



Giordon Stark

Pronouns: he/him/point

Research Statement

Introduction

The Standard Model (SM) of particle physics has been tested for decades, but it is incomplete. The matter/antimatter asymmetry [1], the fine tuning needed to stabilize the Higgs mass at the electroweak scale [2], the lack of dark matter candidates [3, 4], and the hierarchy problem [5] all require physics beyond the Standard Model. My research asks: *where is new physics hiding?* Two possibilities are that new particles have compressed mass spectra, producing soft signatures at the energy frontier, or that they are light, weakly coupled states at the intensity frontier.

These scenarios motivate a dual-experiment research program at the CMS detector at the Large Hadron Collider (LHC) and at the Belle II detector at SuperKEKB. Over eight years at ATLAS, I developed the techniques for compressed electroweak SUSY searches that produced the first LHC limits to surpass LEP constraints on compressed sleptons, and I led the first statistical combination of ATLAS electroweak SUSY searches. I plan to bring this program to CMS for Run 3 and the High-Luminosity LHC (HL-LHC). At Belle II, I propose to search for a Z' gauge boson coupling to heavy-flavor leptons via the $L_\mu - L_\tau$ symmetry, a minimal dark sector model that addresses rare B -decay anomalies and provides a viable dark matter candidate [6]. These programs are supported by a software and statistical infrastructure I helped build: pyhf [7], the HEP Statistics Serialization Standard (HS3), the BMBF-funded DEMOS project [8], and the HEP Packaging Community [9], which together make dual-experiment physics reproducible.

At KIT, I would bring compressed electroweak SUSY and pMSSM global fits to CMS while strengthening the Belle II physics and hardware effort. I am already an institutional partner on DEMOS, a direct connection to the German research landscape. The GridKA computing center, the CMS and Belle II heritage at ETP, and the ML-HEQUPP and Next Generation Triggers communities all fit this program.

Dark Matter Searches at the LHC

Compressed Electroweak Supersymmetry

In R -parity conserving (RPC) SUSY, a stable lightest neutralino $\tilde{\chi}_1^0$ is a dark matter candidate. Scenarios where the lightest electroweakinos are nearly degenerate in mass—compressed spectra—are motivated by naturalness arguments and are weakly constrained by direct detection experiments. At ATLAS, using the full Run 2 dataset (139 fb^{-1}), I led analyses that pushed sensitivity down to mass splittings $\Delta m \sim 1 \text{ GeV}$ using a Vector Boson Fusion topology with virtual electroweak bosons [10]. These were the first LHC limits that surpassed LEP constraints on compressed sleptons, reaching mass splittings as small as 550 MeV.

The techniques I developed at ATLAS (VBF tagging for invisible final states, BDT-based signal extraction in the highly compressed regime, systematic treatments for ISR modeling) are applicable to CMS. The CMS silicon tracker provides superior low- p_T lepton reconstruction, and the particle-flow algorithm improves MET resolution in VBF environments, both of which matter for compressed searches. I plan to lead compressed electroweak SUSY searches in CMS using Run 3 data ($\sim 300 \text{ fb}^{-1}$) and the full HL-LHC dataset (3000 fb^{-1}), targeting the coverage gap I identified in the smuon-chargino plane for models compatible with the muon $g - 2$ anomaly [11]. I also initiated a search for radiative neutralino decays ($\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \gamma$), which become dominant when tree-level channels are kinematically suppressed in compressed scenarios. This photon-plus-MET signature extends LHC sensitivity in the region that connects to the $g - 2$ anomaly and would make for good thesis projects at KIT.

SUSY Combinations and pMSSM Global Fits

As SUSY Combinations Team Contact and Run 2 Summaries Subconvener at ATLAS, I led the statistical combination of 14 electroweak SUSY analyses across four decay channels (WW , WZ , Wh , GGM) [12]. This five-year effort, built on harmonized object definitions, serialized likelihoods,

and preserved analysis workflows, extended mass reach by 30–100 GeV beyond individual searches. The pMSSM interpretation probed a 19-dimensional SUSY parameter space and identified uncovered regions. The program produced four publications: the EWK combination, EWK pMSSM, RPV-RPC reinterpretation, and a general pMSSM paper.

CMS has its own pMSSM program, but the combination and reinterpretation infrastructure is less mature than what I built at ATLAS. I plan to bring the tooling and workflows I developed to help advance this effort. The pMSSM is also a good proving ground for agentic analysis: a 19-dimensional parameter scan involves orchestrating signal generation, detector simulation, likelihood evaluation, and exclusion tracking across dozens of analyses. I built the reinterpretation pipeline [13] and the statistical tools such a program needs, and the bookkeeping complexity of the pMSSM is exactly the kind of problem where AI agents with domain-specific knowledge can make a real difference.

Physics Searches and Hardware at Belle II

Dark Matter via Z' Coupling to Heavy-Flavor Leptons

Belle II provides sensitivity to light dark sector particles at the intensity frontier. I propose to search for Z' bosons in the $L_\mu - L_\tau$ gauge model, where the Z' couples to muons and taus and can mediate interactions with a dark matter candidate in the $\mathcal{O}(\text{GeV})$ mass range [6]. Belle II published the first search for an invisibly decaying Z' in $\mu^+\mu^- + \text{missing energy}$ [14], using pyhf for its statistical analysis, which shows the software bridge I built between the two experiments. With 575 fb^{-1} collected and the dataset growing towards 50 ab^{-1} , the sensitivity will improve by 100–1000 \times from the first publication. I will bring my dark matter search experience from the LHC and my statistical tooling to build a reinterpretation pipeline at Belle II like what exists at ATLAS, in collaboration with phenomenologists such as Wolfgang Altmannshofer and Stefania Gori at SCIPP. ETP's established Belle II group is the right place to do this.

Hardware: The Inner Tracking and Timing Detector

The Belle II collaboration is planning a major detector upgrade for Long Shutdown 2 (~ 2032), with the upgrade baseline to be finalized in July 2026 and a Technical Design Report by 2027. The B-factory Programme Advisory Committee (BPAC) [15] considers the introduction of an Inner Tracking and Timing (ITT) system between the new VTX and the central drift chamber ($17 < R < 34 \text{ cm}$) to be mandatory for sustainable operation in the radiation environment expected after LS2. The BPAC notes the ITT should help attract new resources to the collaboration.

I plan to contribute to the ITT through board-level electronics design: the concentrator or ASIC logic that aggregates data from the front-end sensors and feeds the trigger system within Belle II's 30 kHz maximum trigger rate. This applies my expertise from designing the ATLAS global Feature EXtractor (gFEX) [16, 17], the first hardware trigger to process the entire ATLAS calorimeter on a single board, and from my work on ITk front-end chips at SCIPP. The ITT's sensor technology will likely use AC-LGADs for precision timing, with the LGAD community and fast-timing readout chips maturing from HL-LHC developments. My contribution would focus on the data path from sensors to trigger: designing the board-level logic that performs hit concentration, timing alignment, and trigger primitive generation within the latency budget. This hardware-software codesign approach, informed by the ML-HEQUPP community whitepaper [18] on deploying machine learning inference on FPGAs and ASICs in real-time trigger systems, fits with the existing Belle II expertise at ETP.

Next-Generation Triggers and Instrumentation

The HL-LHC will increase the instantaneous luminosity to $7 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ with up to 200 simultaneous proton-proton collisions per bunch crossing. The CMS Level-1 Trigger must be upgraded from 100 kHz to 750 kHz [19] to handle this environment. The upgraded trigger will, for the first time, incorporate tracking information and high-granularity calorimeter data from HGCal, enabling particle-flow reconstruction at the trigger level with ML-based inference on FPGAs.

My gFEX experience (designing algorithms for pile-up estimation, jet multiplicity, and MET within a 25 ns decision window on a single board processing the entire ATLAS calorimeter) is directly relevant to the CMS L1T upgrade. Recent work in the ML-HEQUPP community [18] and the CERN Next Generation Triggers (NGT) project has shown that neural-network-based anomaly detection triggers like AXOL1TL and CICADA can identify rare events that conventional thresholds miss. CMS already deployed these during Run 3, recording significant datasets with AXOL1TL seeds. I plan to contribute to the next iteration of ML-based trigger algorithms for the HL-LHC, focusing on real-time inference

for compressed SUSY signatures (soft leptons, low MET, VBF topologies) where trigger efficiency limits physics reach.

Looking further ahead, the CERN Future Circular Collider (FCC) feasibility study was confirmed by CERN Council in 2025, with the European Strategy recommendations expected in 2026. A muon collider, the top recommendation of the 2025 National Academies report, would complement both the FCC program and the Belle II physics reach. The software, statistical infrastructure, and trigger expertise I develop is designed to be experiment-agnostic, so the tools can serve the field regardless of which machines ultimately operate.

Software and Statistical Infrastructure

From `pyhf` to the Next Generation of Statistical Tooling

I am a core codeveloper of `pyhf` [7], the pure-Python HistFactory implementation that is the standard for statistical analysis preservation in HEP. Over 50 ATLAS analyses now provide machine-readable likelihoods via `pyhf`'s JSON format; it is the primary statistical tool for Belle II and is used by LHCb, EIC, and future collider efforts. I am currently working on the next-generation tool that builds on lessons from `pyhf` and the HEP Statistics Serialization Standard (HS3). The gap is support for non-binned, non-HistFactory-based analyses: arbitrary statistical models that arise in unbinned fits, simulation-based inference (SBI), and differentiable analysis pipelines. I am developing `pyhs3`, the Python engine for HS3, as an institutional partner in the BMBF-funded DEMOS (Democratizing Models) project [8], whose scope extends beyond particle physics to nuclear physics, hadron physics, laser-plasma physics, and astrophysics. DEMOS builds computational engines in C++, Python, and Julia, with a model-sharing platform and registry that embeds open-science practices across the ErUM research community.

I am currently an affiliated but unfunded partner on DEMOS. At KIT, I would be in-country to strengthen this collaboration with dedicated personpower and to pursue funded participation through BMBF or DFG channels. The Python engine development, building on the JAX-based tooling ecosystem (evermore, `optimistix`, and the emerging `pyhs3`), is a good match for student projects and connects to the differentiable analysis pipeline vision shaped by the IRIS-HEP Blueprint workshops I participated in during February–March 2026.

Analysis Systems, Facilities, and Reproducibility

The HL-LHC will produce an order of magnitude more data than Run 2, requiring a shift in how CMS analyses are performed. I have led efforts within ATLAS to transition from C++-based analysis to modern Python columnar workflows using `Coffea` and `ServiceX`, evaluating ROOT `RNTuple` as the next-generation event data format, and integrating distributed workflow systems such as `TaskVine`. At KIT, I would connect this work to the European computing ecosystem: the GridKA Tier-1 center on campus provides the backbone for CMS computing in Germany, and DESY's analysis infrastructure is a further resource. I plan to help adopt columnar analysis paradigms within the German CMS community, preparing for HL-LHC luminosities.

My end-to-end reinterpretation pipeline [13], combining `mapyde`, `RECAST`, and `REANA`, has enabled systematic reuse of analyses for new physics models. I have also started the HEP Packaging Community [9], a cross-cutting open-source effort to modernize the HEP software stack through secure binary conda packages for deployment on emerging ARM architectures and future HPC environments while strengthening supply-chain security and reproducibility. This is the infrastructure layer that makes dual-experiment research practical: a single physicist should be able to install and run analysis software for CMS and Belle II through a unified, reproducible packaging system.

Agentic AI for HEP Analysis

A new direction in my research program applies AI agents with domain-specific HEP knowledge to accelerate physics analysis. Recent proof-of-concept studies [20, 21] have shown that LLM-based agents can autonomously execute substantial portions of a collider analysis pipeline, from event selection through statistical inference. The tools work; what is missing is the integration with real analysis infrastructure and the domain-specific tooling that would make the results trustworthy.

I am developing agentic frameworks that encode expert knowledge of collider analysis procedures, statistical inference via `pyhf`, fitting strategies, and systematic uncertainty evaluation. The CMS pMSSM program is one concrete target: a 19-dimensional parameter scan involves orchestrating signal generation with `SPheno` and `MicrOMEGAs`, cross-section calculations with `Prospino`, detector simulation, likelihood patching, and exclusion evaluation across dozens of analyses. This is structured,

repetitive, and error-prone work. An agent equipped with access to the reinterpretation pipeline and `pyhf` could explore the pMSSM landscape much faster than a human, identifying uncovered regions and flagging promising search strategies. The open question is what combination of foundational models, domain-specific skills, and tool-use protocols makes agents reliable enough for physics-grade results.

At KIT, this direction connects to the German funding landscape. An ERC Starting Grant or DFG Forschungsgruppe could support a focused effort on agentic analysis infrastructure, while the ErUM-Data Hub is a home for the cross-disciplinary aspects. The statistical tooling I built, the physics analyses I lead, and the computational infrastructure at KIT together make this a viable independent research program.

Summary

My research program is built on two experimental pillars, CMS and Belle II, connected by a software and statistical ecosystem that enables reproducible, cross-experiment physics. At CMS, I bring ATLAS expertise in compressed electroweak SUSY, statistical combinations, and pMSSM global fits to establish a dark matter search program for Run 3 and the HL-LHC. At Belle II, I propose to search for Z' -mediated dark matter candidates in multilepton final states and contribute to the ITT detector upgrade through board-level trigger electronics. The software infrastructure (`pyhf`, `pyhs3/DEMOS`, the HEP Packaging Community, and the differentiable analysis ecosystem) connects these programs, while agentic AI, with pMSSM reinterpretation as one target, provides a new approach to the computational challenges of global fits. ML-based triggers for CMS and Belle II draw on my gFEX heritage and connect to the ML-HEQUPP and NGT communities. At KIT, I would strengthen the CMS and Belle II programs at ETP, grow my research group through DEMOS and the German funding landscape, and contribute to the department's computing, software, and instrumentation work for the HL-LHC era.

References

- [1] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov. "Matter and Antimatter in the Universe". In: *New Journal of Physics* 14.9 (2012), p. 095012.
- [2] Howard Baer et al. "Naturalness, Supersymmetry and Light Higgsinos: A Snowmass Whitepaper". In: *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*. 2013. arXiv: 1306.2926 [hep-ph]. URL: <https://inspirehep.net/record/1238302/files/arXiv:1306.2926.pdf>.
- [3] Gianfranco Bertone, Dan Hooper, and Joseph Silk. "Particle dark matter: Evidence, candidates and constraints". In: *Physics Reports* 405.5-6 (2005), pp. 279–390.
- [4] P. A. R. Ade et al. "Planck 2015 results. XIII. Cosmological parameters". In: *Astron. Astrophys.* 594 (2016), A13. DOI: 10.1051/0004-6361/201525830. arXiv: 1502.01589 [astro-ph.CO].
- [5] Gerard 't Hooft et al. "Recent Developments in Gauge Theories. Proceedings, Nato Advanced Study Institute, Cargese, France, August 26 - September 8, 1979". In: *NATO Sci. Ser. B* 59 (1980), pp.1–438. DOI: 10.1007/978-1-4684-7571-5.
- [6] Wolfgang Altmannshofer et al. "Explaining dark matter and B decay anomalies with an $L_\mu - L_\tau$ model". In: *JHEP* 12 (2016), p. 106. DOI: 10.1007/JHEP12(2016)106. arXiv: 1609.04026 [hep-ph].
- [7] Lukas Heinrich et al. "pyhf: pure-Python implementation of HistFactory statistical models". In: *Journal of Open Source Software* 6.58 (2021), p. 2823. DOI: 10.21105/joss.02823. URL: <https://doi.org/10.21105/joss.02823>.
- [8] DEMOS Consortium. *DEMOS: Democratizing Models*. 2025. URL: <https://democratizing-models.github.io/>.
- [9] Giordon Stark and Matthew Feickert. *HEP Packaging Coordination*. 2025. URL: <https://github.com/hep-packaging-coordination/>.
- [10] Georges Aad et al. "Searches for electroweak production of supersymmetric particles with compressed mass spectra in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector". In: *Phys. Rev. D* 101.5 (2020), p. 052005. DOI: 10.1103/PhysRevD.101.052005. arXiv: 1911.12606 [hep-ex].
- [11] Manimala Chakraborti, Sven Heinemeyer, and Ipsita Saha. "The new "MUON G-2" Result and Supersymmetry". In: (Apr. 2021). arXiv: 2104.03287 [hep-ph].

- [12] Georges Aad et al. "Statistical Combination of ATLAS Run 2 Searches for Charginos and Neutralinos at the LHC". In: *Phys. Rev. Lett.* 133.3 (2024), p. 031802. DOI: 10.1103/PhysRevLett.133.031802. arXiv: 2402.08347 [hep-ex].
- [13] Giordon Stark, Camila Aristimuno Ots, and Mike Hance. "Reduce, reuse, reinterpret: An end-to-end pipeline for recycling particle physics results". In: *SciPost Phys. Codebases* (2024), p. 27. DOI: 10.21468/SciPostPhysCodeb.27. URL: <https://scipost.org/10.21468/SciPostPhysCodeb.27>.
- [14] I. Adachi et al. "Search for an Invisibly Decaying Z' Boson at Belle II in $e^+e^- \rightarrow \mu^+\mu^-(e^\pm\mu^\mp)$ Plus Missing Energy Final States". In: *Phys. Rev. Lett.* 124.14 (2020), p. 141801. DOI: 10.1103/PhysRevLett.124.141801. arXiv: 1912.11276 [hep-ex].
- [15] Belle II Collaboration. *B-factory Programme Advisory Committee Reports*. 2026. URL: https://www.belle2.org/info/bpac_reports/.
- [16] ATLAS Collaboration. *Global Feature Extractor of the Level-1 Calorimeter Trigger: ATLAS TDAQ Phase-I Upgrade gFEX Final Design Report*. Geneva, Nov. 2016. URL: <https://cds.cern.ch/record/2233958>.
- [17] ATLAS Collaboration. *gFEX, the ATLAS Calorimeter Level-1 Real Time Processor*. Tech. rep. ATL-DAQ-PROC-2015-059. Geneva: CERN, Nov. 2015. URL: <https://cds.cern.ch/record/2104248>.
- [18] Julia Gonski et al. "Machine Learning on Heterogeneous, Edge, and Quantum Hardware for Particle Physics (ML-HEQUPP)". In: (Feb. 2026). arXiv: 2602.22248 [physics.ins-det].
- [19] Thiago R. F. P. Tomei. "The CMS Trigger Upgrade for the HL-LHC". In: *EPJ Web Conf.* 245 (2020). Ed. by C. Doglioni et al., p. 01031. DOI: 10.1051/epjconf/202024501031. arXiv: 2003.06460 [physics.ins-det].
- [20] Eli Gendreau-Distler et al. "Automating High Energy Physics Data Analysis with LLM-Powered Agents". In: (2025). arXiv: 2512.07785 [physics.data-an].
- [21] Eric A. Moreno et al. "AI Agents Can Already Autonomously Perform Experimental High Energy Physics". In: (2026). arXiv: 2603.20179 [hep-ex].