Learning Gaussian-Like Deposition to Reduce PIC Noise in HEXAPIC at Scale

Ayushi Agrawal*[1], Shan-Chang Lin[1], Ezhilmathi Krishnasamy[1,2], and Pascal Bouvry[1]

[1] University of Luxembourg, Luxembourg
[2] University of Ljubljana, Slovenia

*Email: ayushi.agrawal@uni.lu

CONNECTING THE DOTS
ISC High Performance
JUNE 22–26, 2026 | HAMBURG, GERMANY

Background

Particle-in-Cell (PIC) simulations are widely used for modeling fusion edge physics, electric propulsion, and space plasmas, but large PIC runs are expensive, often limited by memory and across nodes, not just by raw compute. A common source of error is numerical noise because we represent a huge number of real particles with a finite number of sample particles, which 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 with multicore CPU and single-GPU implementation, and training is performed using a data-parallel approach in PyTorch.

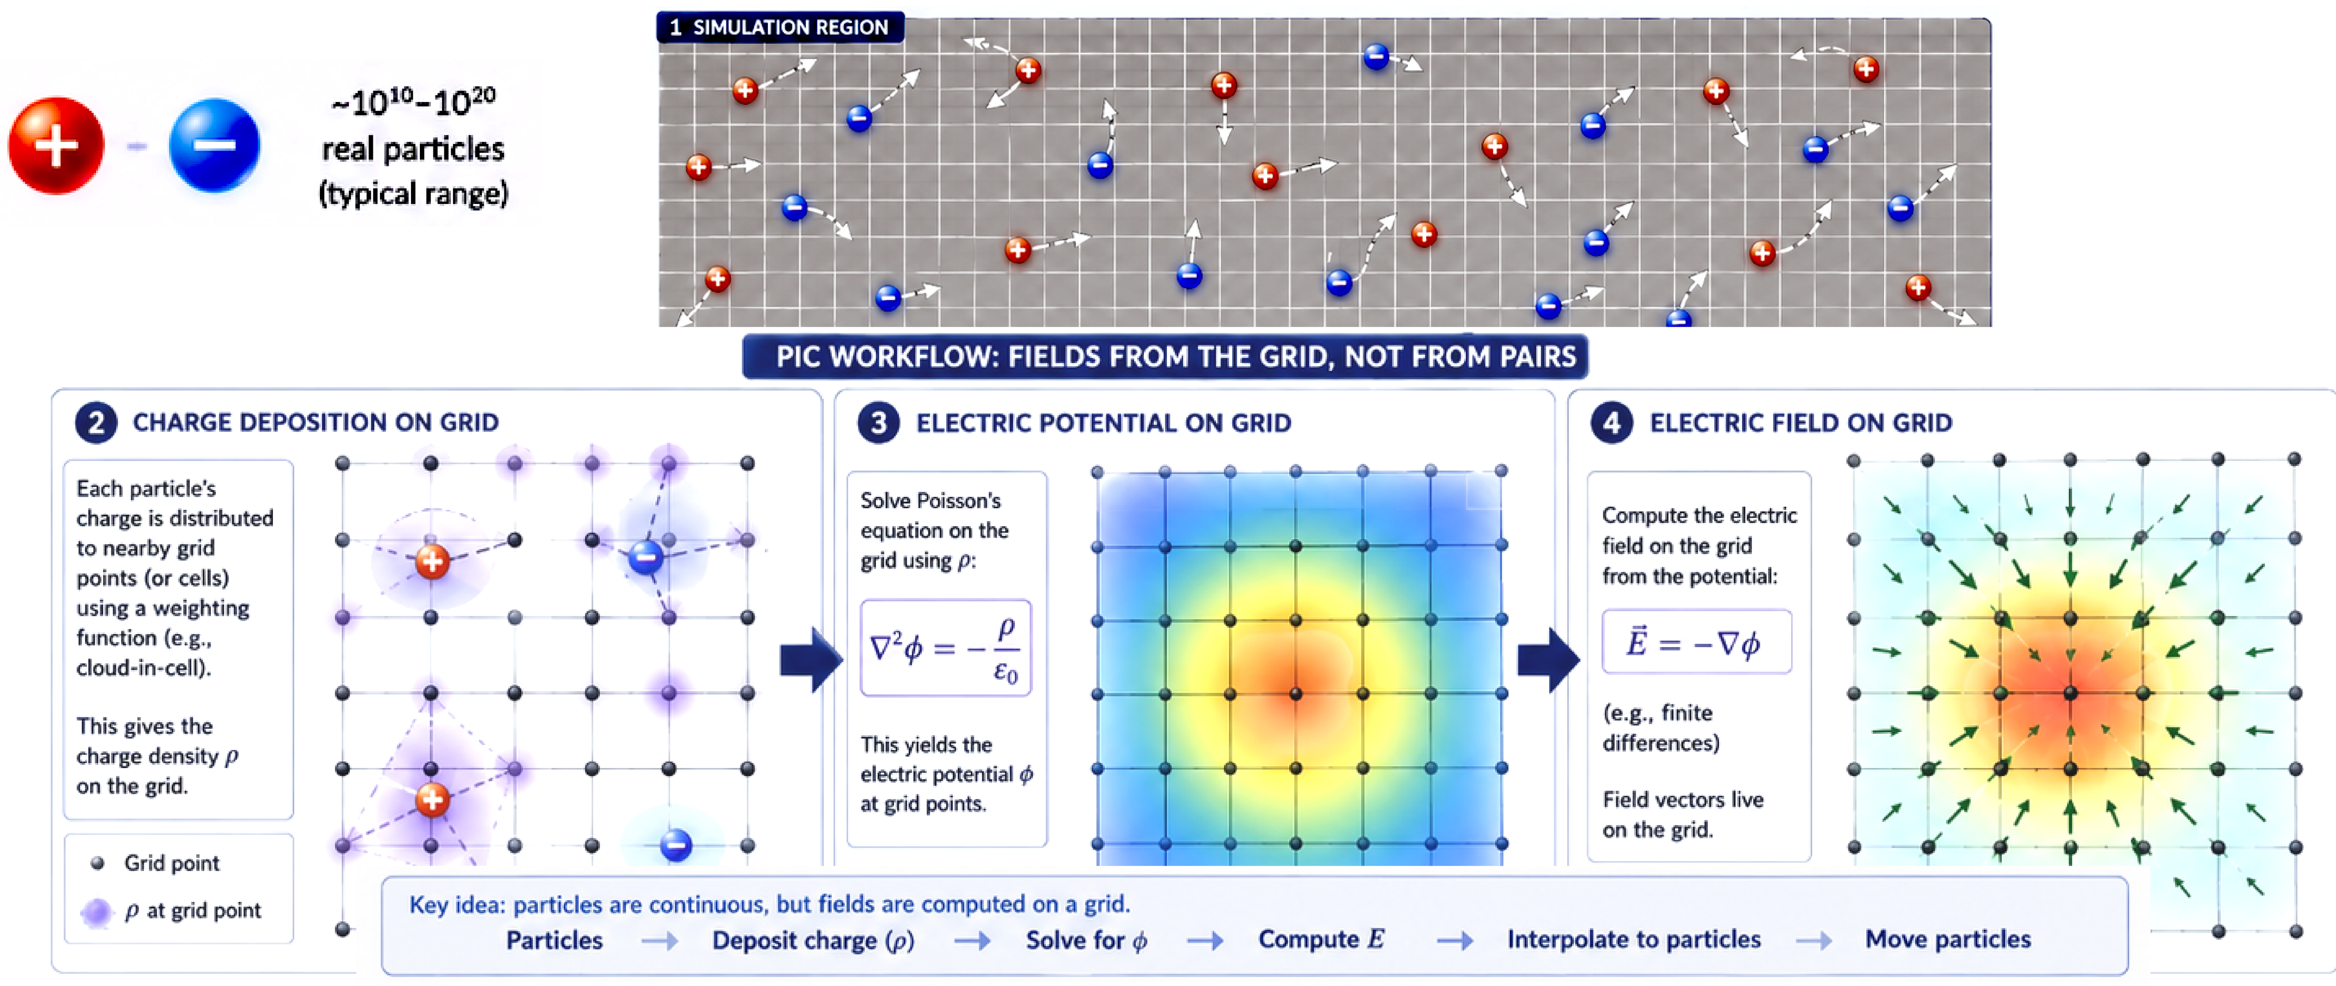

Motivation and Goals

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 signals [1].

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.)

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

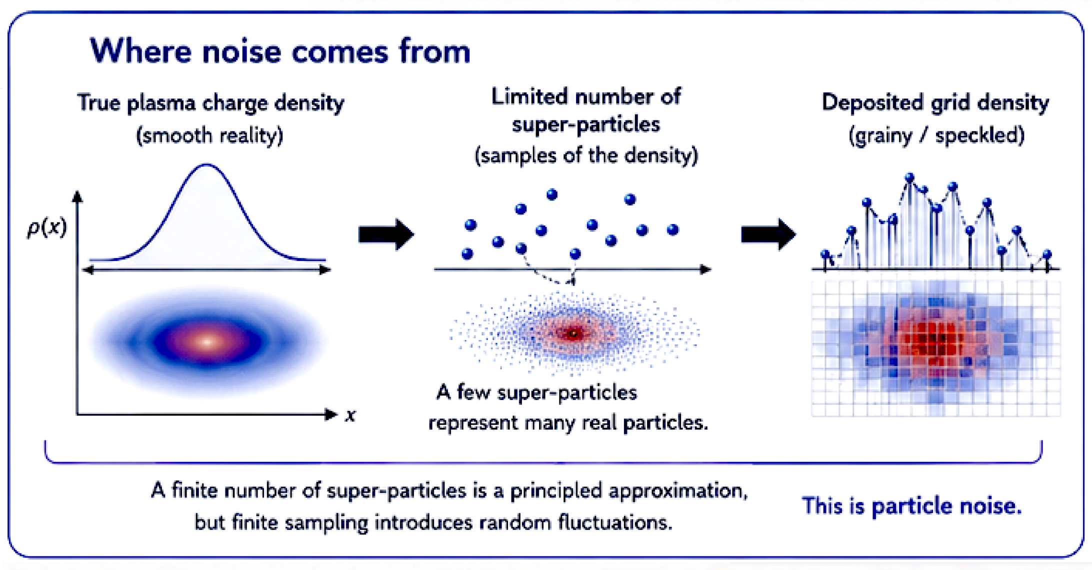

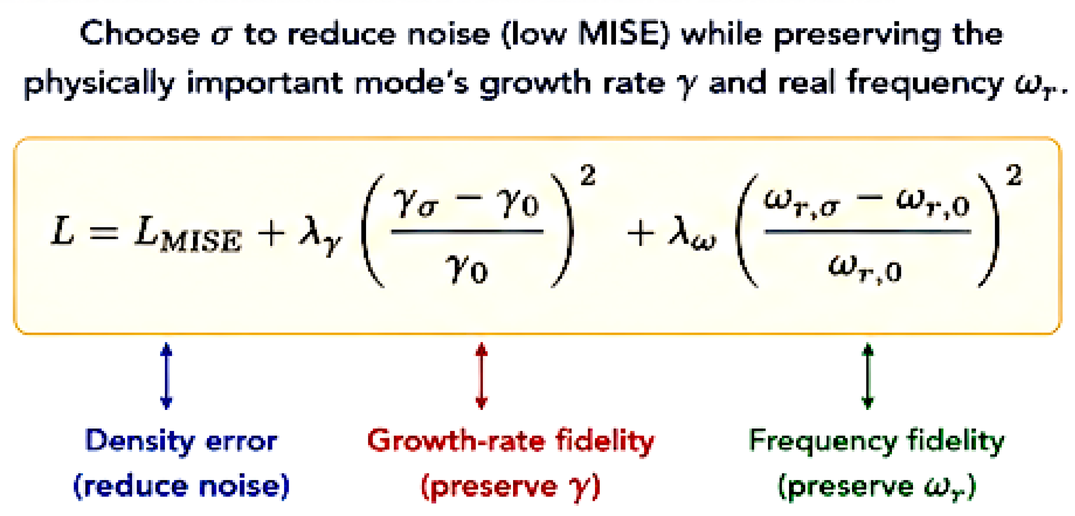

$$L = L_{MISE} + \lambda_\gamma \left(\frac{\gamma_\sigma - \gamma_0}{\gamma_0}\right)^2 + \lambda_\omega \left(\frac{\omega_{r,\sigma} - \omega_{r,0}}{\omega_{r,0}}\right)^2$$

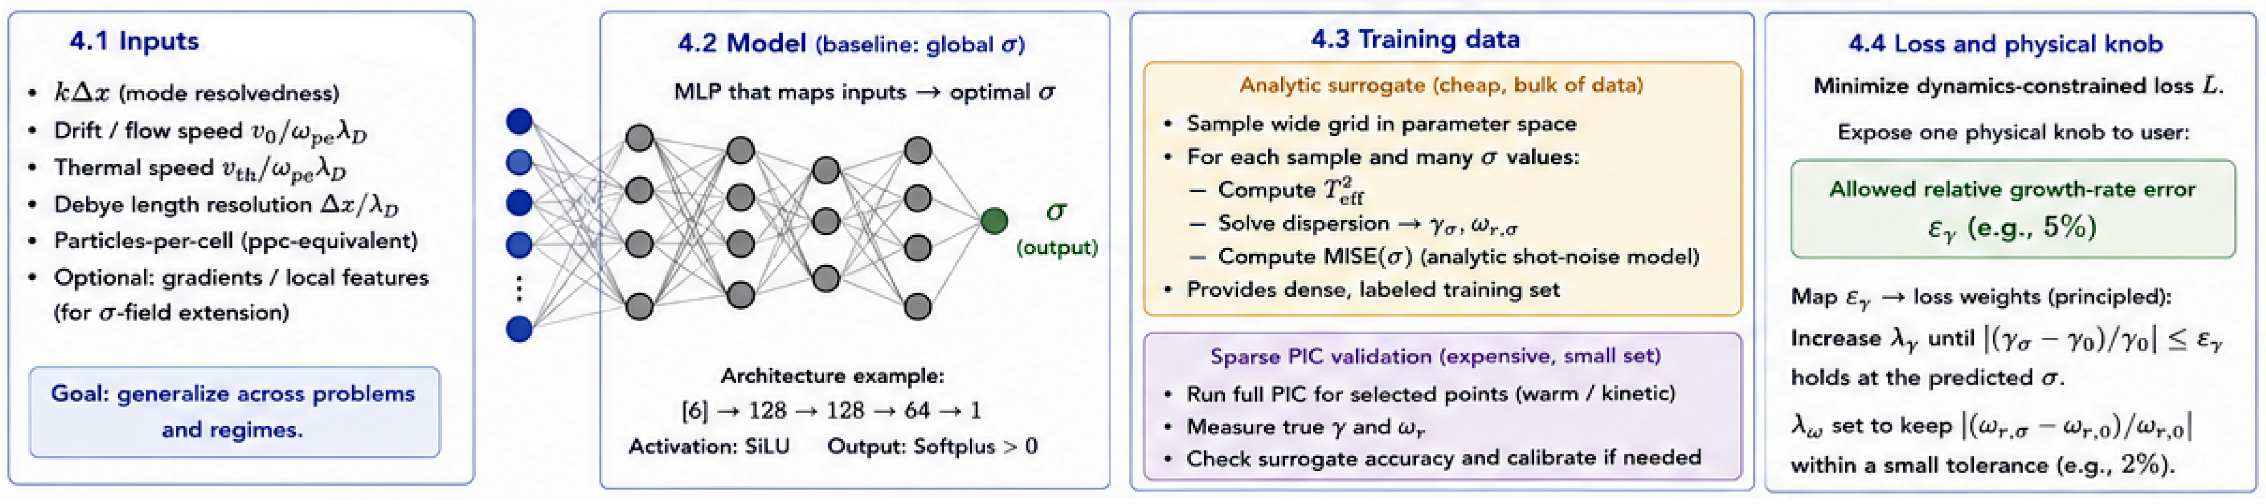

1. Simulation parameter sweep — throughput (task) parallelism

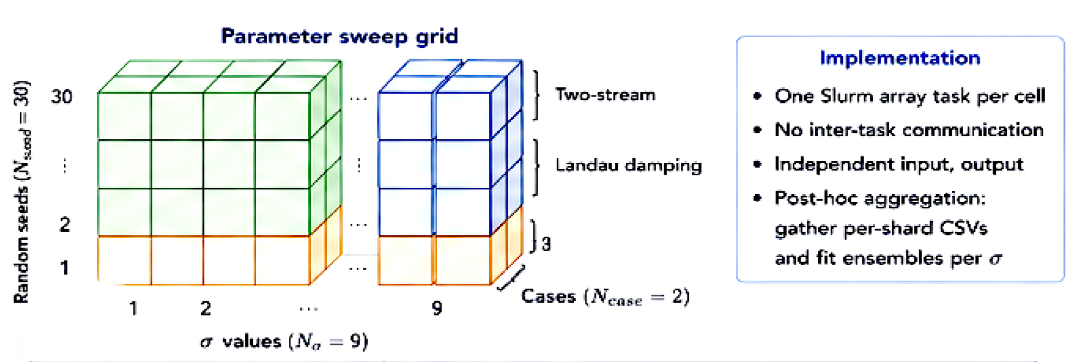

Per-simulation wall time	$t_{sim} \approx 300$ s (5 min)
Campaign size (cells)	$N_\sigma \times N_{seed} \times N_{case} = 9 \times 30 \times 2 = 540$
Concurrent array tasks	$M = 256$ (array concurrency on MeluXina)
Projected campaign wall time	$(\text{cells} / M) \times t_{sim} = (540 / 256) \times 300 \text{ s} \approx 632 \text{ s} \approx 10.5$ min
Speedup vs. serial	$\approx \min(M, \text{cells}) = 256\times$ (linear until $M = 540$)

Measured variance reduction (Landau damping growth rate)

δf fraction α	Variance reduction (vs full-f)	Equivalent particle count increase ($\approx V$)	α^2-scaling ratio (measured)
$\alpha = 1.0$ (full-f)	1×	1×	—
$\alpha = 0.1$	119×	~119×	11.9 (theory = 10)
$\alpha = 0.01$	1,916 ×	~1,916 ×	191.6 (theory = 100)
$\alpha = 0.001$	11,957 ×	~11,957 ×	1,195.7 (theory = 1000)

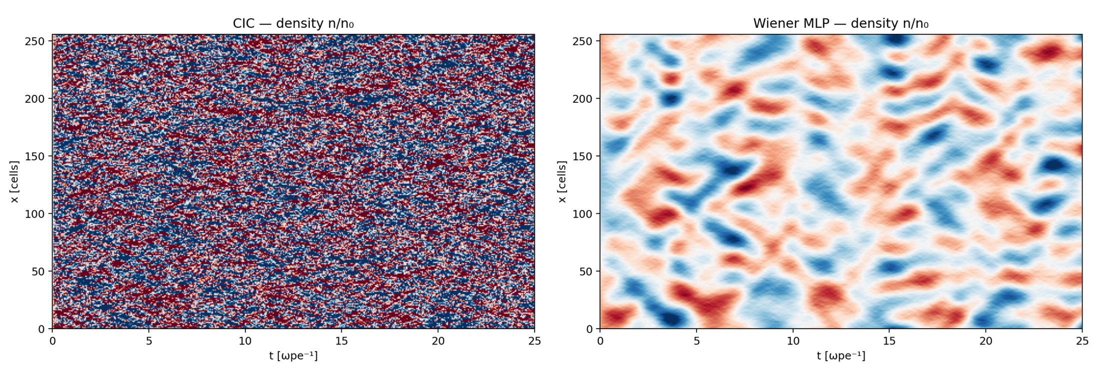

Results

- **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}.

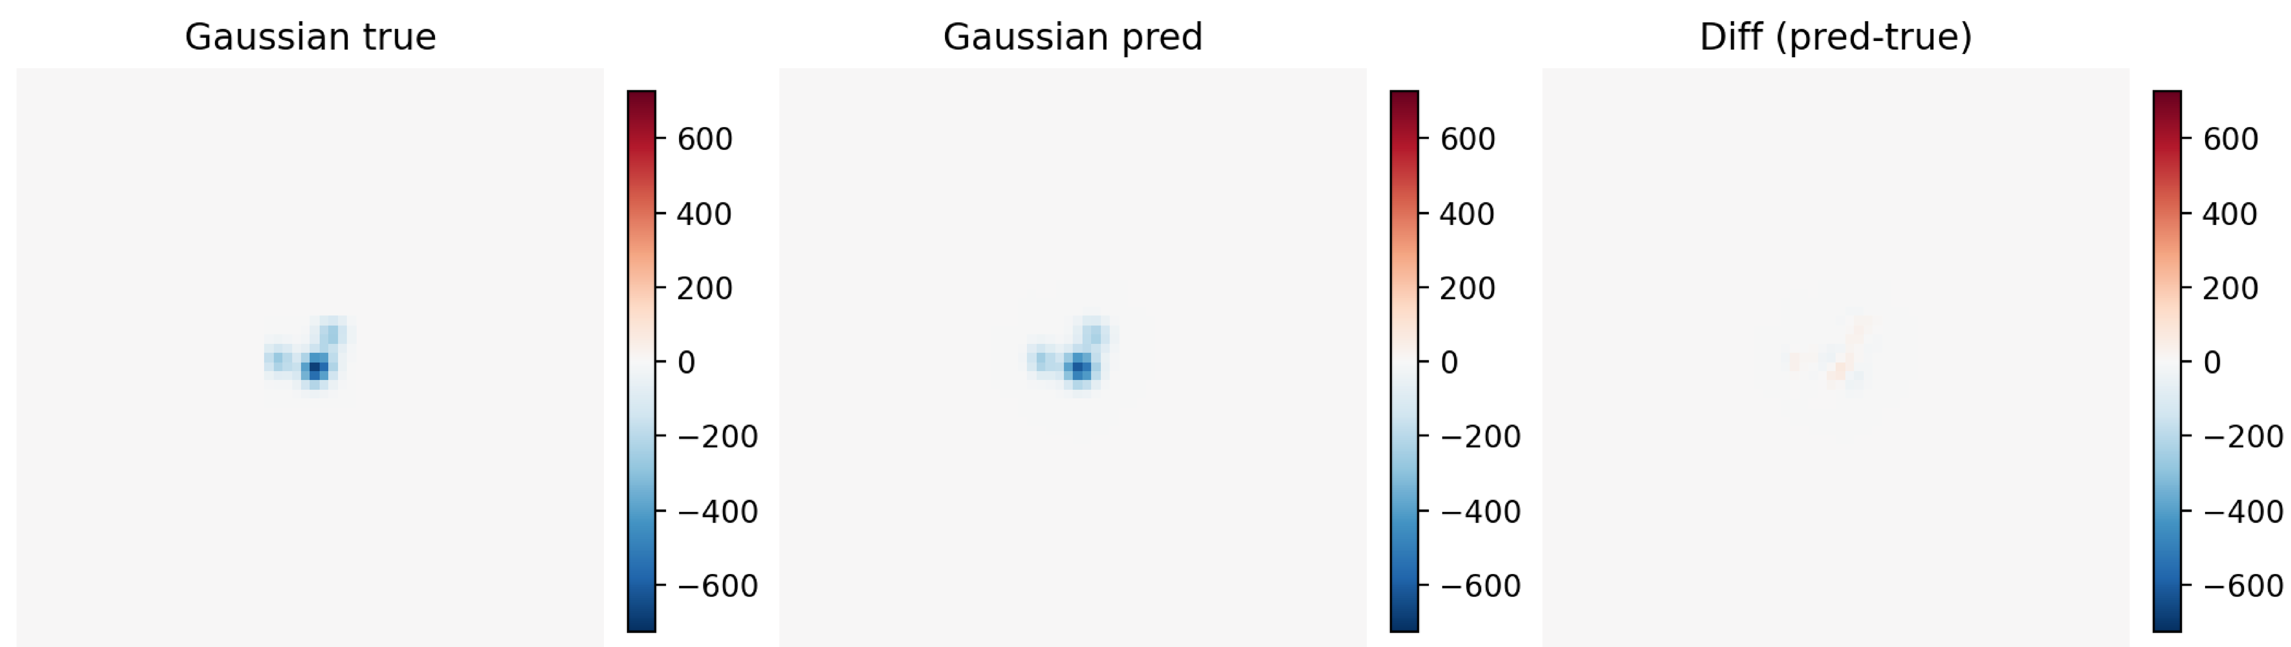

References

[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).

Slovenian Research and Innovation Agency
Luxembourg National Research Fund
UNIVERZA V LJUBLJANI

This work was supported by the Fonds National de la Recherche (FNR), Luxembourg, under the **FNR CORE** project grant number **C23/IS/18105668/HEXAPIC**, and the Slovenian Research and Innovation Agency **(ARIS)**, Slovenia.