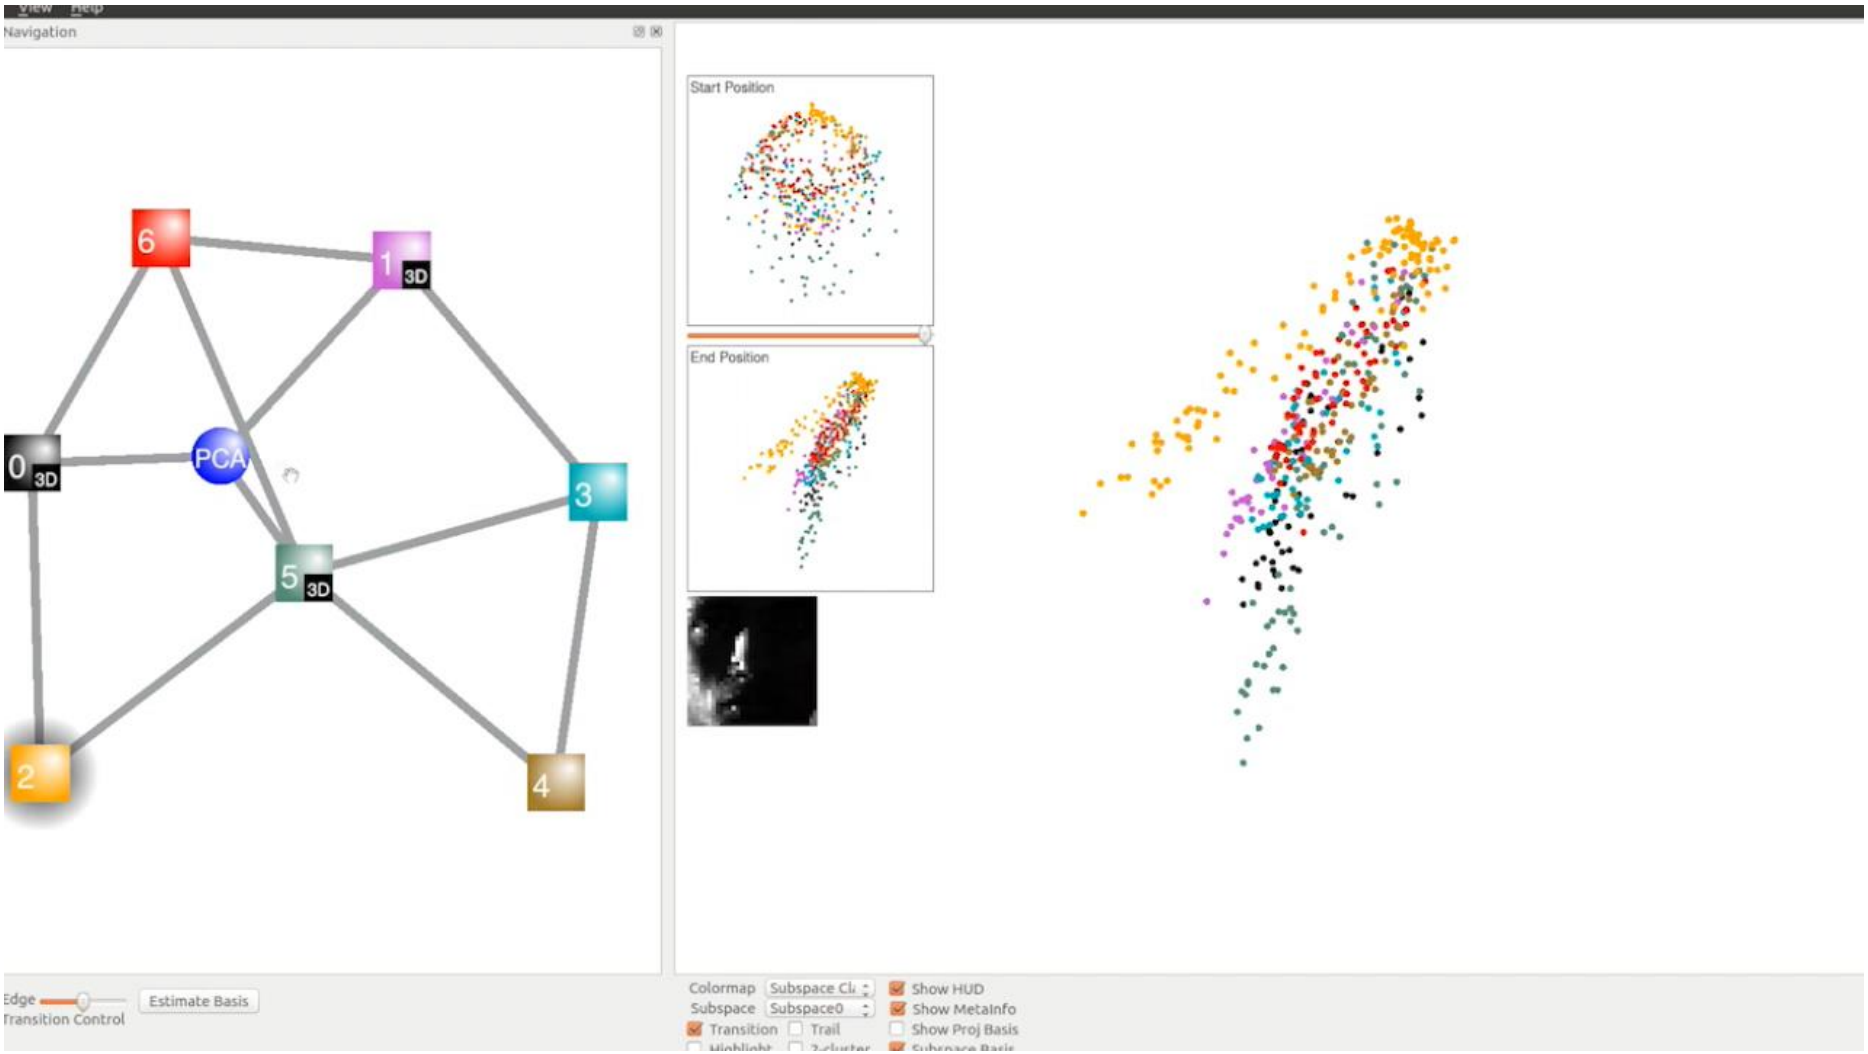

The Lighting Direction Corresponds to the Dominating Trend

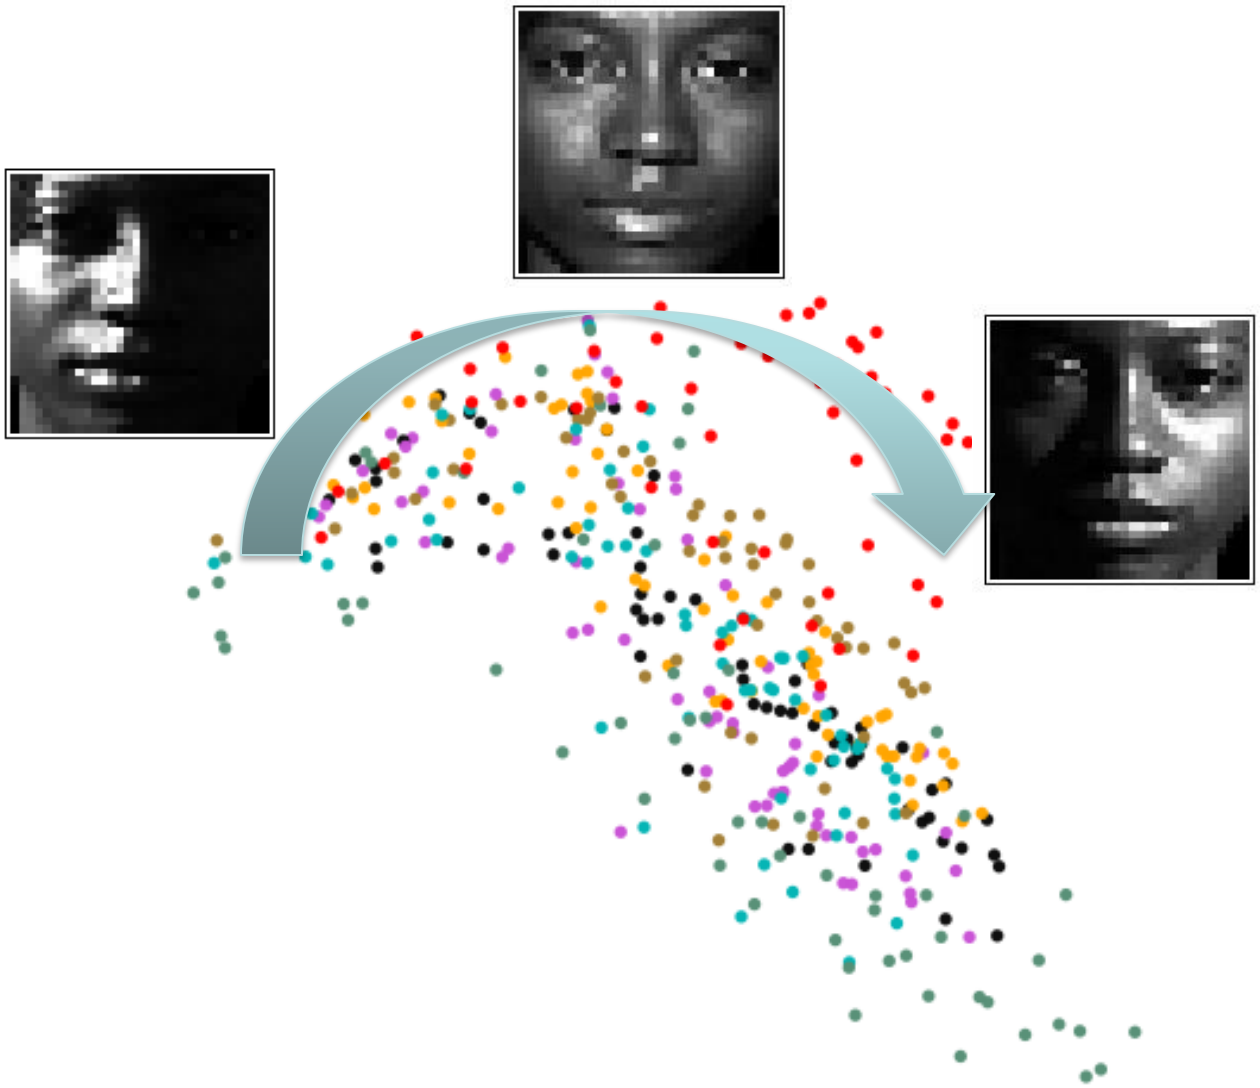

Navigating High-Dimensional Space Can Use Subspace Analysis and Dynamic Projections

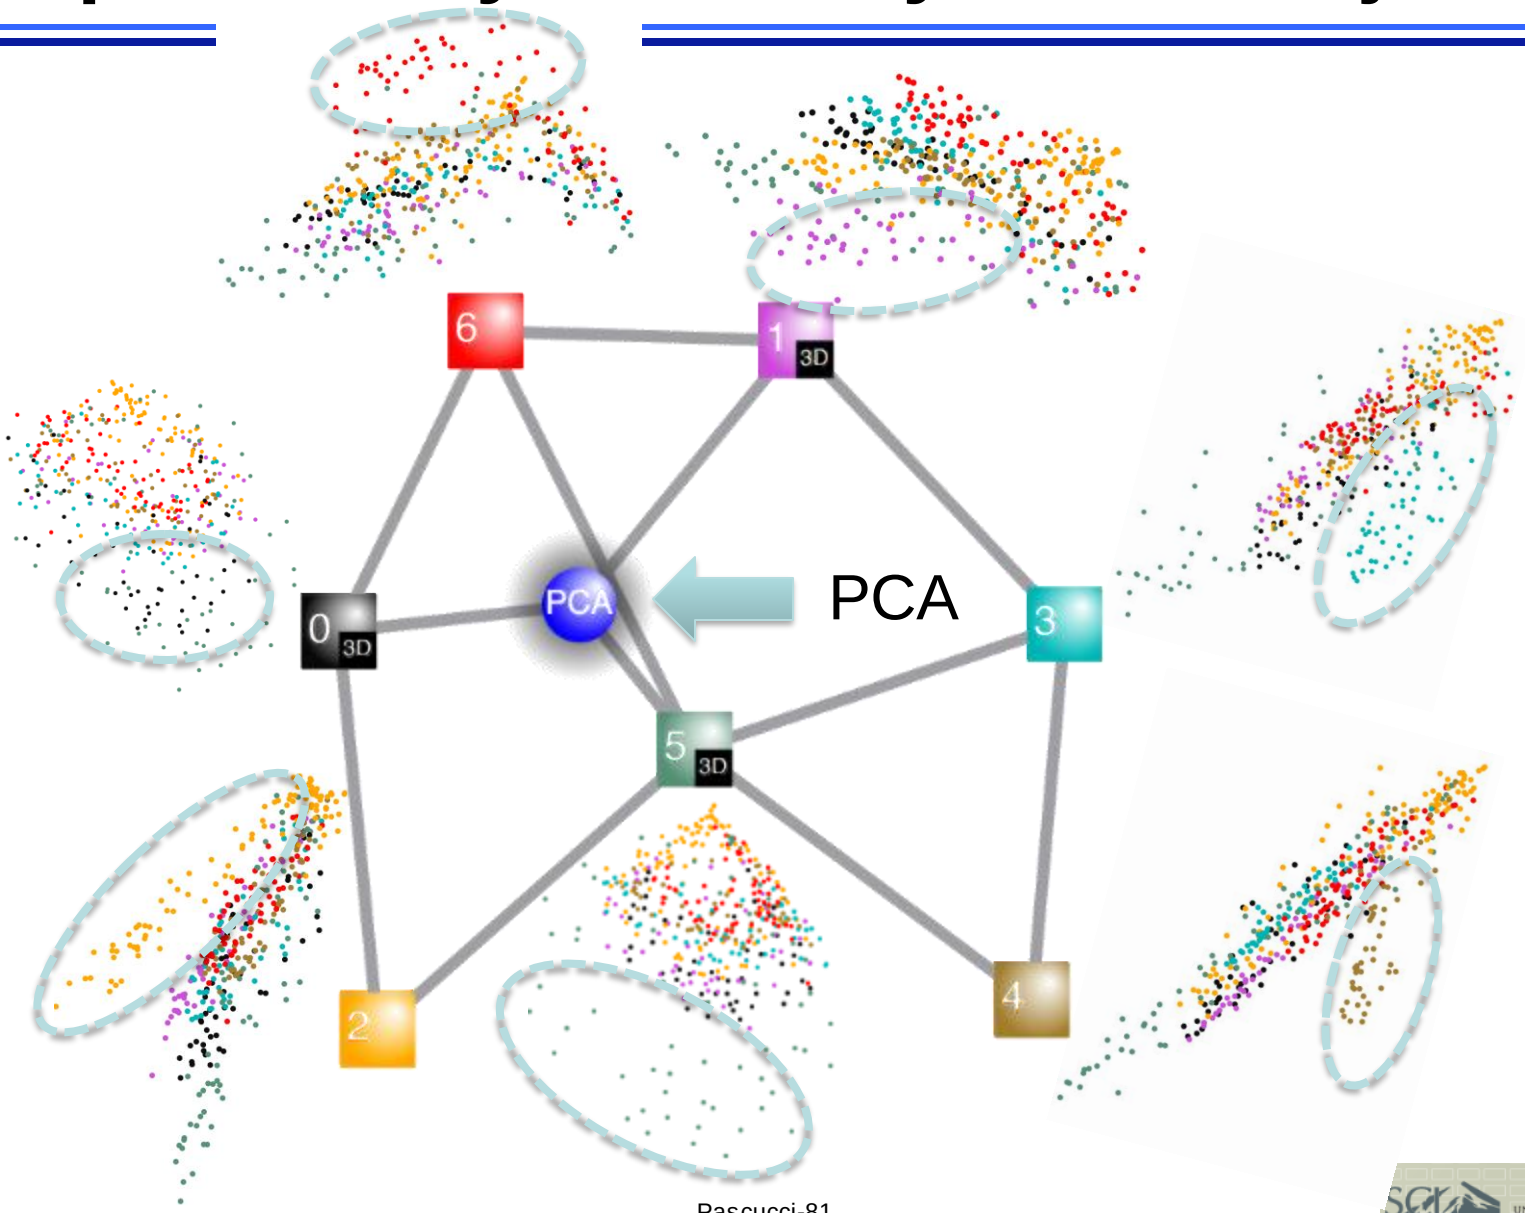

GoogleNews Word Embedding Dataset

- GoogleNews dataset, trained from 100 billion words
- 900 commonly used words and phrases are picked
- PCA is applied to the word vectors to reduce the sampling cost

A Typical Exploration Session of the Word Embedding Data

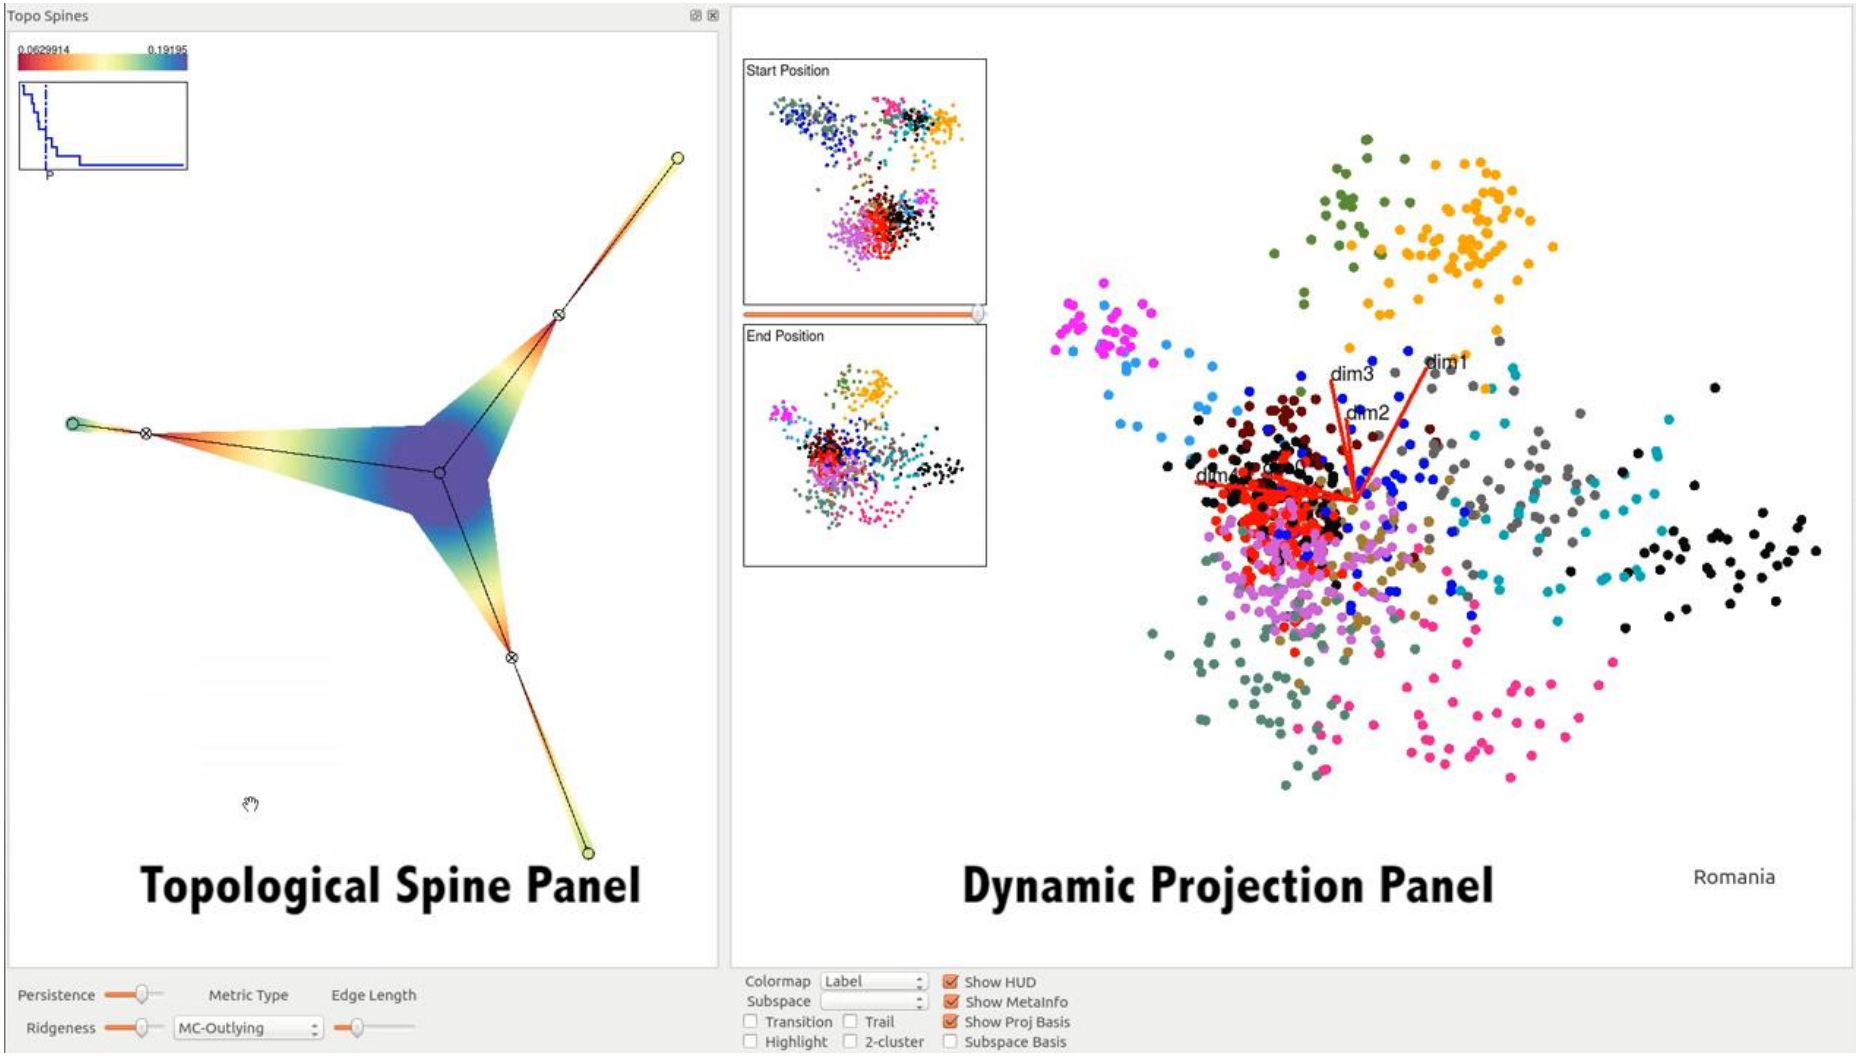

The Set of Local Maxima Provide More Informative Projections than the Global Maximum

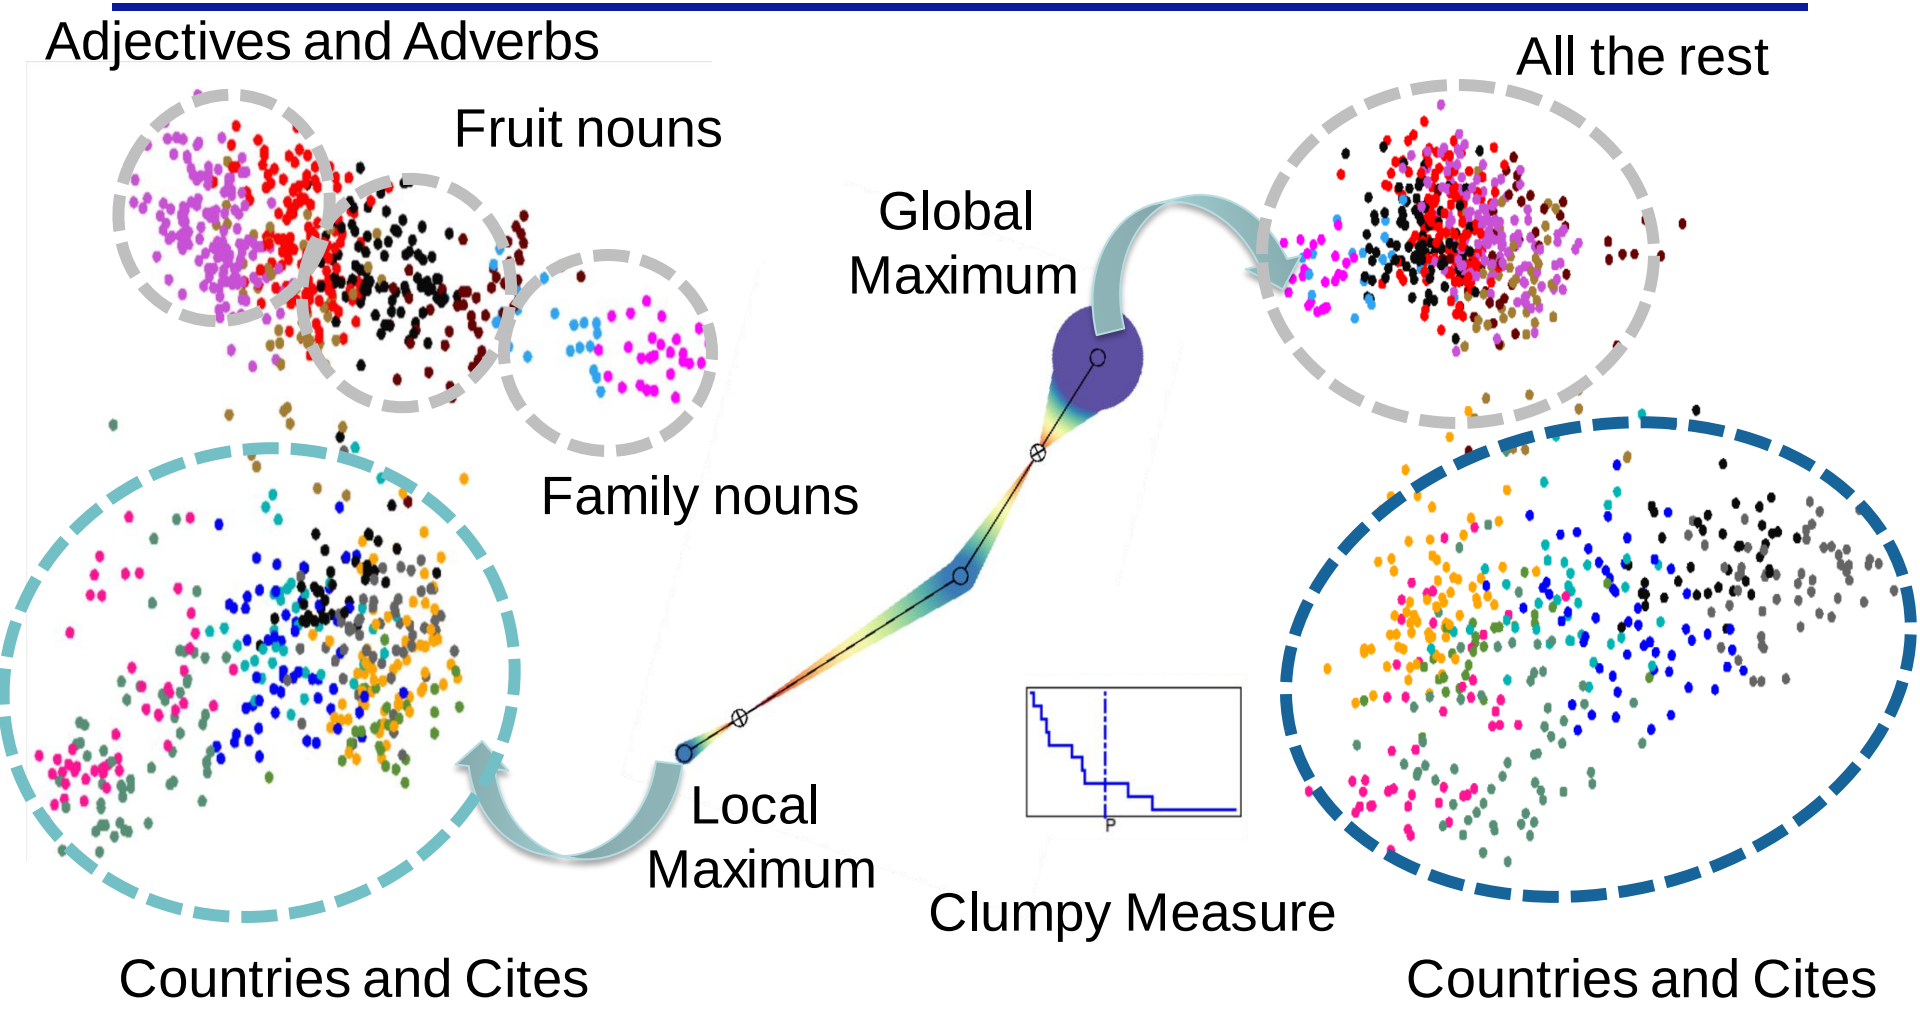

Color corresponds to word type label

Path to Commercialization

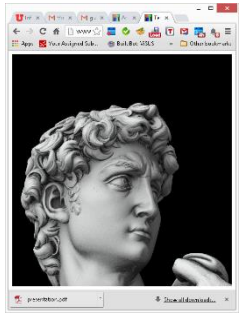

- Mathematical research: space-filling curves
- Algorithmic research: Progressive computing
- System research: scaling from iPhone to HPC
- Application research: digital photography, medical imaging, …
- Commercialization:
  - Training
  - Oil and gas
  - Histological exams
  - Imaging labs
  - Precisions agriculture
  - Microscopy

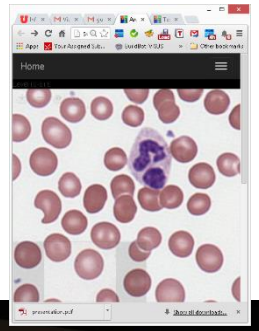

ViSOAR: a Unified Solution for Distribution of Imaging Data in Medical and Geology

Example of Visualization and Annotation of outcrop and medical data

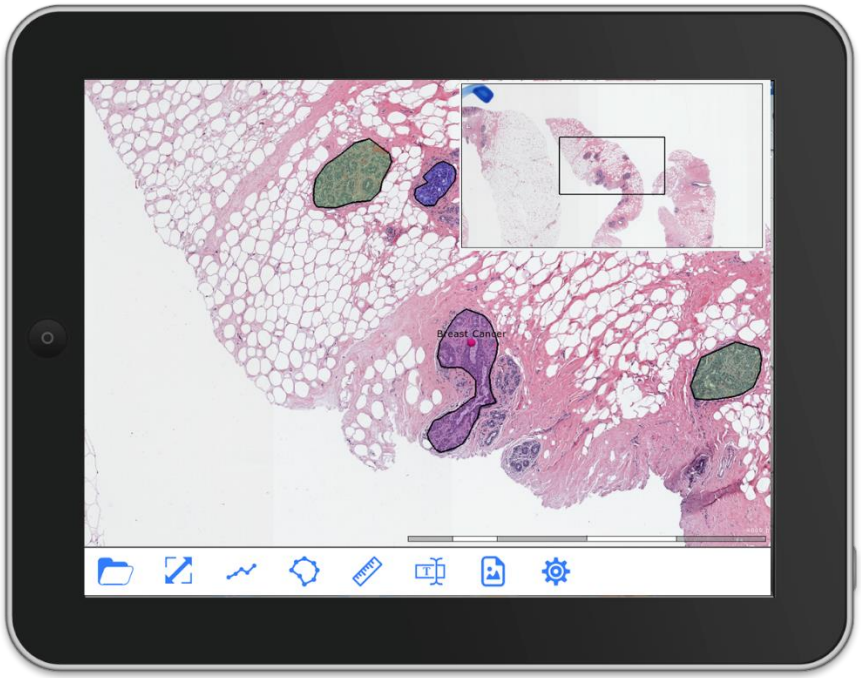

Healthcare: remote access and diagnostics for Doctor and Patients on commodity devices

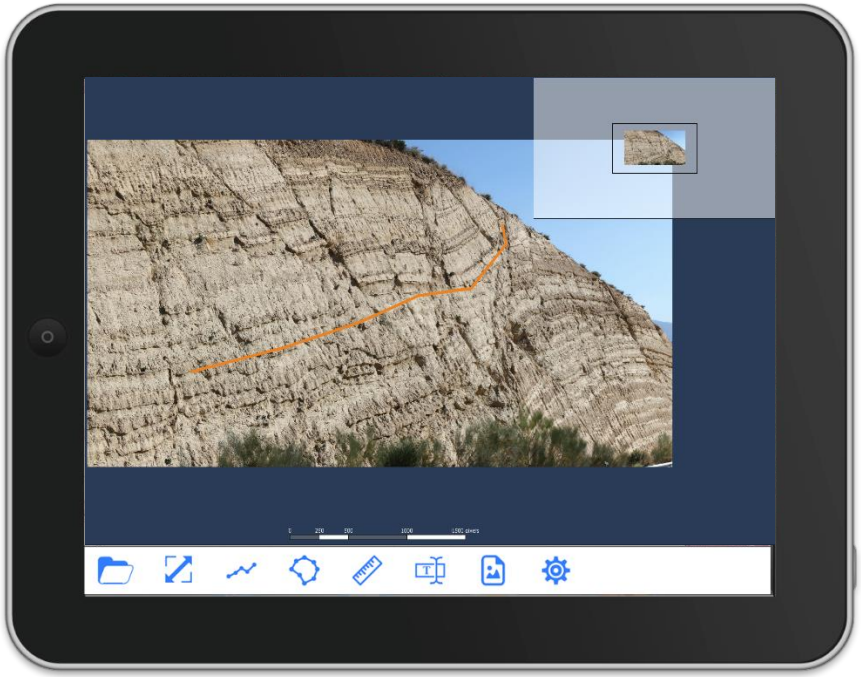

Oil and Gas: support of geologists for fast and reliable identification of Faults and Horizons in reservoirs

Rural Healthcare Needs that the Technology Addresses

Real-Time Access to Diagnostic Tools for Radiology and Histopathology on a Variety of Platforms

Mobile Data Acquisition and Distribution for Providers and Patients

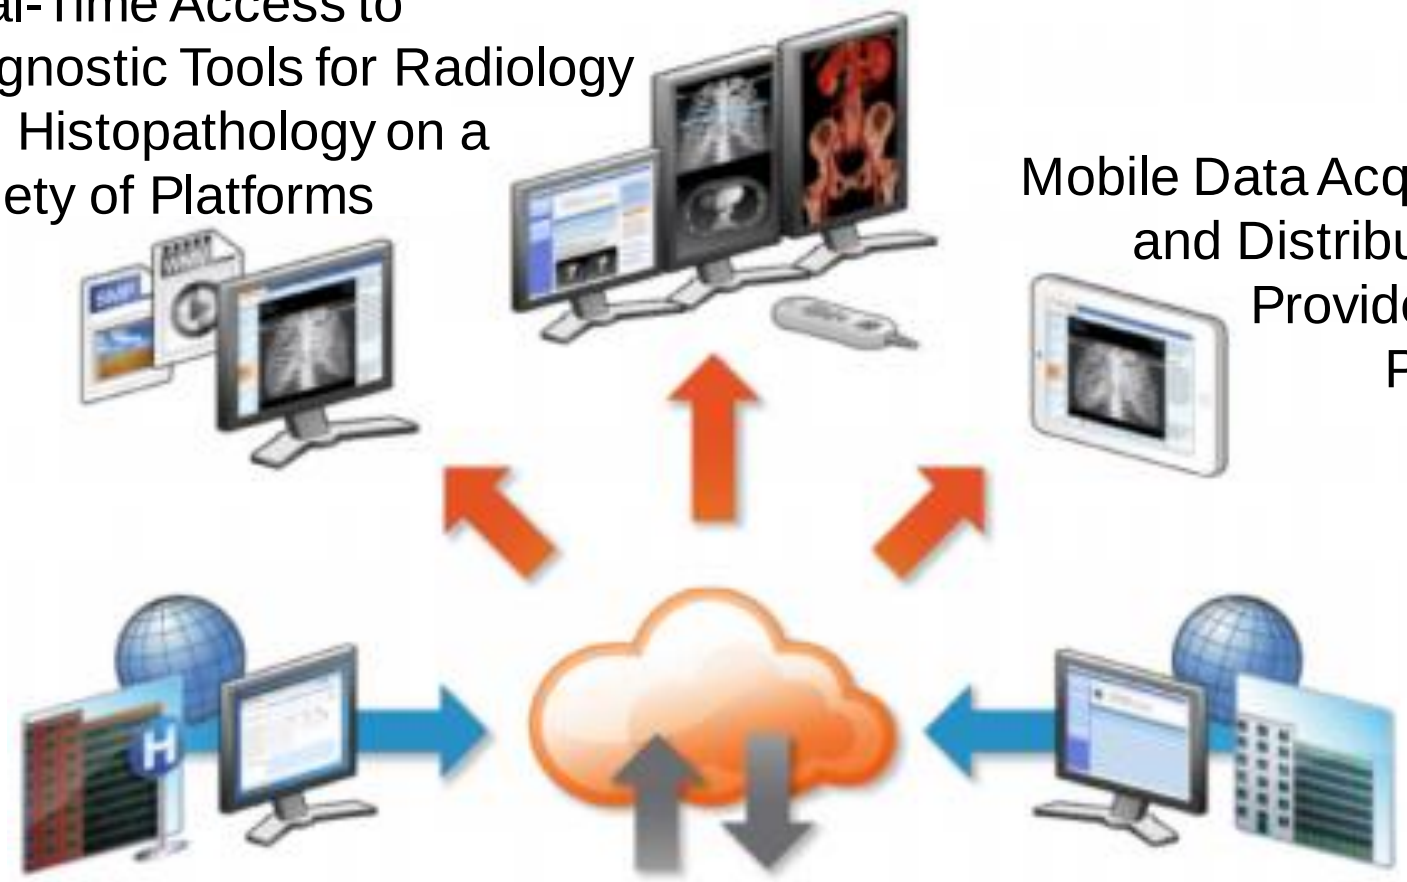

Simplified, purely Web-based access

Simple and Scalable Storage Solutions

Distributed Data Access Services

https://visoar.org/mamografia/

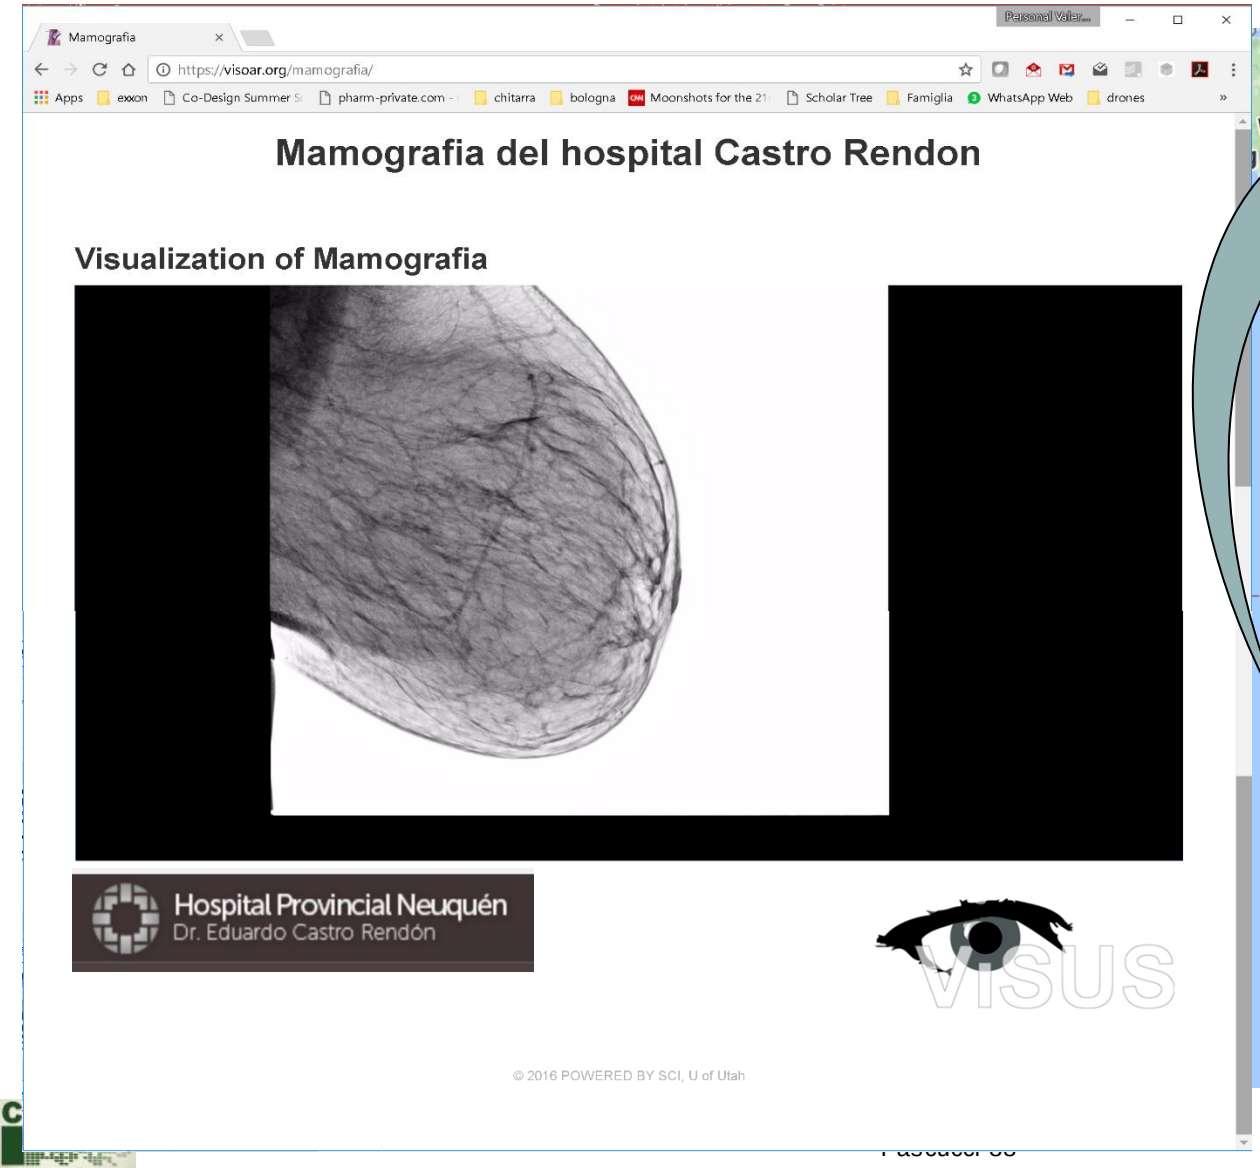

ViSOAR: a Unified Solution for Distribution of Imaging Data in Medical and Geology

Example of Visualization and Annotation of outcrop and medical data

Healthcare: remote access and diagnostics for Doctor and Patients on commodity devices

Mohamed E. Salama, MD, Chief of Hematopathology, Professor of Pathology

ARUP LABORATORIES

Poisson Solver for Image Cloning in Massive Image Collections

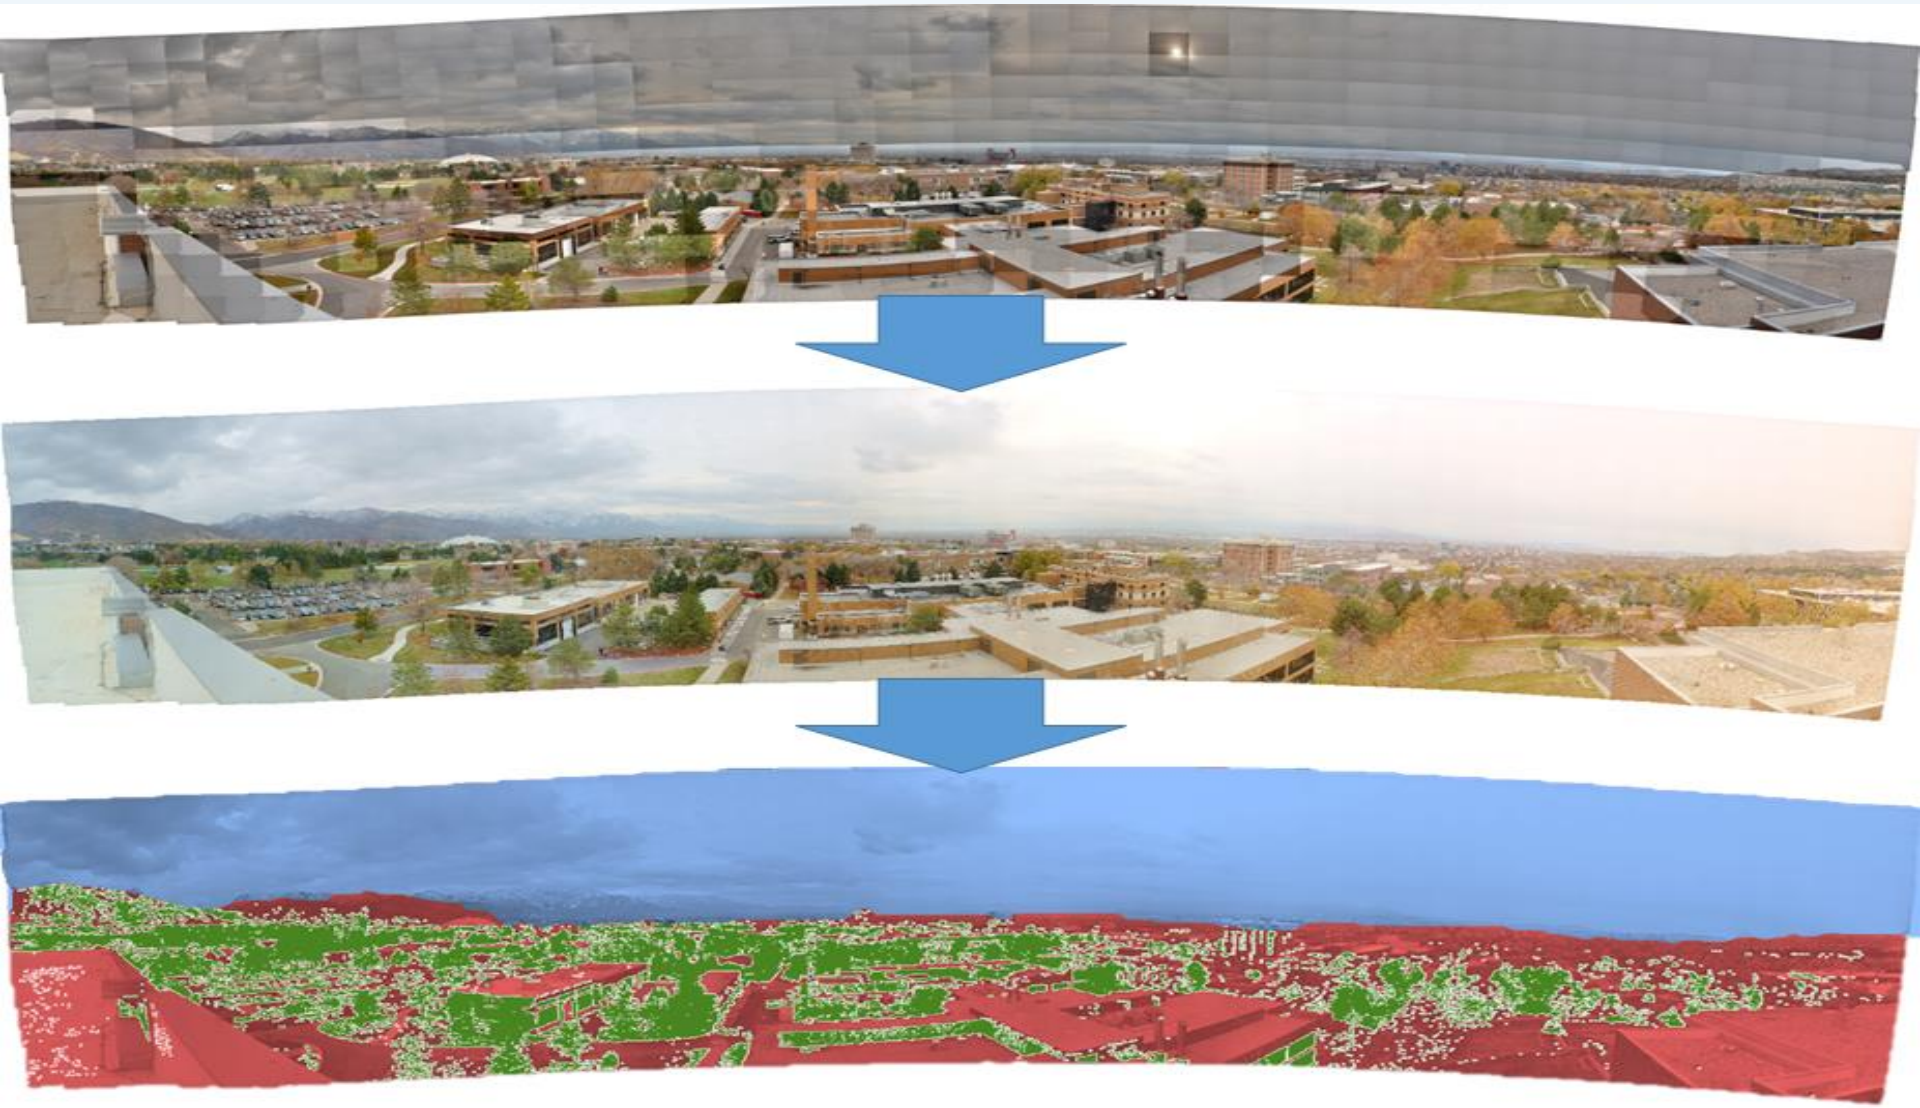

Precision Agriculture Application

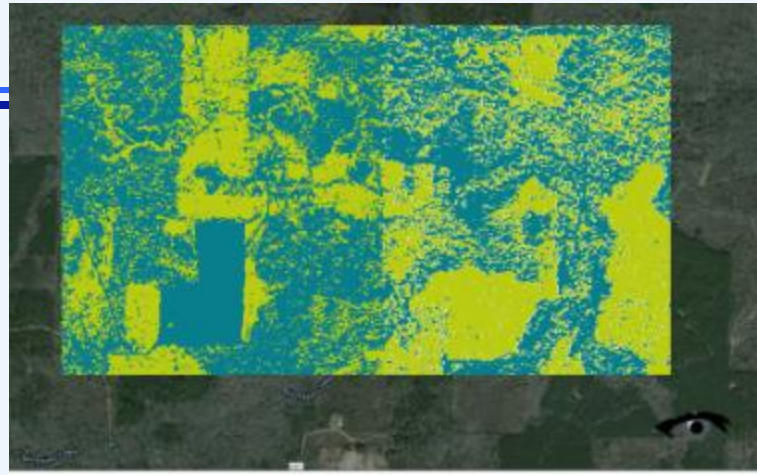

Automation in agriculture propels the need for high-quality data that enables farmers to reduce costs and increase yields

UAV capability to deliver cheap, reliable, high-resolution and *frequent* imagery

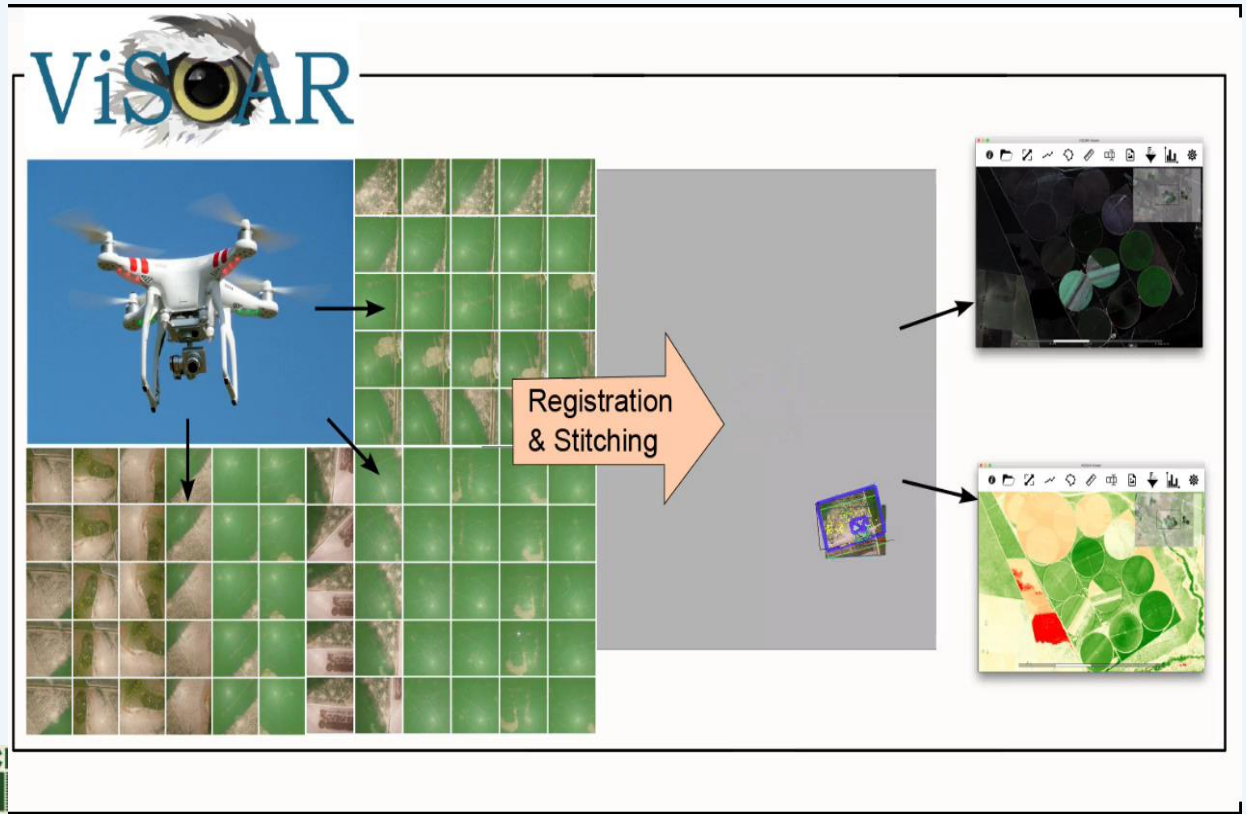

Cyberinfrastructure for Data Management, Analysis, and Visualization Can be a Catalyst for a Virtuous Cycle of Collaborative Activities

- Tight cycle of :
  - basic research,
  - software deployment
  - user support
  - commercialization
- Coordination among many projects:
  - unified techniques for several applications
- University-Lab-Industry collaboration
- Focused technical approaches:
  - performance tools for fast data access
  - general purpose data exploration
  - error bounded quantitative analysis
  - feature extraction and tracking
  - ……..
- Wide Spectrum of Interdisciplinary collaborations:
  - motivating the work
  - formal theoretical approaches
  - feedback to specific disciplines