Sandia's Journey toward Trustworthy and Efficient Foundation Models

Justin Newcomer, PhD

Deputy Executive, Advanced Simulation and Computing Program

Wednesday, March 19, 2025

SOS27 Workshop

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2025-02972C

Challenges for Trustworthy and Efficient Foundation Models

- Data Challenges
  - Diverse and Representative Data
  - Data Privacy and Security
- Model Design Challenges
  - Complexity, Scalability, Robustness
- Evaluation Challenges
  - Comprehensive Metrics and Benchmarking
  - Real-world Testing
- Deployment Challenges
  - Model security
  - Model awareness and guardrails
  - Continuous learning and adaption
  - Assessing Uncertainty

Challenges for Trustworthy and Efficient Foundation Models

- Model Architecture Optimization
  - Pruning, Quantization, Knowledge Distillation
  - Model Compression
- Efficient Training Algorithms and Techniques
  - Dropout and Early Stopping
  - Mixed Precision Training
  - Batch Normalization and Layer Normalization
- Hardware Considerations
  - Dataflow architectures and accelerators
  - Distributed training
- Environment and Framework Optimization
  - Efficient Libraries
  - Shared Infrastructure and Knowledge Base

Trust, Trustworthy, and Credibility

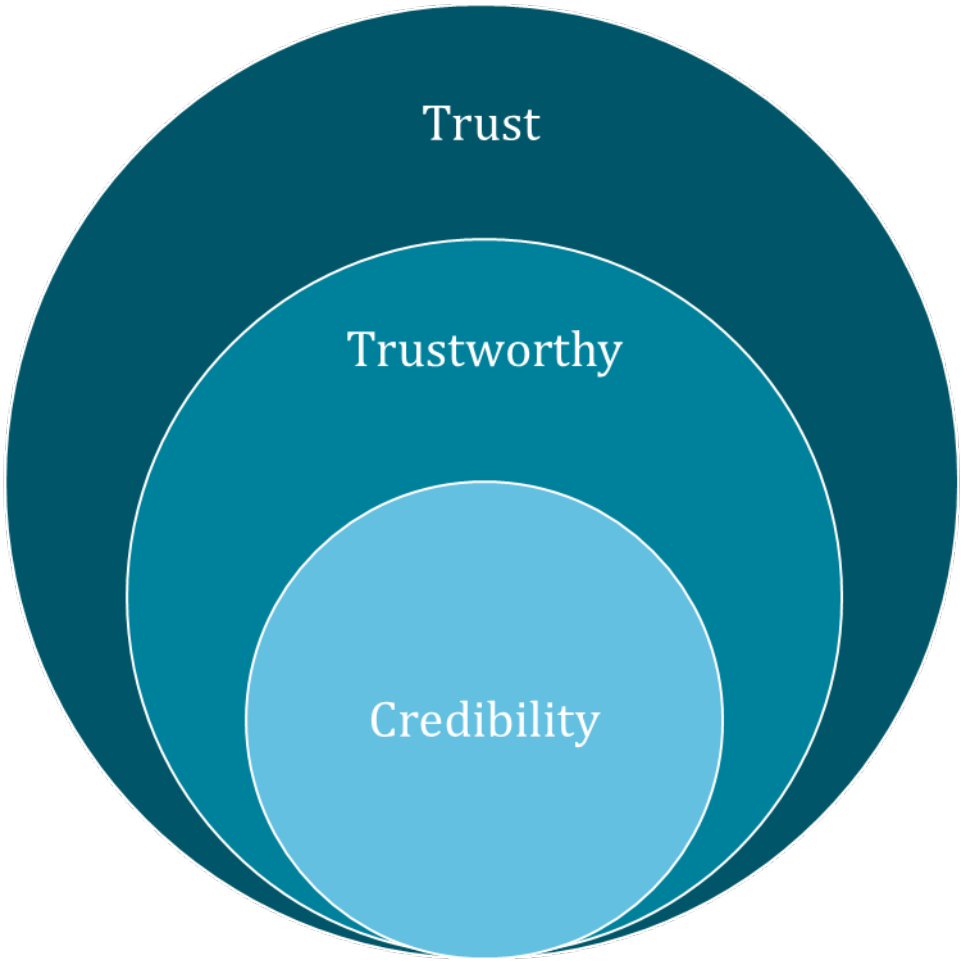

Trust: Relationship between AI and decision maker

- Model inference and/or predictions support decision-making process.
- Appropriate trust and use

Trustworthy: Objective measures of correctness, reliability, and security

- Data properties and biases addressed
- Domain information incorporated
- Interpretability / explainability where appropriate

Credibility: identifies the technical basis of the model.

- Model selection
- Verification and validation
- Uncertainty quantification

Credibility supports model trustworthiness, which in turn supports (does not guarantee) trust.
Trust in a model does not guarantee that credibility has been established.

BANYAN – An Institute for Generative AI @ Sandia

Bring together GenAI projects and use the synergy across the lab
Share datasets, models, software stacks, knowledge, & industry interactions

"roots"
ASC Federated Models development & several smaller projects

ASC Federated Models with LANL (PI: Dan O'Malley) and LLNL (PI: Josh Kallman):
Federated model training across three labs, Build large AI-ready data sets for ND, Benchmarks that require reasoning

ASC Industry Collaboration
Joint effort on AI/ML with Cerebras and NVIDIA

ASC Co-pilots for code
Kokkos Copilot for parallel code generation

"trunk"
Two Large projects

BANYAN LDRD
Multimodal model training for mission critical components

PARADIGM LDRD
Neurosymbolic / GraphRAG approach for multimodal data

~ 50 member team

BANYAN - Multimodal model training for critical components

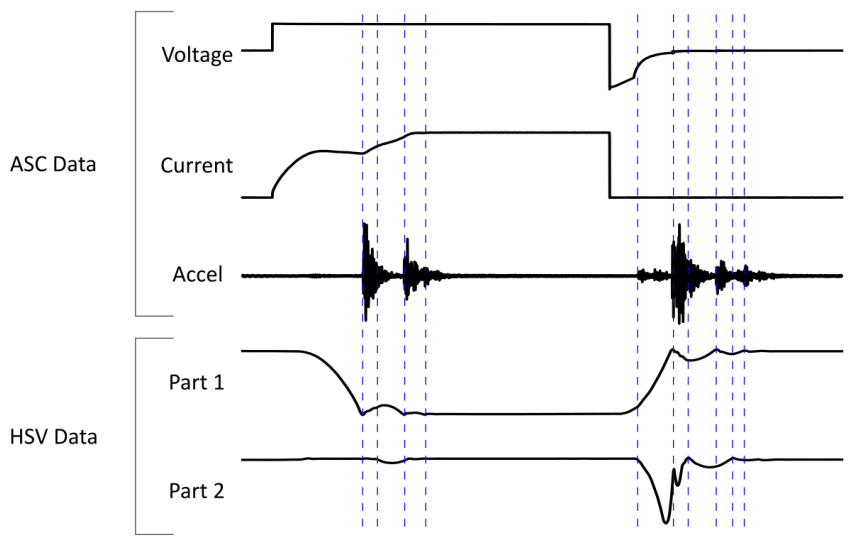

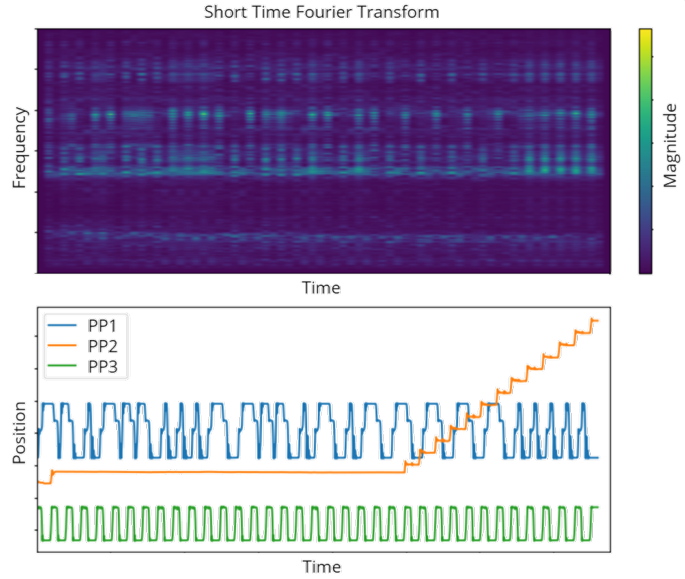

- Phenomenology:
  - Part movements are largely autocorrelated, but show abrupt changes at “events”
  - “Events” also seen in input data (as inflection points in electricals & “spurts” of sound/vibrations)
- Hypothesis:
  - Part movements can be replicated using a combination of auto-correlation-based models and input from exogenous inputs
  - The spectral info in acoustics can be used to detect changes in part movement and the timing of the changes

Producing critical metrics to influence design, identifying precursor behavior indicative of possible failures, and providing more insight and faster conclusions for critical anomaly resolution

PARADIGM - Increase Trustworthiness of Generative AI Models

- Can trustworthiness be achieved by integrating symbolic reasoning with GenAI models?
- Will augmenting the GenAI or constraining the context be sufficient to engender trust?
- Test case: Estimate the yield of the Beirut blast (August 2020) using open-source data
- Thrust areas:
  - Neural augmentation: Configure generative models to estimate blast yield, using seismic, imagery, infrasound, and cellphone video data
  - Knowledge graphs: Design knowledge graphs mimicking the thought processes of SME to filter out bogus output or constrain an LLM via RAG
  - Trust: Devise metrics that allow conformal predictions to quantify/detect shortcomings in the workings of LLMs
  - System: Investigate, configure & deploy new architectures needed to run GenAI models on Sandia premises

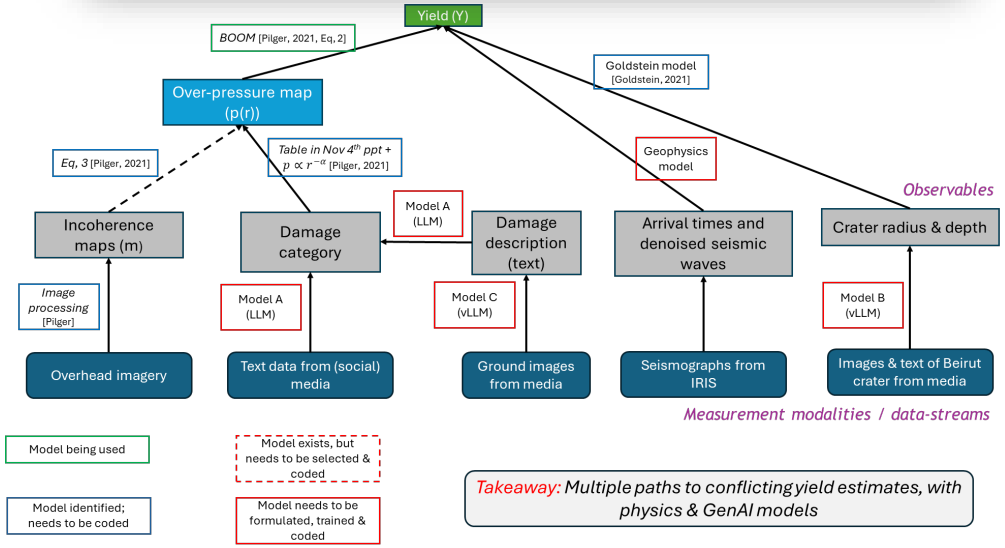

Demonstrating the necessary and sufficient conditions for combining GenAI and physics models will significantly increase the trustworthiness of GenAI models

Kokkos Copilot: Can we use GenAI to be a copilot for programmers to write parallel portable C++ code?

Motivation: Writing parallel code is hard; Writing portable parallel code with C++ is harder

- Kokkos – performance portable programming model, close to bare metal programming
- Needs knowledge of the domain, C++, GPUs and portability issues

Data Collection and Data Generation

- Three levels of Trusted data sources
- 30 repositories from various Kokkos-based open-source projects
  - Adoption of Kokkos in multiple software projects helped us with trusted training data
- 500 code and prompt pairs automatically generated from comments and basic building blocks following the comments (Automatically generated, but human curated)
  - Assumption: Comments actually reflect what is in the next basic building block
- 138 human generated code and prompt pairs as "gold standard"
  - Not released publicly, used to do our quick evaluation

Three different data sets with three levels of trustworthiness are vital to overall success

PCMM: Predictive Capability Maturity Model

The computational simulation (CompSim) credibility process assembles and documents evidence to ascertain and communicate the believability of predictions that are produced from computational simulations.

Evidence Basis

- Plan
- Execute
- Organize & Analyze

Elements

- Categories for collecting evidence
- Dependent on model paradigm

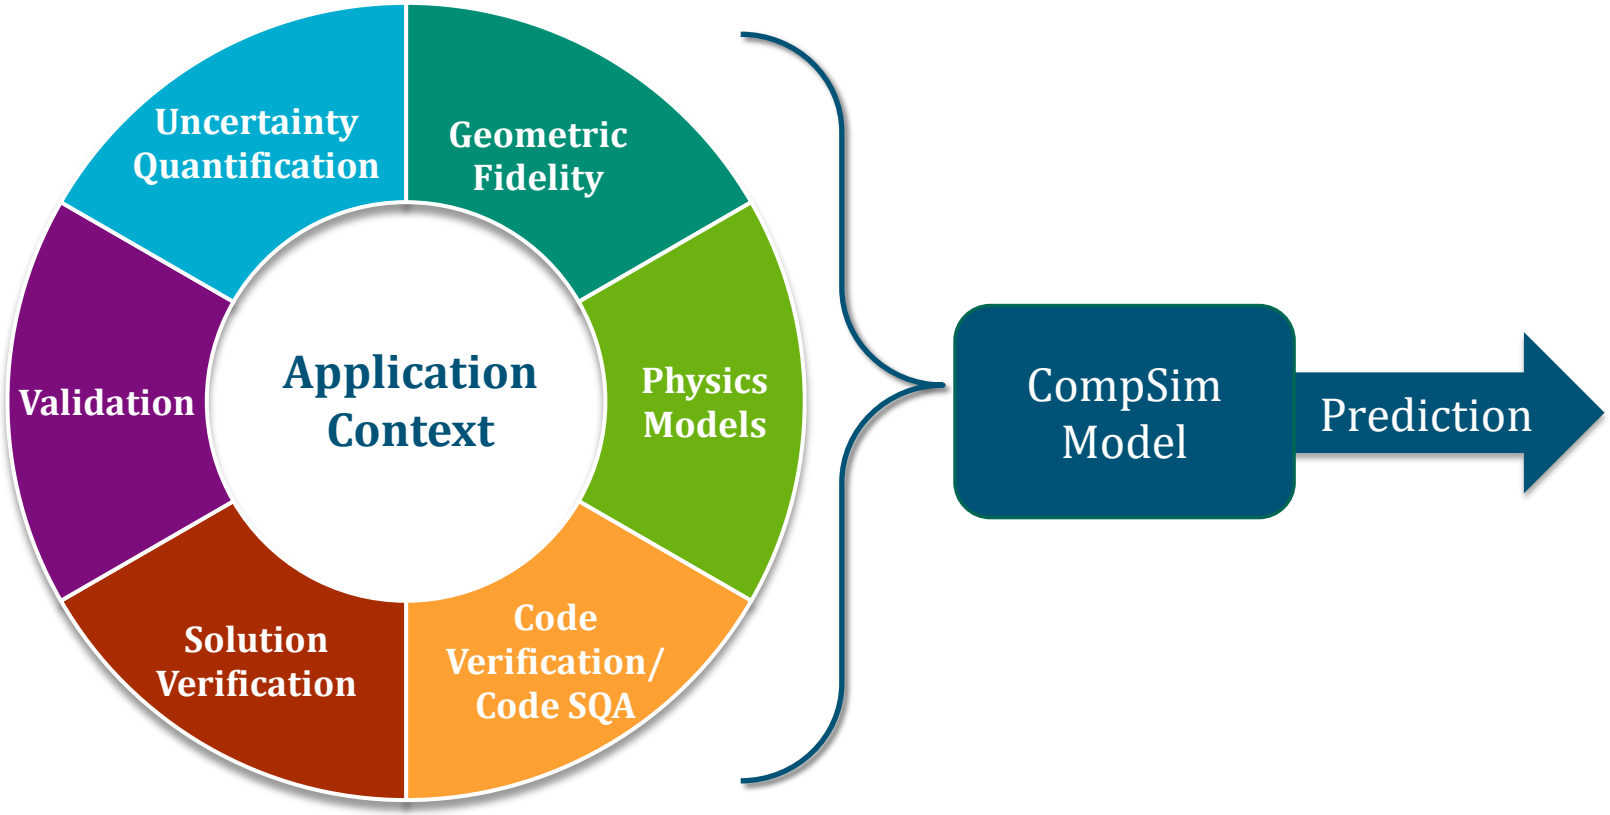

PCMM evolved from industry standards and lab/academic collaboration, forming 30+ years of experience in verification, validation, and uncertainty quantification (VV/UQ) for complex problems with limited data.

"Predictive Capability Maturity Model for Computational Modeling and Simulation" by Oberkampf, W.L., Pilch, M., and Trucano, T.G., SAND2007-5948

Developing a Credibility Model for AI

Credibility process assembles and documents evidence to ascertain and communicate the believability of predictions that are produced from those models.

UQ and Validation are currently the core areas of research that exists for AI credibility.

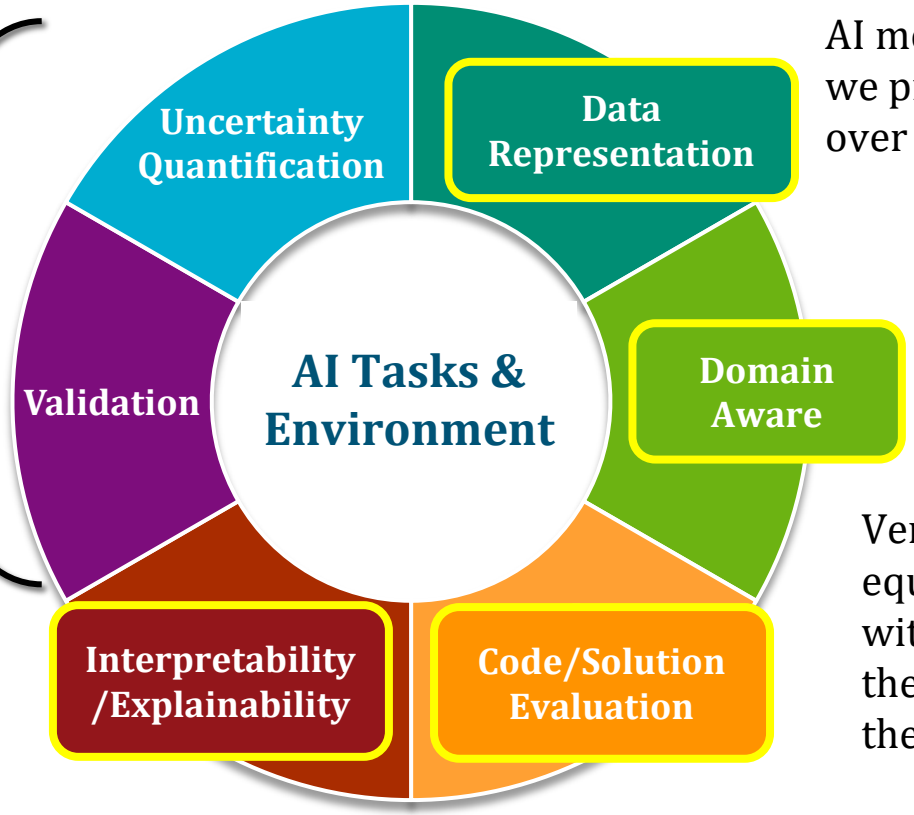

AI models learn patterns from data, so we prioritize data representations over the geometries.

AI is applied more broadly and it is not only physical principles we want to preserve.

Verification asks “are we solving the equations correctly”, AI models do not start with equations. We may want to reconsider the terminology for AI and address more of the community of practice methods.

A credibility model for AI is essential to ensure confidence in deployed AI solutions

Our work builds upon NNSA's 20+ years of experience in verification, validation, and uncertainty quantification (VV/UQ)

Counter Adversarial Machine Learning (CAML)

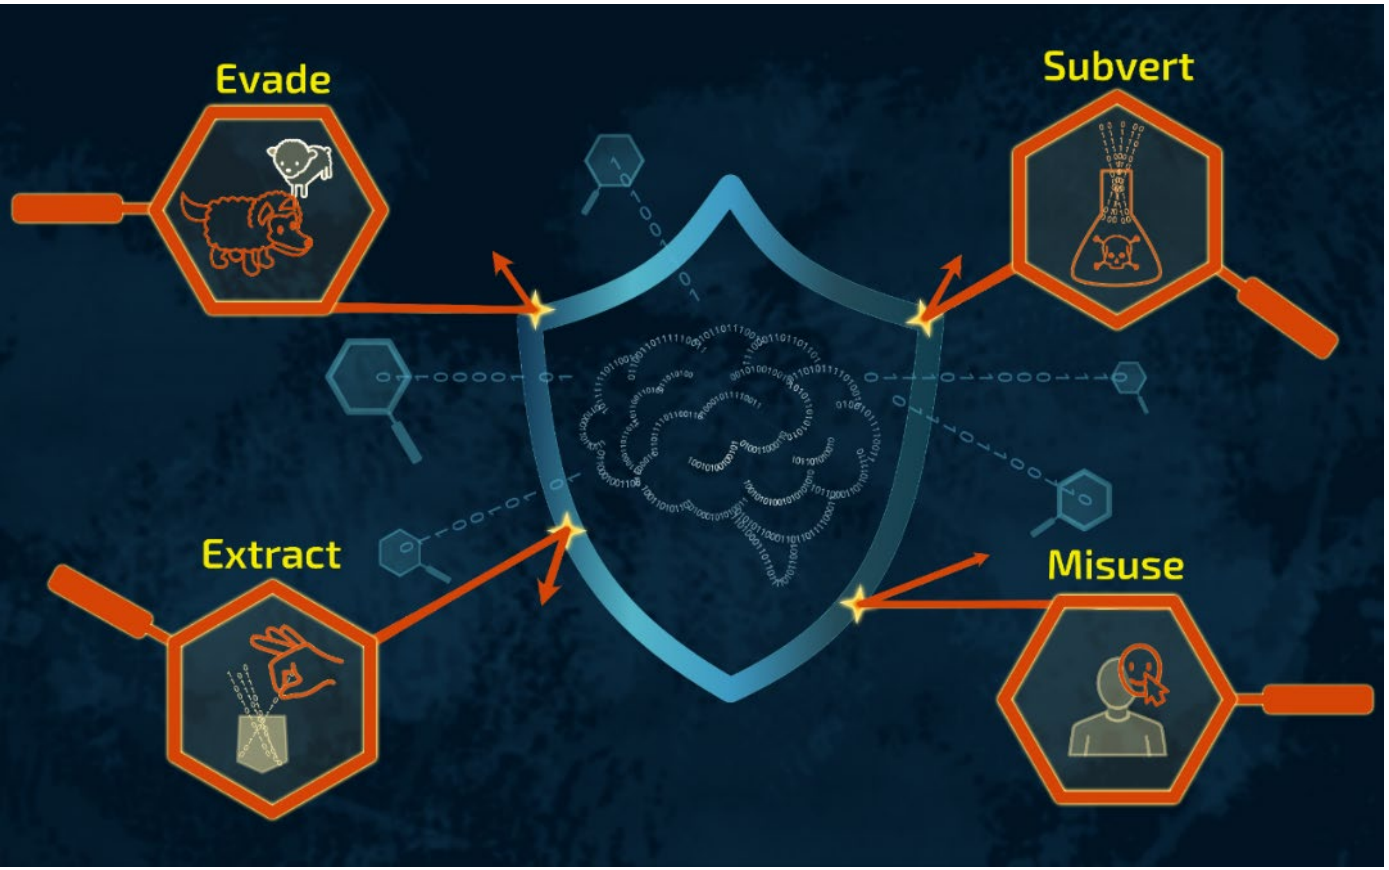

EVADE attacks make minimal changes to the input so it can fool the ML system into misclassifying them, while the underlying input is still essentially its original self

SUBVERT attacks alter some foundation of the model to create a backdoor issue in it, often by "poisoning" the training data or code to create attacker-known backdoors

EXTRACT attacks steal data from the model

MISUSE attacks employ benign ML systems for malicious purposes, such as "deep fake" attacks where any actor can drive another person's facial expressions as they see fit

Defending national security systems against attacks

NNSA Tri-lab Federated Learning Project

- This project aims to develop a large foundation model tailored for security applications across the three National Nuclear Security Agency (NNSA) laboratories (Sandia, LANL, LLNL)
  - No data sharing between labs to maintain trust and security of private and/or sensitive data
  - Individual models will be trained or fine-tuned on local and different computer architectures providing robustness, fault tolerance, and redundancy
  - Collaborating with NVIDIA to explore NVFLARE to exchange model weights between the tri-labs and demonstrate assembly of the weights into a single AI model
    - Several options for federated training – Swarm mode where instances exchange weights with the server
    - Also exploring FLOWER and APPFL frameworks
- Federated learning provides opportunities for more efficient foundation model training
  - Distributed training and decentralized data storage, efficient resource utilization, efficient communication protocols
- Future work will explore use risks of the combined model and the effectiveness of guardrails

The federated learning framework will set a precedent for collaborative and efficient AI model development in environments where data sharing is restricted, positioning the NNSA laboratories at the forefront of AI innovation in national security

1 Trillion Parameter Model Training on a single Cerebras CS-3

- Wafer Scale architecture using on-chip SRAM for reduced data movement and lower energy costs
- The size of the model is not constrained by the memory capacity on chip due to the weight-streaming architecture
- All model weights are stored externally in a unique 55 terabyte MemoryX device
- Installation to Trillion parameter model training in ~5 weeks
- The model was then scaled up seamlessly to 16 CS-3 systems, demonstrating a step-change in the linear scalability and performance of large AI models

Sandia has fielded a four wafer system – Kingfisher

Achieving the above typically requires thousands of GPUs, megawatts of power, and many weeks of hardware and software configuration

Cerebras Wafer Scale Engine outperformed the world's leading supercomputers in molecular dynamics

1,470 time-steps/s

Frontier Supercomputer

37,888 GPUs
21 megawatts
Fastest general purpose supercomputer

980,000 time-steps/s

Anton 3 Supercomputer

512 Custom ASICs
400 kilowatts
Runs only molecular dynamics

1,100,000 time-steps/s

Cerebras CS-2

1x Wafer Scale Engine
27 kilowatts
General purpose AI & HPC Accelerator

Using just a single wafer-scale chip and a tiny fraction of the power of specialized supercomputers, we were able to run MD simulations at over 1 million steps per second - an industry record


SANDIA
NATIONAL
LABORATORIES
AI-info@sandia.gov
WWW.SANDIA.GOV/AI/