SOS27: Abstract for session “Running Machine Learning at Scale”

High-performance computing (HPC) has become an indispensable tool for advancing Machine Learning (ML), but scaling these workloads to the largest systems available presents a host of challenges. From the technical demands of distributed training, through the logistics of storage aspects, to the operational requirements of managing extended runs on large amounts of GPUs, the journey is fraught with hurdles.

By combining insights from leading experts at national laboratories, in academia and in industry, the session aims to foster a holistic discussion about the operational and technical demands of large-scale ML training, as well as the opportunities for innovation that these represent for us.

The session will also consider the disparity between industry and academic institutions in scaling ML applications. While industry benefits from dedicated resources and streamlined processes, academic and lab environments often operate under different constraints. By examining these contrasts, the session will seek to identify common ground, shared solutions, and opportunities for collaboration that could accelerate innovation across domains.

Through a mix of real-world experiences and forward-looking perspectives, this session will delve into the technical and strategic landscape of running ML applications at scale, inspiring new approaches to make these transformative workloads more efficient and impactful for our users.