

CSAID AI Jamboree Flash Talks




January 30, 2025

Smarter Storage - Nick Smith

- Distributed storage is our gold mine: all science insight requires performant, reliable access to petabytes of data
 - Performance: choose right technology for right dataset
 - Smart caching (prior art using deep Q-learning: [2208.06437](#))
 - Reliability: predict and reduce impact of hardware failure
 - All modern storage distributes data across multiple drives/servers/racks such that one loss does not prevent data access. Frequency of failures at scale causes trouble for hash-based data location such as Ceph ([meta](#)). Predicting and pre-handling failure may help? Related: [disk prediction](#)

Optimizing resource requests- K. Herner

- Often a clear peak with a long tail in batch job resource profiles
 - Initial instinct is to provision for the most demanding jobs.
 - The more you request, the less can run concurrently...
- Often also have an initial setting and then increase on retries
 - What are the optimal initial settings for a given task? Can we determine in a programmatic way? Can we dynamically update job resources before exhaustion?
 - Some literature exists on the topic
- Historical data on previous runs can often guide us
 - Use it to so generate (close to) optimal requests ahead of time?
- How fine-grained can we get with resource requests as a function of input data? Per input file? Dynamic adjustments in jobs? Many possibilities?
- **Can we make a generic tool useful to multiple experiments?**

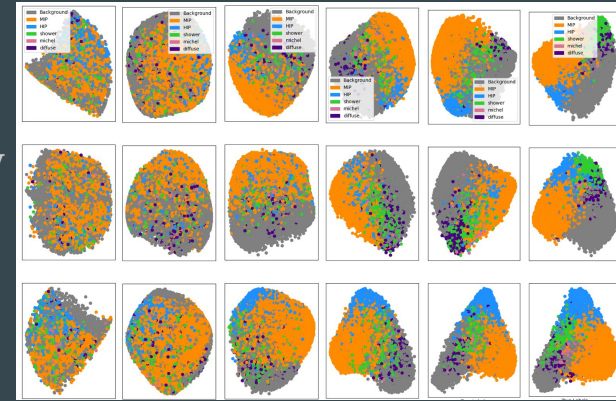


Making Monitoring Data Accessible to AI - K. Retzke



- The Landscape project collects and stores a large volume of data:
 - Batch systems, disk & tape storage, data transfers, system/host metrics, service metrics, traces, logs
 - Metrics in time-series databases Graphite and Prometheus/Mimir
 - Log/event data in semi-structured (JSON) documents in Elasticsearch/OpenSearch
 - More Info: <https://landscape.fnal.gov/docs/>
- How do we make this data more accessible for AI projects?
 - Online data warehouse/data lake - Clickhouse, Iceberg, etc?
 - Offline files/buckets (data swamp?) - Parquet, CSV?
- What tools/interfaces should Landscape provide to enable AI-powered research?

Network Explainability - M. Voetberg & G. Cerati

- Built-in network explainability required to not add uncertainty into analysis of the network
 - Most existing post-hoc methods require either training a method or additional optimization problem
 - Must be better solutions than asking an LLM for reasoning (inherent fraught by non-negligible error bars) or looking at a network latent space (biased by both user interpretation and subset of data used for evaluation and visualization)
- Generic framework for putting together AI methods that can be explained generically (e.g. network blocks) - similar to having saliency maps for CNNs.
- Linear network probes worked well for a classification network with discrete "concepts", but unscalable for different types of networks and problems.
 - How exact to do this? Required to actually look into both internal network components and data structure features.
 - In line with the upcoming FOA
 - Building off previous work with GNNs



Latent Space Clustering study of GNN classifications



“Concept Probes” - decoders measuring how well different combinations of classes are classified at points in the network’s inference process

Dynamic learning systems for storage and network - Lindsey Gray

- Data placement and satisfactory user feedback loop in HEP/Astro data processing systems is an unsolved problem
- Try to use RL and Agents to solve this? i.e. Define **actions** that are the results of **policies**
 - “User initiates a job **J** that requires dataset **D** running at site **X**”
 - Create virtual network circuit between sites?
 - bandwidth guarantee?
 - Transfer data from sites **Y, Z** to **X**?
 - partial xfer to **X** for better throughput?
 - How long do we keep data at **X**?
 - Delay executing **D** until bandwidth between **X,Y,Z** is suitable for xfer or rapid completion?
- Understanding what to do requires a just-in-time global optimization over multiple parameters with constraints
 - Using MDMM (Lagrange Multipliers for NNs), Reinforcement Learning can create a feedback loop that satisfies physical constraints of network / storage
 - Learn policies based on satisfying constraints, user requests, and dataset features
 - Agents can be used as an automated system for creating/submitting requests to lower level tools from high-level user request
 - Agent requests can then be used to close to loop for managing data in a geographic region, subject to the constraints of WAN traffic, job load, storage

Tape Storage Digital Twin - S. Norberg

- We have a lot of data about how tape is used via multiple experiments, from the steady state usage to the bursty experimental usage
 - With this data I want to create a digital twin to run tests on
 - Dev systems do not include the problems we see in the production systems (complexities of user interactions are not included)
 - We have a virtual tape environment available but it does not simulate experimental access
- If we had a digital twin we could then do testing on a more realistic system
- One way to use the digital twin would be with working on how to place and take data off the tape system(It is more or less a first come first serve system now.)
- There are only so many problems a human can find, having a digital twin test stand for tapes can help to find problems before they happen

Backup and additional convo topics

LLM + RAG Ticket Triage?

- System for automatically redirecting tickets/common questions to documentation or correct POC
 - Many tickets can be answered with “Here is the page you’re looking for”. Why don’t we automate that?
 - Free up time of already oversubscribed support staff
- See Gennadiy’s talk from this morning for a Proof of Concept
- Use open source model (Ollama, etc) hosted on site, condition with existing documentation sources and codebases

Smart Heterogeneous Computing Management - Nhan Tran (?)

- Classic question of “How do we know when people need what compute?”
- Predictive management of systems - recommend for up or down scaling on certain resources (e.g. HEP Cloud)