

Bhaskar Agarwal, Ph.D.

AI & HPC for Scientific Discovery | Ph.D. Astrophysicist | Technical Advisor & Educator

Bologna, Italy | EU Resident | a.bhaskar.is@gmail.com | +39 348 592 4404 | [LinkedIn](#) | [Medium](#)

PROFILE

Ph.D. astrophysicist and data science leader with experience in building and deploying HPC enabled AI & ML solutions across R&D projects and industry. My work spans forecasting, optimisation, geospatial analytics and decision systems, combining hands-on modelling with rapid-prototyping and technical leadership. I currently lead multidisciplinary teams, taking concepts from problem definition and experimentation through deployment and adoption.

PROFESSIONAL EXPERIENCE

Group Lead & Senior Data Scientist, [OmegaLambdaTec](#), Germany (Remote). 2025–Present

Leading R&D team of data scientists and software engineers, setting research directions, creating prototypes and turning them into production-ready systems. Hands-on algorithm design, system architecture and technical validation.

- **[LocationAI](#), Applied Geospatial AI Research:** Government-funded research collaboration with WIGeoGIS on AI-driven location intelligence. Designing an adaptive, multi-resolution H3 spatial discretisation for urban, suburban and rural classification, and leading development of an agentic passenger-footfall simulation combining LLM reasoning with deterministic models.
- **Scientific Computing & R&D:** Transferred physics-inspired numerical methods, including adaptive time-stepping from astrophysical modelling, to industrial optimisation, achieving up to 100-1000× faster convergence and reusable methods for energy-management applications.
- **Technical Product Leadership:** AI products for electricity-grid planning ([GrydNova](#)), agentic engineering simulation (SimLumina) and Earth-observation intelligence (TerrAI). Defined architectures, roadmaps and validation criteria with multidisciplinary teams and external partners.

Senior Data Scientist & Technical Lead, [General Logistics Solutions \(GLS\)](#), Italy. 2023–2025

Translated loosely specified operational questions into measurable analyses, modelling experiments and product decisions; defined KPIs and success metrics.

- **Technical lead for the EU AI Hubs**, defining best practices and prototyping solutions for country-level adoption.
- Built and deployed a geospatial optimisation and decision platform across seven European countries, **reducing operational costs by 20% (€1–5M/yr per country)**. Led a team of six while owning modelling choices, validation, performance and stakeholder communication.
- Developed demand-forecasting models achieving **<5% MAPE (€0.5–1M/yr)** while reducing computational runtime from months to minutes.

Data Scientist & HPC Specialist, [CINECA](#) Supercomputing Centre, Italy. 2020–2023

Technical advisor to research and industry teams on Italy's Tier-0/1 supercomputers (Marconi100, Leonardo): scoping AI and data problems, porting and optimising workloads for NVIDIA GPUs, and training users.

- **GPU-accelerated AI on HPC:** Ported and optimised deep-learning and data-science workloads for NVIDIA V100 nodes (Marconi100) and DGX systems using TensorFlow, Horovod (distributed data-parallel training), RAPIDS and Dask. Took part in pre-production testing on A100 GPUs ahead of Leonardo's launch.
- **Multi-GPU debugging:** With colleagues, traced a bug in a Dask-based all-reduce that caused failures when scaling from single to multiple GPUs on an NVIDIA DGX, restoring reliable multi-GPU data processing.
- **AI for weather and climate ([EU H2020 CYBELE](#)):** Built the ML pipeline for a day-ahead hail-prediction system for Valencia agriculture, with GMV and the CACV farmers' federation, delivered as a Docker container for the project's demonstrator. Converted 6 TB of downscaled numerical-weather simulations (2.2 km grid, 68k grid points × 1,000+ atmospheric features, 5 years) into ML-ready data, cutting preprocessing from 5 days to 4 hours with parallel I/O, and achieved a **4–5× training speedup** with Horovod data parallelism on V100 GPUs.
- **Federated learning ([EU ORCHESTRA](#)):** Co-developed federated-learning infrastructure for GDPR-compliant analysis of SARS-CoV-2 cohort data across European partners; [published in Digital Health \(2024\)](#).

- **Training and advocacy:** Taught the Master's course in Data Science at [Bologna Business School](#) and delivered training and consulting on ML, distributed computing and scientific workloads. Wrote a two-part "AI with HPC" tutorial series for Towards Data Science (2022) on parallel data processing and multi-GPU training with Horovod.

Data Scientist, [OmegaLambdaTec](#), Germany. 2018–2019

Built predictive and optimisation solutions for demand and operational decision support.

- Developed a linear-programming energy-management system from first principles, in collaboration with a German University as a part of a joint research project.
- Built a day-ahead gas-price forecasting solution that achieved a podium result in the Ruhr Challenge 2018.

EDUCATION & RESEARCH

Ph.D. Astrophysics, LMU & Max Planck Institute. 2013

M.Sc. Astronomy, University of Sussex. 2008

B.Sc. Physics, University of Pune. 2007

Theoretical Astrophysicist. LMU, Max Planck Institute, Yale and Heidelberg. 2010–2018

Numerical and semi-analytical modelling, computational experiments and statistical comparison with observations; [20+ peer-reviewed publications](#) (h-index: 16, citations: 1,500+). Delivered talks at over 20 international conferences.

SELECTED RESEARCH & PROFESSIONAL ACTIVITIES

- **Proposal Reviewer**, [Fortissimo Plus](#), 2024–25: Independent evaluator for European HPC and AI innovation calls.
- **Research publications and technical writing (post-academia):**
 - "An innovative technological infrastructure for managing SARS-CoV-2 data across different cohorts in compliance with General Data Protection Regulation", Dellacasa et al., [Digital Health](#), 15 May 2024.
 - "Towards an automatic approach to modelling the circumgalactic medium: new tools for mock making and fitting of metal profiles in large surveys", Longobardi et al., [MNRAS](#), 1 Jan 2023.
 - "AI with HPC" series, Towards Data Science (2022): [parallel data processing on supercomputers](#); [multi-GPU training with Horovod](#).
- **Technical Leadership & Innovation:** Provided technical leadership and architectural guidance to early-stage technology ventures, supporting research-driven AI/ML prototyping.

SKILLS

High Performance & Accelerated Computing

- Porting, benchmarking and optimising AI workloads on HPC; Python, SLURM, Docker, NVIDIA NGC.
- Parallel computing & data pipelines for large-scale scientific and AI workloads: multiprocessing, parallel I/O, PyArrow/Parquet, Dask, RAPIDS (cuDF, cuML), Horovod.

Artificial Intelligence & Machine Learning

- PyTorch, TensorFlow, Hugging Face Transformers, scikit-learn, XGBoost; NumPy, pandas, SciPy.
- Distributed training, fine-tuning (LoRA/PEFT), hyperparameter optimisation, model evaluation and benchmarking.
- Forecasting, anomaly detection, classification, numerical optimisation and first-principles modelling.

Generative AI & Agentic Systems

- RAG, embeddings, semantic search, vector databases (Qdrant, pgvector), knowledge graphs (Neo4j).
- LangChain/LangGraph, Pydantic, tool calling, structured outputs and agentic scientific simulations.

Scientific & Geospatial Computing

- GeoPandas, H3, OpenStreetMap; spatial analysis, simulation and mathematical optimisation.

Soft Skills: Curious, problem solving, scientific writing, public speaking, technical training, stakeholder communication.

LANGUAGES: English (Native), Hindi (Native), Italian (B2/C1). German (basic)