Ho Fung Tsoi

Postdoctoral Researcher
Department of Physics and Astronomy, University of Pennsylvania
hftsoi@sas.upenn.edu

Webpage: hftsoi.github.io

Last updated: *Feb 2026*

Research Interests

Experimental particle physics at the CERN Large Hadron Collider (ATLAS experiment) for new physics searches beyond the Standard Model; development of novel machine learning algorithms: self-supervised pretraining, anomaly detection, symbolic regression; low-latency (sub-microsecond) machine learning algorithms on FPGAs.

Education

University of Wisconsin–Madison, Ph.D. Physics 2024
Thesis: Search for exotic Higgs boson decays with CMS and fast machine learning solutions for the LHC
Advisor: Sridhara Dasu

The Chinese University of Hong Kong, B.Sc. Physics 2018
Thesis: First passage time problem of the time-dependent Ornstein-Uhlenbeck process
Advisor: Chi-Fai Lo
Visiting student at University of California Berkeley & Berkeley Lab (Jan – Aug 2017)

Professional Experience

Postdoctoral Researcher, University of Pennsylvania (Supervisor: Dylan Rankin) 2024 – Present

Awards and Honors

CERN, CMS PhD Thesis Award 2024
UWisconsin, Cornelius P. and Cynthia C. Browne Fellowship 2024
CUHK, CN Yang Scholarship 2017
CUHK, Wei Lun Foundation Exchange Scholarship 2017
CUHK, Professor Dennis Yam Kuen Lo Physics Award 2016

Selected Research Experience

New Physics Search at the LHC

- Lead analyst of the CMS Run 2 search for exotic Higgs boson decays into pseudoscalars in the final state of two b quarks and two τ leptons, and of the combination with the $\mu\mu$bb final state [5] (*2020–2024*).

Novel ML Algorithms for Particle Physics

- Self-supervised pretraining methods for jet tagging [14] and mass regression (*2025–Present*).
- Symbolic regression for automating parametric modeling of binned distributions in LHC physics analyses [7] (*2024–2025*).

Low-Latency (Sub-Microsecond) ML Algorithms on FPGAs

- Neuromorphic spiking neural networks (SNNs) for dN/dx clustering counting in drift chambers [12] (*2025–Present*).
- Sparse CNNs for spatially sparse image data on FPGAs [9] (*2024–Present*).
- Compression of neural networks using symbolic regression for fast inference [4,6] (*2022–2024*).

Trigger System at the LHC

- Commissioning the GNN-based tau trigger (GNTau) at the ATLAS High-Level Trigger for Run 3 [10] (*2024–Present*).
- Development of the CICADA anomaly detection trigger algorithm at the CMS Level-1 Trigger [3,11] (*2022–2024*).
- Development of the Data Quality Monitoring (DQM) software for the Calorimeter Layer-1 subsystem at the CMS Level-1 Trigger (*2020–2024*).
- HL-LHC trigger sensitivity projection for the CMS searches for exotic Higgs boson decays [1] (*2020–2021*).

Selected Publications

Listed below are selected publications, conference proceedings, and technical reports to which I was a primary author (*: refereed). I also co-authored papers by CMS and ATLAS as a collaboration member. For the complete list see INSPIRE. Preprints at arXiv.

[14] **jBOT: Semantic Jet Representation Clustering Emerges from Self-Distillation** 2026
Ho Fung Tsoi, Dylan Rankin
submitted to *SciPost Phys.* [arXiv:2601.11719]

[13] **SparsePixels++: Scalable Sparse Convolution on FPGAs** 2026
Ho Fung Tsoi, Dylan Rankin, Vladimir Loncar, Philip Harris
manuscript in preparation

[12] **Neuromorphic Cluster Counting** 2026
Kam Wai Lai, Ho Fung Tsoi, Dylan Rankin
manuscript in preparation

[11] **Anomaly Detection in the CMS Level-1 Trigger in Run 3** 2026
CMS Collaboration
manuscript in preparation

[10] **Performance of the ATLAS Tau Trigger in Run 3** 2026
ATLAS Collaboration
manuscript in preparation.

[9] **SparsePixels: Efficient Convolution for Sparse Data on FPGAs** 2025
Ho Fung Tsoi, Dylan Rankin, Vladimir Loncar, Philip Harris
submitted to *ACM Trans. Reconf. Tech. Syst.* [arXiv:2512.06208]

[8*] **hls4ml: A Flexible, Open-Source Platform for Deep Learning Acceleration on Reconfigurable Hardware** 2025
Jan-Frederik Schulte et al.
submitted to *ACM Trans. Reconf. Tech. Syst.* [arXiv:2512.01463]

[7*] **SymbolFit: Automatic Parametric Modeling with Symbolic Regression** 2024
Ho Fung Tsoi, Dylan Rankin, Cecile Caillol, Miles Cranmer, Sridhara Dasu, Javier Duarte, Philip Harris, Elliot Lipeles, Vladimir Loncar
Comput Softw Big Sci **9**, 12 (2025) [arXiv:2411.09851]

[6*] **SymbolNet: Neural Symbolic Regression with Adaptive Dynamic Pruning for Compression** 2024
Ho Fung Tsoi, Vladimir Loncar, Sridhara Dasu, Philip Harris
Mach. Learn.: Sci. Technol. **6**, 015021 (2025) [arXiv:2401.09949]

[5*] **Search for Exotic Decays of the Higgs Boson to a Pair of Pseudoscalars in the $\mu\mu$bb and $\tau\tau$bb Final States** 2024
CMS Collaboration
Eur. Phys. J. C **84**, 493 (2024) [arXiv:2402.13358]

[4*] **Symbolic Regression on FPGAs for Fast Machine Learning Inference** 2023
Ho Fung Tsoi, Adrian Alan Pol, Vladimir Loncar, Ekaterina Govorkova, Miles Cranmer, Sridhara Dasu, Peter Elmer, Philip Harris, Isobel Ojalvo, Maurizio Pierini
EPJ Web of Conferences **295**, 09036 (2024) [arXiv:2305.04099]

[3] **Level-1 Trigger Calorimeter Image Convolutional Anomaly Detection Algorithm** 2023
CMS Collaboration
CMS Detector Performance Summaries (2023), CERN-CMS-DP-2023-086

[2] **Searches for exotic Higgs boson decays with the CMS experiment** 2023
Ho Fung Tsoi *on behalf of the CMS Collaboration*
Proceedings of the European Physical Society Conference on High Energy Physics (*PoS* EPS-HEP2023 402)

[1] **The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger** 2021
CMS Collaboration
CMS Technical Design Report (2021), CERN-LHCC-2021-007, CMS-TDR-022

Conference, Workshop, Seminar Talks

[10] Coordinating Panel on Advanced Detectors (CPAD) — UPenn, USA — Oct 2025
SparsePixels: Efficient Convolution for Sparse Data on FPGAs

[9] 2024 CMS Thesis Award Ceremony Talk (CMS Week) — CERN, Switzerland — Oct 2025
Search for Exotic Higgs Boson Decays with CMS and Fast Machine Learning Solutions for the LHC

[8] Fast Machine Learning for Science Conference (FastML) — ETH Zurich, Switzerland — Sep 2025
SparsePixels: Efficient Convolution for Sparse Data on FPGAs

[7] High-energy physics seminar — UPenn; U.Washington; Fermilab, USA — 2023, 2024
Search for Exotic Higgs Boson Decays with CMS and Fast Machine Learning Solutions for the LHC

[6] US LHC Users Association Meeting (US LUA) — Fermilab, USA — Dec 2023
CICADA: Anomaly Detection for New Physics Searches at the CMS Level-1 Trigger

[5] Machine Learning at Level-1 Trigger Workshop (ML@L1) — CERN, Switzerland — Dec 2023
Anomaly Detection – CICADA: Status, Plans, and Prospects for Phase-2

[4] CMS Machine Learning Town Hall — CERN, Switzerland — Sep 2023
L1 Anomaly Detection with Calorimeter Inputs: Status and Opportunities

[3] European Physical Society Conference on High Energy Physics (EPS-HEP) — Hamburg, Germany — Aug 2023
Searches for exotic Higgs boson decays with the CMS experiment

[2] International Conference on Computing in High Energy & Nuclear Physics (CHEP) — Norfolk VA, USA — May 2023
Symbolic Regression on FPGAs for Fast Machine Learning Inference

[1] International Conference on Applied Mathematics — CityU, Hong Kong — May 2016
First Passage Time Problem of the Time-Dependent Ornstein-Uhlenbeck Process: a Model for Stochastic Decision-Making Process

Professional Service

Membership

- Member, ATLAS Collaboration, CERN — 2024 – Present
- Affiliated member, Accelerated AI Algorithms for Data-Driven Discovery (A3D3) Institute — 2024 – Present
- Member, CMS Collaboration, CERN — 2020 – 2024

Collaboration Role

- **Co-convener**, SUSY Monte Carlo & Interpretation subgroup [CMS] — 2023 – 2024
- **Contact person**, Higgs boson Monte Carlo group [CMS] — 2021 – 2023
- **Software developer**, Data Quality Monitoring system for CaloLayer-1 trigger [CMS] — 2020 – 2024

Journal Referee

IEEE Transactions on Evolutionary Computation (2026); *Advanced Theory and Simulations* (2025).

Teaching and Mentoring

Teaching

- Guest speaker, PHYS 3359 [UPenn]: Data Analysis for the Natural Sciences II: Machine Learning (Spring 2025)
- Teaching assistant, undergraduate physics courses [U.Wisconsin] (Fall 2018 – Spring 2020)

Master students

- Kam Wai Lai [UPenn]: Neuromorphic AI with spiking neural networks on FPGAs (2025 – Present)
- Yuyan Wang [UPenn]: Self-supervised learning with JEPA-based pretraining for LHC jet tagging (2025 – Present)

Undergraduate students

- Abhay Agarwal [UPenn]: Scalable sparse convolution on FPGAs (2026 – Present)

- Alex Yang [UPenn]: Self-supervised learning with SimCLR/VICReg for reconstructing the mass of supersymmetric particles (2025 – Present)
- Ashni Kumar [Drexel]: Acceleration of reinforcement learning agent inference on FPGAs for astronomical observations (2025)