

NILOTPAL KAKATI

Machine Learning Researcher · Weizmann Institute of Science, Israel

[Personal website](#) · nilotpalkakati09@gmail.com · [LinkedIn](#) · [ORCID](#)

PROFILE

Machine-learning researcher with a PhD from the ATLAS experiment at CERN, where several methods I led are in production — including the primary jet flavour tagger for Run 3 (Nature Communications, 2026) and a CUDA graph-construction pipeline delivering orders-of-magnitude speedup over CPU baselines. Co-developed top-scored research proposals with ETH Zürich (MSCA 98/100 on scientific excellence) on flow-matching latent-dynamics world models for soft-robotic control — extending my generative-physics work toward physical/robotic systems.

EXPERIENCE

Researcher (PhD then Postdoc), Weizmann Institute of Science

Feb 2020 – present

Department of Particle Physics and Astrophysics · Member of the ATLAS Collaboration, CERN

- **GN2 jet flavor tagger:** Led implementation and software development of GN2, the state-of-the-art ATLAS jet flavor tagger now in production for Run 3 analyses (Nature Communications, 2026). Co-designed the architecture's auxiliary-task structure: a shared transformer backbone with multiple prediction heads supervising physically-motivated intermediate quantities alongside the primary classification target. Ablation studies show the auxiliary supervision is responsible for substantial performance gains over single-task baselines. Analysis contact for the GN2 publication.
- **HGPflow particle reconstruction:** Reformulated end-to-end particle reconstruction from a black-box set-to-set task into a physics-motivated energy-assignment problem, implemented as hypergraph learning. Energy conservation between detector clusters and reconstructed particles is enforced architecturally in the incidence matrix rather than via loss penalty. Outperforms unconstrained baselines with significantly fewer parameters. Eur. Phys. J. C, 2025.
- **Multimodal foundation model:** Proposed and leading development of a multimodal foundation model for particle physics — a 4M-style masked-modeling framework integrating up to seven heterogeneous physics modalities (calorimeter cells, tracks, jets, final-state particles, etc.) in a single set-compatible encoder-decoder architecture. Modality-specific tokenizers map each input stream into a shared token space; cross-modal masked prediction yields a single backbone that supports both generative and discriminative downstream tasks.
- **Self-supervised jet representations:** Co-developed a contrastive pretraining framework where the augmentation is *re-simulation itself* — share upstream physics, vary downstream stochastic stages. Invariances are inherited from the data-generating process rather than synthesized by hand-designed transforms. Demonstrated downstream gains across multiple tasks (MLST, 2025).
- **Flow matching super-resolution:** Proposed and leading a flow-matching model with a transformer backbone for post-readout calorimeter super-resolution: improves effective detector granularity from learned priors over realistic shower physics, avoiding costly hardware upgrades. Demonstrates stable long-horizon sampling and physically consistent outputs (energy conservation, plausible shower morphology). NeurIPS ML4PS 2025; full paper submitted to Phys. Rev. D.
- **Generative simulation surrogates:** Built physics-guided generative surrogates replacing Geant4-level detector simulation at distribution-level fidelity (Phys. Rev. D 110, 092013, 2024). Compared diffusion and flow-matching variants; flow-matching emerged as the winner on stability and conditioning.
- **CUDA graph-construction pipeline:** Designed and shipped a fully GPU-resident graph-construction pipeline for the GN2 inference path, removing the CPU bottleneck. Orders-of-magnitude wall-clock speedup; supports current ATLAS Run 3 production.

SKILLS

Machine Learning: Generative models (diffusion, flow matching, GANs, normalizing flows), transformers, GNNs/hypergraph networks, masked multimodal modeling (4M-style), self-supervised learning (SimCLR, CLIP-style alignment), set-conditional generation.

Engineering: Python, PyTorch, CUDA, C++, SLURM/HPC workflows, Docker, distributed training, ONNX deployment, TB-scale dataset pipelines (HDF5/Parquet/ROOT). Reproducible experiment infrastructure.

Domain: World models for physical and robotic systems, physics-informed ML, learned simulators, surrogate modeling for stochastic dynamics.

SELECTED PUBLICATIONS

Full list (20+ publications) available on request. Kakati is an author on every entry below; lead and corresponding-author roles are noted where applicable. Particle-physics author lists are alphabetical by convention.

- ATLAS Collaboration (analysis contact: Kakati). “Transforming jet flavor tagging at ATLAS”. **Nature Communications, 2026.**
- N. Soybelman et al., “Topology-agnostic, hierarchical graph network for event reconstruction” — **Phys. Rev. D, 2026**
- Kakati et al. (corresponding author). “HGPflow: Extending Hypergraph Particle Flow to Collider Event Reconstruction”. **Eur. Phys. J. C, 2025.**
- Rieck et al. “Self-Supervised Learning Strategies for Jet Physics.” **Machine Learning: Science and Technology, 2025.**
- Arguin et al., “Automatizing the search for mass resonances using BumpNet” — **JHEP, 2025**
- Kakati et al. (corresponding author). “Denoising Graph Super-Resolution towards Improved Collider Event Reconstruction”. Submitted to Phys. Rev. D; arXiv:2409.16052.
- Kobylanskii et al. “Advancing set-conditional set generation: Diffusion models for fast simulation of reconstructed particles”. **Phys. Rev. D 110, 092013, 2024.**
- di Bello et al. (Kakati: corresponding author). “Reconstructing particles in jets using set transformer and hypergraph prediction networks.” **Eur. Phys. J. C, 2023.**

EDUCATION

PhD, Weizmann Institute of Science 2020 – 2025

Department of Particle Physics and Astrophysics. Defended September 2025.

BS-MS in Physics, Indian Institute of Science Education and Research (IISER) Mohali 2014 – 2019

AWARDS

- *John F. Kennedy Prize (2026)* — Weizmann Institute of Science. For outstanding Ph.D. research.
- *ATLAS Outstanding Achievement Award (2024)* — For contributions to jet flavor tagging with GNNs.
- *FGS Dean Award of Excellence (2023–2024)* — In Physics, for international doctoral students.
- *DST-INSPIRE Fellowship (2014–2019)* — Government of India.

TEACHING

- Practical Deep Learning for Science (lectures + tutorials), Weizmann Institute of Science — 2021, 2023, 2024, 2025.
- AI Hub at Weizmann (tutor) — Sep 2023 – Feb 2024.