A Data-Centric Perspective on Scientific Workflows in the Computing Continuum

Silvina Caíno-Lores, PhD
scainolo@utk.edu

THE UNIVERSITY OF TENNESSEE
KNOXVILLE

BIG ORANGE. BIG IDEAS.®

Data Management Across the Computing Continuum Minisymposium, PASC, 26th of June of 2023

Dr. Silvina Caino-Lores

Research Assistant Professor
scaino-lores@acm.org

Research interests

- Cloud computing
- HPC systems and applications
- Big Data and data science
- Workflow systems
- Computational reproducibility
- <u>Converged environments (HPC+BD+AI)</u>
- <u>Scientific computing</u>

Interdisciplinary and collaborative work

- Railway and electric grid infrastructure
- Hydrogeology
- Molecular dynamics
- Protein crystallography
- Urban traffic planning
- Astrophysics
- ML/AI (especially DL)

Acknowledgements

- **Argonne National Laboratory** (B. Nicolae, T. Peterka, O. Yildiz)
- **Cornell University** (M. Cuendet, E. Kots, A. Plante, H. Weinstein)
- **IBM** (Global University Program)
- **OLCF** (Allocation CSC427)
- **NSF** (NSF awards 1741057 and 2103845)
- **University Carlos III of Madrid** (J. Carretero)
- **University of Neuchatel** (P. Kropf, A. Lapin)
- **University of New Mexico** (H. Carrillo-Cabada, T. Estrada, H. Sahni)
- **University of North Texas** (S. Bhowmick)
- **University of Southern California** (E. Deelman, R. Ferreira da Silva, T. Anh Do, L. Pottier, K. Vahi)
- **Syracuse University** (D. Brown)
- **University of Tennessee-Knoxville** (J. Luettgau, I. Lumdsen, J. Marquez, R. Patel, A. Rorabaugh, C. Schuman, M. Taufer, M. Wyatt)

And many more!

Scientific Workflows

Scientific workflows define the sequence of tasks needed for the creation, collection, exploration, exploitation, and preservation of data

- Enable *in silico* science in all disciplines
- Capture complex interactions of tasks and data dependencies
- Typically depend on high-performance computing systems

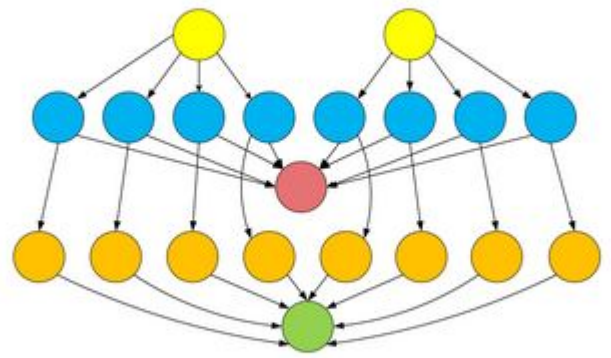

The CyberShake workflow. Credit: Anwar N, Deng H. *A Hybrid Metaheuristic for Multi-Objective Scientific Workflow Scheduling in a Cloud Environment.*

BIG ORANGE BIG IDEAS

Structure of Modern Scientific Workflows

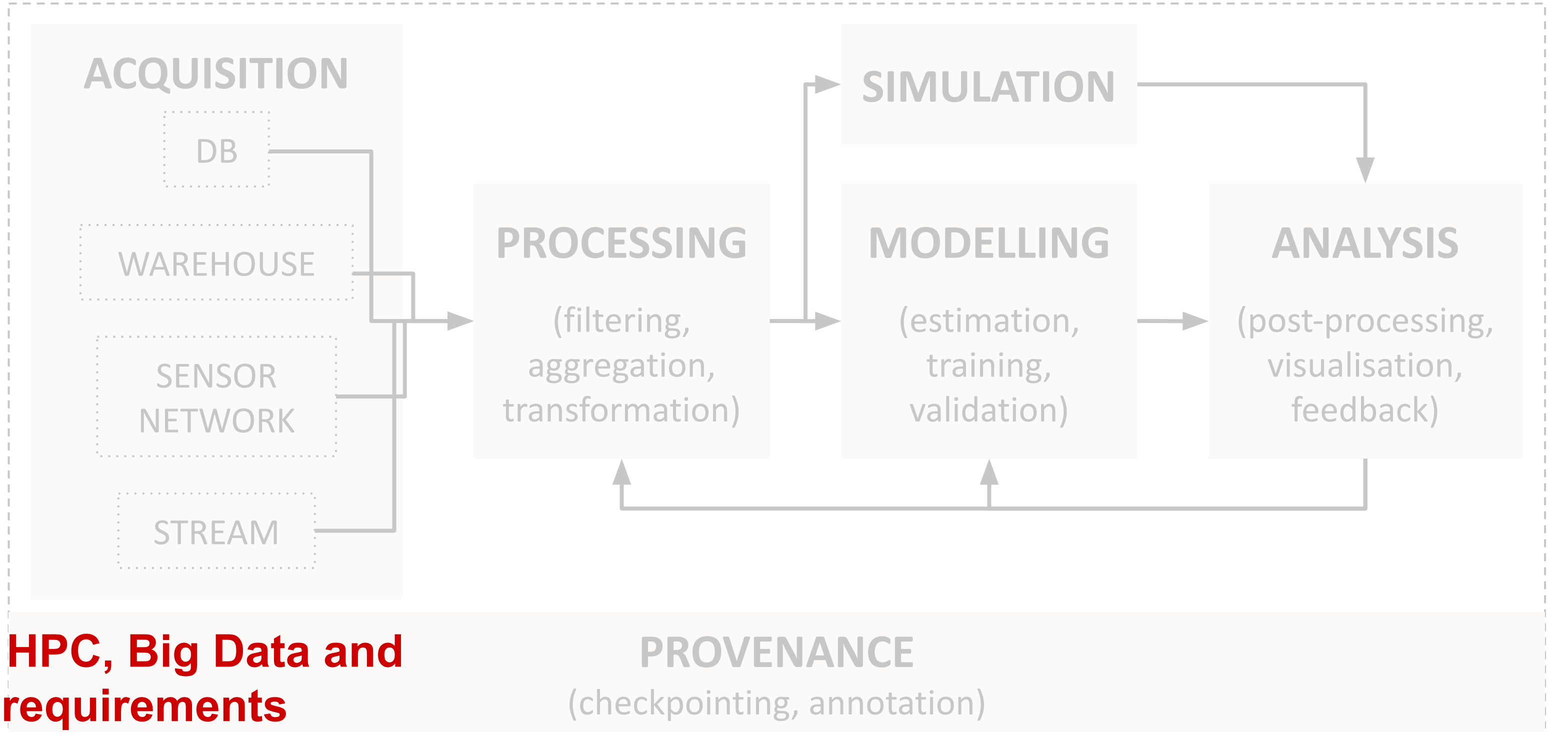

Mixed HPC, Big Data and ML/AI requirements

Structure of Modern Scientific Workflows

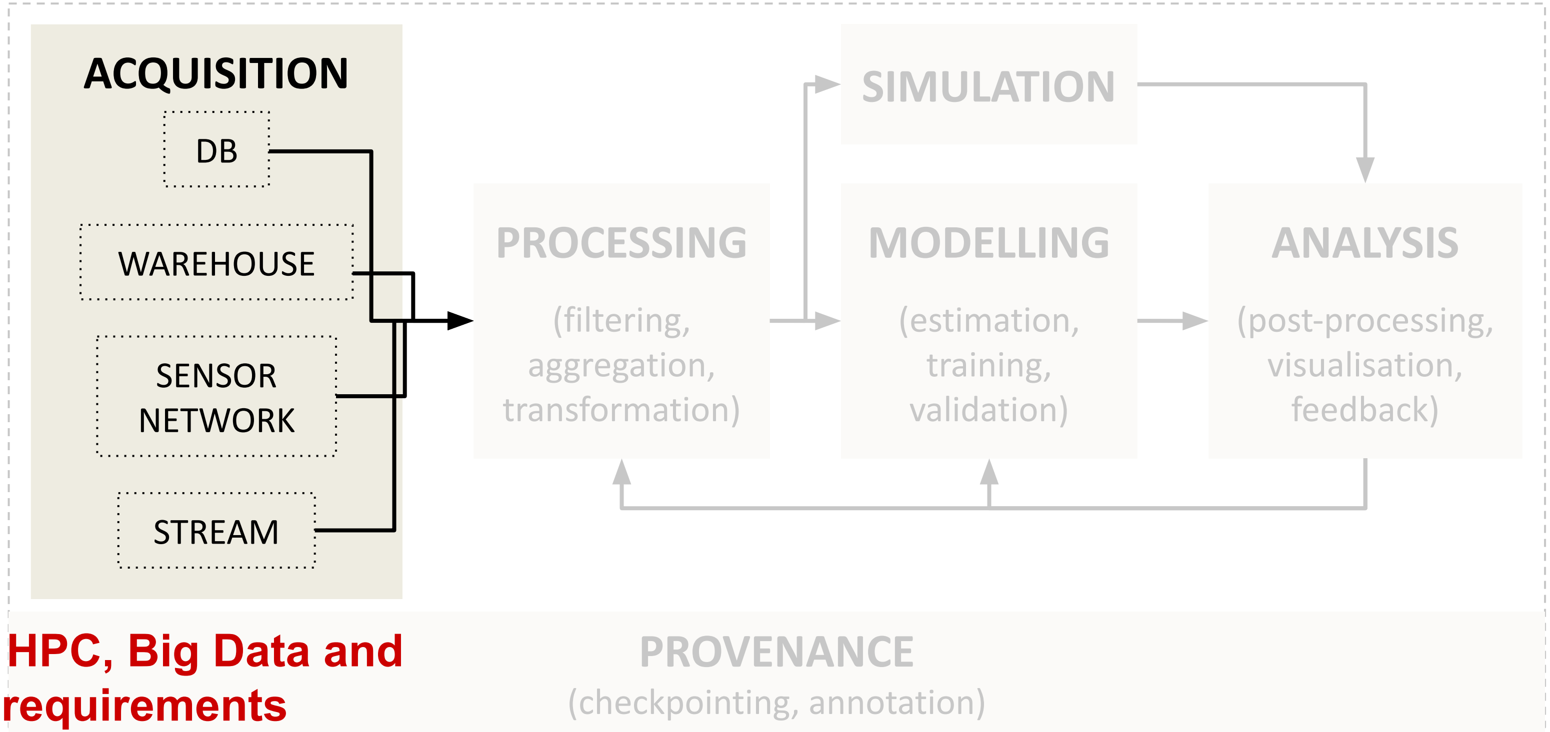

Mixed HPC, Big Data and ML/AI requirements

Structure of Modern Scientific Workflows

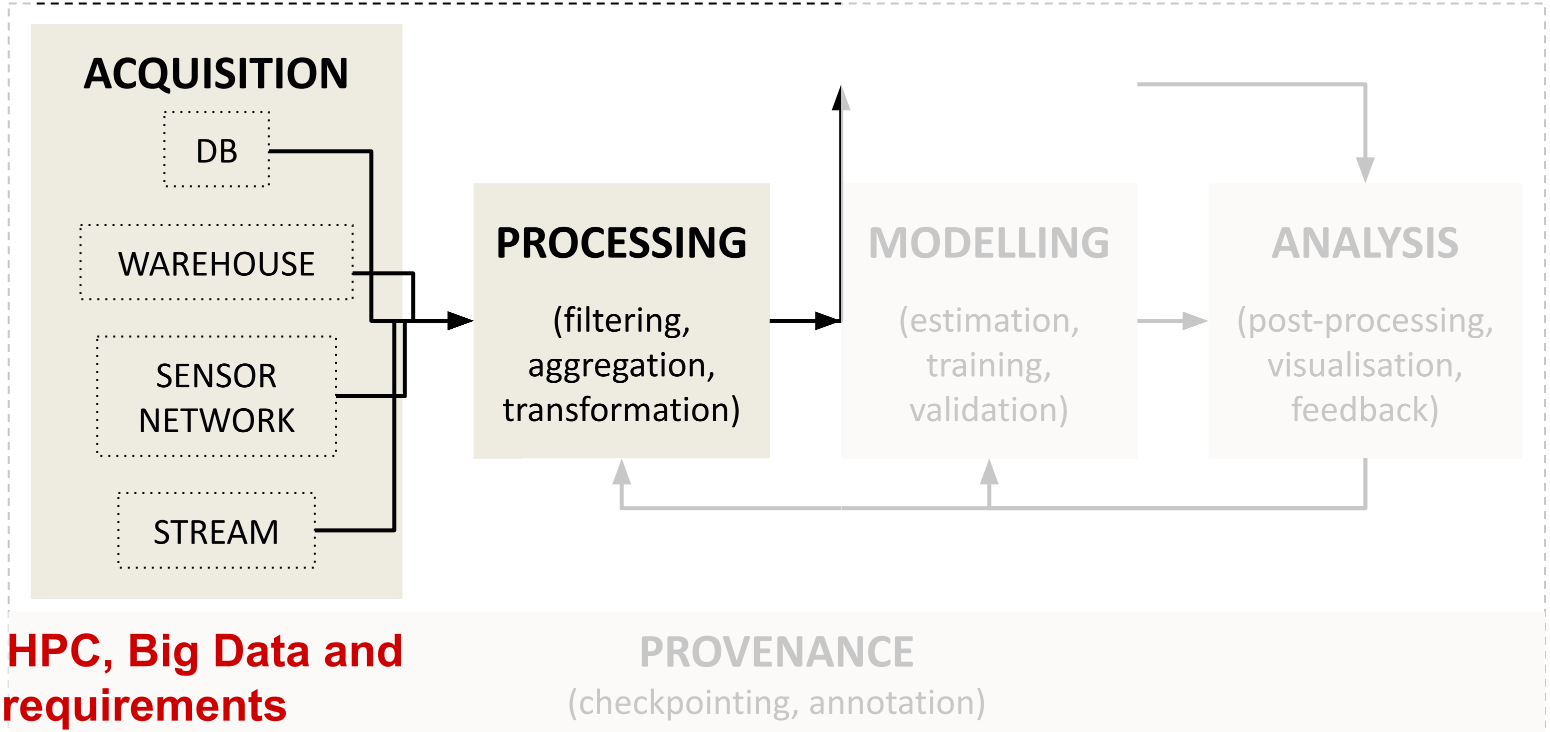

Mixed HPC, Big Data and ML/AI requirements

Structure of Modern Scientific Workflows

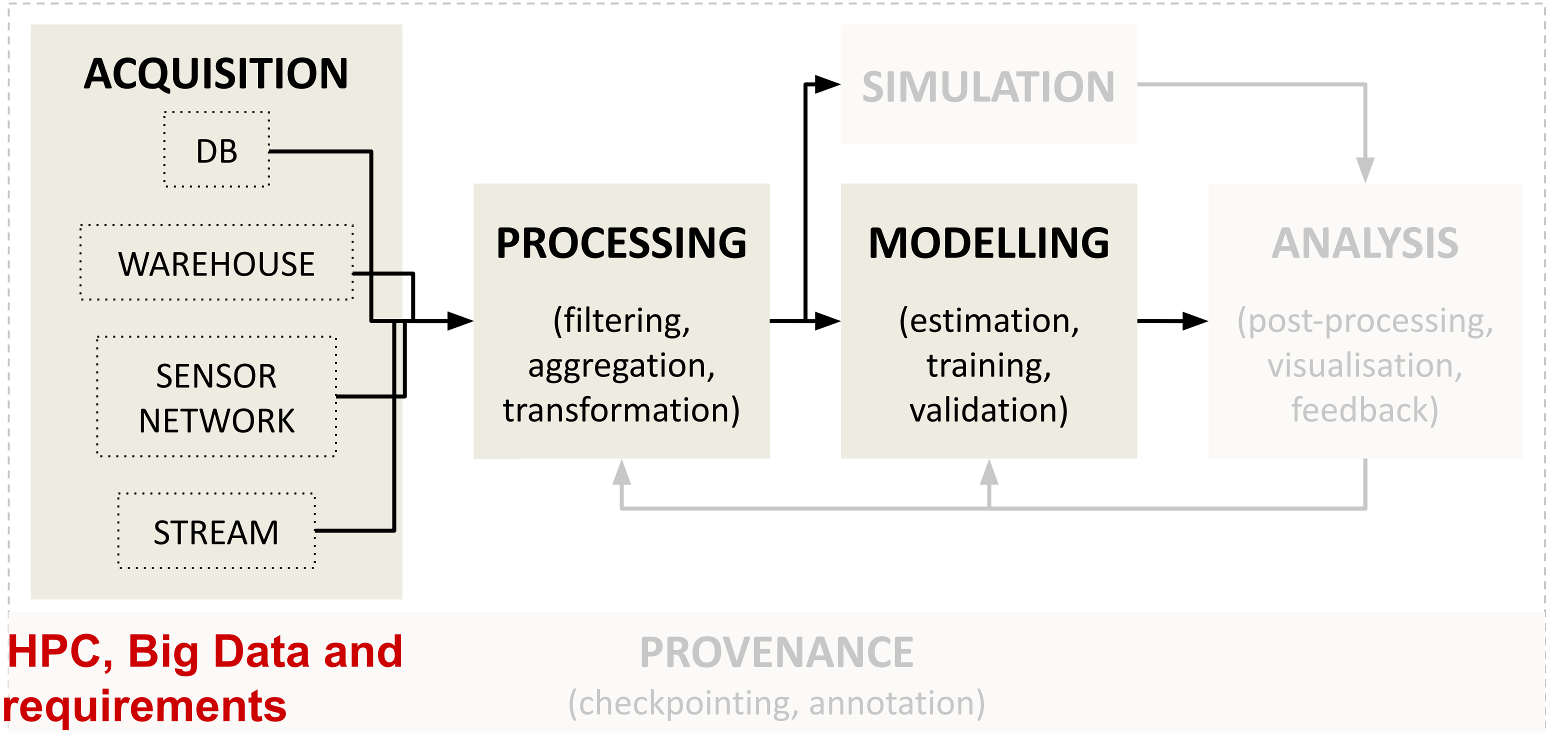

Mixed HPC, Big Data and ML/AI requirements

BIG ORANGE BIG IDEAS

Structure of Modern Scientific Workflows

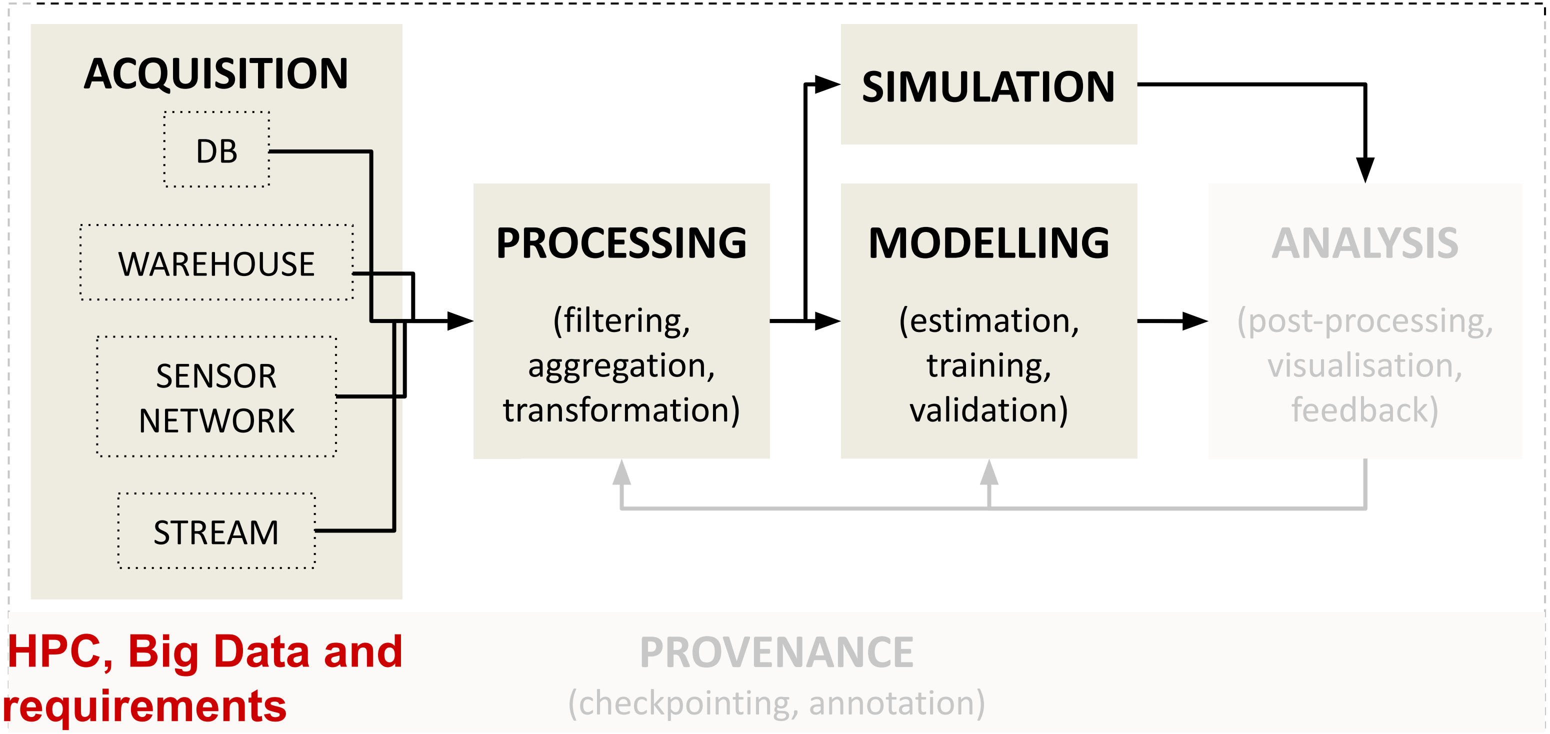

Mixed HPC, Big Data and ML/AI requirements

Structure of Modern Scientific Workflows

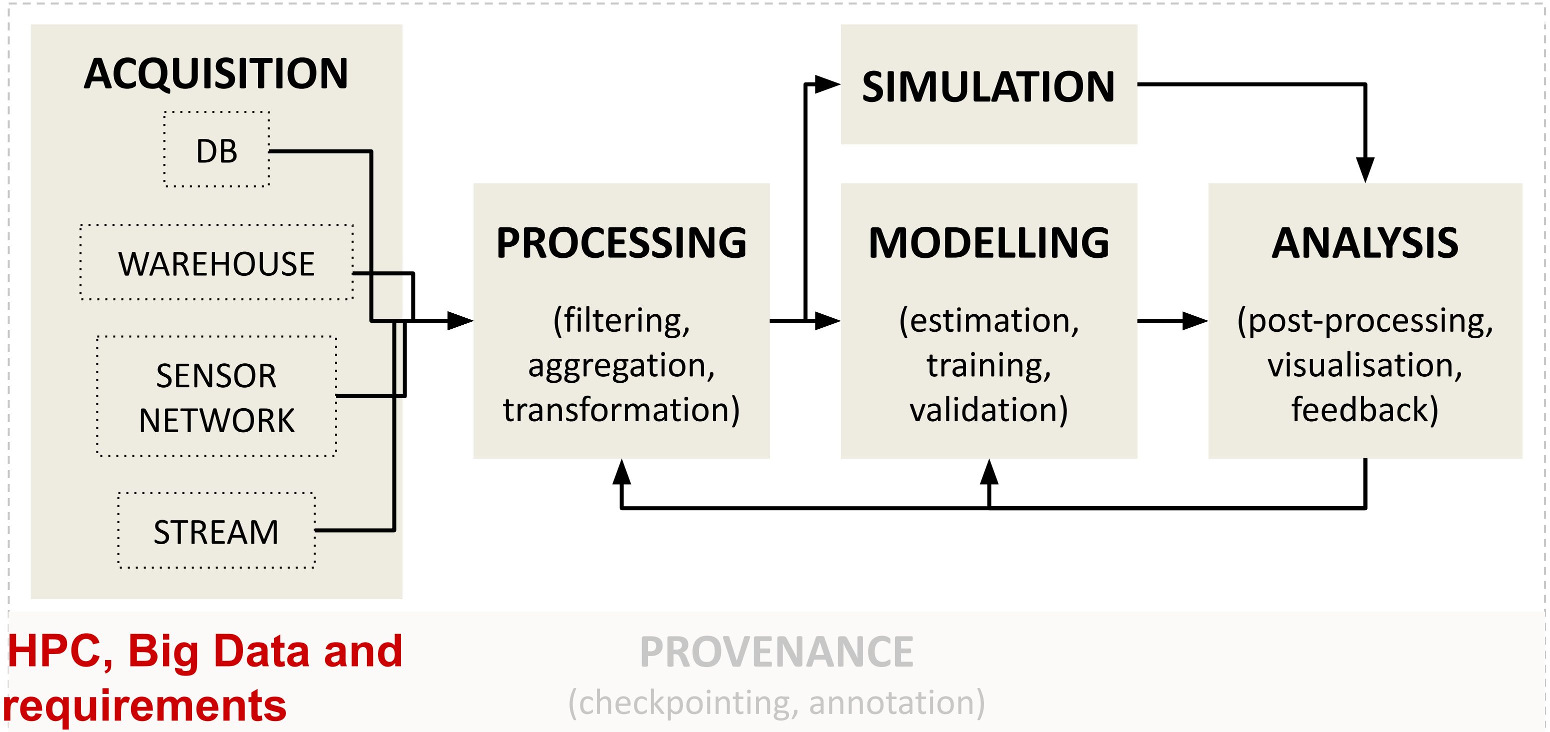

Mixed HPC, Big Data and ML/AI requirements

Structure of Modern Scientific Workflows

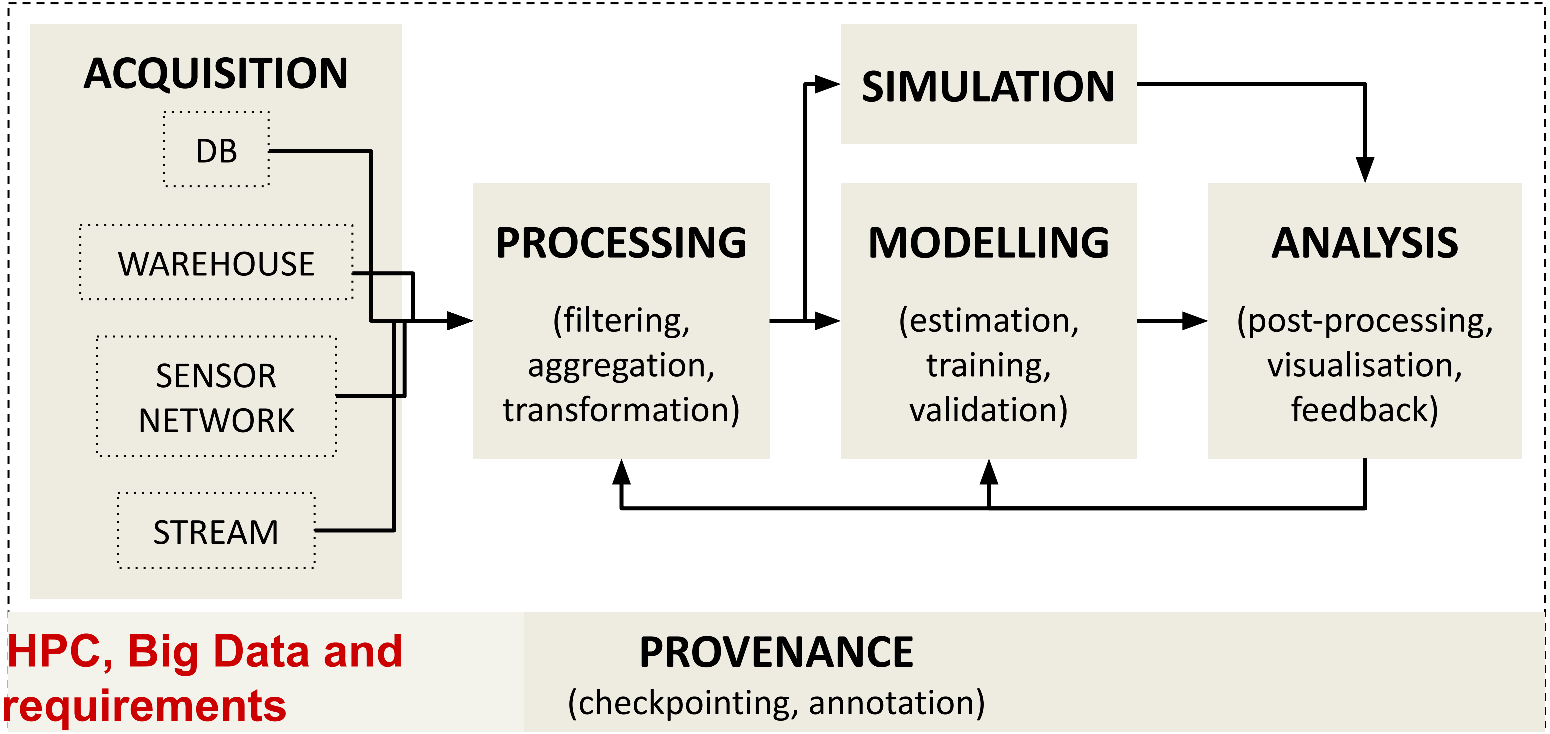

Mixed HPC, Big Data and ML/AI requirements

Scientific Workflows in the Computing Continuum

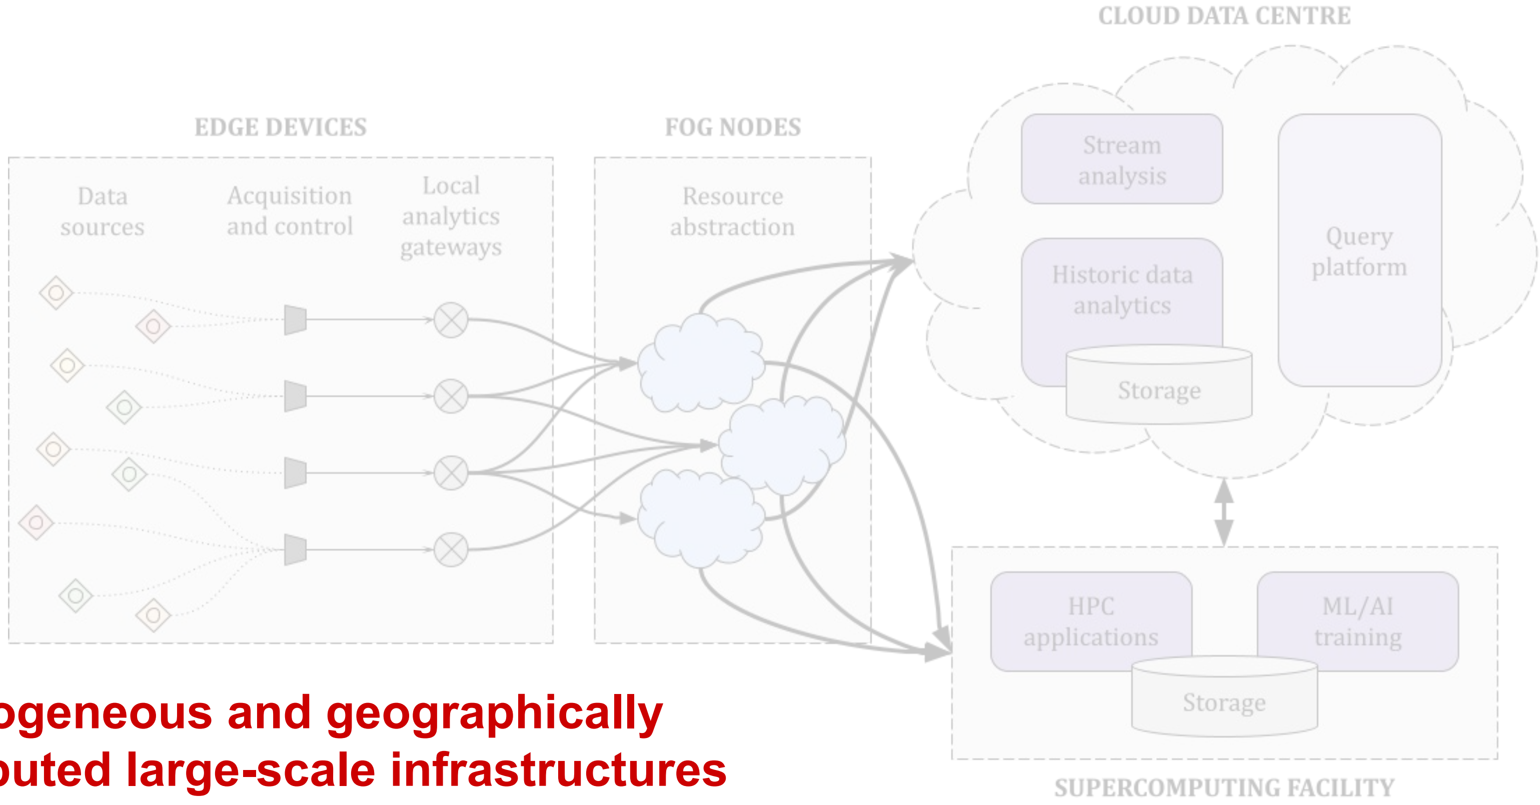

Heterogeneous and geographically distributed large-scale infrastructures

Scientific Workflows in the Computing Continuum

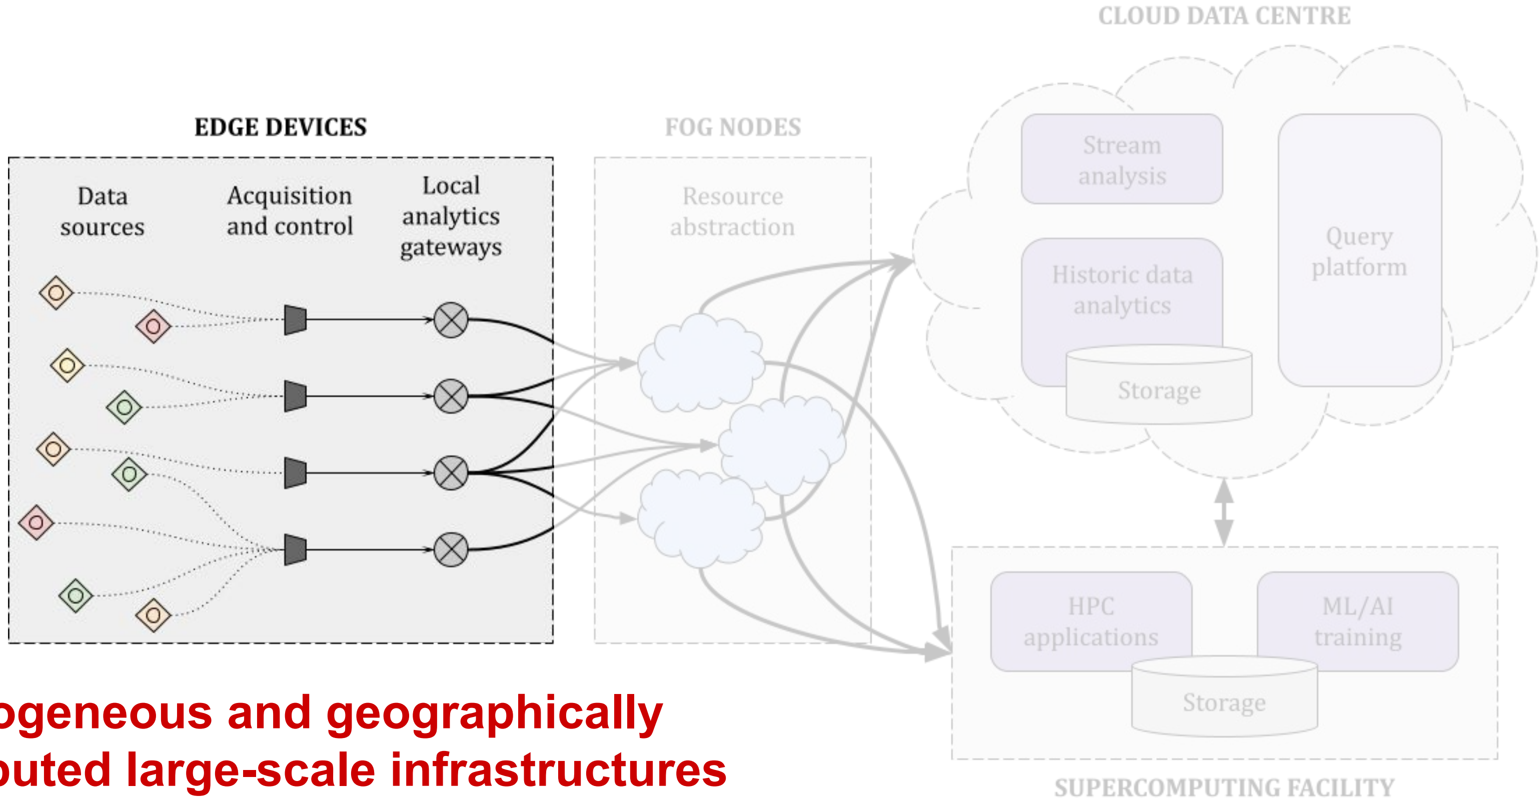

Heterogeneous and geographically distributed large-scale infrastructures

Scientific Workflows in the Computing Continuum

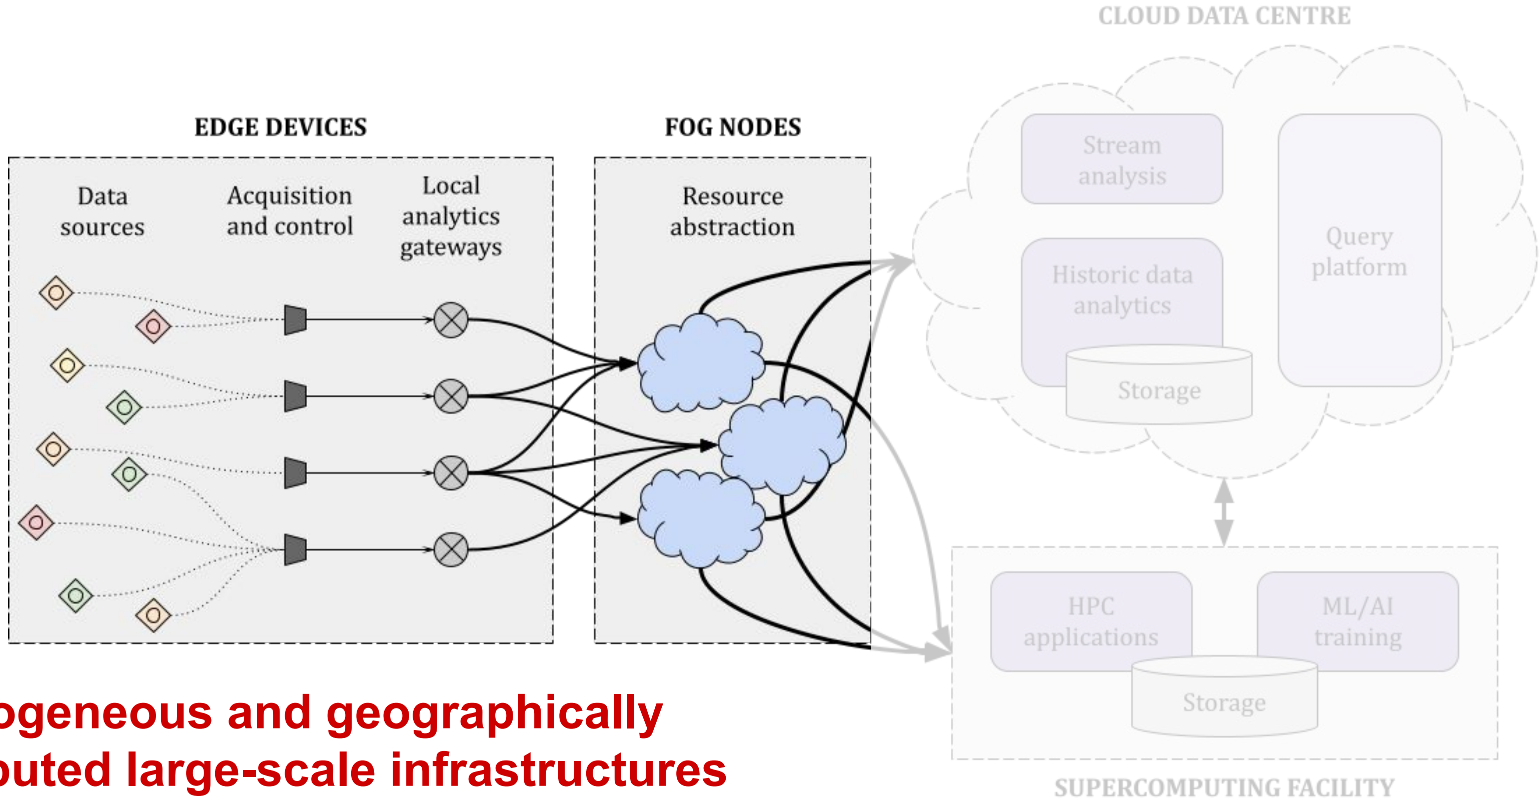

Heterogeneous and geographically distributed large-scale infrastructures

Scientific Workflows in the Computing Continuum

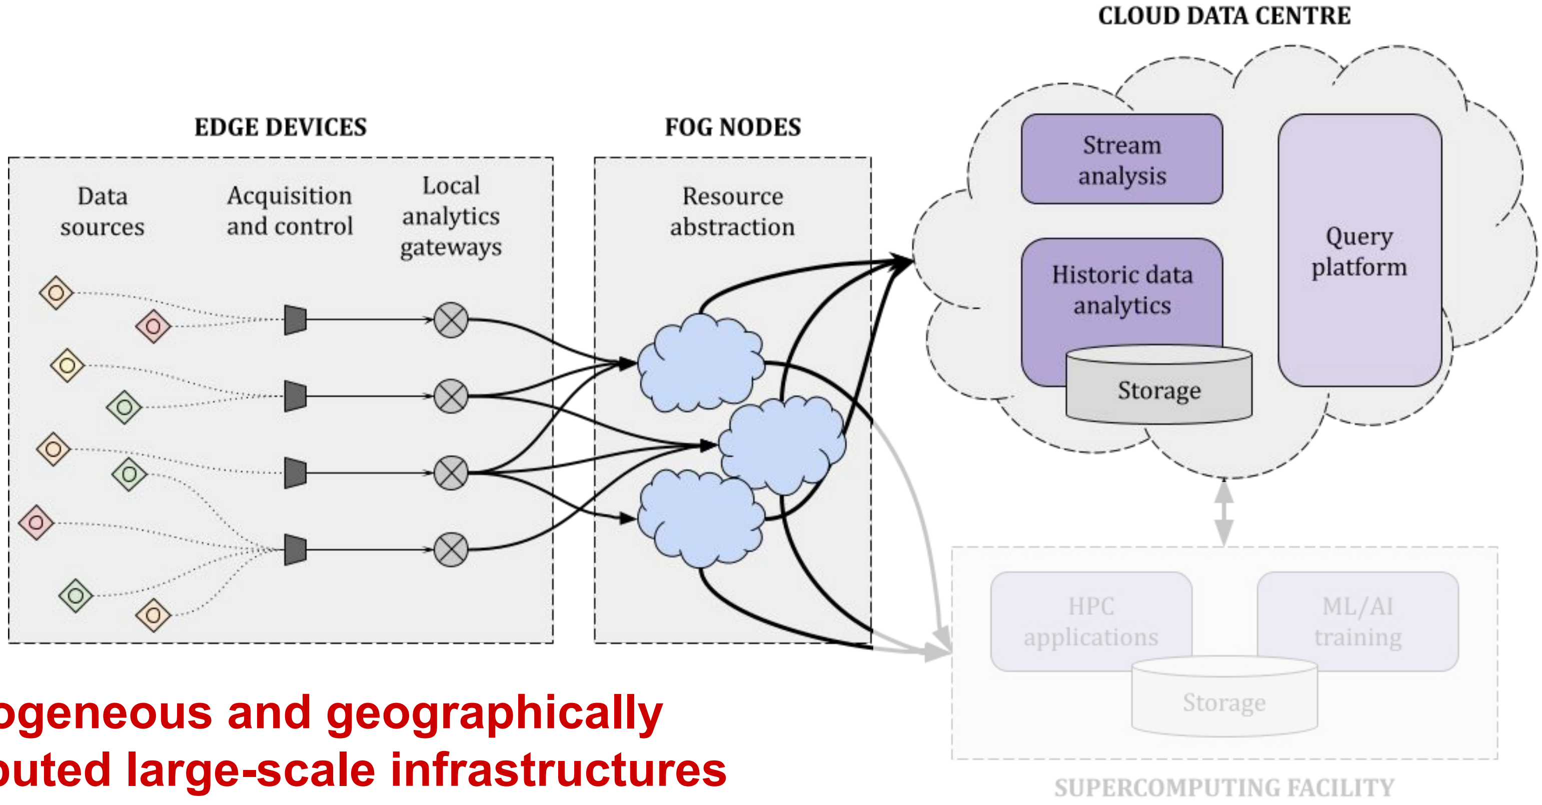

Heterogeneous and geographically distributed large-scale infrastructures

Scientific Workflows in the Computing Continuum

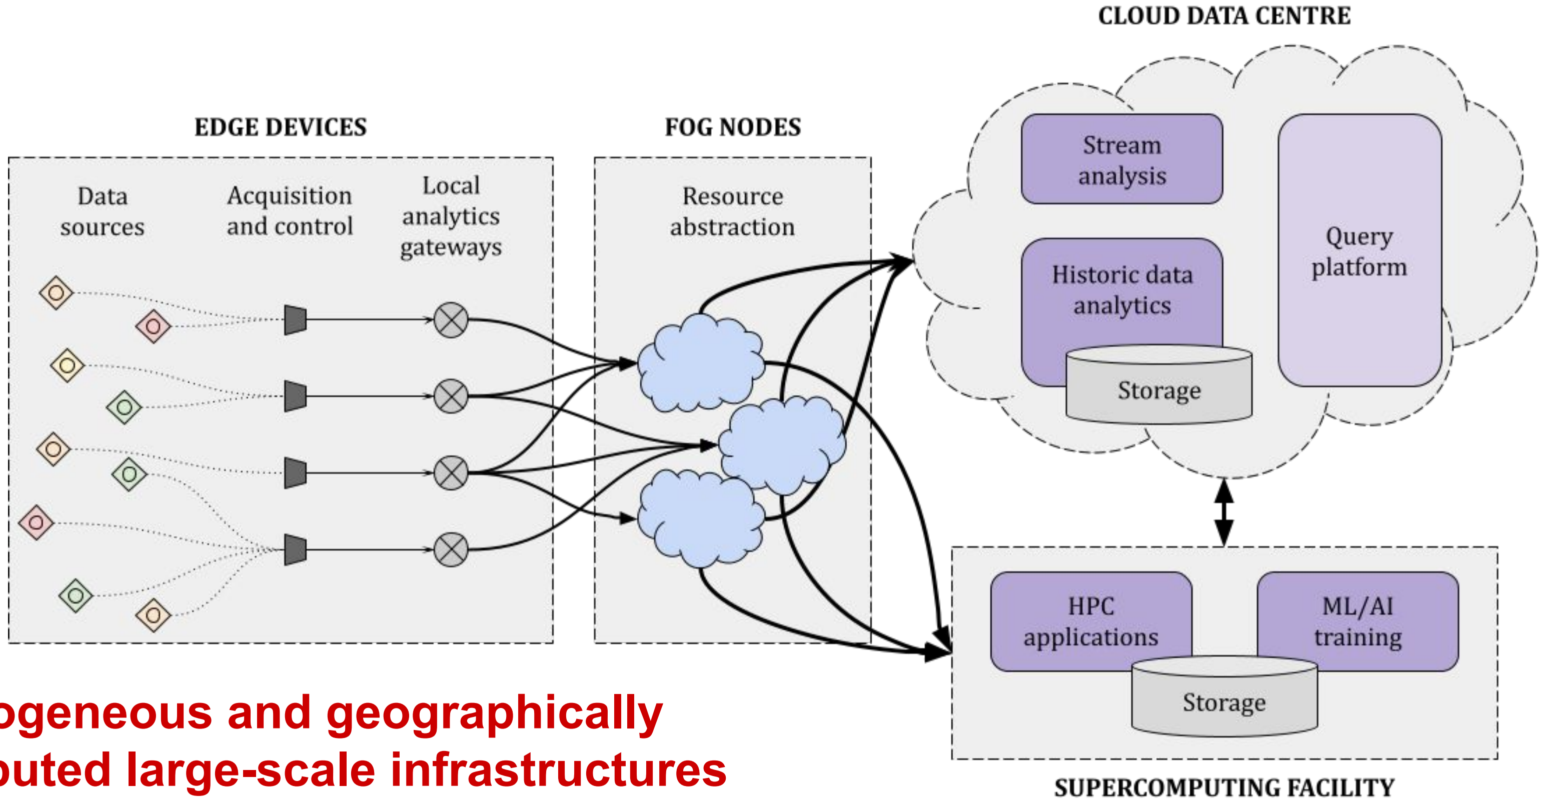

Heterogeneous and geographically distributed large-scale infrastructures

Data Oriented Challenges in the Continuum

Heterogeneity in hardware, platforms and applications lead to challenges in:

- Interoperating different programming and data models
- Representing data in a unified manner
- Placing and transferring data efficiently
- Supporting provenance and reproducibility

Data Oriented Challenges in the Continuum

Heterogeneity in hardware, platforms and applications lead to challenges in:

- Interoperating different programming models
- Representing data in a unified manner
- Placing and transferring data efficiently
- Supporting provenance and reproducibility

Not a exhaustive list!

Data Oriented Challenges in the Continuum

Heterogeneity in hardware, platforms and applications lead to challenges in:

- Interoperating different programming models
- Representing data in a unified manner
- Placing and transferring data efficiently
- Supporting provenance and reproducibility

Not a exhaustive list!

How can we converge these diverse ecosystems without losing their respective benefits?

BIG ORANGE BIG IDEAS

Case Studies:
Geographically Distributed Water Forecasting

Hydrogeological Simulations for Water Forecasting

- Water resource management requires fast and reliable knowledge of very complex hydrogeological systems
- Relies on real-time stochastic simulation of water profiles
  - Multiscale nonlinear processes and matrix operations from multi-physics models
  - Input data from several geographically distributed sensors
  - Severe deadlines (forecasts, status reports for emergencies)

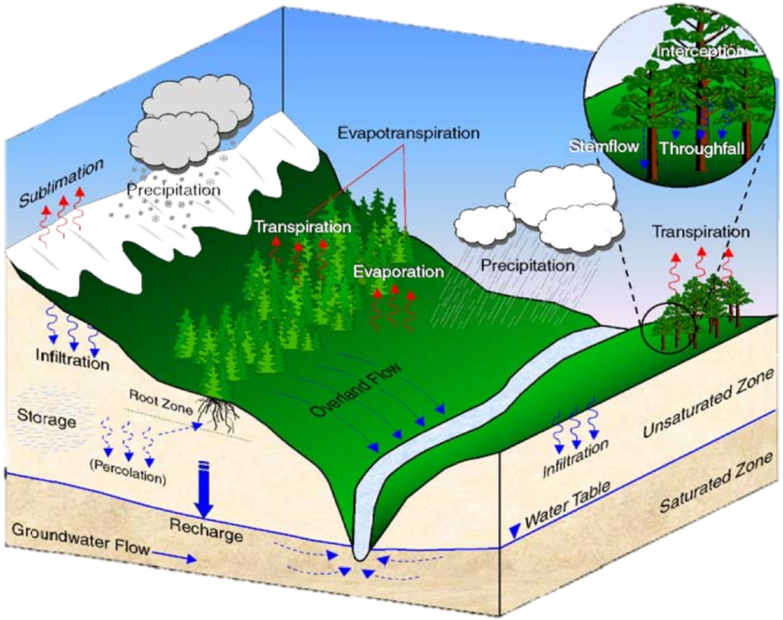

Hydrogeological Simulations for Water Forecasting

HPC has increased the scale and complexity of the simulations:

- Multicore systems
- Distributed cluster computing
- Grid-like technologies

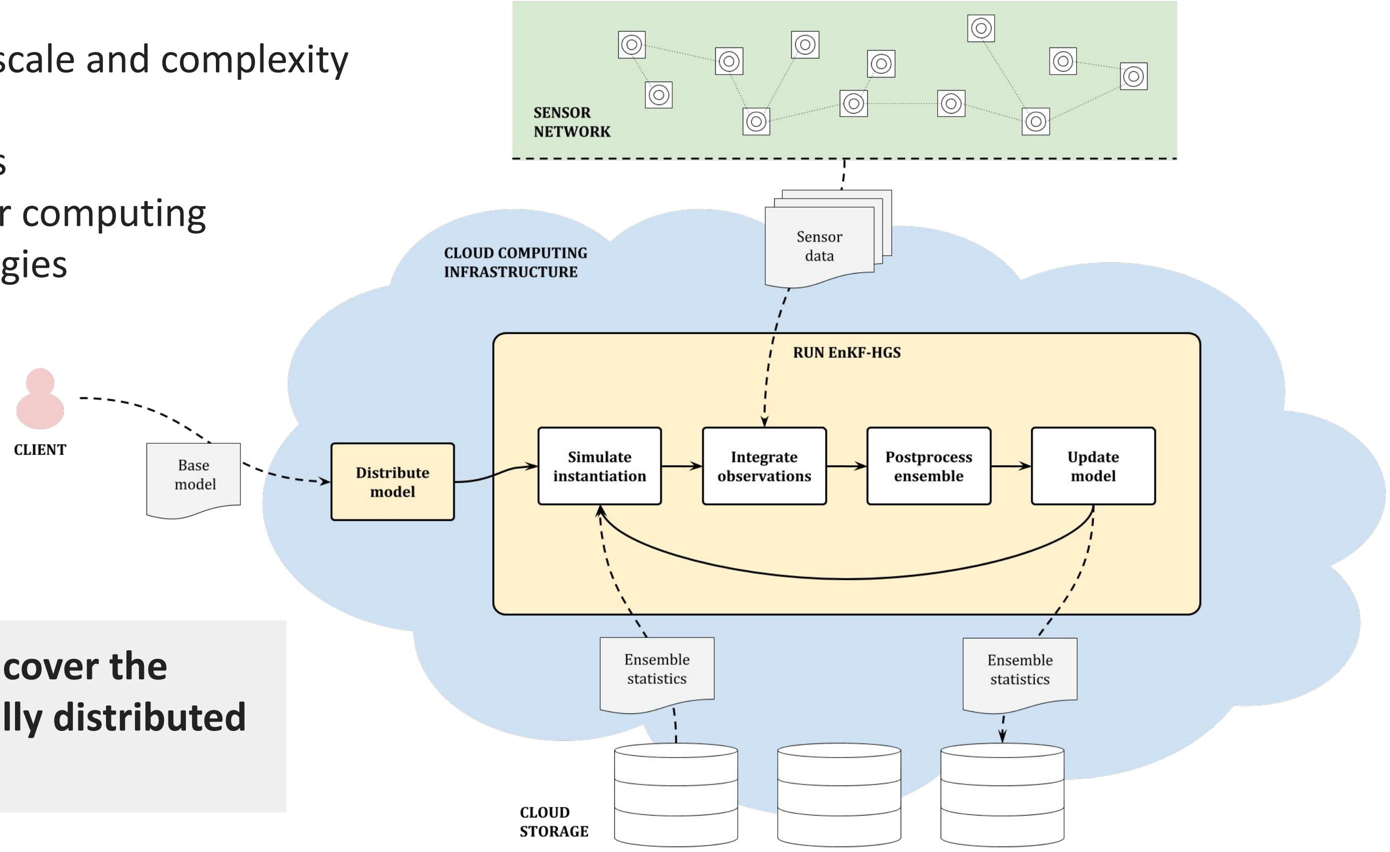

But HPC insufficient to cover the needs of a geographically distributed sensor network

Data-Centric Transformation of an HPC Application

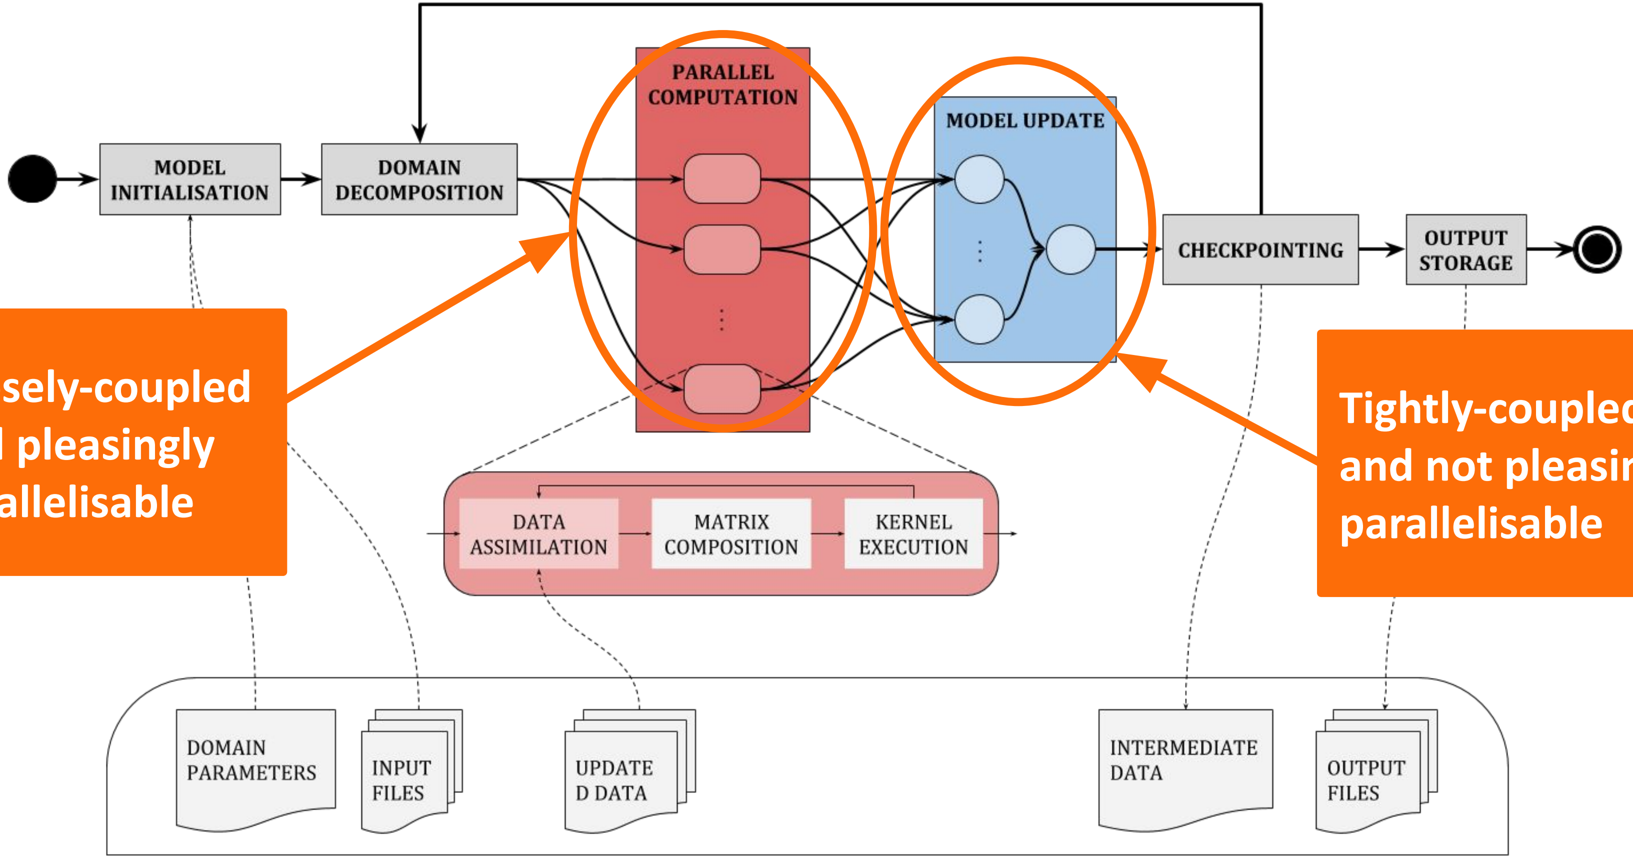

Loosely-coupled and pleasingly parallelisable

Tightly-coupled and not pleasingly parallelisable

Caíno-Lores, S., Carretero, J., Nicolae, B., Yildiz, O., & Peterka, T. (2019). Toward High-Performance Computing and Big Data Analytics Convergence: The Case of Spark-DIY. *IEEE Access*, *7*, 156929-156955.

A Generalist Architecture for HPC-BDA Applications

- Enable transparent access to existing BDA and HPC features
- Expose a unified data abstraction and operational model
- Build on existing runtimes
- Allow process-centric workloads to interact with BDA platforms and infrastructures

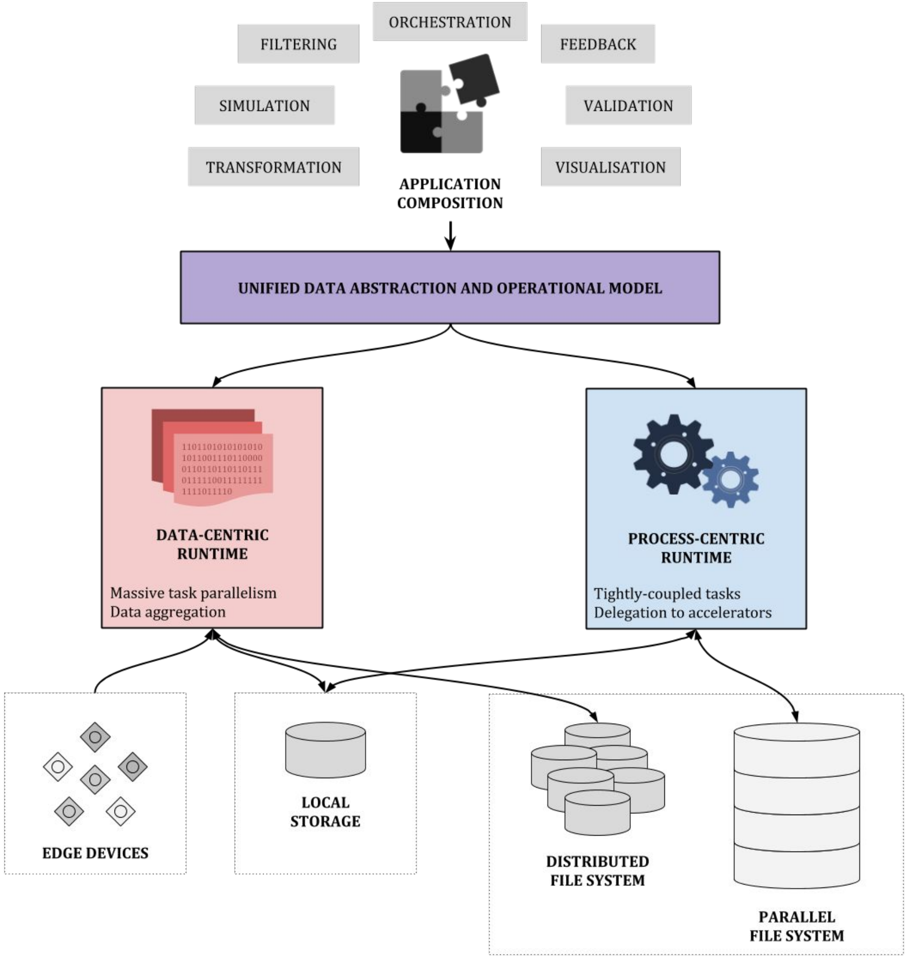

A Generalist Architecture for HPC-BDA Applications

- **Enable transparent access to existing BDA and HPC features**
- Expose a unified data abstraction and operational model
- Build on existing runtimes
- Allow process-centric workloads to interact with BDA platforms and infrastructures

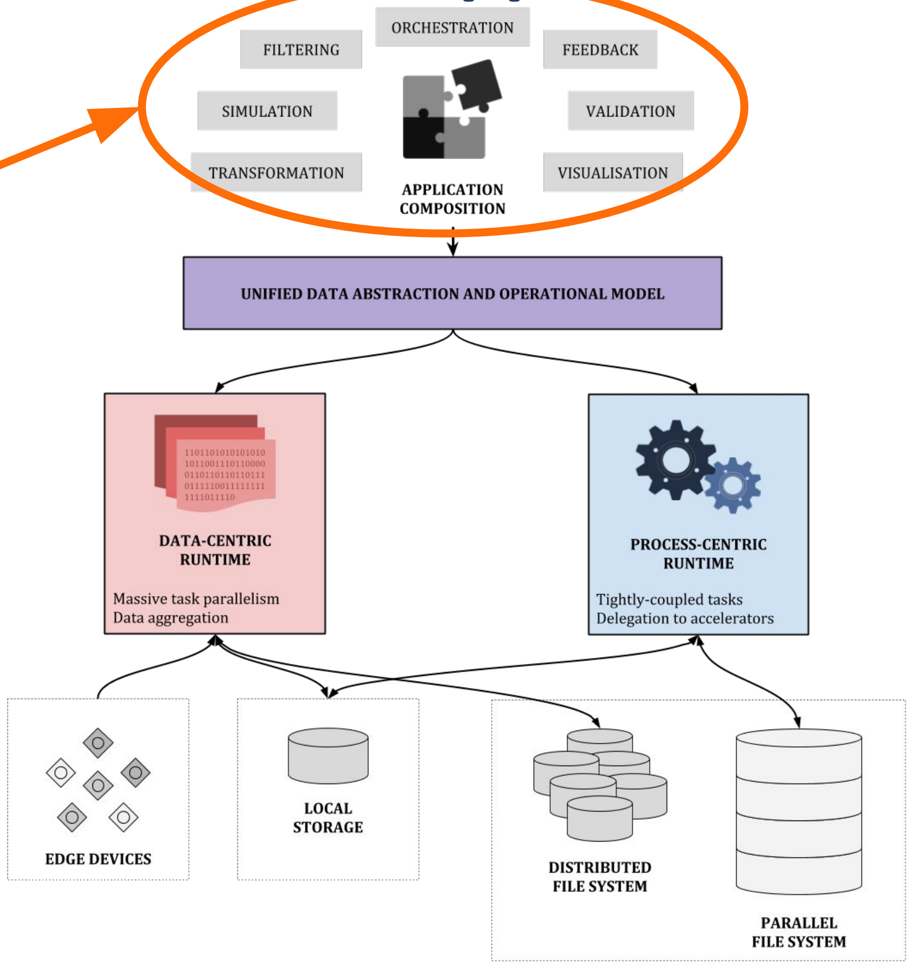

A Generalist Architecture for HPC-BDA Applications

- Enable transparent access to existing BDA and HPC features
- **Expose a unified data abstraction and operational model**
- Build on existing runtimes
- Allow process-centric workloads to interact with BDA platforms and infrastructures

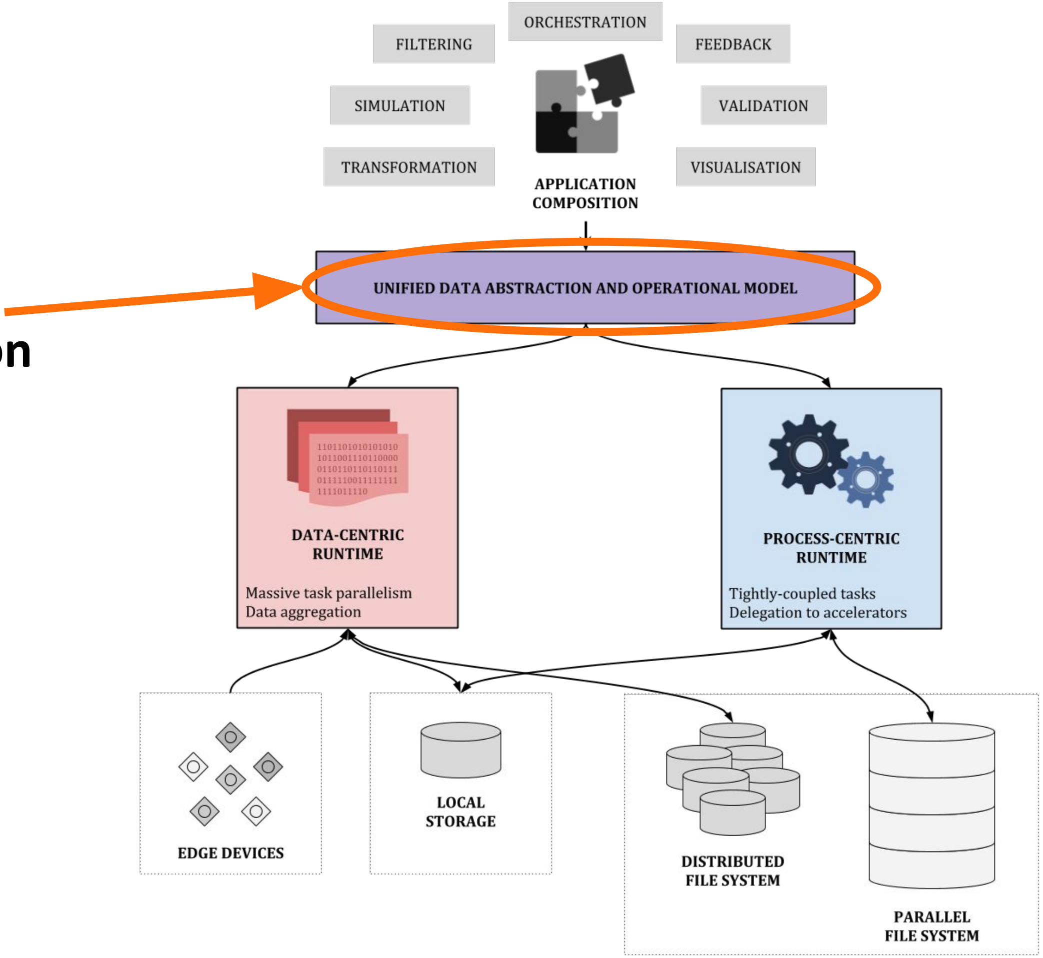

A Generalist Architecture for HPC-BDA Applications

- Enable transparent access to existing BDA and HPC features
- Expose a unified data abstraction and operational model
- **Build on existing runtimes**
- Allow process-centric workloads to interact with BDA platforms and infrastructures

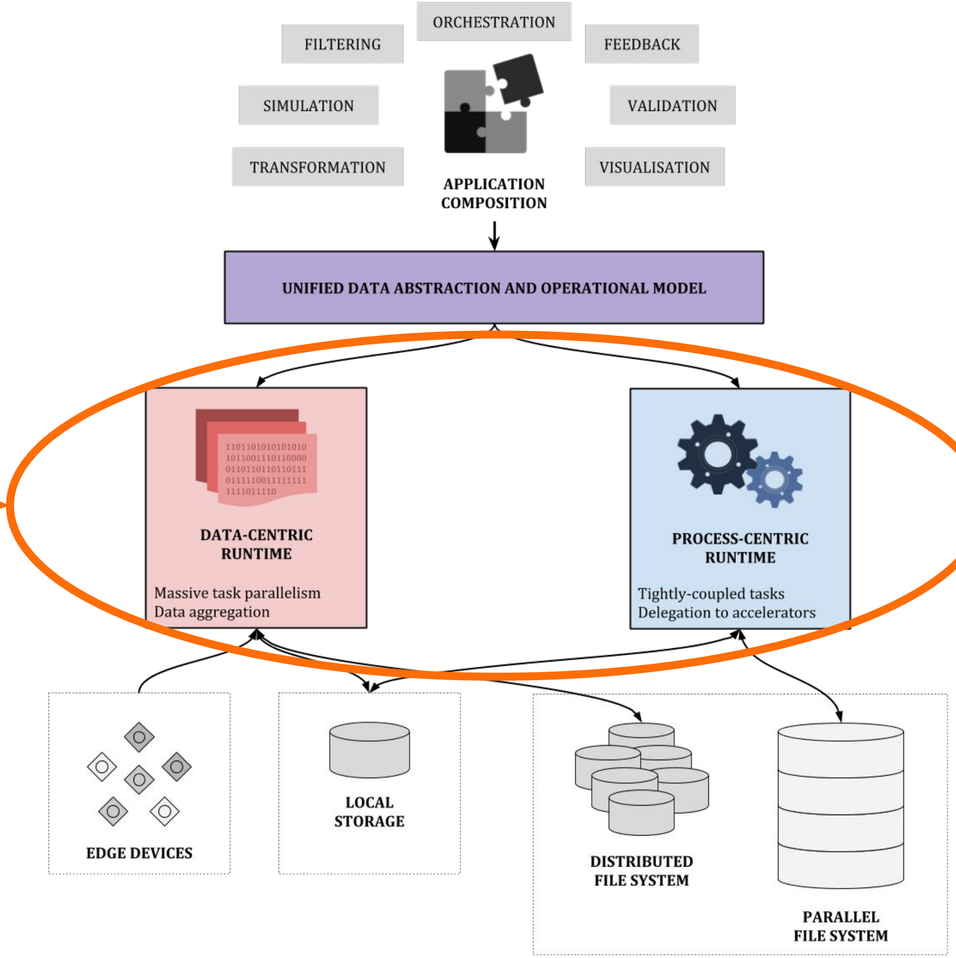

A Generalist Architecture for HPC-BDA Applications

- Enable transparent access to existing BDA and HPC features
- Expose a unified data abstraction and operational model
- Build on existing runtimes
- **Allow process-centric workloads to interact with BDA platforms and infrastructures**

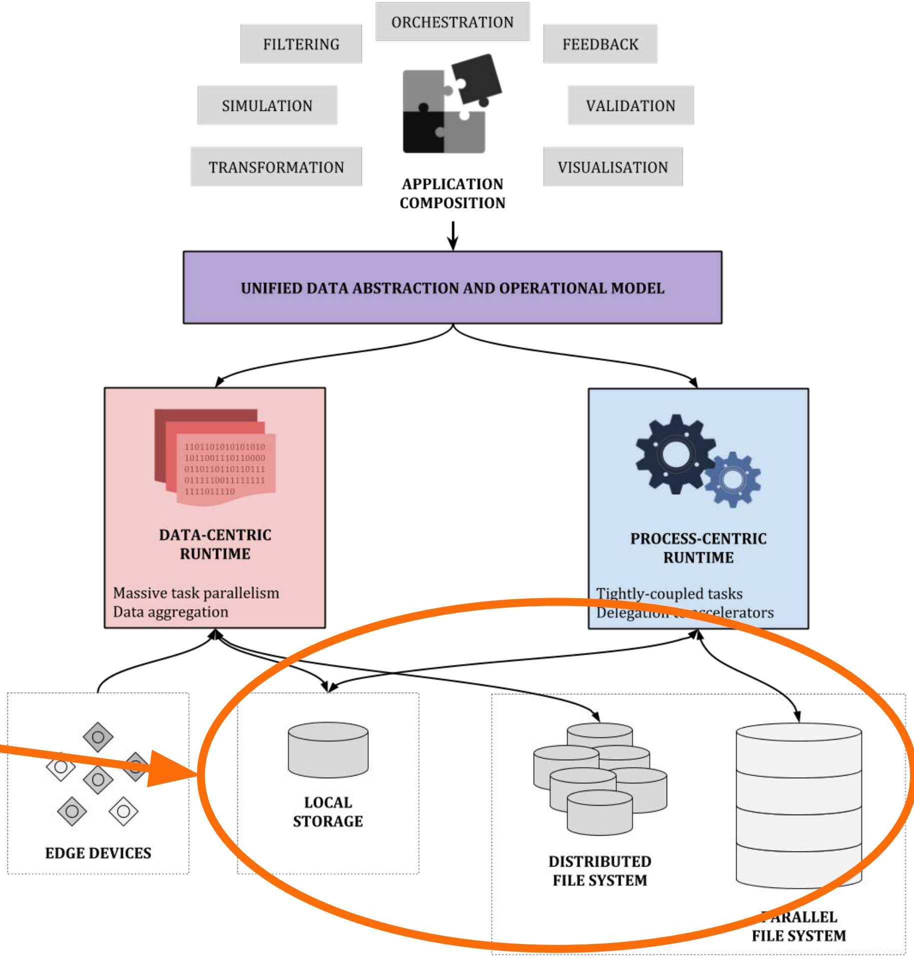

Case Study:
Massively Parallel In Situ Analysis of Molecular Dynamics Simulations

Classical Molecular Dynamics (MD) Simulations

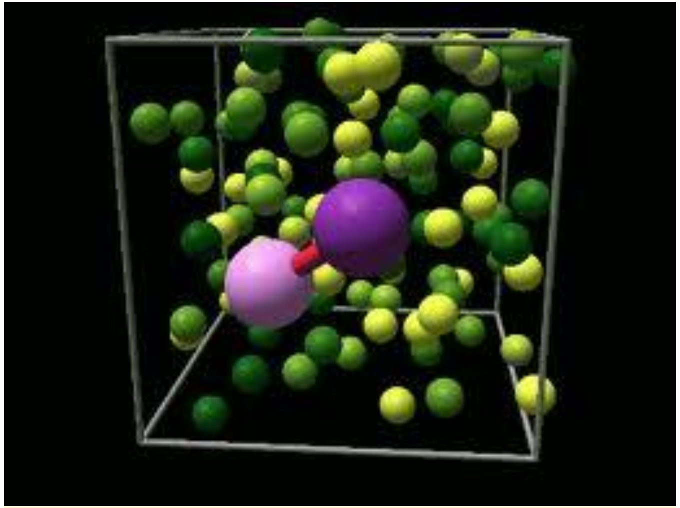

Forces on single atoms
↳ Acceleration
↳ Velocity
↳ Position

1. MD step computes **forces** on single atoms (e.g., bond, dihedrals, nonbond)
2. Forces are added to compute **acceleration**
3. Acceleration is used to update **velocities**
4. Velocities are used to update the **atom positions**
5. Every *N* steps (stride)
   ☐ ***Store 3D snapshot or frame***

Classical Molecular Dynamics (MD) Simulations

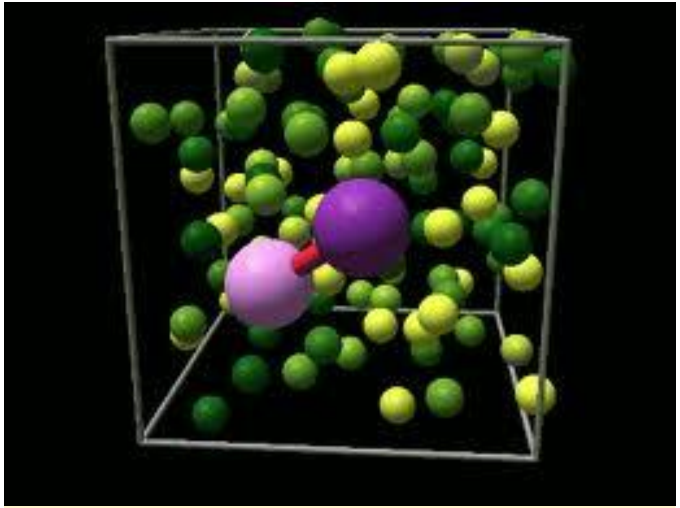

Forces on single atoms
↳ Acceleration
↳ Velocity
↳ Position

1. MD step computes **forces** on single atoms (e.g., bond, dihedrals, nonbond)
2. Forces are added to compute **acceleration**
3. Acceleration is used to update **velocities**
4. Velocities are used to update the **atom positions**

Data generation

5. Every *N* steps (stride)
   □ ***Store 3D snapshot or frame***

Scientist-Driven Analysis of MD Trajectories

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

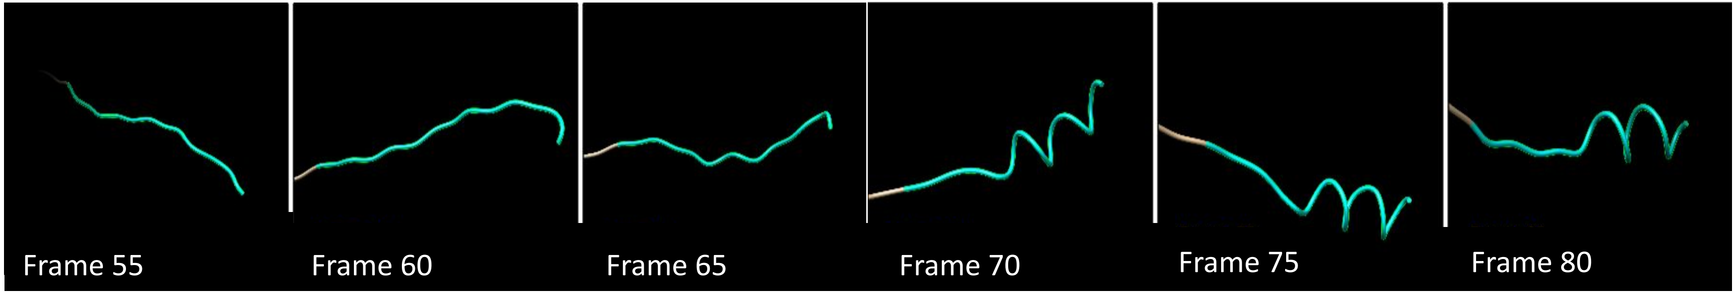

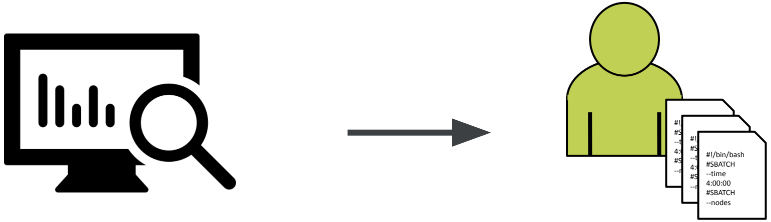

Visualization

Scientist

[1]Travis Johnston, Buyu Zhang, Adam Liwo, Silvia Crivelli, and Michela Taufer. *"In-Situ Data Analysis and Indexing of Protein Trajectories,"* JCC 2017.

Scientist-Driven Analysis of MD Trajectories

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

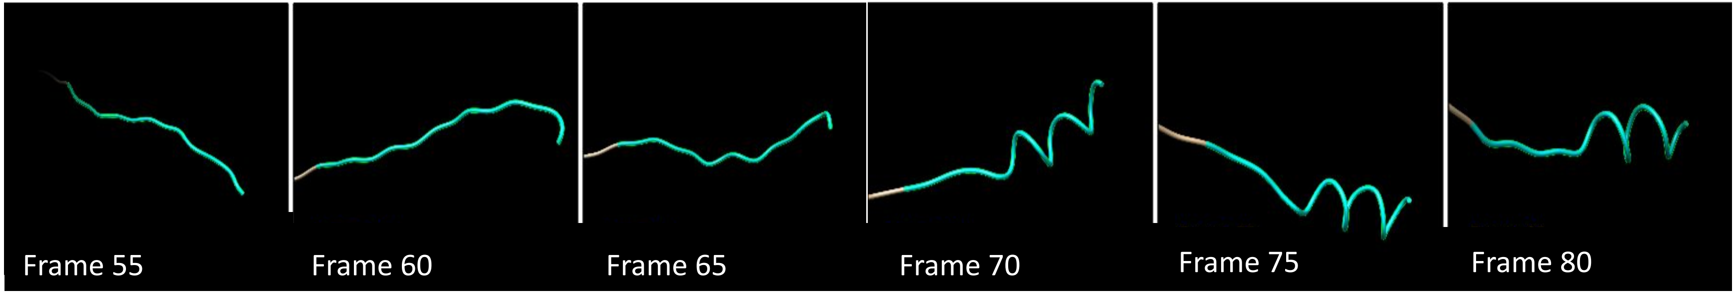

Simulation and analysis are isolated!

Visualization

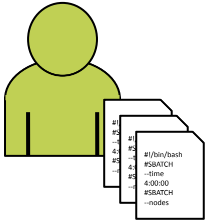

Scientist

[1]Travis Johnston, Buyu Zhang, Adam Liwo, Silvia Crivelli, and Michela Taufer. *"In-Situ Data Analysis and Indexing of Protein Trajectories,"* JCC 2017.

From HPC to BDA MD Simulations

An holistic approach that co-locates simulation and analysis can benefit from

- Natural integration with data streaming
- Massive task parallelism
- In-memory storage

However, we must find a way to integrate MD simulations with the analytics

In Situ Analysis of MD Trajectories: A4MD

See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

Run n-Stride simulation steps

MD code
(e.g., GROMACS)

Plumed

See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

Run n-Stride simulation steps

MD code
(e.g., GROMACS)

Plumed

ML-inferred algorithms Collective variables

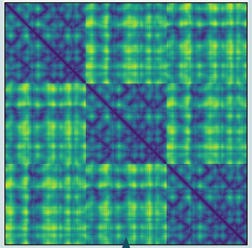

Analytics representations and algorithms

See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

Run n-Stride simulation steps

MD code
(e.g., GROMACS)

Plumed

In-memory Staging Area
(DataSpaces)

ML-inferred algorithms
Collective variables

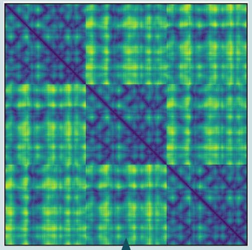

Analytics representations and algorithms

 See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

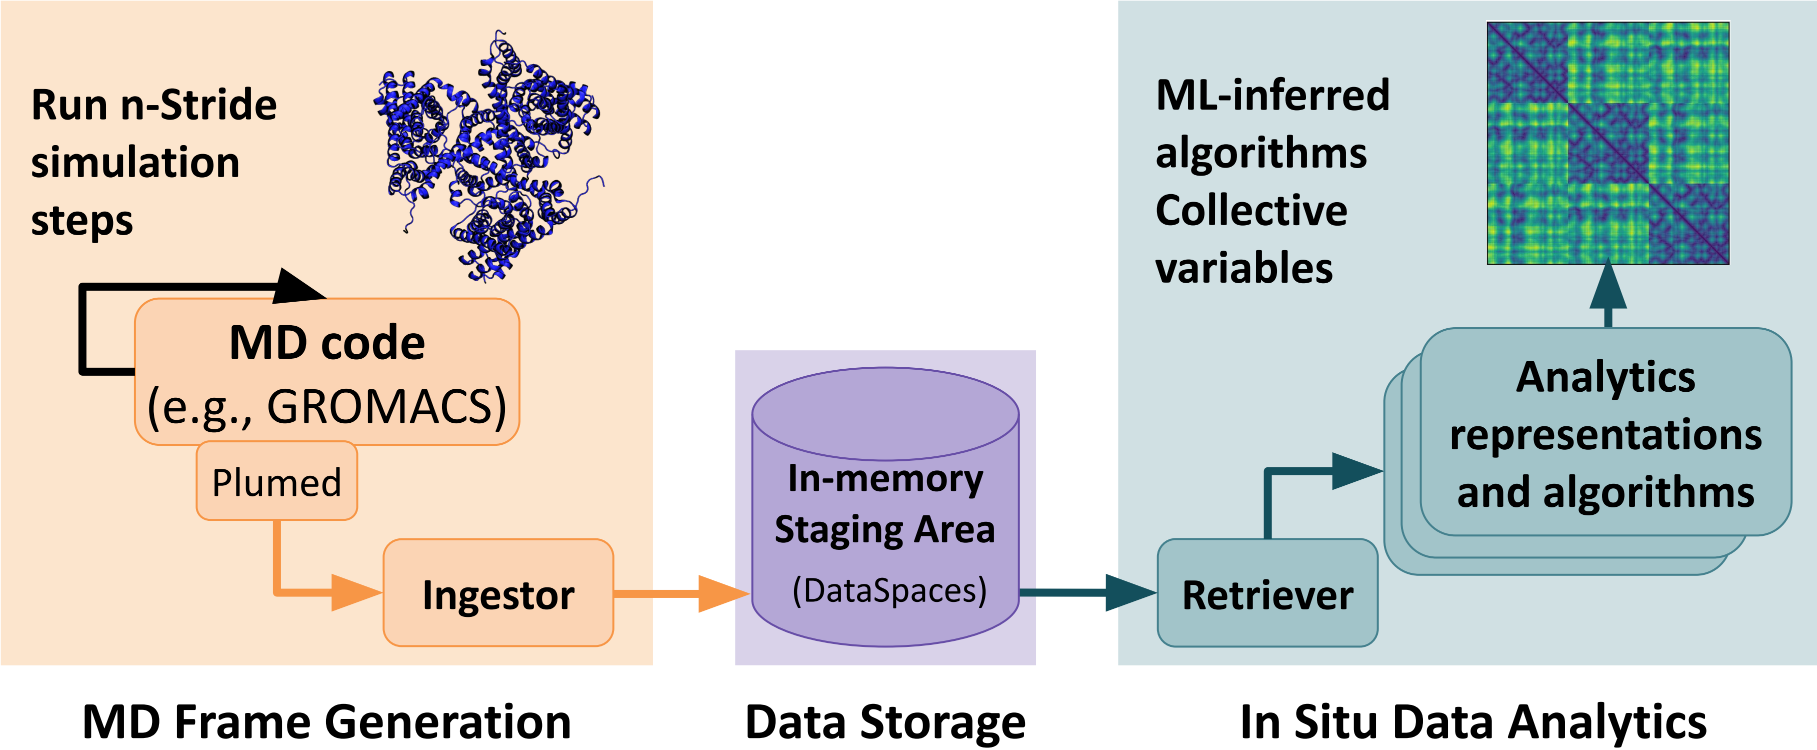

See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

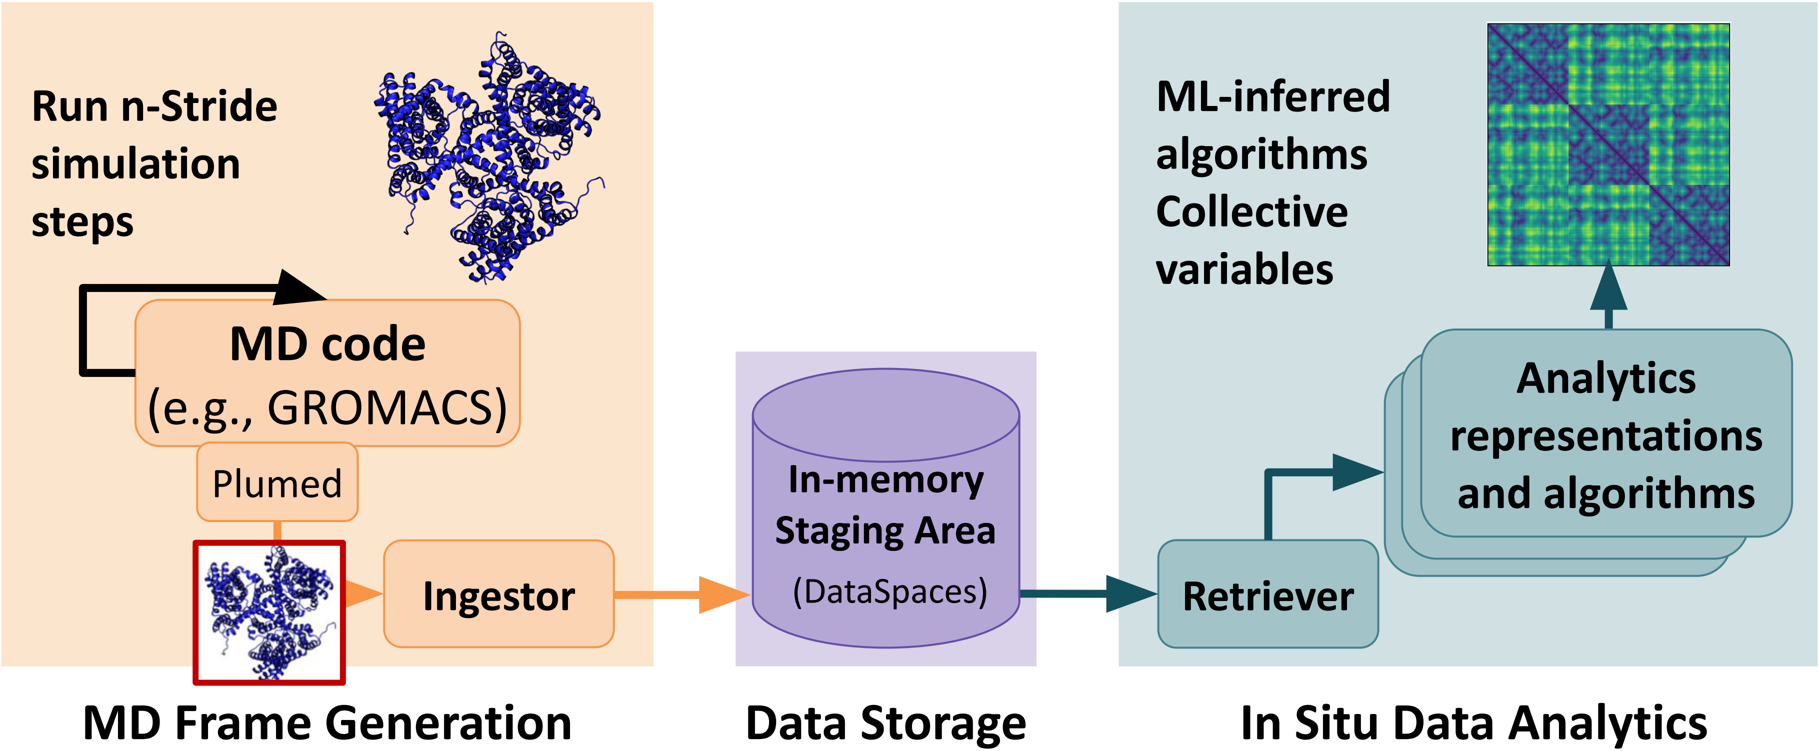

See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

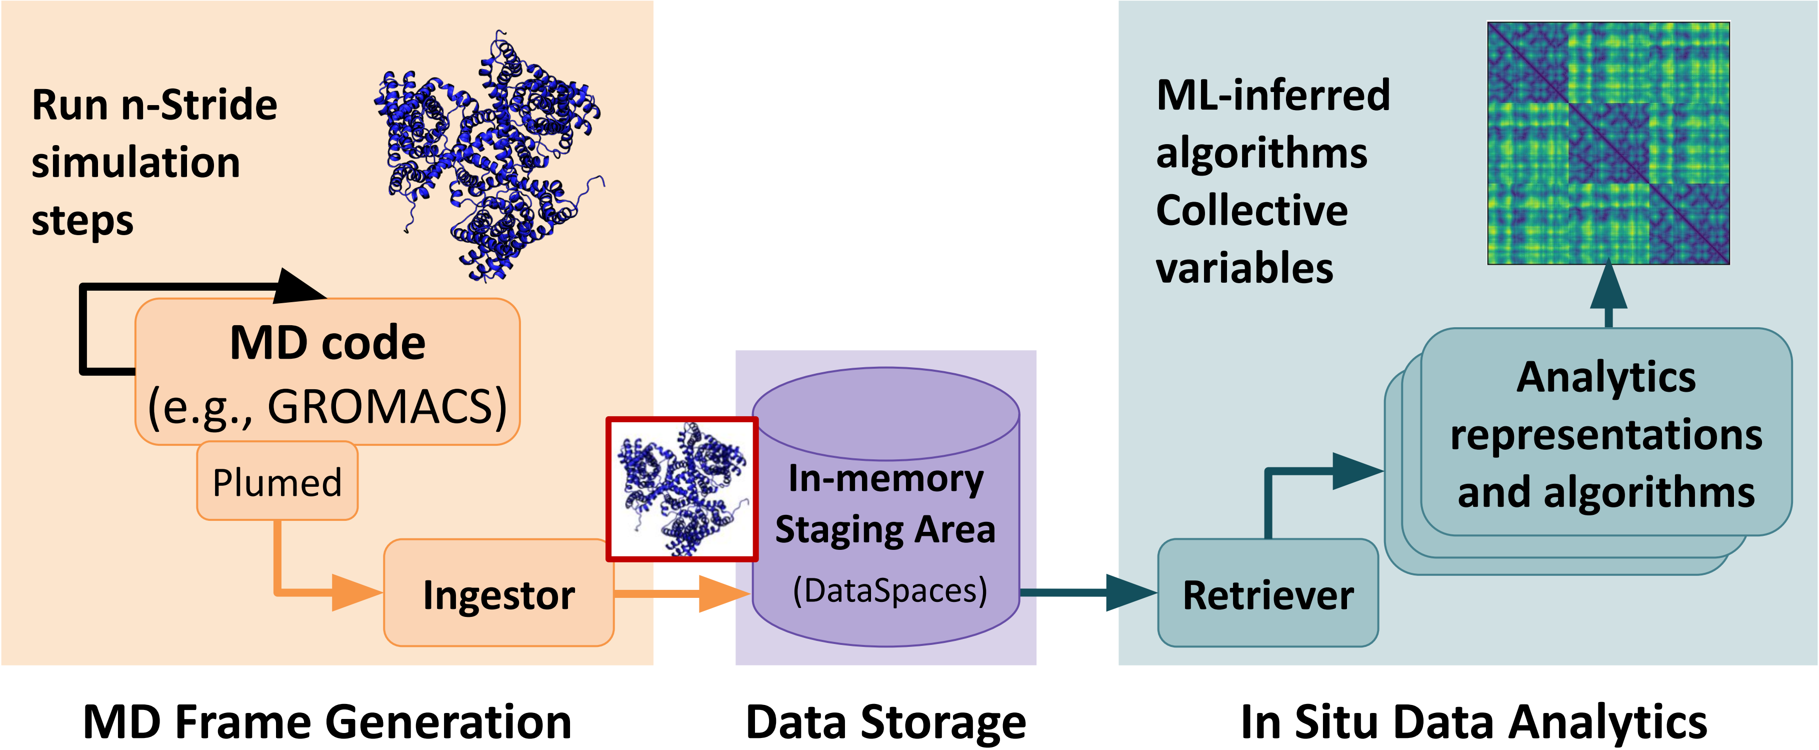

See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

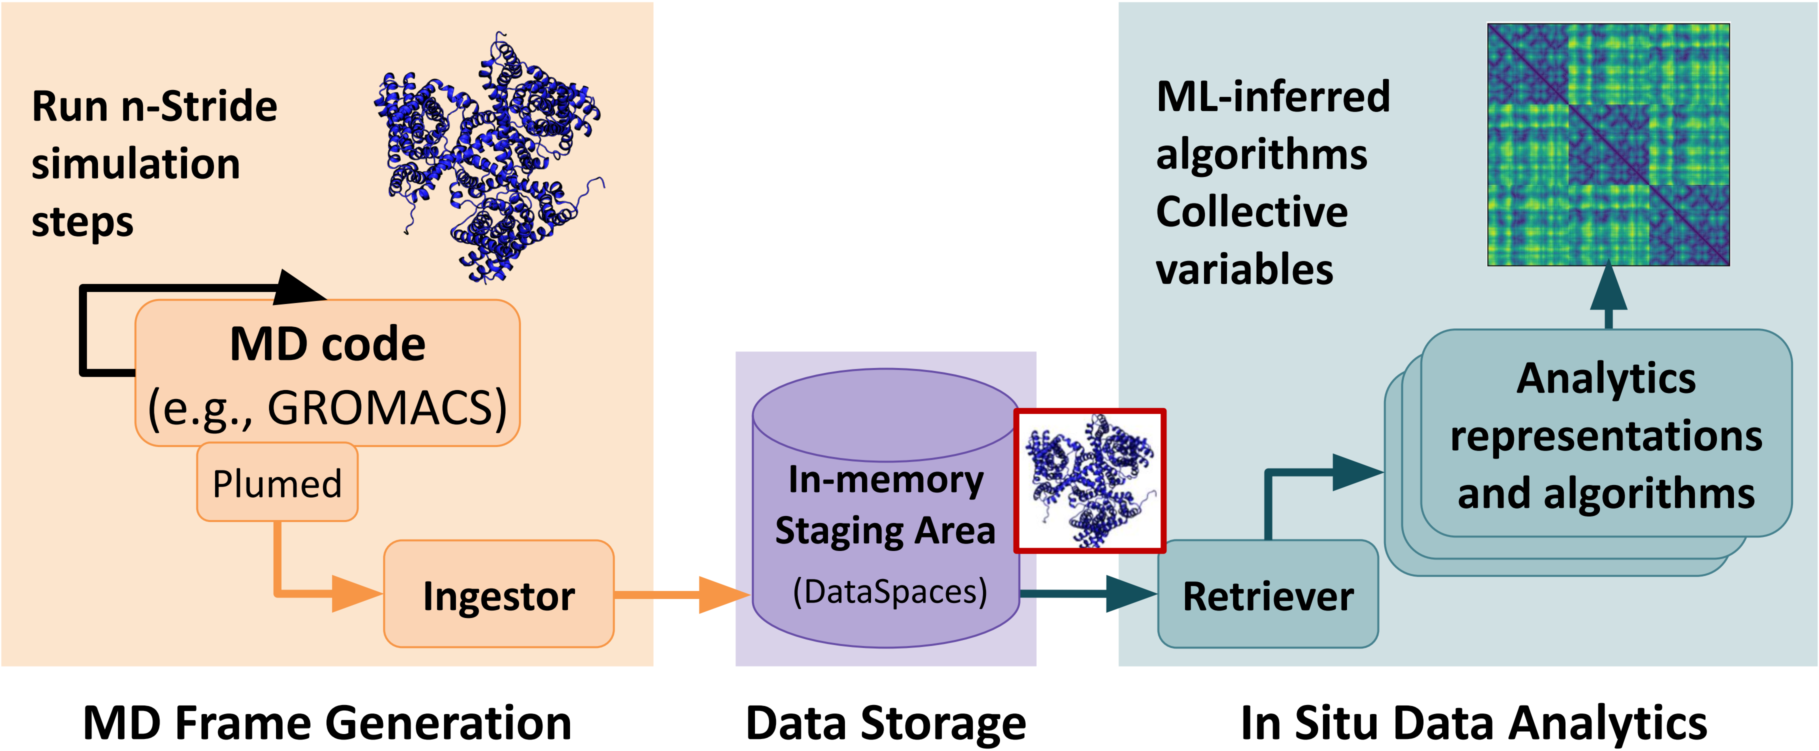

See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

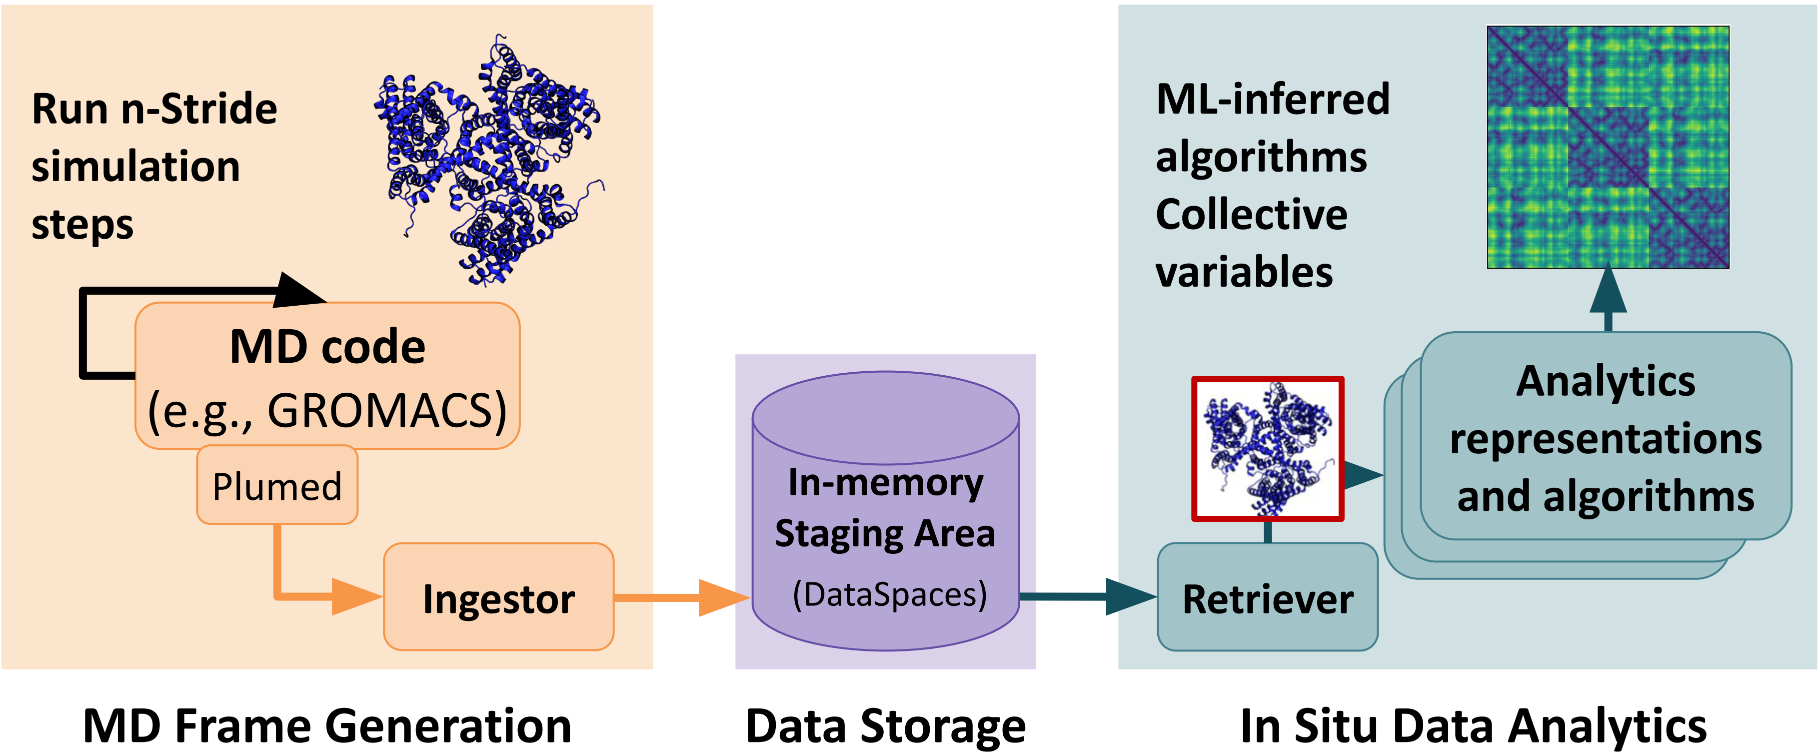

See https://analytics4md.org/ for more information regarding A4MD

In Situ Analysis of MD Trajectories: A4MD

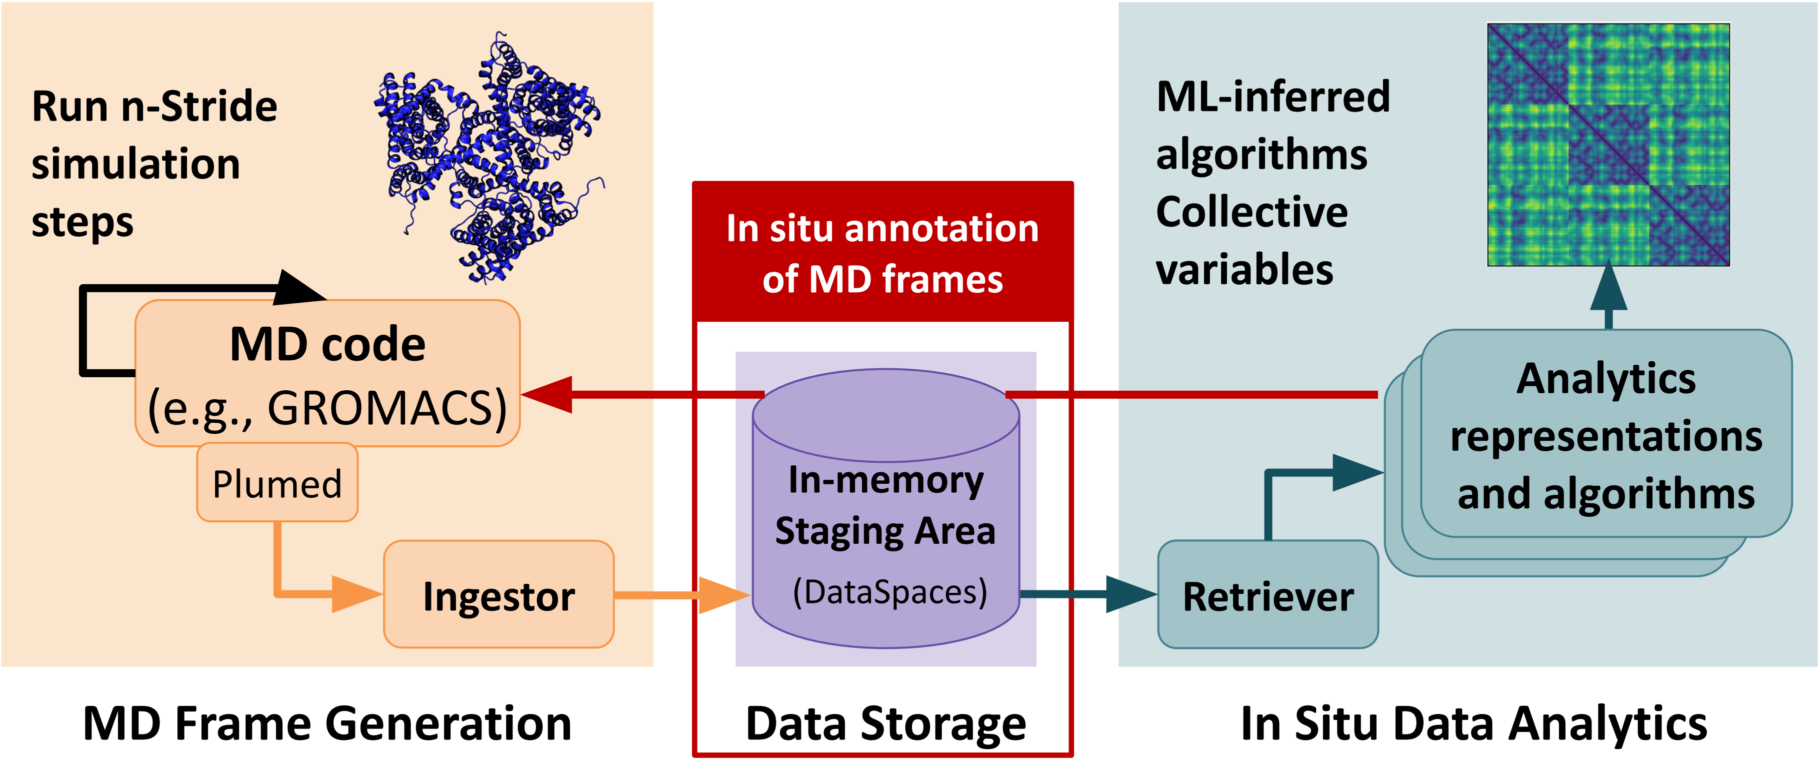

 See https://analytics4md.org/ for more information regarding A4MD

Application: Folding Runtime Detection

- The LEV collective variable (CV) can detect the folded state of an alpha helix with high accuracy using just one frame.

- The CV can be analysed in situ!

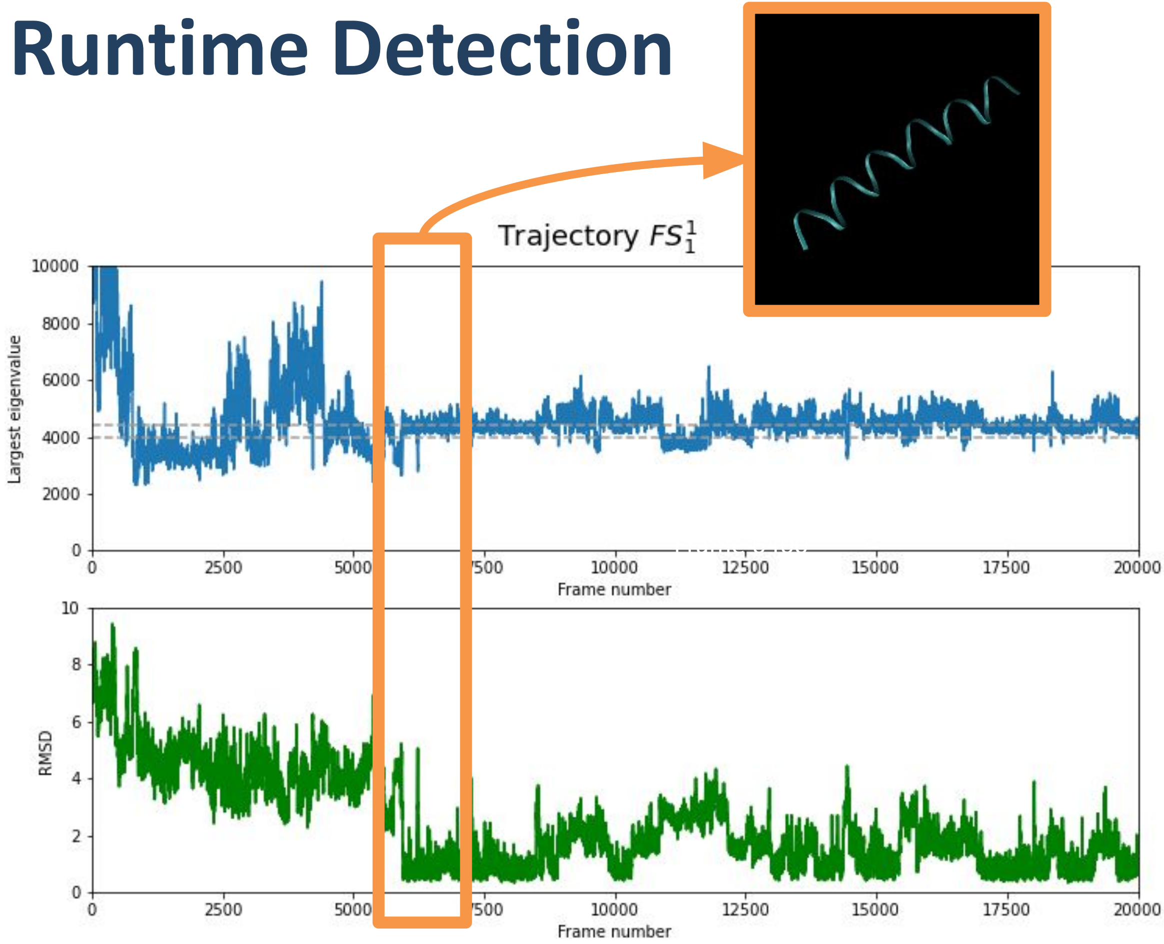

S. Caino-Lores *et al.*, "Runtime Steering of Molecular Dynamics Simulations Through In Situ Analysis and Annotation of Collective Variables," *Platform for Advanced Scientific Computing*, 2023.

Application: ML-Based Runtime Event Detection

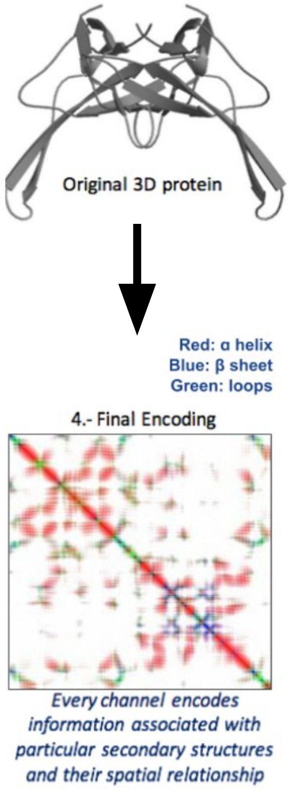

Non-negative matrix factorization

- An increase in reconstruction loss indicates that a NMF model trained for several frames is not suitable for the new observations
- A new model is trained for each event

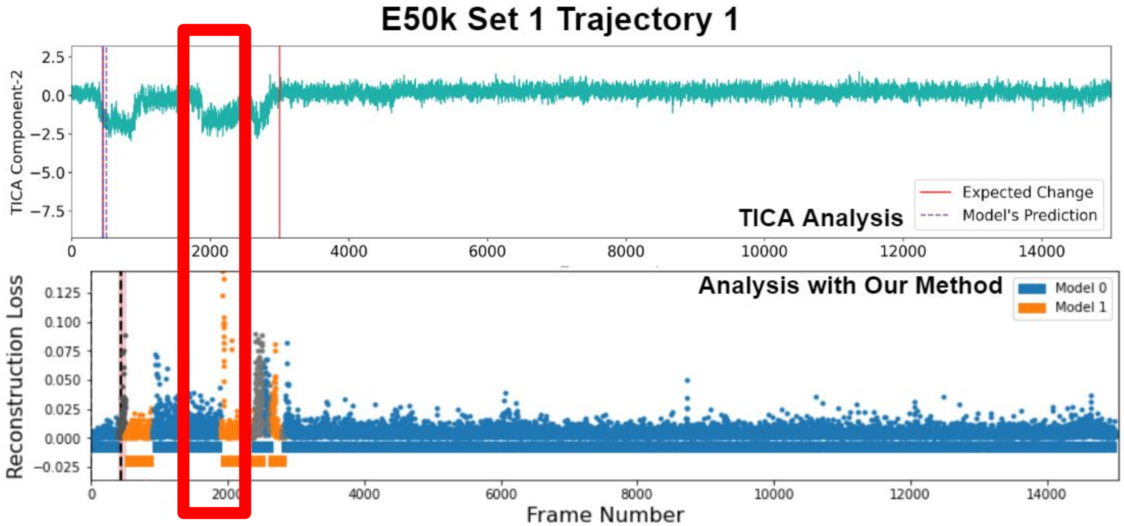

H. Carrillo-Cabada *et al.*, "A Graphic Encoding Method for Quantitative Classification of Protein Structure and Representation of Conformational Changes," in *IEEE/ACM Transactions on Computational Biology and Bioinformatics*

Case Studies:

Runtime Neural Network Fitness Prediction for Neural Architecture Search

Neural Networks and Neural Architecture Search

Neural Networks (NN) can be utilized to extract information from scientific data

Protein images

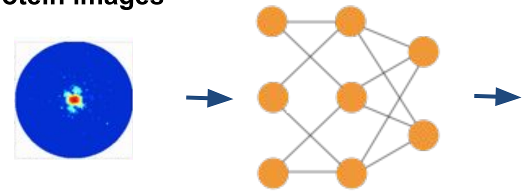

Protein properties:

- Protein type
- Orientation
- Structure

Olaya *et al*, "Identifying Structural Properties of Proteins from X- ray Free Electron Laser Diffraction Patterns". *In Proceedings of the 18th IEEE International Conference on e-Science (eScience), pages 1–10, Salt Lake City, Utah, USA, October 2022. IEEE Computer Society.*

Neural Networks and Neural Architecture Search

Neural Networks (NN) can be utilized to extract information from scientific data

Protein images

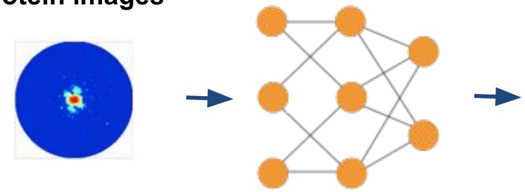

Protein properties:
- Protein type
- Orientation
- Structure

Custom NNs are needed for each dataset

Olaya *et al*, "Identifying Structural Properties of Proteins from X- ray Free Electron Laser Diffraction Patterns". *In Proceedings of the 18th IEEE International Conference on e-Science (eScience), pages 1–10, Salt Lake City, Utah, USA, October 2022. IEEE Computer Society.*

Neural Networks and Neural Architecture Search

Neural Networks (NN) can be utilized to extract information from scientific data

Protein images

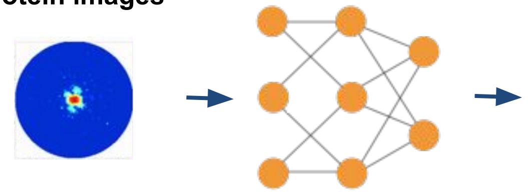

Protein properties:
- Protein type
- Orientation
- Structure

Custom NNs are needed for each dataset

Neural Architecture Search (NAS) can automatically find an optimal NN for a given dataset.

NAS workflow

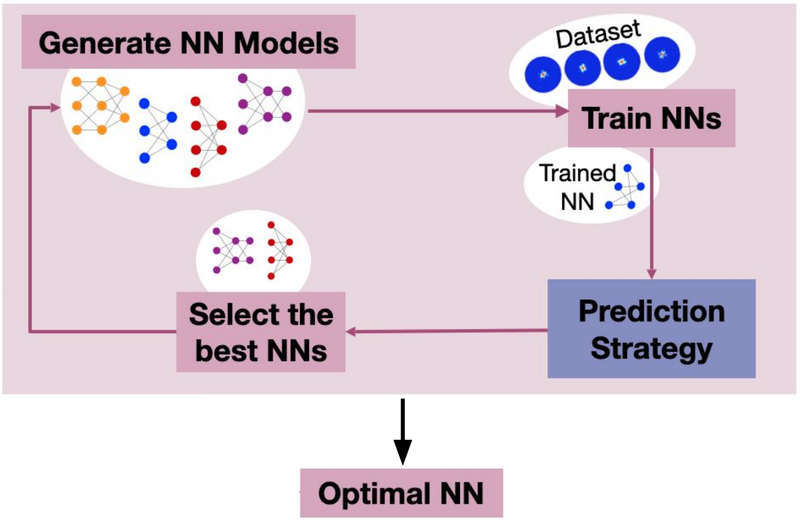

BIG ORANGE BIG IDEAS

Transforming NAS Workflows

NAS workflow

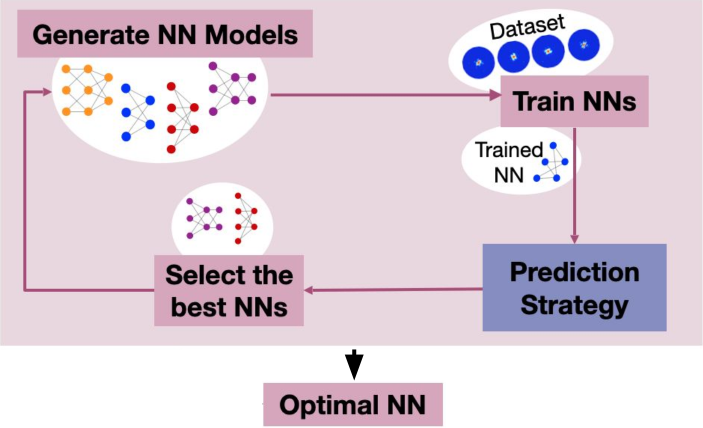

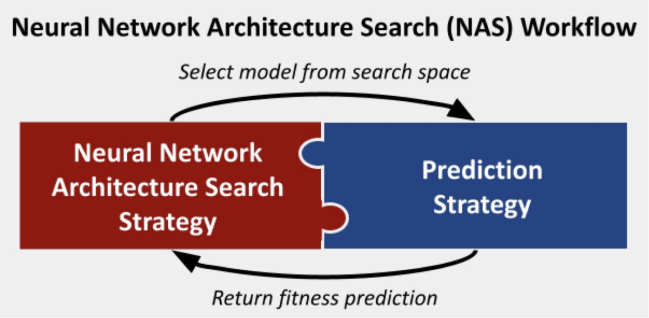

Transforming NAS Workflows

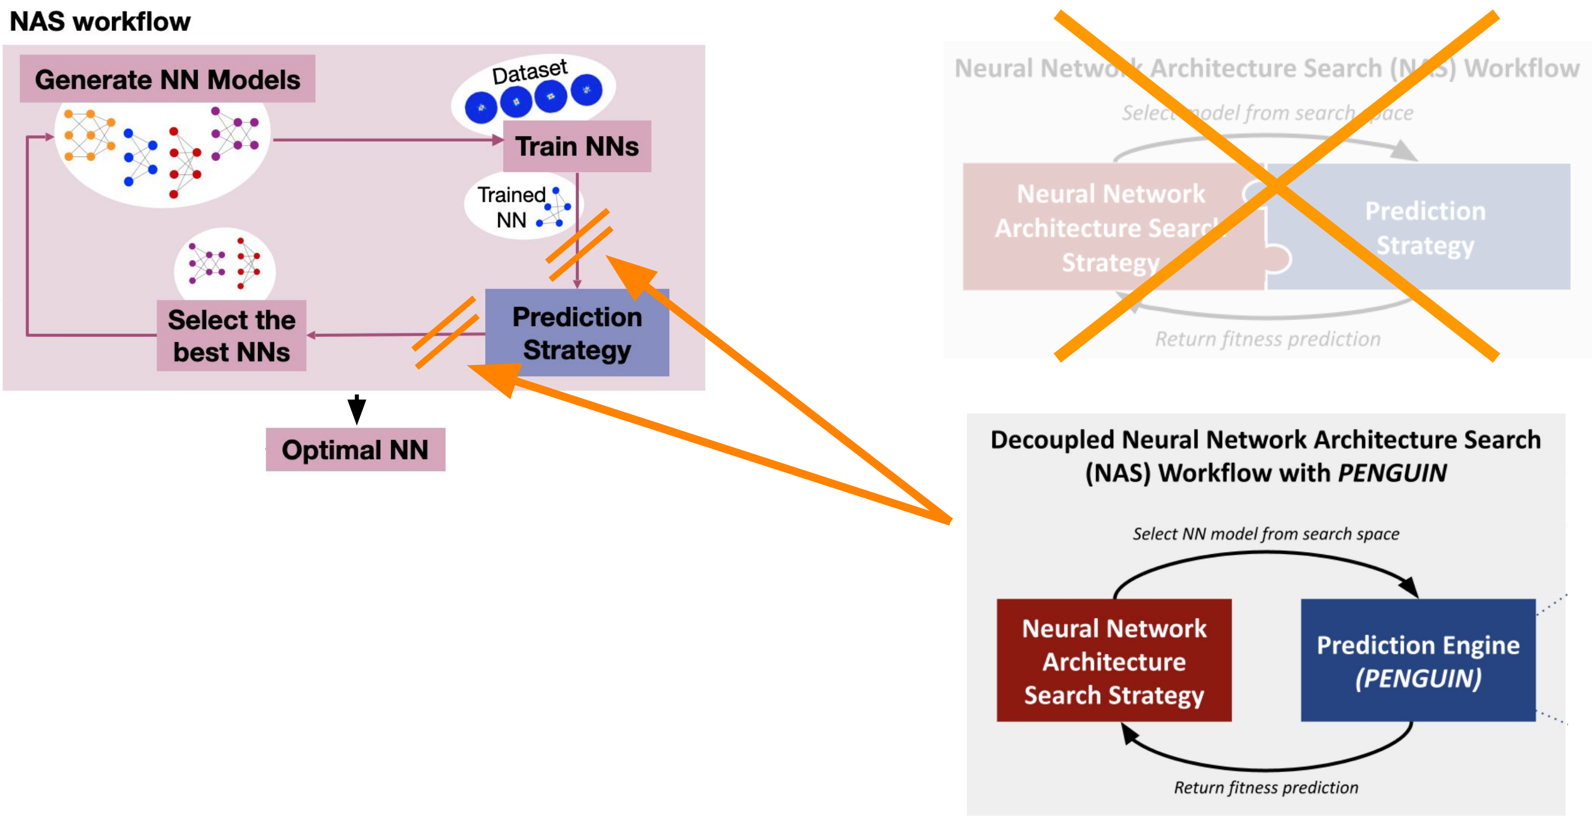

Rorabaugh, A. K., Caino-Lores, S., Johnston, T., & Taufer, M. (2022). Building high-throughput neural architecture search workflows via a decoupled fitness prediction engine. *IEEE Transactions on Parallel and Distributed Systems*, *33*(11), 2913-2926.

Transforming NAS Workflows

1. Build the necessary interfaces to decouple existing NAS from the prediction strategy
2. Enable runtime analysis of generated NNs
3. Inform the NAS about NN predicted performance
4. Support the generation of NN record trail

Users can plug in any NAS and tailor the prediction engine to their problem and dataset

Modular NAS Workflow

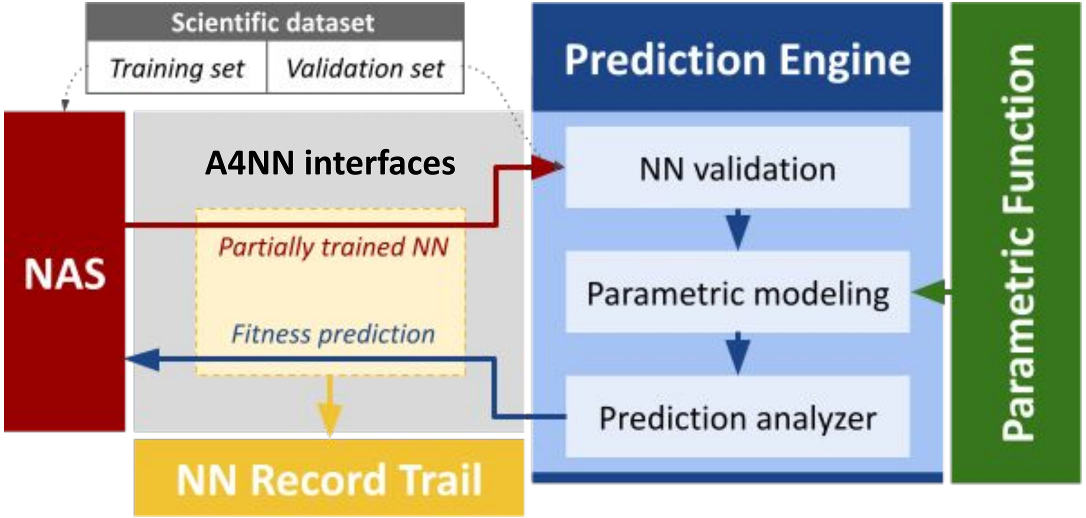

Open Access NN Data Commons

1. Extract metadata from NAS workflow executions
2. Track record trail of each NN
3. Classify NNs according to taxonomies
4. Build data commons containing record trails, scripts, tutorials, and tools to ensure FAIR data

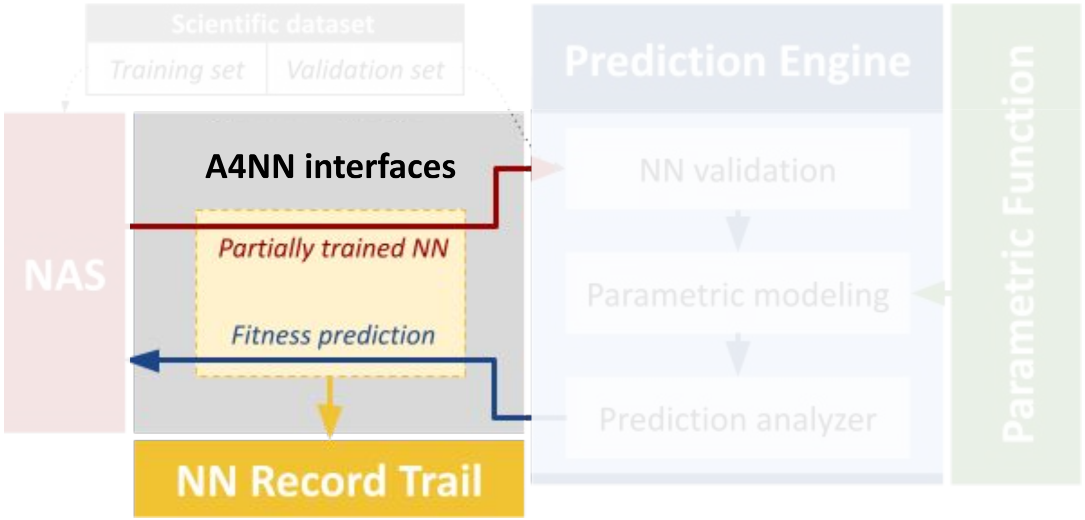

Use Case Summary

Use Case Summary

1. Simulations and AI workloads can act as “streaming devices”

Use Case Summary

1. Simulations and AI workloads can act as “streaming devices”
2. Data exchange components must cover multiple roles
   - Within the workflow: unification, staging, communication, annotation, provenance
   - Within the system: workload balancing, fault tolerance, performance

Use Case Summary

1. Simulations and AI workloads can act as “streaming devices”
2. Data exchange components must cover multiple roles
   - Within the workflow: unification, staging, communication, annotation, provenance
   - Within the system: workload balancing, fault tolerance, performance
3. Orchestration requires introducing elasticity (i.e. dynamic task allocation according to workload needs)

Use Case Summary

1. Simulations and AI workloads can act as “streaming devices”

2. Data exchange components must cover multiple roles
   - Within the workflow: unification, staging, communication, annotation, provenance
   - Within the system: workload balancing, fault tolerance, performance

3. Orchestration requires introducing elasticity (i.e. dynamic task allocation according to workload needs)

4. Programmability is key
   - Workflows require interfaces between different programming languages and data models

The Future of Scientific Workflows in the Computing Continuum

What’s Next?

Next steps towards interoperability of computing ecosystems for scientific workflows

What’s Next?

Next steps towards interoperability of computing ecosystems for scientific workflows

- Shift towards data- and user-centric workflow composition
  - Focus on data-resource mapping, metadata and user insights
  - Make the data management layer the central component of the workflow

What's Next?

Next steps towards interoperability of computing ecosystems for scientific workflows

- Shift towards data- and user-centric workflow composition
  - Focus on data-resource mapping, metadata and user insights
  - Make the data management layer the central component of the workflow
- Better data characterisation
  - Understand data volumes, formats, generation/consumption rates, metadata, etc.

What's Next?

Next steps towards interoperability of computing ecosystems for scientific workflows

- Shift towards data- and user-centric workflow composition
  - Focus on data-resource mapping, metadata and user insights
  - Make the data management layer the central component of the workflow
- Better data characterisation
  - Understand data volumes, formats, generation/consumption rates, metadata, etc.
- New data abstractions and management approaches
  - Find common representations and interfaces
  - Abstract the interaction with the "data lake"

What’s Next?

Next steps towards interoperability of computing ecosystems for scientific workflows

- Shift towards data- and user-centric workflow composition
  - Focus on data-resource mapping, metadata and user insights
  - Make the data management layer the central component of the workflow
- Better data characterisation
  - Understand data volumes, formats, generation/consumption rates, metadata, etc.
- New data abstractions and management approaches
  - Find common representations and interfaces
  - Abstract the interaction with the “data lake”
- Sustainable SW stack for scientific workflows in a changing HW landscape
  - Leverage existing solutions as building blocks
  - Promote modular and extensible workflow composition

Assimilation of Emerging Technologies

Emerging workflows are bringing unconventional HW into the picture (e.g., neuromorphic, quantum) and additional challenges

- Very limited resources
- Immature interoperation capabilities outside of commercial environments
- Heavy data transformation overhead
- Limited SW stack
- No unified data abstractions

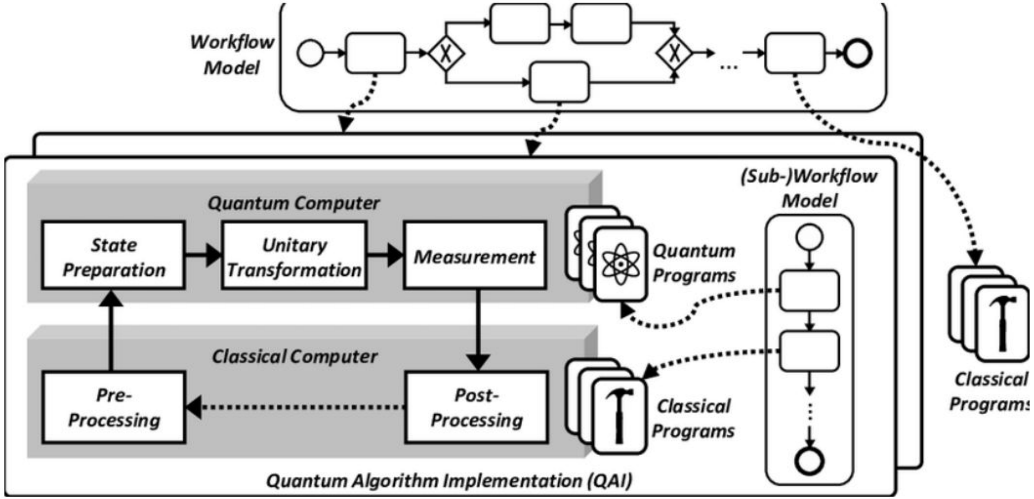

General Structure of a Hybrid Quantum Application.

Weder, B., Barzen, J., Leymann, F., & Vietz, D. (2022). Quantum software development lifecycle. In *Quantum Software Engineering* (pp. 61-83). Cham: Springer International Publishing.

A Data-Centric Perspective on Scientific Workflows in the Computing Continuum

Silvina Caíno-Lores, PhD
scainolo@utk.edu

THE UNIVERSITY OF TENNESSEE KNOXVILLE

BIG ORANGE. BIG IDEAS.®