

Hewlett Packard Enterprise



# SMARTSIM: ENHANCING HPC WORKFLOWS WITH MACHINE LEARNING

Andrew Shao, PhD | Senior HPC&AI Research Scientist, HPE 13 July 2023 andrew.shao@hpe.com

►SEARCH►TR/01►03

# ARCHETYPES OF MACHINE LEARNING

- Supervised learning:
  - Given these inputs and known outputs, derive a relationship
  - Linear regression falls under supervised learning
  - Artificial neural network:
    - Linear and non-linear transformations with many free parameters
    - Useful especially when the basis is unknown
- Reinforcement learning
  - Train model to take actions within a given ruleset based on predefined reward function
    - Play a video game with these inputs that gives you a high score
- Unsupervised Learning (not discussed today)
  - Find relationships in unlabeled data

All archetypes rely on data to learn, generalize, and predict



### MACHINE LEARNING TASKS



#### • Forward pass

- Pass inputs through a model with a given set of weights (i.e. values) for every free parameter
- Backward pass
  - Calculate the gradient of the error w.r.t. the weights
- Training Loop (similar to optimization)
  - Perform a forward pass with given weights and features
  - Calculate error compared to expected output
  - Use **backward pass** to estimate gradient
  - Update the weights
- Trained model:
  - Architecture (layers, layer types, etc.)
  - Weights (values of the free parameters)
- Inference: Forward pass of a trained model to Trained Model make a new prediction

### CHALLENGES IN EMBEDDING ML INSIDE AND OUTSIDE OF SIMULATIONS

#### Technical

- How can we call an ML model from a Fortran/C/C++ code?
- How can we scale this to the needed compute scales?
  - ML models in physics simulations tend to be small
  - Want to efficiently allocate GPU resources

#### These technical challenges impede scientific progress!

- Naturally focuses the community on how to replace simulations with AI
- Al to solve physics problems tends to be in idealized contexts or static data
- Scientists are asking the questions, but do not have the tools to test this practically

### SMARTSIM ENABLES AI-ENHANCED SIMULATIONS

### Lower the barriers to entry

- Enable Fortran, C, C++ simulations to interact with ML packages efficiently
- Rapidly prototype and iterate on new ML models
- Allow scientists to focus on building applications *not* code or infrastructure

# Introduce a new workflow paradigm

- Simulations as producers and consumers of data and ML models
- Data is kept in-memory and available in-flight



# SMARTSIM OVERVIEW



### **ABOUT SMARTSIM**

The SmartSim open-source library bridges the divide between traditional numerical simulation and data science **SmartSim** enables simulations to be used as engines within a system, producing data, consumed by other services to add value and create **new applications** 

- Use Machine Learning (ML) models in existing Fortran/C/C++ simulations
- Communicate data between C, C++, Fortran, and Python applications
- Train ML models and make predictions using TensorFlow, PyTorch, and ONNX
- Analyze data streamed from HPC applications while they are running



### SMARTSIM AND SIMULATION INTEGRATIONS

**Computational Fluid Dynamics**:

OpenFOAM, FLEXI, PHASTA, libCEED, NekRS (in progress)

**Climate and Weather**: MOM6, NEMO, CESM

Molecular Dynamics LAMMPS, OPENMM

Easy to instrument an existing codebase to enable new AI/ML/data science applications

- Very little new code needs to be added



### **HIGH-LEVEL ARCHITECTURE**





### **CLUSTERED DATABASE**





### CO-LOCATED DATABASE





# PROTOTYPICAL USE CASES



### ONLINE VISUALIZATION AND ANALYSIS

I want to perform visualization and analysis while my simulation is running

 Stream data from C/C++/Fortran simulations for analysis, and visualization in real time.



jupyter

ndel for the Gent-McWilliams coefficient) to an xarray dataset. The array at timester 1 is plotted in the first figure: the

Extract SmartRedis to an varray datase

### CONTINUOUS ONLINE TRAINING

I want to create machine learning models on simulation data without performing application collectives or writing to the filesystem

- SmartRedis Client can be used inside DataLoaders for TensorFlow and PyTorch to perform Stream Training
- Trained models can be checkpointed, saved and later sent to the database for online inference.
- Can be set from any language and called from any language (e.g., model set from Python, called from Fortran)



# ONLINE MACHINE LEARNING TRAINING/INFERENCE

• I want to use continuously train a machine-learning model from my simulation and use the updated model for inference inside my solver



- Task: Need to use velocities to understand how to optimally deform the mesh
- ML Workflow
  - 1. Extract displacements
  - 2. Train model on displacement between boundary and mesh boundary
  - Use ML model to predict mesh displacement for unstructured mesh

Collaboration with OpenFOAM Data-Driven Modelling Special Interest Group (Maric and Weiner) Molecular Dynamics with DeepDriveMD



### Molecular Dynamics with DeepDriveMD

- Goal: predict the folded configuration of a protein starting from its atomic structure
  - The original paper: <u>DeepDriveMD: Deep-Learning Driven Adaptive Molecular Simulations for Protein Folding</u>
- Problems:
  - Each protein is composed of hundreds or thousands of atoms
  - Protein folding happens in discrete steps which are driven by energy minimization, but with random fluctuations
    - Molecular Dynamics simulation compute evolution of protein and store steps at regular time intervals: a collection of steps is called a trajectory
  - How can one efficiently explore the space of all possible conformations?
    - Most trajectories will end up in suboptimal states
    - Trajectories far apart may collapse on the same one after a sufficiently large number of steps
- Solution:
  - Run short simulations, store steps
  - Project conformations on low-dimensional space (latent space of a Convolutional Variational Autoencoder)
  - Cluster discovered conformations
  - Explore less sampled regions of the low-dimensional space

### DEEPDRIVEMD DATA FLOW



# RELEXI – RL+CFD+SMARTSIM



### **REINFORCEMENT LEARNING IN CFD**

#### CFD problems that can be posed as RL

- Training a neural network to behave like a turbulence model
- Design and test an active flow control system

#### **Modeling Assumption:**

Full Solution ~ Coarse Solution + Turbulence Model

- Supervised Learning
  - Inherent risk because ML model does not interact with the simulation -> instability, crashes, etc.
- Reinforcement Learning solution:
  - Train a model of turbulence in-situ to ensure stability
  - Agent: predict Smagorinsky coefficient based on local u, v in LES
  - Reward function: Compare LES+ML turbulent spectra to DNS solution

Work done by Kurz, Offenhauser, and Beck [2022], https://doi.org/10.1016/j.ijheatfluidflow.2022.109094



### ARCHITECTURE OF RELEXI



### RL MODEL REPRODUCES "TRUE" SPECTRA



32 DOF

- RL Model convergence takes ~1 day
  - 1,024 CPUs (8 nodes), 8 Nvidia A100s (1 node)
- Closely reproduces energy spectra down to resolved wavenumbers
  - Skill beyond even the 'resolvable' limit
  - Some generalizability to different Reynolds numbers
- Incurs about ~10% performance cost
  - Small in comparison to Implict LES
- Flexibility in defining "reward" function allows different/multiple objectives

### **RL FOR ACTIVE FLOW CONTROL**

Problem Statement:

• Reduce drag by controlling 'jets' at top and bottom of cylinder that inject momentum RL Solution

• Find an optimal policy to control 'jets' based on information at 5 probe locations that minimizes

$$r_t = -\langle c_D \rangle_{avg} - 0.2 \left| \langle c_L \rangle_{avg} \right|$$



# TURBULENCE MODELLING WITH ONLINE INFERENCE IN CLIMATE/WEATHER

Recent experiment done

# TURBULENCE IN CLIMATE/WEATHER

- Resolving turbulence in the ocean requires excessive compute resources O(>10,000 CPUs)
- Hypothesis: Including turbulence models can improve low-resolution models
- Problem: Mathematically deriving new turbulence models is hard
- Solution: Use deep learning to predict the effect of turbulence from coarse data

CNN

- HPE: Partee et al. [2022]
- Courant Institute: Guillaumin and Zanna [2021]





Collaboration between HPE/National Center for Atmospheric Research Partee et al. [2022], doi.org/10.1016/j.jocs.2022.101707

CRAYLABS@HPE.COM

26

# DEPLOYING A SECOND TURBULENCE MODEL ON FRONTIER

#### • Prologue

- Collaborators had a small, one-node ocean simulation with CNN parameterization
- Successfully integrated SmartSim into MOM6
- Cluster available to them did not have enough nodes to support a 'realistic' simulation
- June 12:
  - SmartSim engineers get access to allocation on Frontier
  - Install SmartSim from scratch
  - Run base case
- June 13:
  - Start doing scaling studies using SmartSim ensembles to understand how to efficiently use Frontier Resources
- June 14:
  - Run one year of realistic simulation on 50 nodes of Frontier
- June 15
  - Run application on 672 nodes of Frontier (~20,000 CPUs and 5,000 AMD GPUs)



# **RESOURCES AND CONCLUSIONS**



### DOCUMENTATION

#### www.craylabs.org



GETTING STARTED	=	r: 0	+	I≡ Contents
Introduction	_			Experiment
Installation	SmartSim ADI			Settings
Community	SmartSim API			Orchestrator
Contributing Examples				Model
TUTORIALS	Experiment			Ensemble Machine Learning
Getting Started	Emericant date (name) are noth launcherly	Initializa an Eveneriment instance		Slurm
Online Analysis	Experimentinit(name[, exp_path, launcher])	Initialize an Experiment Instance		Ray
Online Inference	<pre>Experiment.start(*args[, block, summary,])</pre>	Start passed instances using Experiment launcher		
Online Training				
Ray Integration	Experiment.stop(*args)	Stop specific instances launched by this Experiment		
SMARTSIM	<pre>Experiment.create_ensemble(name[, params,])</pre>	Create an Ensemble of Model instances		
Experiments	Experiment create model(name run settings)	Create a general purpose Model		
Orchestrator				
Launchers	<pre>Experiment.create_database([port, db_nodes,])</pre>	Initialize an Orchestrator database		
SmartSim API	Evneriment create run sattings(eve[ ])	Create a RunSettings instance		
SMARTREDIS		oreate a nanoeccingo instance.		
SmartRedis	<pre>Experiment.create_batch_settings([nodes,])</pre>	Create a BatchSettings instance		
Integrating into a Simulation	Experiment generate(*args[ tag_overwrite])	Generate the file structure for an Experiment		
Python				
C++	<pre>Experiment.poll([interval, verbose,])</pre>	Monitor jobs through logging to stdout.		
Fortran				
Data Structures	Experiment.finished(entity)	Query if a job has completed.		
Runtime Requirements	<pre>Experiment.get_status(*args)</pre>	Query the status of launched instances		
SmartRedis API				
REFERENCE	Experiment.reconnect_orchestrator(checkpoint)	Reconnect to a running Orchestrator		
Changelog	<pre>Experiment.summary([format])</pre>	Return a summary of the Experiment		
Code of Conduct				
Developer	<pre>class Experiment(name, exp_path=None, launcher='local') [source]</pre>			
	Bases: object			
Theme by the Executable Book Project	Experiments are the Python user interface for SmartSim.			
	Experiment is a factory class that creates stages of a workflow and manages their execution.			

### WHERE CAN I FIND MORE?

Other SmartSim Information

- Contact: <u>CrayLabs@hpe.com</u>
- SmartSim Repository: <u>https://github.com/CrayLabs/SmartSp</u>
  - Includes docker container with tutorials!
- SmartRedis Repository: <u>https://github.com/CrayLabs/SmartRedis</u>
- Seminar given at NCAR: <a href="https://www.youtube.com/watch?v=2e-5j427AS0">https://www.youtube.com/watch?v=2e-5j427AS0</a>
- SmartSim Case studies: <u>https://github.com/CrayLabs/SmartSim-Zoo</u>
- SmartSim Slack workspace (link on repo)



### **CLOSING WORDS**

#### The value in combining AI and Simulations with SmartSim

- SmartSim can be used to create novel applications that involve both HPC and AI
- The lightweight SmartRedis client is easy to embed into simulation code
- Using a data-in-motion philosophy is key to creating complex AI/HPC workflows
- HPC and AI workflows have the potential to open up new branches of scientific discovery

#### **Opportunities for Collaboration**

Our team is open and eager to work with scientists across many domains

- Use simulations to generate data and symbolic regression for automatic equation discovery
- Building RL or blackbox optimizer to tune simulations for particular goals
- Improve accuracy of simulations by integrating ML models
- Adapt generative AI techniques for ML-assisted engineering design

Questions, comments, other new ideas?

# HANDS-ON WORKSHOP

RUNNING AN AI-ENABLED OCEAN SIMULATION

# OVERVIEW OF THE WORKSHOP EXERCISE

- By the end of the workshop:
  - Understand how to instrument a simulation with SmartRedis
  - Write a SmartSim driver script to configure and launch a simulation
  - Be aware of Frontier-specific best practices
- Exercise
  - Run a toy version of the ocean example
    - Idealized North Pacific ('double gyre')
  - Components:
    - SmartSim modulefile
    - MOM6 ocean executable
    - Pre-trained ML model
    - SmartSim driver
  - Phase 1: Overview of implementation
  - Phase 2: Run the example
  - Phase 3: Scale the simulation
- Resources
  - <u>https://github.com/CrayLabs/OLCF\_SmartSim2023.git</u>
  - Ask SmartSim Team Members in the Slack Channel for help





Courtesy of Cheng Zhang [Princeton]

### SMARTREDIS IMPLEMENTATION IN MOM6

```
input(1) = "input"//CS%key suffix
    MOM6 is written in pure Fortraput(1) = "model_input"//CS%key_suffix
•
                                      model output(1) = "model output"//CS%key suffix
                                      output(1) = "output"//CS%key suffix
                                      ! script 1
                                      db return code = CS%client%run script(CS%script key, "pre process", input, model input)
                                     if (CS%client%SR error parser(db return code)) call MOM error(FATAL, "Run script1 in the database failed")
                                      ! ML model
                                      call cpu_clock_begin(CS%id_run_model)
                                      db return code = CS%client%run model(CS%model key, model input, model output)
                                  if (CS%client%SR_error_parser(db_return_code)) call MOM_error(FATAL, "Run model in the database failed")
                                      call cpu clock end(CS%id run model)
                                      ! script 2
                                      call cpu clock_begin(CS%id_run_script2)
                                      db return code = CS%client%run script(CS%script key, "post process", model output, output)
                                      if (CS%client%SR_error_parser(db_return_code)) call MOM_error(FATAL, "Run script2 in the database failed")
                                      call cpu_clock_end(CS%id_run_script2)
                                    ! extract the output from Python
                                      call cpu_clock_begin(CS%id_unpack_tensor)
                                      db_return_code = CS%client%unpack_tensor(output(1), out_for, shape(out_for))
                                      if (CS%client%SR error parser(db return code)) call MOM error(FATAL, "unpack tensor from the database failed")
                                      call cpu clock_end(CS%id_unpack_tensor)
```

# THE MOM6 SMARTSIM DRIVER

• Link to driver:

https://github.com/CrayLabs/OLCF\_SmartSim2023/blob/main/smartsim\_drivers/double\_gyre/call\_MOM6.p

- Main steps:
  - Create a SmartSim Experiment
  - Set the resources needed to launch and run MOM6
  - Attach and configure a MOM6 simulation
  - Configure a 'colocated' database to run alongside the simulation
  - Add ML model and pre/post-processing scripts
  - Launch the model
- Variation 1:
  - Scale the number of CPUs used by the database and the simulation
- Variation 2:
  - Enable CPU 'pinning' according to Frontier's reserved system CPUs

# THE IMPORTANCE OF PINNING ON FRONTIER

- Colocated database and simulation code are slightly asynchronous
  - Simulation code waits for ML inference to complete
  - Pinning prevents processes from changing cores while waiting, reducing overall performance
- Frontier reserves cores for system processes
  - Pinning those CPUs is not allowed
- SmartSim provides custom pinning for collocated database, Slurm allows cpu pinning of the application
- Rerun:
  - python call\_MOM6.py --enable-pinning



### **CLOSING AND TAKEAWAYS**

- Use this SmartSim driver as a starting point for your own simulations
- Other examples can be found in the tutorials: <u>https://github.com/CrayLabs/SmartSim/pkgs/container/smartsim-tutorials</u>
- Please contact us at our Slack channel: <u>https://join.slack.com/t/craylabs/shared\_invite/zt-1z9cmf1yp-PrzmFadoofJ0hp5XHtDcMQ</u>