

# Distributed statistical inference with pyhf

*Tuesday, July 6, 2021 3:00 PM (30 minutes)*

pyhf is a pure-python implementation of the HistFactory statistical model for multi-bin histogram-based analysis with interval estimation based on asymptotic formulas. pyhf supports modern computational graph libraries such as JAX, PyTorch, and TensorFlow to leverage features such as auto-differentiation and hardware acceleration on GPUs to reduce the time to inference. pyhf is also well adapted to performing distributed statistical inference across heterogeneous computing resources (clusters, clouds, and supercomputers) and task execution providers (HTCondor, Slurm, Torque, and Kubernetes) when paired with high-performance Function as a Service platforms like funcX or highly scalable systems like Google Cloud Platform that allow for resource bursting. In this notebook talk we will give interactive examples of performing statistical inference on public probability models from ATLAS analyses published to HEPData in which we leverage these resources to reproduce the analyses results with wall times of minutes.

**Primary authors:** STARK, Giordon Holtsberg (University of California, Santa Cruz (US)); FEICKERT, Matthew (Univ. Illinois at Urbana Champaign (US)); HEINRICH, Lukas Alexander (CERN)

**Presenter:** FEICKERT, Matthew (Univ. Illinois at Urbana Champaign (US))

**Session Classification:** Plenary session Tuesday