

Gabriel Pedde Ungureanu | Physics

PhD | AI/ML Engineer

Brescia, Italy

Email: gpemdeun@sisssa.it • Web: gabriel-pedde.github.io/Gabriel-Pedde
GitHub: [Gabriel-Pedde](https://github.com/Gabriel-Pedde) • LinkedIn: [gabriel-pedde-ungureanu-4388371ab](https://www.linkedin.com/in/gabriel-pedde-ungureanu-4388371ab)

Profile

PhD particle physicist (6 papers, ~60 citations) with hands-on AI/ML engineering at EuroHPC scale: RLHF fine-tuning of Llama-3.2 (+11 pp on GSM8K), LLM multi-agent systems for scientific discovery, and distributed GPU training on Leonardo (CINECA). I build systems that pair LLM creativity with deterministic, verifiable ground truth.

Research & Engineering Experience

SISSA

Doctoral Researcher in Particle Physics and AI for Science

Trieste, Italy

2022–2026

- Non-perturbative QFT via symbolic computation on Ulysses (SISSA EuroHPC); 6 published papers, ~60 citations.
- Building an LLM “**AI co-scientist**” (LangGraph, arXiv-MCP + RAG): multi-agent generation for General Relativity discovery, refereed by a **deterministic symbolic/numerical evaluator**—the LLM never grades itself.
- Recovers known vacuum metrics (Schwarzschild, Kerr, Taub–NUT) symbolically; Israel–Khan two-body solution to **machine precision** via self-improving Co-Scientist evolution.

Education

SISSA / ICTP Joint Programme

Master in High Performance Computing and AI (MHPC)

Trieste, Italy

2025–2026

Concurrent with the last year of PhD. Machine learning, deep learning, reinforcement learning, LLM agents, and distributed GPU computing.

Università Cattolica del Sacro Cuore

BSc Mathematics & MSc Physics (Laurea Triennale & Laurea Magistrale)

Brescia, Italy

2017–2022

Theoretical physics, QFT, condensed matter; strong analytical foundations. I did a master degree thesis in **quantum complexity theory** at the interface of holography and quantum information

Awards & Publications

6 peer-reviewed papers in theoretical physics (~60 citations) — [INSPIRE-HEP](#) | [Google Scholar](#)

Best graduate of the year, MSc Physics — Università Cattolica del Sacro Cuore

110/110 *cum laude* — MSc Physics | 110/110 *cum laude* — BSc Mathematics

Istituto Toniolo merit fellowship — BSc, Università Cattolica del Sacro Cuore

Selected Projects — AI & ML

LLM RL fine-tuning: Implemented RLOO/PPO from scratch to fine-tune **Llama-3.2-1B** on GSM8K with a rule-based reward signal and KL regularisation; **+11 pp answer accuracy** (35 % → 46 %), format compliance from 51 % to 100 %. Distributed on **Leonardo** (CINECA). Stack: PyTorch, TRL, SLURM | [GitHub](#)

Multi-Agent LLM System: Built a multi-agent debate framework (CrewAI, Ollama) across **5 Qwen2.5 sizes** (0.5B–14B), $N = 4$ agents, 10 rounds. Models $\geq 3B$ converge with $< 10^{-3}$ **variance at round 0** — a sharp capacity-threshold phase transition independent of training. Stack: CrewAI, Ollama, Python | [GitHub](#)

Distributed GPU: CUDA kernel development through to cuBLAS; distributed dense linear algebra with **MPI + CUDA**. Lattice Boltzmann on a 2048×2048 grid: **~120× GPU speedup** over CPU (0.75 s vs 91.5 s, 806 GB/s bandwidth). MPI + OpenMP atmospheric simulation: near-linear scaling $2 \rightarrow 32$ nodes (3087 s → 317 s) | [GitHub](#)

Technical Skills

AI / ML: PyTorch, Hugging Face Transformers & TRL, scikit-learn, cuML

LLM / Agents: RLHF, REINFORCE/RLOO, PPO, GRPO, fine-tuning, RAG, MCP, LangGraph, LangChain, CrewAI, Ollama, Claude Code

ML methods: CNNs, autoencoders, representation learning, dimensionality reduction, Bayesian inference, clustering

Programming: Python, C++, C, Fortran, Mathematica, $\LaTeX$

Scale / HPC: CUDA, cuBLAS, OpenACC, MPI, OpenMP, SLURM, PETSc/SLEPc

MLOps: Docker, Singularity/Apptainer, Git, HDF5/NetCDF

Clusters: Leonardo (CINECA), Ulysses (SISSA), Argo (ICTP)

Languages: Italian (native) | English (C1) | Romanian (A2)

I authorize the processing of my personal data pursuant to Italian Legislative Decree 196/2003 and art. 13 GDPR 679/16.