Giordon Stark — Cover Letter

SCIPP, NS2, Room #337 – 1156 High Street – Santa Cruz, CA 95064

gstark@cern.ch • giordonstark.com • Built May 10, 2026 from @486c31

Nikhef
Science Park 105
1098 XG Amsterdam
The Netherlands

May 10th, 2026

Dear Professor Verkerke, Dr. Caron, and the Search Committee,

I am writing to apply for the Staff Scientist position in AI and Machine Learning in Physics at Nikhef. I obtained my Ph.D. in Physics from the University of Chicago in 2018 and am currently a Project Scientist at the University of California, Santa Cruz, working on the ATLAS Experiment. My research program combines active physics analysis (dark matter searches through compressed electroweak supersymmetry and intensity-frontier BSM searches), statistical inference tooling built on automatic differentiation, and a new direction in agentic AI for physics workflows. I am excited by the opportunity to help define the research agenda of Nikhef's new AI/ML group and to deepen the impact of AI/ML across Nikhef's programmes, from LHC physics to astroparticle experiments.

My physics analysis program at ATLAS targets dark matter candidates in compressed electroweak SUSY spectra, where the signal signatures (soft leptons, low missing transverse momentum, VBF topologies) are subtle enough that machine learning methods for background estimation and signal discrimination are essential. Using the full Run 2 dataset, I led analyses that pushed sensitivity down to mass splittings $\Delta m \sim 1$ GeV using BDT-based signal extraction in a Vector Boson Fusion topology, producing the first LHC limits to surpass LEP constraints on compressed sleptons. As SUSY Combinations Team Contact, I then led the statistical combination of 14 electroweak SUSY analyses, extending mass exclusion limits by up to 100 GeV and published in PRL. The pMSSM interpretation of these results probed a 19-dimensional SUSY parameter space and identified uncovered regions, including a gap in the smuon-chargino plane for models compatible with the muon $g-2$ anomaly. This combination required building the infrastructure to serialize, validate, and jointly fit probability models across a dozen independent analysis teams. I also initiated a search for radiative neutralino decays ($\tilde{\chi}_2^0 \to \tilde{\chi}_1^0\gamma$), a photon-plus-MET signature that extends sensitivity in the compressed regime. At Belle II, I propose to search for Z' bosons in the $L_\mu - L_\tau$ gauge model, where the Z' couples to muons and taus and mediates interactions with a dark matter candidate in the GeV mass range. Belle II published its first search in this channel using `pyhf` for the statistical analysis, showing the software bridge I built between the two experiments.

I am a core developer of `pyhf`, the Python statistical modeling framework built on automatic differentiation backends (PyTorch, JAX, TensorFlow) that has become the standard for likelihood publication in ATLAS and the official statistical tool for Belle II. The choice of autodiff backends was deliberate: it makes `pyhf` natively compatible with gradient-based optimization, differentiable analysis pipelines, and simulation-based inference (SBI) workflows that are central to the next generation of AI/ML in HEP. Over 50 ATLAS analyses now publish machine-readable likelihoods through `pyhf`, and the framework is used by LHCb, EIC, and future collider efforts. I am now developing `pyhs3`, the Python engine for the HEP Statistics Serialization Standard, in partnership with the BMBF-funded DEMOS project. This next-generation tool extends beyond HistFactory to support arbitrary statistical models (unbinned fits, SBI, differentiable pipelines) and connects to the JAX-based tooling ecosystem (evermore, optimistix) shaped by the IRIS-HEP Blueprint workshops on statistical methods, simulation-based inference, and differentiable analysis. The vision is end-to-end differentiable physics analysis, where gradients flow from event selection through histogram construction to inference, enabling gradient-based optimization of analysis strategies. This is AI/ML infrastructure for physics at the foundational level.

A new direction in my program applies AI agents with domain-specific HEP knowledge to accelerate physics analysis. Recent proof-of-concept studies have shown that LLM-based agents can autonomously execute substantial portions of a collider analysis pipeline. I am developing agentic frameworks that encode expert knowledge of analysis procedures, statistical inference via `pyhf`, and systematic uncertainty evaluation. The pMSSM is one concrete target: a 19-dimensional parameter scan involves orchestrating signal generation, detector simulation, likelihood evaluation, and exclusion tracking across dozens of analyses. An agent equipped with access to the reinterpretation pipeline and `pyhf` could explore this landscape faster than a human, identifying uncovered regions and flagging promising search strategies. I have also developed MCP tool servers that allow AI agents to programmatically search experimental databases, check data replicas, retrieve cross-sections, and validate sample metadata. The cross-experiment applicability of this approach is what makes it interesting for Nikhef: agentic tools developed for ATLAS can serve LHCb, and the statistical tooling already supports Belle II. Extending to KM3NeT or XENONnT is a question of encoding domain knowledge, not rebuilding infrastructure.

I have supervised undergraduate and graduate students across hardware, software, and physics projects at both UChicago and UC Santa Cruz. At UCSC, I built the full ITk pixel module test bench infrastructure, created the Python library for the ATLAS ITk Production Database, and developed QC/QA software used across $\sim$20 international testing clusters for $\sim$12,000 pixel modules. My mentoring includes students working on this ITk testing programme, `pyhs3` development, compressed SUSY analyses, and analysis facility infrastructure. I currently serve on the Executive Committee of the DPF Coordinating Panel for Software and Computing, and I have organized and contributed to over 40 workshops and bootcamps through the HEP Software Foundation, Software Carpentry, and IRIS-HEP. I am experienced in funding procurement: I have active NSF awards, a pending DOE proposal for the HEP Software Packaging Initiative, and a proposal in development for agentic AI systems for HEP. I am familiar with European funding mechanisms through my partnership with the BMBF-funded DEMOS project. I also have experience with ML-based real-time trigger systems through my design of the ATLAS global Feature EXtractor (gFEX), the first hardware trigger to process the entire ATLAS calorimeter on a single board, and through my engagement with the ML-HEQUPP community and the CERN Next Generation Triggers project.

I am enthusiastic about joining the Nikhef community beyond the scope of the position itself. I value the collaborative, welcoming environment described in the posting. As a Deaf physicist, I have spent my career working to make scientific tools and training materials accessible to people of all backgrounds, from developing ASL content for physics education to contributing to the Snowmass whitepaper on accessibility in high-energy physics. I am excited by the prospect of helping build Nikhef's AI/ML group from the ground up, defining a research agenda that has real impact on the physics, and contributing to the European AI-in-science initiatives that Nikhef is joining. I am happy to provide any additional material upon request.

Sincerely,

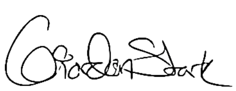

Giordon Stark (pronouns: he/him/point)

Enclosed: curriculum vitæ, List of Publications, contact details for three referees