

Session Program

31 August 2026 to 4 September 2026



Fast Machine Learning for Science Conference 2026

Reception & Poster Session

Monday 31 August

17:36

Reception & Poster Session

17:36–19:06

Residual Geometry as a Diagnostic for Hidden Failure Modes in Fast Language Model Compression

Speaker

Jun-Sik Yoo

17:36–19:06

Deploying a Diffusion model on Voyager for cosmological applications

Speaker

Javier Hernandez-Nicolau

17:36–19:06

Machine Learning-Based Large-R Jet Tagging for the ATLAS Level-0 Global Trigger

Speaker

Maryam Bayat Makou

17:36–19:06

Jet Tagger for Real-Time Triggering at the LHC

Speaker

Kalib McEuen

17:36–19:06

TensorFlow-to-ONNX Migration of CMS DeepTau Models for Efficient Inference in CMSSW and SONIC

Speaker

Pritam Palit

17:36–19:06

Accelerating Lattice Gauge Simulations with Gauge-Equivariant Graph Neural Networks

Speaker

Gia-Wei Chern

17:36–19:06

Generative Cosmological Emulators as Fluid Flows in Data Space

Speaker

Sandip Roy

17:36–19:06

The CMS ML-based Data Quality Monitoring Infrastructure

Speaker

Dr Antonio Vagnerini

17:36–19:06

A Mamba-2 based Selective State Space Model for Improving MET Estimation in CMS Run 3 and the HL-LHC

Speaker

Aleesha Kallil Tharayil

17:36–19:06

TriggerFlow: A Reproducible End-to-End Framework for Machine Learning-Driven Level-1 Trigger Systems in CMS**Speaker**

Marius Köppel

17:36–19:06

Using Variational Autoencoders to Generate Atmospheric Models for Brown Dwarfs**Speaker**

Julia Haynes

17:36–19:06

Proxy-guided DSE for Neural Networks on FPGAs**Speaker**

Gabriele Trotta

17:36–19:06

Nereid: an Inference Server for HPC**Speakers**

Ethan Colbert, Noah Paladino

17:36–19:06

Direct Parameter Estimation of Binary Neutron Star Mergers with State Space Models**Speaker**

Kyungseop Yoon

17:36–19:06

Design and Implementation of Tiny Real-Time ML for FPGAs with Agentic Workflows**Speaker**

Qibin Liu

17:36–19:06

A Controlled Comparison of FINN and Vitis AI for Small Neural Networks**Speaker**

Ruthwik Reddy Sunketa

17:36–19:06

Bridging the Gap Between Enterprise Memory-Intensive Workloads and Real-Time Distributed ML Systems: A Scheduling Reality Check**Speaker**

Atul Garg

17:36–19:06

SNAC-Pack: Surrogate Neural Architecture Codesign Package for Hardware-Aware Machine Learning**Speaker**

Jason Weitz

17:36–19:06

Pruning LUT-Based Neural Network Ensembles for Efficient FPGA Inference**Speaker**

Sijan Shrestha

17:36–19:06

Multi-SG Parameter Estimation for Anomalous Gravitational Wave Signals**Speaker**

Jack Redepinning

17:36–19:06

A Controlled Benchmark for Memory Mechanisms in Efficient-Attention Transformers**Speaker**

Robert Sneiderman

17:36–19:06

Sub-100ns On-Probe Spike Sorting via Resource-Constrained FPGA**Speaker**

Aarav Gaur

17:36–19:06

ML Classifiers for the search of new post-AGB candidates in all-sky surveys.**Speaker**

Diego Osvaldo Ochoa de la Cruz

17:36–19:06

Lightweight Transformers for High-Energy Physics Experiments**Speaker**

Erdem Yigit Ertorer

17:36–19:06

Reject or Compress. Event streaming and triggering system using Byte Latent Transformer.**Speaker**

Leonid Didukh

17:36–19:06

SUNet: A Spectral UNet for Axion Dark Matter Signal Denoising**Speaker**

Akbot Assan

17:36–19:06

A Binary-Weight Transformer for Real-Time Particle Jet Tagging at the LHC**Speaker**

Kai Yamaguchi

17:36–19:06

To Chop or Not To Chop: the Question for Variational Autoencoders as Anomaly Detectors in High Energy Physics**Speaker**

Sterre Hoogendoorn

17:36–19:06

HACE: A Surrogate-Guided Agent for hls4ml Configuration on Scientific FPGAs**Speaker**

Mr Venkata Sai Prathyush Turaga

17:36–19:06

SparsePixels++: Scalable Sparse Convolution on FPGAs

Speaker

Ho-Fung Tsoi

17:36–19:06

PLANT: Conditional Generative Models for Fast Gravitational-Wave Population Synthesis**Speaker**

Rhea Senthil Kumar

17:36–19:06

An extreme edge ML-based electron/photon trigger for the ATLAS Global Trigger**Speaker**

Adrian Chitan

17:36–19:06

Benchmarking BitNet-Style Low-Precision Networks for Fast Jet Classification Toward FPGA Deployment**Speaker**

Dorian Sloom

17:36–19:06

CNN-based event classification for first level triggers at HL-LHC**Speaker**

Jim Brooke

17:36–19:06

Memory-Efficient Activation Checkpointing with Sliding Window and Hirschberg's Algorithm for 0/1 Knapsack Solving in PyTorch**Speaker**

Jędrzej Maczan

17:36–19:06

PermuGini-RF: A Hybrid Feature Selection Model of Permutation and Gini Importance for Mobile Phone Price Classification**Speaker**

FEMI JOHNSON

17:36–19:06

TSCC as an AI-for-Science Platform: Enabling Research with LLMs, Agentic Tools, and Campus Cyberinfrastructure**Speakers**

Subhashini Sivagnanam, Yuwu Chen

17:36–19:06

Machine-Learning-Based Self-Trigger for the COHERENT Cryogenic CSI Detector**Speaker**

Zepeng Li

17:36–19:06

FlowFI: Dataflow and Quantization Aware Fault Injection for Resilient Accelerators**Speaker**

Raymond Duenas

17:36–19:06

FQTree: Fine-Grained Quantization and Hardware Generation of Boosted Decision Trees**Speakers**

Zhiqiang Que, Chang Sun

17:36–19:06

Predicting grain-size-aware flow behavior via a physics-informed neural network elasto-viscoplastic model**Speaker**

Adnan Eghtesad

17:36–19:06

Real-Time Multi-Messenger Trigger Fusion Using Machine Learning for Sub-Threshold Event Detection**Speakers**

Mr Kaamesh Chandrasekaran, Ms Madhushree Naga

17:36–19:06

Fast machine learning trigger system for real-time event identification for Ba-tagging in a Liquid Xenon TPC**Speaker**

Melissa Medina Peregrina

17:36–19:06

ErwinParT: A Radially-Partitioned Hierarchical Transformer for Ultra-Low Latency Jet Tagging at the CMS Level-1 Trigger**Speaker**

Vaibhav Lohia

17:36–19:06

Applications of machine learning to classification of cosmic X-ray sources: results and challenges.**Speaker**

Yichao Lin

17:36–19:06

From Keras 2 to Keras 3: QKerasV3 Integration with the hls4ml FPGA Flow**Speaker**

Marius Köppel

19:06