TRAINING PHYSICS-INFORMED NEURAL NETWORKS FOR ADDITIVE MANUFACTURING PDES USING AWS CLOUD COMPUTING

FAMU-FSU College of Engineering

NIKKI ZAHEDI, SOUMIK SARKAR, RAGHAV GNANASAMBANDAM

M²M Math Helps Manufacturing

FSU FLORIDA STATE UNIVERSITY

BACKGROUND

- Laser-based additive manufacturing involves complex heat transfer governed by PDEs, where accurate temperature modeling is critical for predicting melt pool behavior and material properties.
- Traditional methods like FEM are effective but computationally expensive and difficult to scale for rapidly changing thermal fields [2].
- Physics-Informed Neural Networks (PINNs) incorporate PDE constraints directly into the training loss, enabling mesh-free approximation of temperature fields without labeled simulation data [1].

OBJECTIVE

- Apply Physics-Informed Neural Networks (PINNs) to model heat transfer dynamics in additive manufacturing.
- Deployed and optimized the PINN training workflow on AWS for scalable, cost-efficient execution.
- Configured cloud compute environments (EC2, SageMaker).
- Implemented checkpointing and GPU/CPU switching to reduce runtime and cost.

REFERENCES

PHYSICS-INFORMED NEURAL NETWORKS

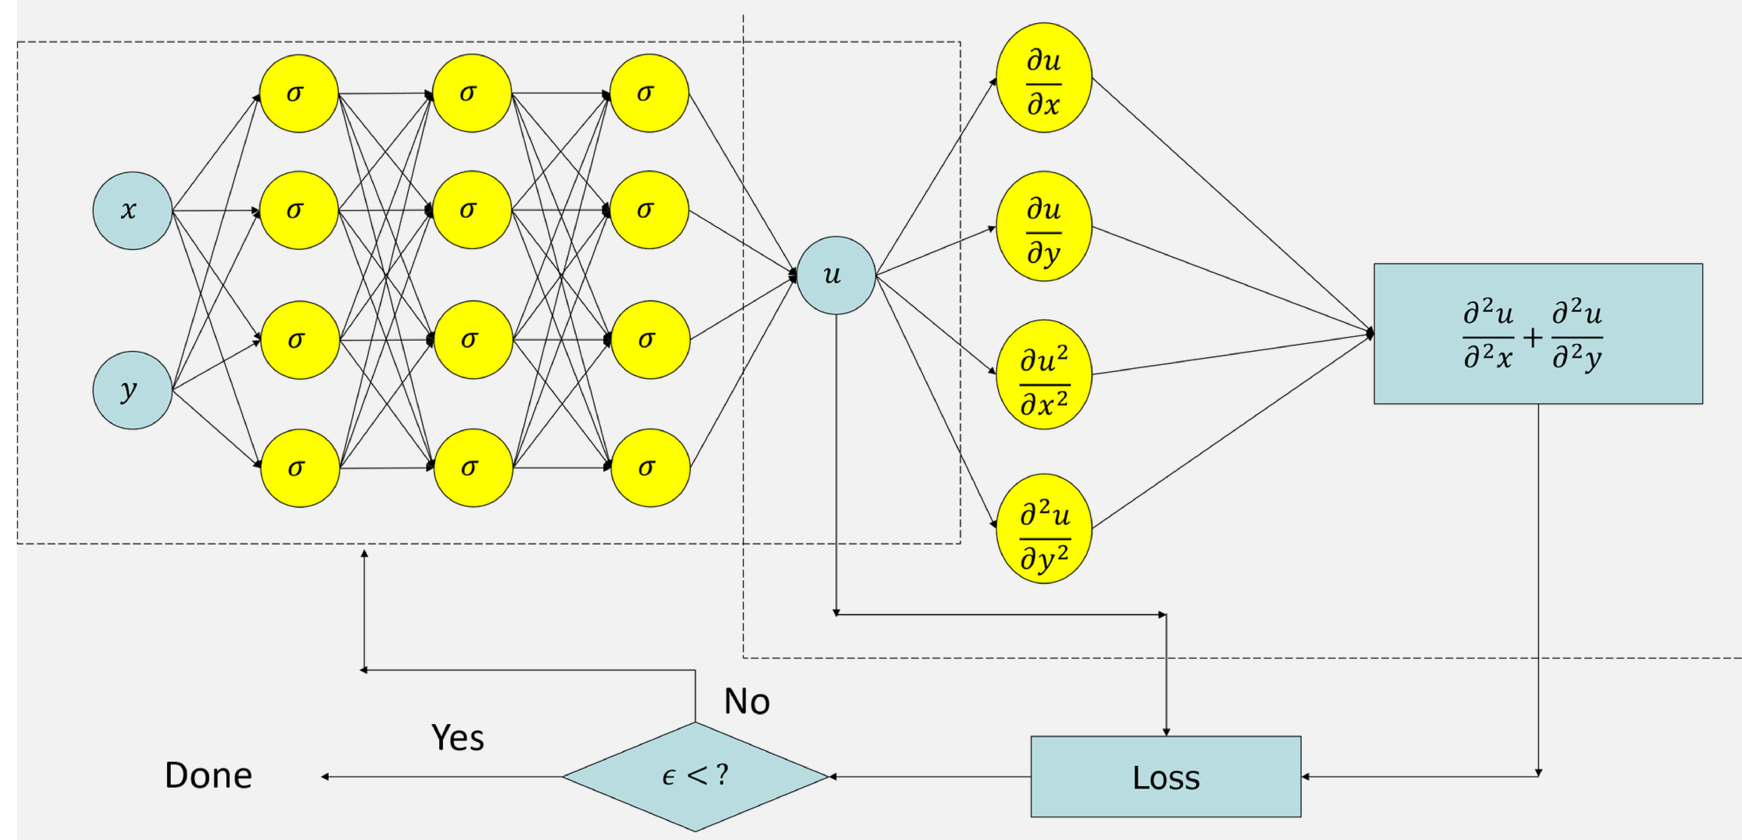

FIGURE 1: PHYSICS-INFORMED NEURAL NETWORK (PINN) ARCHITECTURE ENFORCING THE HEAT EQUATION THROUGH AUTOMATIC DIFFERENTIATION AND PHYSICS-BASED LOSS TERMS.

PINN MODELING FRAMEWORK

- The neural network represents the temperature field as a function of space and time.
- The heat equation, boundary and initial conditions are enforced through the loss function.
- Training minimizes physics-based residuals, enabling approximation of thermal behavior without labeled data [1].

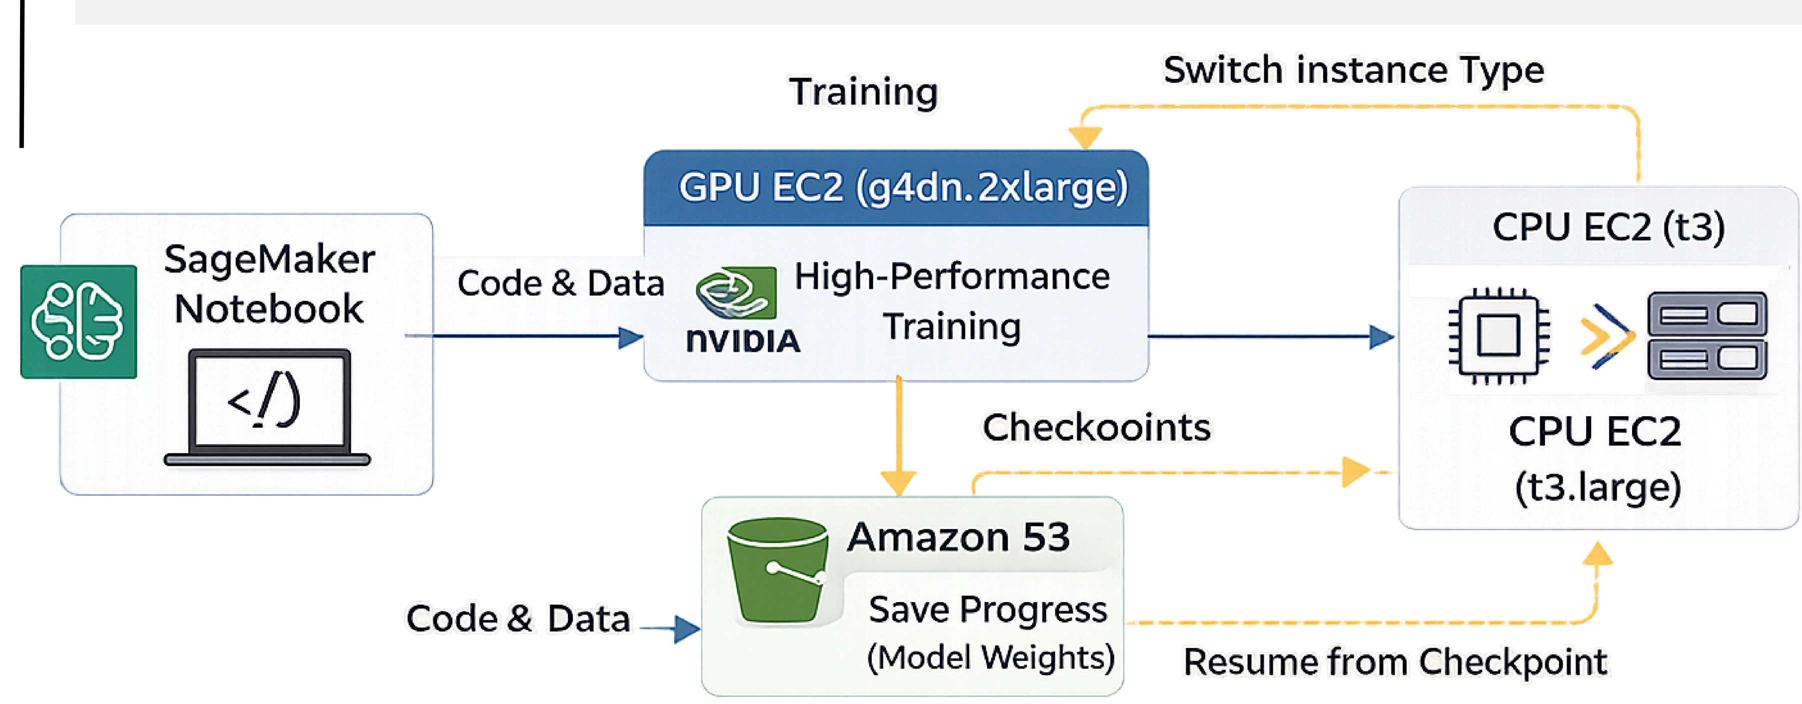

FIGURE 2: AWS CLOUD WORKFLOW FOR SCALABLE PINN TRAINING USING GPU ACCELERATION, CHECKPOINTING IN AMAZON S3, AND COST-EFFICIENT CPU CONTINUATION.

AWS DEVELOPMENT STRATEGY

- Deployed training workflows with EC2 and SageMaker.
- Used GPU instances for accelerated training and CPU instances for lower-cost evaluation.
- Stored checkpoints in Amazon S3 to enable seamless job resumption across instance types.

TOTAL COMPUTE COST FOR PINN TRAINING ON AWS

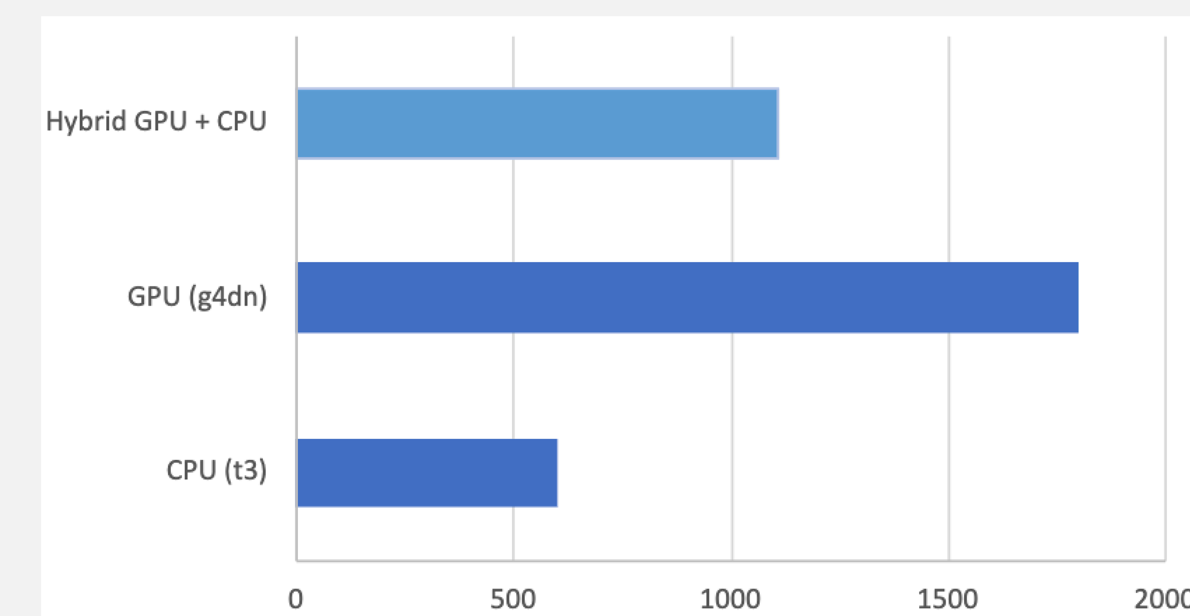

FIGURE 3: TOTAL COMPUTE COST ACROSS AWS TRAINING ENVIRONMENTS, HIGHLIGHTING COST SAVINGS FROM THE HYBRID GPU–CPU WORKFLOW..

PERFORMANCE AND COST OPTIMIZATION

- Benchmarked multiple instance and storage configurations.
- GPU training significantly reduced convergence time compared to CPU-only runs.
- Checkpointing and instance switching reduced overall compute cost by ~30%.

CONCLUSION

- PINNs provide a scalable framework for modeling heat transfer in additive manufacturing.
- Cloud infrastructure enables efficient large-scale training.
- Cost-aware AWS optimization improved runtime, reproducibility, and resource efficiency.