

Background

Particle-in-Cell (PIC) simulations are widely used to study plasmas by tracking many charged “sample particle” while computing the electric and magnetic fields on a grid. They are important for problems such as fusion edge physics, electric propulsion, and space plasmas, but large PIC runs are expensive. At scale, performance is often limited by moving data through memory and across nodes, not just by raw compute. A common source of error in PIC is numerical noise: because we represent a huge number of real particles with a finite number of sample particles, the particle-to-grid step often done with the standard cloud-in-cell method can introduce fluctuations that hide weak physical signals and reduce diagnostic quality. Using smoother, higher-order deposition can reduce this noise, but it usually increases memory traffic and scalability. In this work, we explore a ML approach that learns a smooth, Gaussian-like deposition behavior to reduce high-frequency noise while staying compatible with the PIC loop. We currently have a baseline multicore CPU and single-GPU implementation, and training is performed using a data-parallel approach in PyTorch.

Motivation and Goals

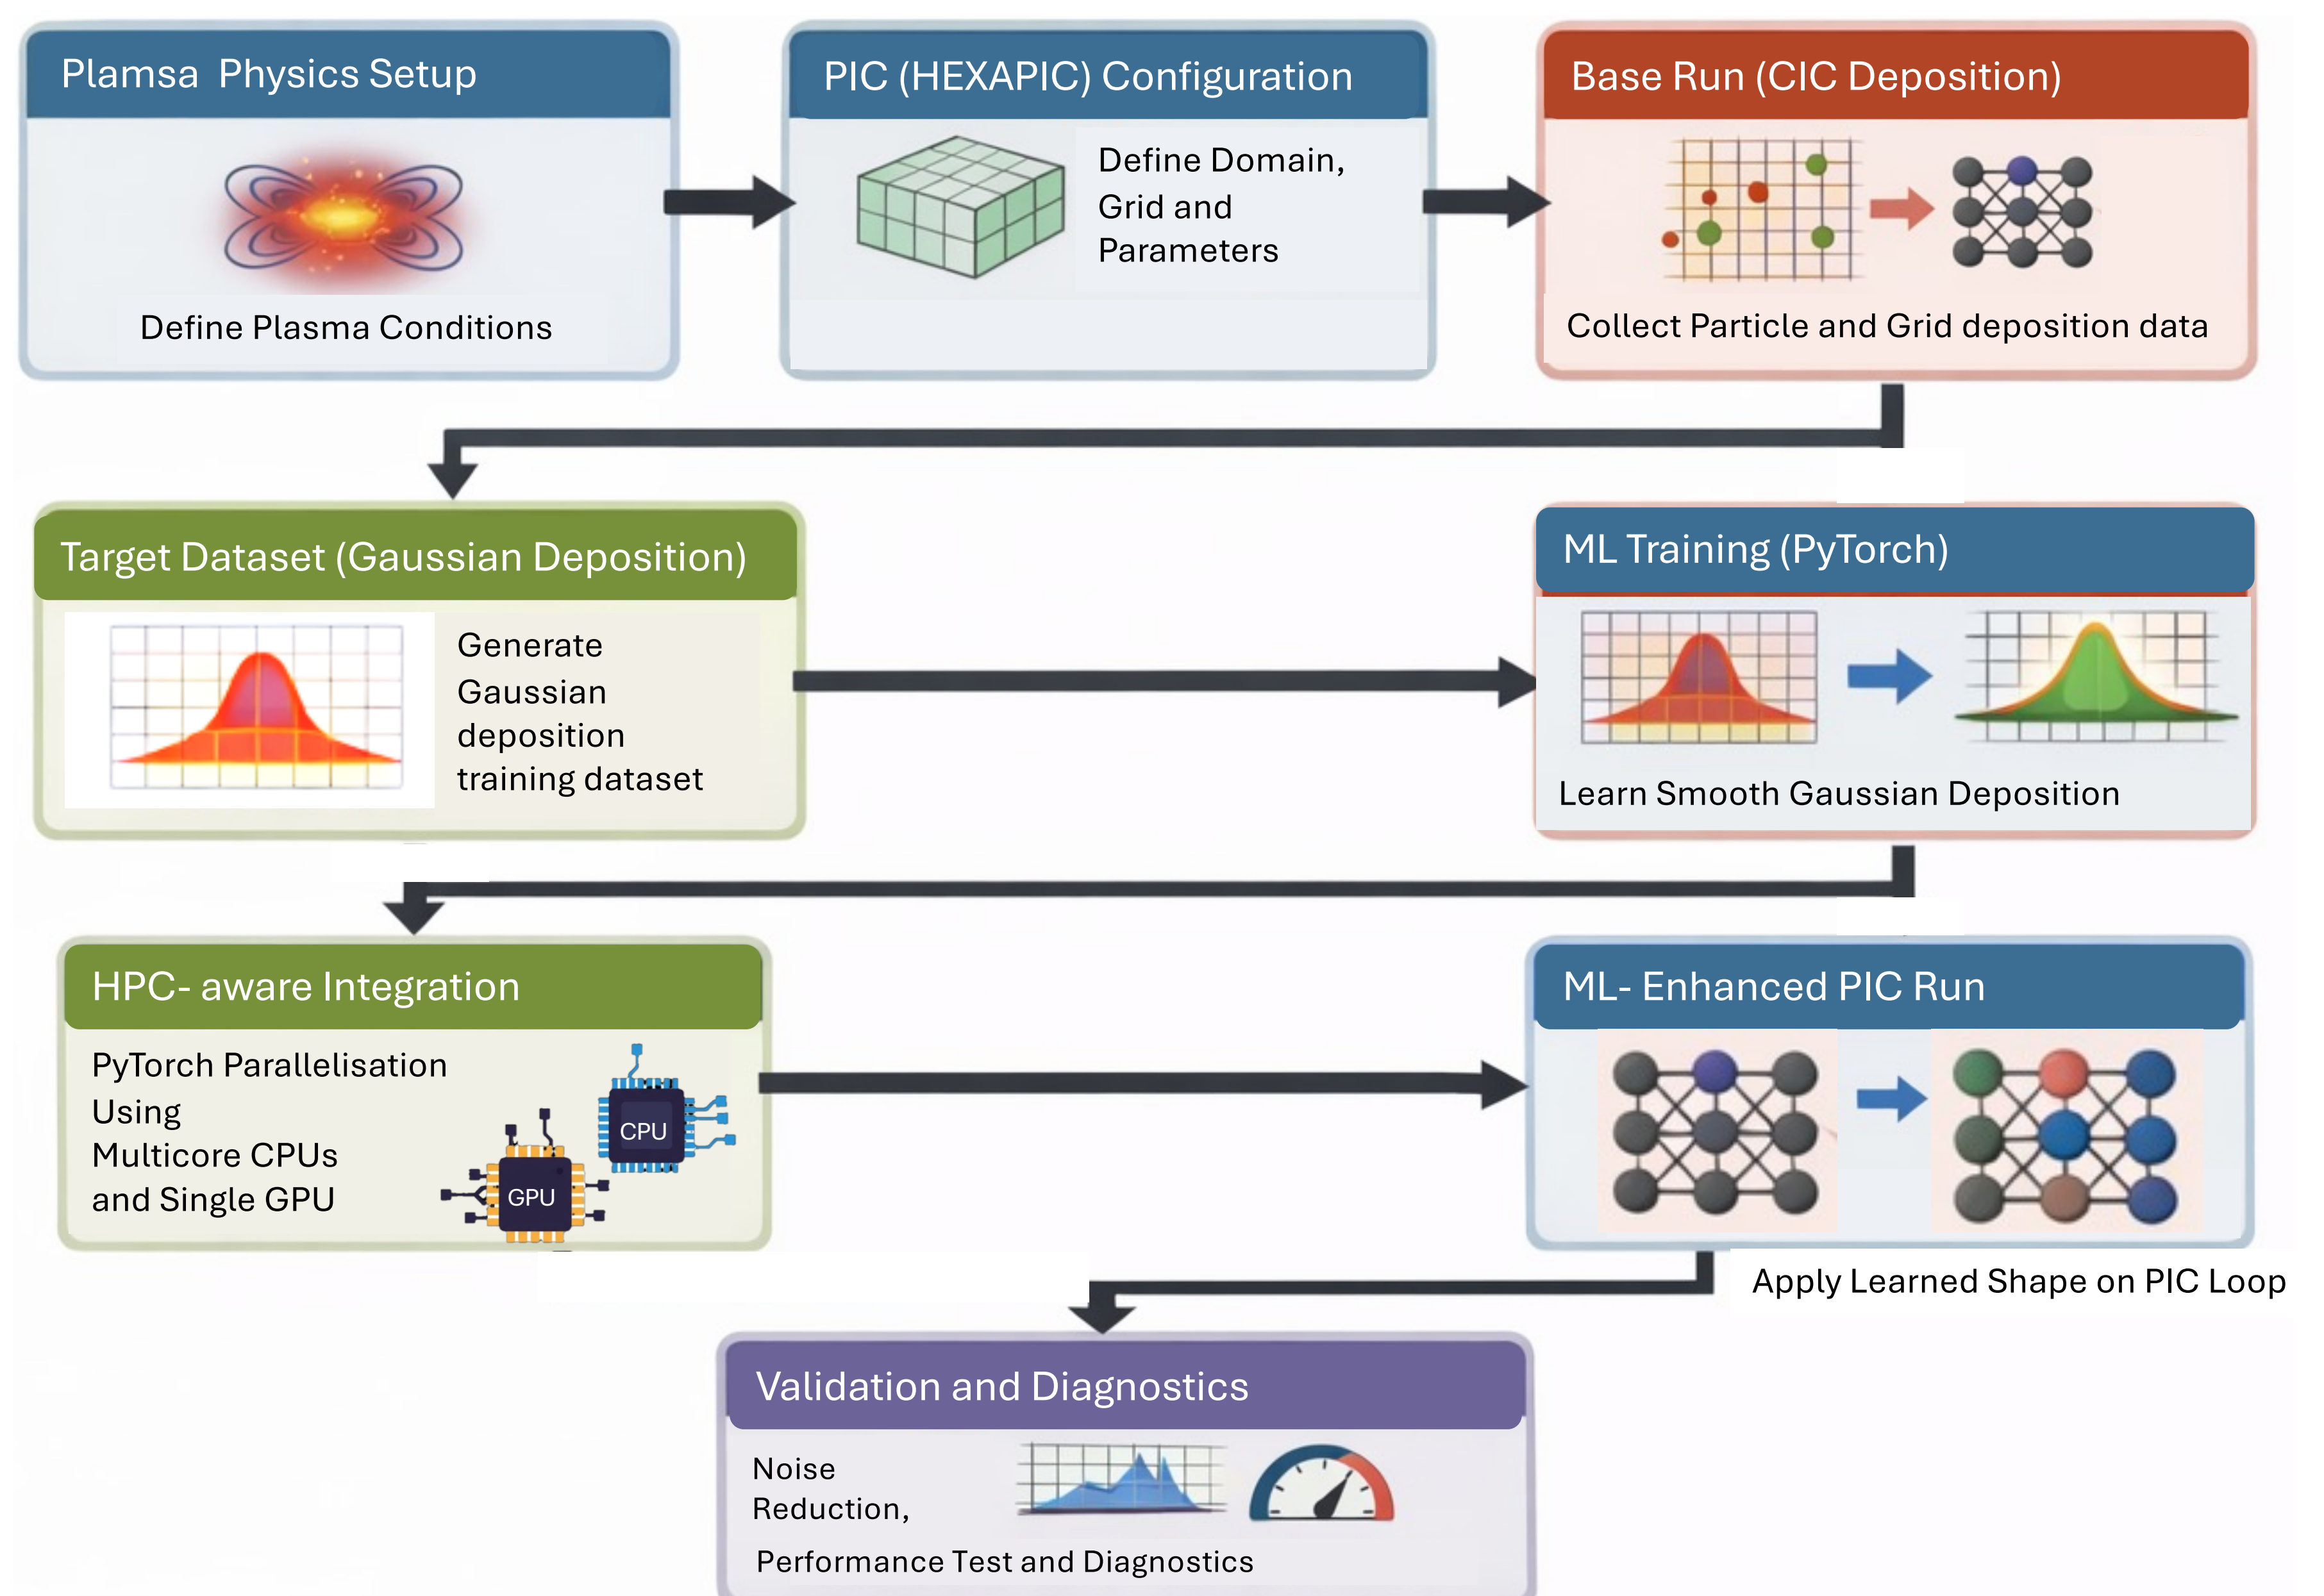
Motivation

- Large PIC studies are expensive, and at scale the main limits are often memory and node-to-node communication, not just compute speed.
- The standard cloud-in-cell (CIC) particle-grid coupling is fast and scalable, but its noise can reduce diagnostic quality and hide weak physical signals¹.

Goals

- Reduce high-frequency noise introduced during particle-to-grid deposition while keeping the PIC loop unchanged in its main structure.
- Integrate a Machine Learning (ML)-based, Gaussian-like smooth deposition approach into the Heterogeneous Exascale PIC (HEXAPIC) workflow.
- Support scalable training and testing using distributed PyTorch (data parallelism, and model parallelism using multicore CPUs and a single GPU.)

Workflow



ML Model

Model form (learned convolution): We assume the mapping is well-approximated by a linear, shift-invariant operator:

- **Separable kernel:** $K(x, y) = k_x \otimes k_y$.
- **Full 2D kernel:** $K \in \mathbb{R}^{K \times K}$.
- **Total charge conservation:** nonnegative and sum to 1.

Training data:

- CIC field patch \rightarrow Gaussian-deposited field patch.
- **Boundary:** Padding selectable to match the domain circular for periodic domains; `reflect/replicate/zeros` otherwise.
- **Loss:** MSE between predicted and target Gaussian fields.


We used the AMD EPYC Rome CPU and NVIDIA A100 GPU for training on the MeluXina HPC machine.

HPC Integration

PyTorch supports various types of parallelization, with data and model parallelization being the most common. In data parallelization, the training model’s data is split across multiple CPUs and GPUs, allowing each computing resource to compute its own gradients, which are then summed to update the model’s weights for the next training iteration. Conversely, model parallelization distributes the layers of the model across different computing resources, requiring the entire dataset to be available for any layer processing. It’s crucial to choose the appropriate parallelization strategy within the HPC framework to accelerate computations; otherwise, even with simple datasets, processing times can increase significantly, despite using HPC capabilities across multicore CPUs and GPUs².

Preliminary Results (Regression)

- **Training setup:** Trained on 51,833 supervised samples on a 64×64 grid using a 2D learned kernel with $K = 9$, covering **mixed-sign charge** (positive and negative plasma charge density).
- **Qualitative agreement:** For an exemplary case, the predicted Gaussian-deposited field closely matches the true Gaussian deposition. The difference map (pred – true) remains near zero, with small residuals localized around the charge peaks.
- **Accuracy:** $MAE = 7.53 \times 10^{-4}$. Mean abs. charge error: 1.246913×10^{-2} .
- Further ML tuning and HPC parallelization are ongoing. Future work will focus on multi-GPU scaling across multiple compute nodes.



References

- ¹ Touati *et al.*, *Plasma Phys. Control. Fusion* **64**, 115014 (2022)
- ² Krishnasamy *et al.*, “Large-Scale Deep Learning Medical Image Methodology and Applications Using Multiple GPUs,” in *ICABME* (2023).