A Performance Characterization of Scientific Machine Learning Workflows

Patrycja Krawczuk, **George Papadimitriou**, Ryan Tanaka, Tu Mai Anh Do,
Srujana Subramanya, Shubham Nagarkar, Aditi Jain, Kelsie Lam,
Anirban Mandal, Loïc Pottier, Ewa Deelman

Workshop on Workflows in Support of Large-Scale Science (WORKS)
November 15th, 2021

This work was funded by the US Department of Energy under contracts **DE-SC0012636M** and **DE-SC0022328,** and by NSF under contract **1664162.**

CONTENT

- Motivation
- Scientific Machine Learning Workflows (SciMLW)
  - Galaxy Classification Workflow
  - Lung Segmentation Workflow
  - Crisis Computing Workflow
- Experimental Setup
- Executable Workflow Characteristics
- Characterization of the Individual ML Workflow Stages
- Workflow Level and Execution Environment Optimizations
- Summary and Future Work

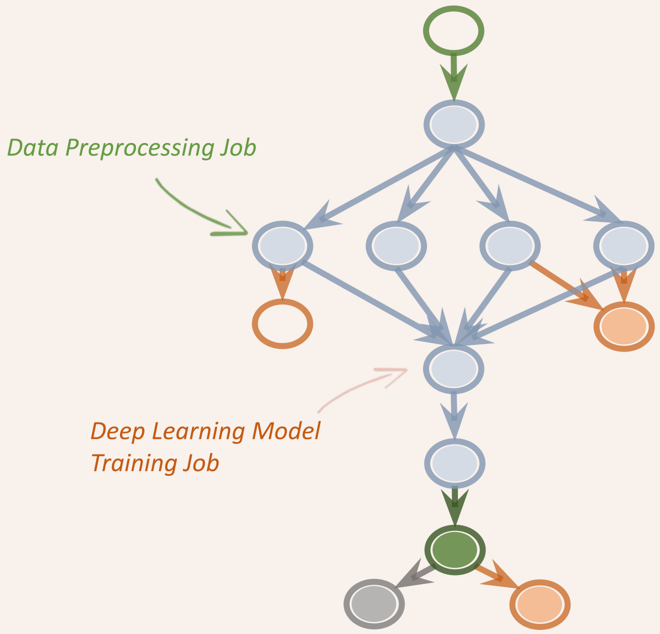

MOTIVATION: SCIENTIFIC MACHINE LEARNING WORKFLOWS (SciMLW)

The landscape of scientific workflows is changing as researchers employ machine learning techniques in their experiments

Examples of SciML experiments:

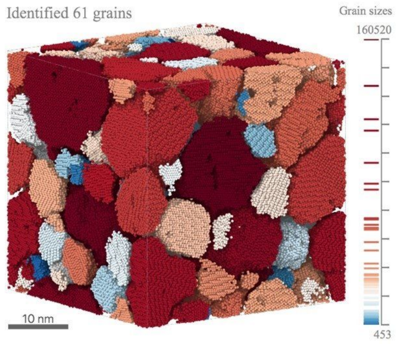

Machine learning can quickly analyze complex phenomena like this simulation of ice crystals." (DOE)

Image: Argonne National Laboratory

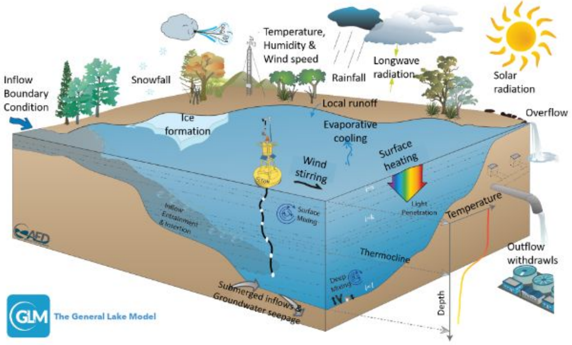

Physics-Guided ML models are used to simulate Lake Temperature Profiles. (X.Jia et. al)

Image: GLM

Machine learning techniques are used to remove instrumental artifacts from big astronomical data

Image: CHBD

SciMLW: STEPS IN SCIENTIFIC MACHINE LEARNING WORKFLOWS

ML Pipelines are composed by **3 main steps**:

- Data Transformations
- Learning
- Validation and Analysis

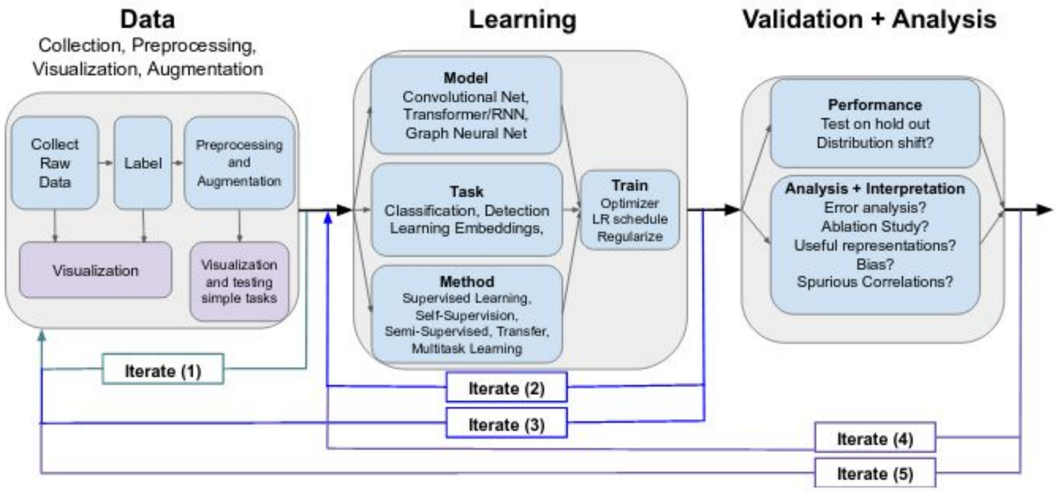

Image: "A Survey of Deep Learning for Scientific Discovery" M.Raghu et. al

Iterative Nature of Design Process
Results from the different stages informing the redesign and re-running of other stages.

SciMLW: GALAXY CLASSIFICATION WORKFLOW

Workflow Context:

- The galaxy morphology classification is a critical step towards understanding how galaxies form and evolve.
- The workflow utilizes the Galaxy Zoo 2 dataset (61,578 RGB images, each 424x424x3 pixels, 1.9 GB of compressed data)

Workflow Steps Overview:

- ***Dataset Generation and Split*** filters out galaxies based on their feature score
- ***Pre-process Images*** jobs where data transformations are applied
- ***Augment Images*** jobs generate additional images of galaxies
- ***VGG16 HPO*** job finds a good set of hyperparameters
- **Train VGG16** job where the model is trained with the chosen hyper-parameters.
- ***Inference and Evaluation*** job runs predictions on the test set, generates statistics and plots that provide insights into the quality of the trained model

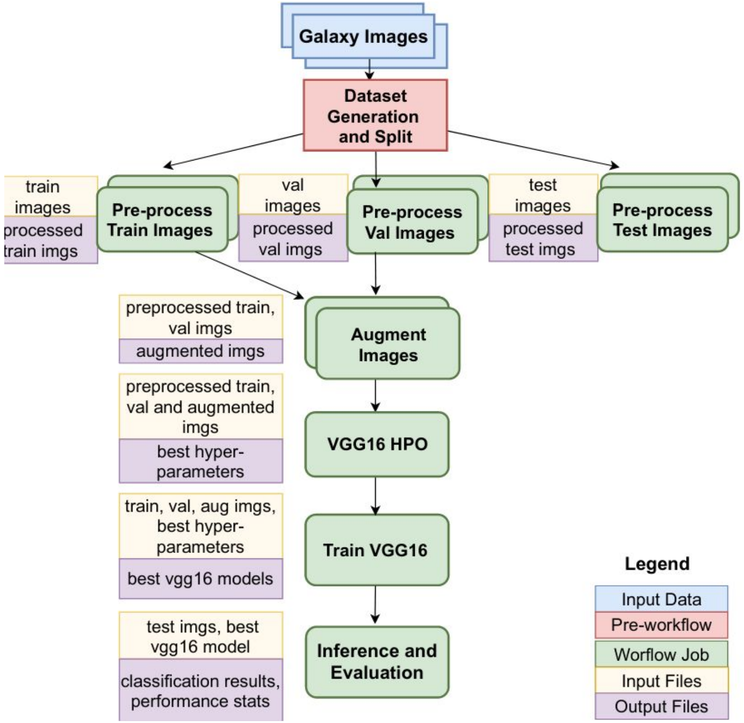

Model	Size	#Params
VGG-16	528 MB	138 M

USC Viterbi School of Engineering renci

SC21

SciMLW: LUNG SEGMENTATION WORKFLOW

Workflow Context:

- Precise detection of the borders of organs and lesions in medical images such as X-rays, CT, or MRI scans is an essential step towards correct diagnosis and treatment planning.
- The Lung Segmentation Workflow uses a the *Chest X-ray Masks and Labels dataset* (800 high-resolution X-ray images and masks, 5.4 GB)

Workflow Steps Overview:

- ***Dataset Split*** images and masks are split into train, test and validation sets
- ***Pre-process Augment Images*** job where data transformations are applied
- ***UNet HPO*** job finds a good set of hyperparameters
- ***Train UNet*** job where the model is trained with the chosen hyper-parameters
- ***Inference on UNet*** job generates masks for test data
- ***Evaluation*** job generates statistics and plots that provide insights into the quality of the trained model

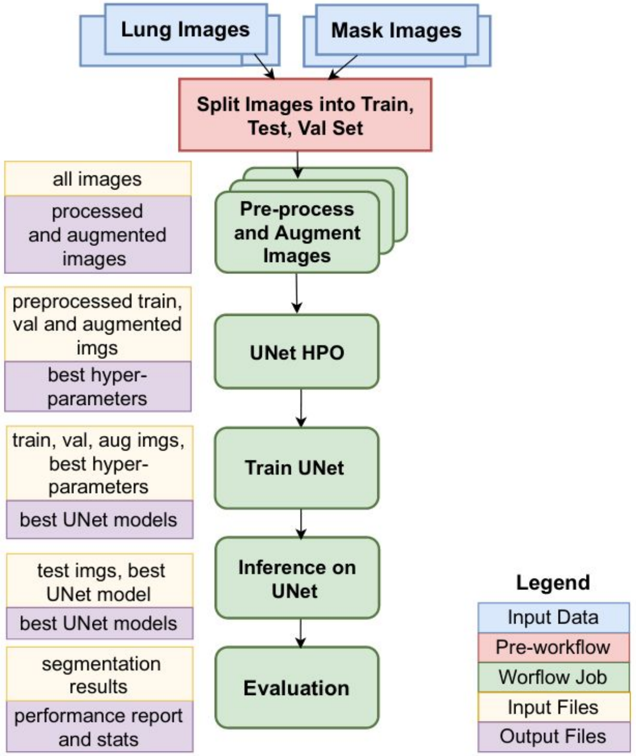

Model	Size	#Params
UNet	81.6 MB	24.4 M

USC Viterbi School of Engineering

renci

SC21

SciMLW: CRISIS COMPUTING WORKFLOW

Workflow Context:

- Social Media (SM) platforms like Twitter have proven to be valuable sources of critical information during disaster events. The published multi-modal content can provide timely and actionable information to local officials.
- The workflow consists of the two pipelines that ingest, respectively, pictorial and textual parts of SM posts. We use the CrisisMMD v2.0 datasets (18,082 images and 16,058 texts, about 2 GB of data

Workflow Steps Overview:

- *Extract and Split* divides data into train, test and validation sets
- *Pre-process Tweet Text/Images* jobs where data transformations are applied
- *BiLSTM/ ResNet50 HPO* jobs finds a good set of hyperparameters
- Train/ Fine-tune BiLSTM/ResNet50 jobs where the models are trained with the chosen hyper-parameters.
- *Inference on BiLSTM/ ResNet50* jobs runs predictions on the test sets
- *Late Fusion Classification and Evaluation* job combines predictions from both pipelines, generates final predictions

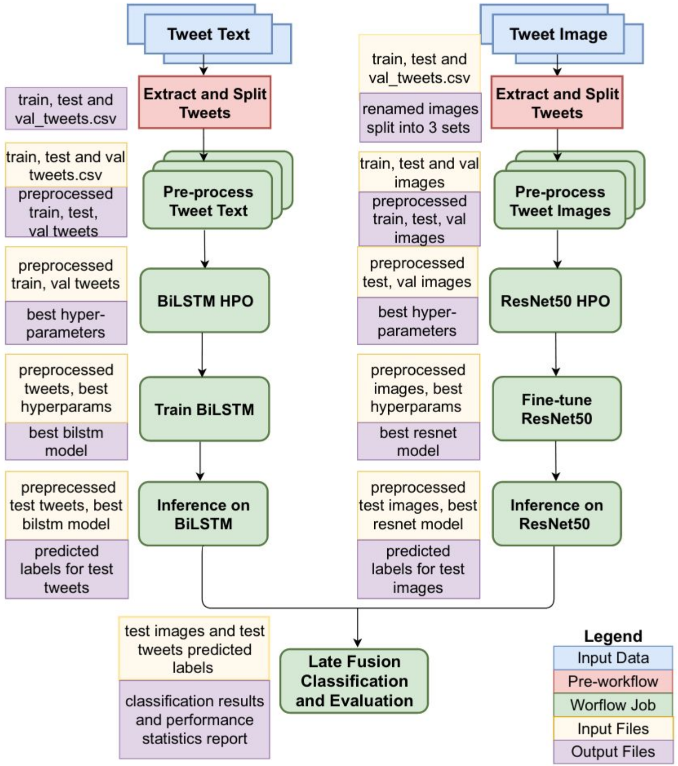

Model	Size	#Params
ResNet50	98 MB	25 M
BiLSTM	9 MB	1 M

USC Viterbi School of Engineering

renci

SC21

EXPERIMENTAL SETUP: Pegasus, Panorama and Chameleon

Pegasus (https://pegasus.isi.edu) is a popular workflow management system that enables users to design workflows at a high-level of abstraction. The workflow descriptions developed by the users are independent of available execution resources and locations of data and executables.

Pegasus relies on HTCondor to run and manage the generated workflows.

Panorama enables **end-to-end online workflow monitoring** and provides execution traces of the computational tasks (CPU and GPU), statistics for individual transfers and infrastructure-related metrics

The NSF Chameleon Cloud is a large-scale, deeply programmable testbed designed for systems and networking experiments.

Chameleon leverages OpenStack to deploy isolated instances of cloud resources for user experiments.

Pegasus

panorama360

Chameleon

EXPERIMENTAL SETUP: CHAMELEON Non SharedFS

- 1 Node in CHI@TACC
  - Served as submit node
- 3 Nodes in CHI@Chicago
  - Served as worker nodes
  - 2 Nodes with GPUs Nvidia RTX6000 (24GB VRAM)
- All nodes were equipped with
  - 24 physical cores
  - 192GB of RAM
  - 10Gbps network cards
- Intersite connectivity
  - 100Gbps dedicated link
  - ~31ms RTT

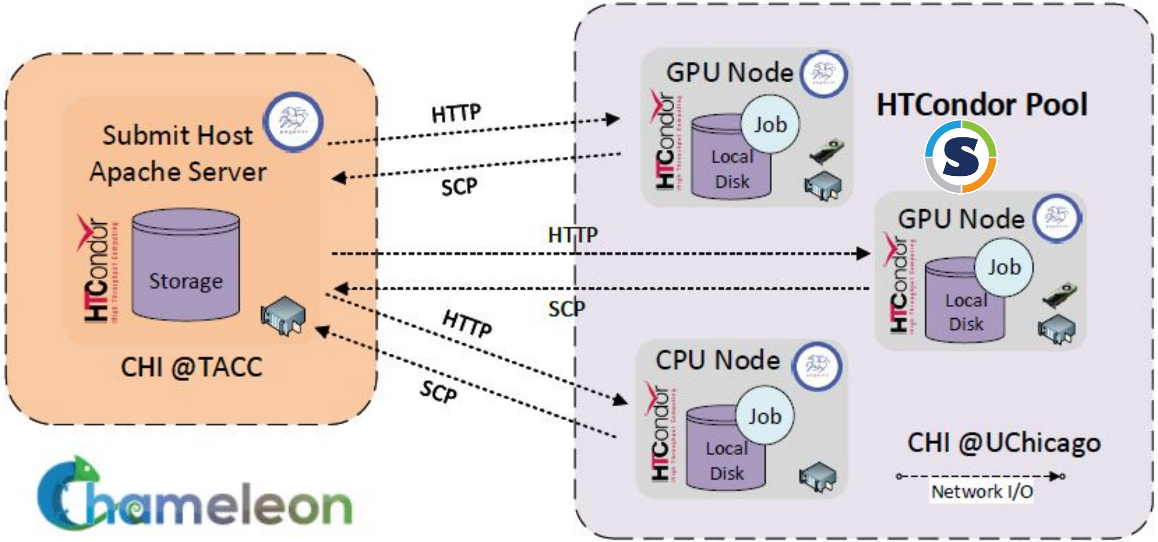

Our 3 SciMLW: EXECUTABLE WORKFLOWS CHARACTERISTICS

Pegasus Executable Workflows

Workflow Name	#Jobs	#Aux. Jobs	#Input Files	Input Size (GB)	Container Size (GB)	#Trials	#Epochs	Batch Size
Galaxy Classification Workflow	11	12	28793	0.374	2.4	2	5	32
Lung Segmentation Workflow	6	12	1408	3.6	4.1	10	25	32
Crisis Computing Workflow	16	12	12747	3.2	4.6	2	5 (ResNet) 10 (BiLSTM)	8 (ResNet) 128 (BiLSTM)

I/O collected using Darshan

Workflow Name	I/O Read (GB)	I/O Write (GB)	Peak Memory (GB)	Peak GPU Memory (GB)	CPU Hours	GPU Hours
Galaxy Classification Workflow	2.88	2.66	41.19	4.42	33.23	1.93
Lung Segmentation Workflow	5.45	7.45	5.96	22.99	37.63	1.64
Crisis Computing Workflow	8.29	3.31	16.74	22.64	48.76	2.45

Our 3 SciMLW: DATA PREPROCESSING

Characteristics

- Good cpu utilization
- Low memory usage
- No GPU usage

- Relatively short running tasks
  - Due to the size of input
- The workflows preprocess almost all their input data

Workflow Name	Data Type	I/O Read (MB)	I/O Write (MB)	Avg. CPU(%)	Peak Memory (GB)	Avg. GPU(%)	Peak GPU Memory (GB)	Avg. Exec. Time (sec)
Galaxy Classification Workflow	Images	435.47	180.1	99.85	0.08	-	-	135.21
Lung Segmentation Workflow	Images	3404.66	48.36	120	0.37	-	-	337.85
Crisis Computing Workflow	Images	1320.19	1697.21	660.74	0.07	-	-	212.48
	Text	11.84	1.12	121.99	0.133	-	-	8.8

Our 3 SciMLW: HYPER PARAMETER TUNING

Characteristics

- GPU enabled
- Galaxy Class. → High RSS - Low GPU Mem. Usage
- Lung Seg. → Low RSS - High GPU Mem. Usage
- Crisis Comp. → Low RSS - High GPU Mem. Usage

- Long running tasks
- High reads

Workflow Name	Data Type	I/O Read (MB)	I/O Write (MB)	Avg. CPU(%)	Peak Memory (GB)	Avg. GPU(%)	Peak GPU Memory (GB)	Avg. Exec. Time (sec)
Galaxy Classification Workflow	Images	1087.6	527.84	1426.22	41.19	53.22	4.13	3421.48
Lung Segmentation Workflow	Images	1955.36	82.93	100.75	5.96	68.19	22.99	4947.24
Crisis Computing Workflow	Images	391.15	763.69	186.22	6.44	68.73	1.76	1424.91
	Text	3926.37	0.53	269.51	4.13	20.92	22.64	1031.6

*Partial I/O captured with Darshan due to #records limit

Our 3 SciMLW: TRAINING

Characteristics

- GPU enabled
- 1 iteration of the HPO
- Peak GPU memory doesn't change
- Max RSS is smaller (significantly for the Galaxy workflow)
- High writes

Workflow Name	Data Type	I/O Read (MB)	I/O Write (MB)	Avg. CPU(%)	Peak Memory (GB)	Avg. GPU(%)	Peak GPU Memory (GB)	Avg. Exec. Time (sec)
Galaxy Classification Workflow	Images	543.62	1489.55	789.51	18.8	68.92	4.13	1453.53
Lung Segmentation Workflow	Images	119.3	7494.12	104.83	5.12	62.48	22.99	557.4
Crisis Computing Workflow	Images	391.15	724.55	179.75	5.94	63.73	1.76	871.99
	Text	1963.2	5.84	272.56	4.03	20.9	22.64	625.09

*Partial I/O captured with Darshan due to #records limit

Our 3 SciMLW: EVALUATION/INFERENCE

Characteristics

- Mix of GPU and non-GPU enabled tasks
- Process fewer data (testing set)
- Lower GPU utilization
- GPU memory remains fairly the same
- Crisis Comp. ⟶ Inference can be GPU accelerated, but wasn't supported at the time of the runs.

Workflow Name	Step	I/O Read (MB)	I/O Write (MB)	Avg. CPU(%)	Peak Memory (GB)	Avg. GPU(%)	Peak GPU Memory (GB)	Avg. Exec. Time (sec)
Galaxy Classification Workflow	Inference & Evaluation	889.86	527.82	404.01	3.78	26.76	4.42	51.84
Lung Segmentation Workflow	Inference	104.74	0.86	106.89	3.1	10.39	22.99	22.51
	Evaluation	0.39	0.37	129.44	0.43	-	-	6.13
Crisis Computing Workflow	Inference	478.24	196.41	2386.63	16.64	-	-	3003.4
		6.41	0.58	111.57	4.13	42.41	22.64	65.95
	Evaluation	0.95	0.14	161.01	11.18	3.72	0.83	211.76

*Partial I/O captured with Darshan due to #records limit

WORKFLOW OPTIMIZATIONS: CAN WE IMPROVE?

- A **workflow management system** can provide optimization opportunities **without changing** your workflow design!
- Workflows can be mapped to different execution environments
- Data can be collocated and shared between the tasks using them

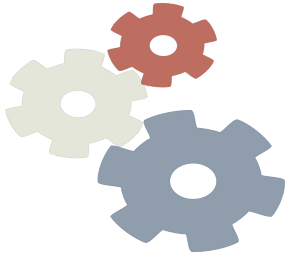

WORKFLOW OPTIMIZATIONS: STRATEGIES

Data Placement Strategies

- Baseline
- Container installed on the workers
- Tasks clustered into larger jobs
- Use of shared file system
  - New execution environment
  - NFS storage on the submit host
  - Shared with the workers

Non Shared FS

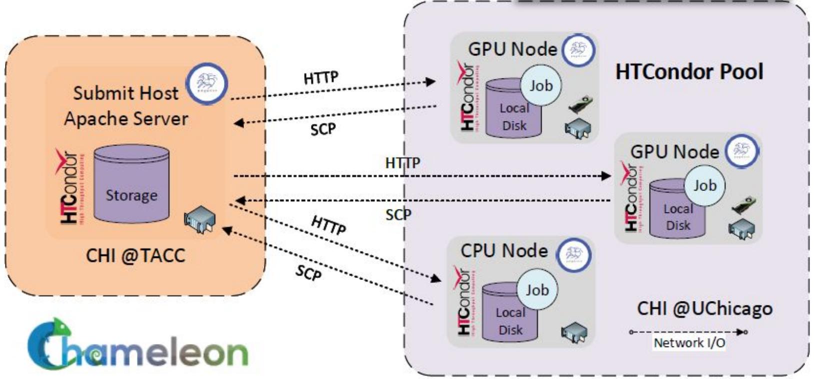

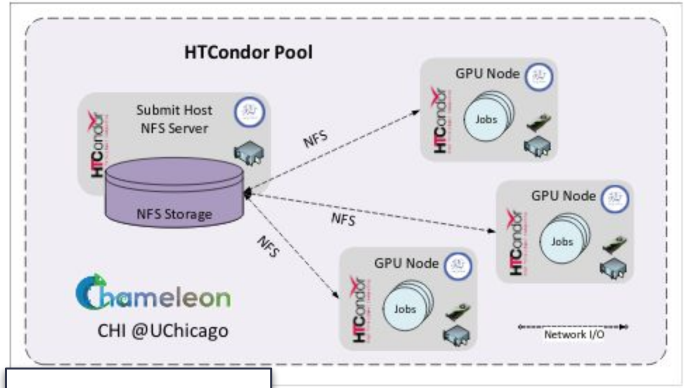

Shared FS

WORKFLOW OPTIMIZATIONS: REDUCING DATA MOVEMENT

- Approx. same amount of files staged in for Baseline, Container Installed and NFS
  - Baseline, Container Installed copy files
  - NFS symlinks files
- Clustering reduces the number of files need to be staged in
  - Data reuse throughout the ML pipelines
  - Fewer aux. jobs than the rest

EXECUTABLE WORKFLOW SCENARIOS AND TRANSFERS SETTINGS

Workflow	Scenario	Jobs	Aux. Jobs[1]	Transfer Threads[2]	Files Staged In
Galaxy Classification	Baseline	11	12	8	111950
	Container Inst.	11	12	8	111950
	NFS	11	12	8	111938
	Clustering	1	4	24	28803
Lung Segmentation	Baseline	7	12	8	9446
	Container Inst.	7	12	8	9446
	NFS	7	12	8	9438
	Clustering	1	4	24	1417
Crisis Computing	Baseline	16	12	8	51006
	Container Inst.	16	12	8	51006
	NFS	16	12	8	50987
	Clustering	3	6	24	12758

[1] The number of auxiliary jobs generated by Pegasus during planning.
[2] Number of threads used by Pegasus to transfer (stage in/out) job input files.

WORKFLOW OPTIMIZATIONS: EFFECT ON COMPUTE TIME

- Compute time remains the same across Baseline, Container Installed and Clustering
  - Compute scratch location is on the worker's local disc.
- Compute time increases for the NFS case
  - Compute scratch location is on disk hosting the NFS.
  - I/O and Network limitations

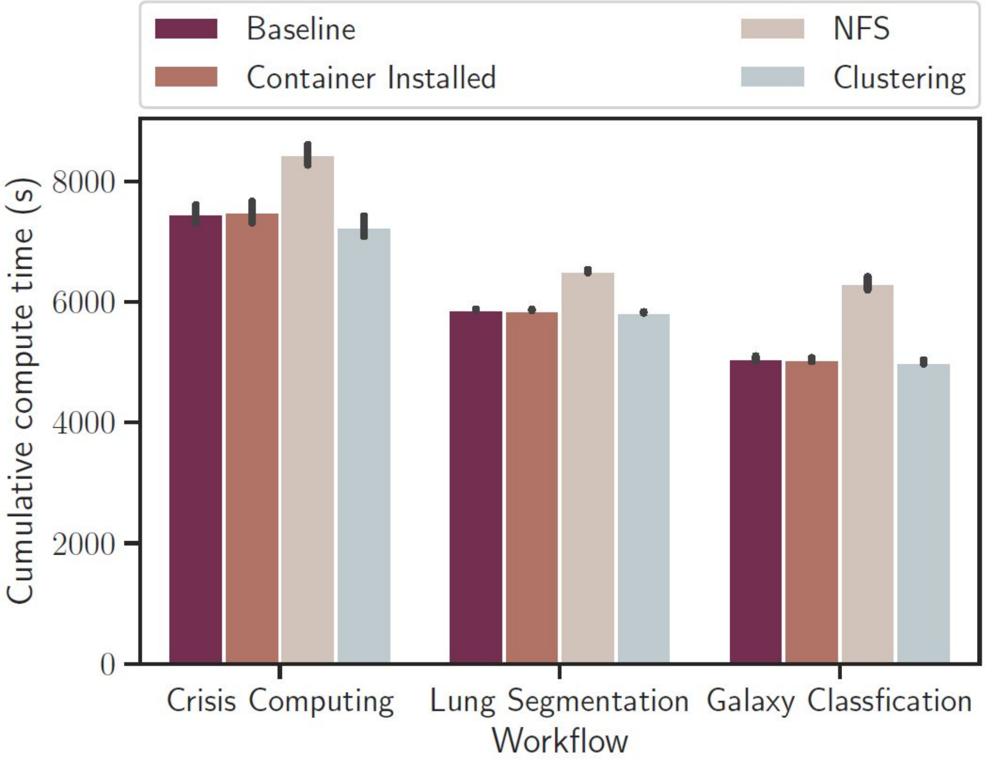

WORKFLOW OPTIMIZATIONS: EFFECT ON TIME SPENT ON TRANSFERS

- Consistently
  stage in time > stage out time
- Container Installed, NFS, Clustering improve the time spent on moving data from the Baseline
  - ML enabled container size is comparable with input data size
  - NFS strategy avoids copying files by symlinks them
  - Clustering reduces the total number of files transferred

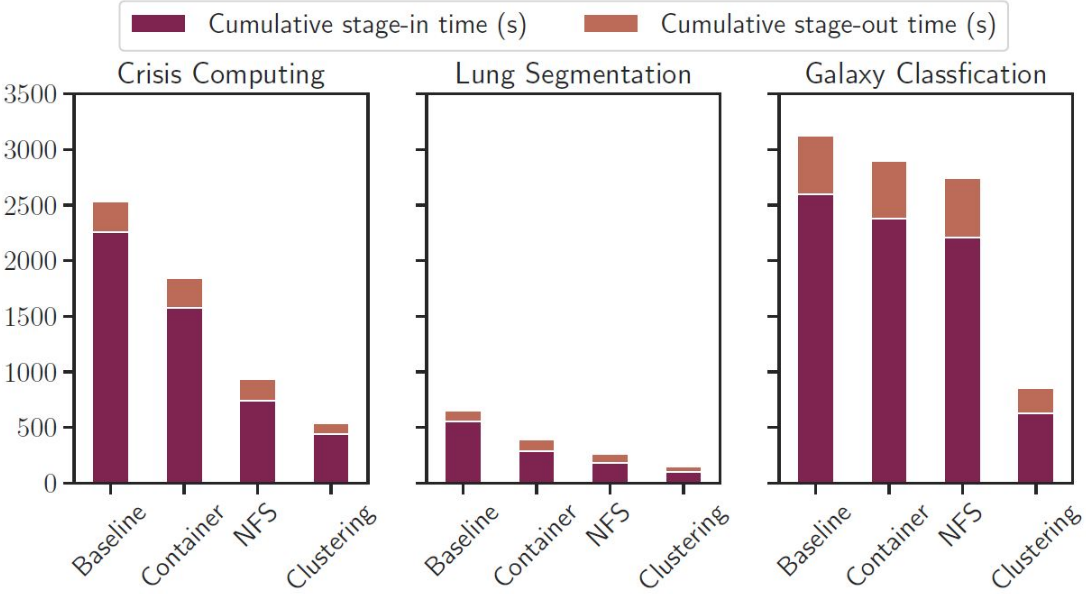

WORKFLOW OPTIMIZATIONS: EFFECT ON WORKFLOW MAKESPAN

- Installing the container on the workers and clustering affect positively the workflow makespan
  - Reducing the data moved
- But, the NFS increases the workflow makespan
  - The execution overhead overshadows the data movement improvements
- Most of the time is spent on compute!

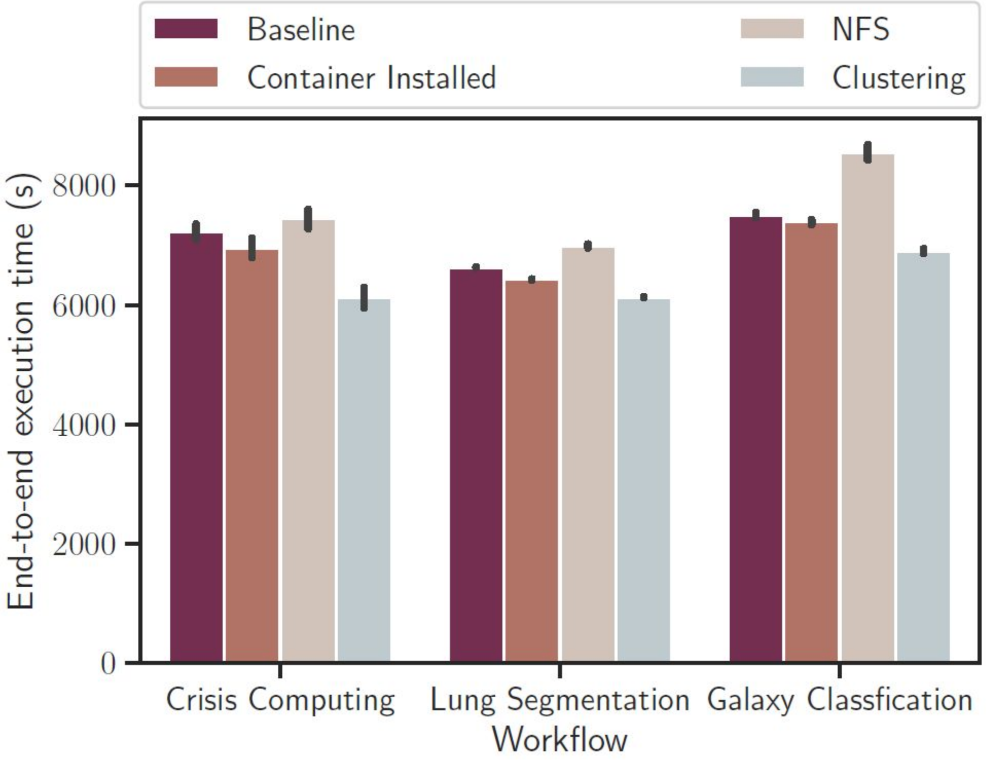

REPRODUCIBILITY: ACCESSING THE WORKFLOWS

- The SciMLW are available online
  - Galaxy Classification Workflow: https://zenodo.org/record/5297663#.YXWXCp7MKUk
  - Lung Segmentation Workflow: https://zenodo.org/record/5297480#.YXWXDZ7MKUk
  - Crisis Computing Workflow: https://zenodo.org/record/5298197#.YXWXEZ7MKUk
- The datasets used in this work are publicly accessible
- Instructions to reproduce our Chameleon setups can be found inside the workflow folders

CONCLUSION

- Scientific Machine Learning Workflows (SciMLW)
  - Galaxy Classification Workflow
  - Lung Segmentation Workflow
  - Crisis Computing Workflow
- We presented their execution characteristics
  - Implemented in Pegasus WMS
  - Executed on Chameleon cloud
- Performed workflow and execution environment optimizations
  - Reduce data transfers and optimize data access
  - Reduction in total makespan of the workflows

QUESTIONS ?

Patrycja Krawczuk
University of Southern California
krawczuk@usc.edu

George Papadimitriou
University of Southern California
georgpap@isi.edu

This work was funded by the US Department of Energy under contracts **DE-SC0012636M** and **DE-SC0022328,** and by NSF under contract **1664162.**