

# Decode the Workload: Training Deep Learning Models for Efficient Compute Cluster Representation

October 19 - 25, 2024

CHEP  
2024



Ahmed Mohammed, Mark Jones, Diana McSpadden, Malachi Schram, Bryan Hess, Kishansingh Rajput, Zhenyu Dai, Laura Hild, Wesley Moore

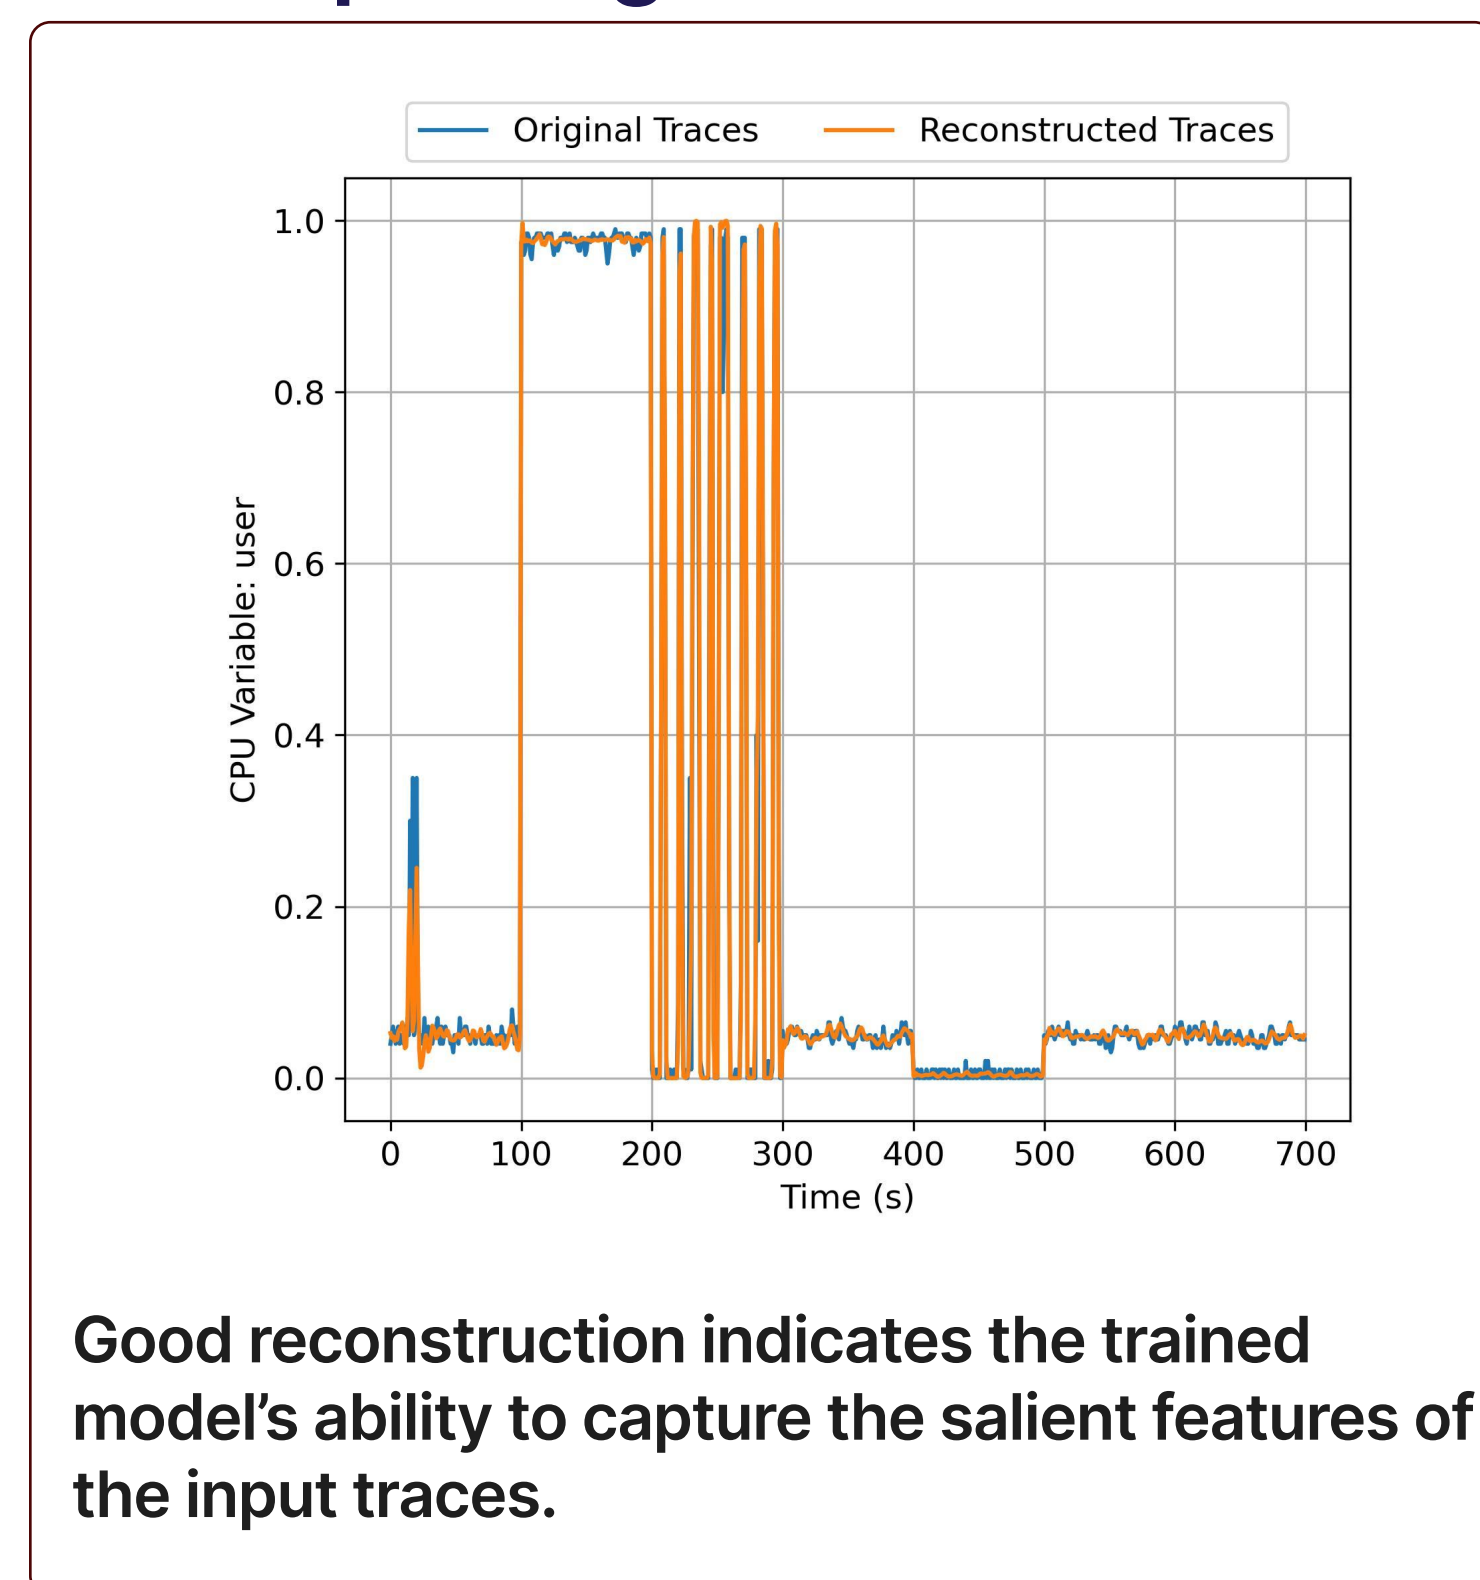
Thomas Jefferson National Accelerator Facility  
ahmedm@jlab.org

Monitoring compute cluster activity at higher sampling fidelity enhances job differentiation, while Graph Neural Network (GNN) assisted autoencoders produce higher reconstruction error for anomalous jobs.

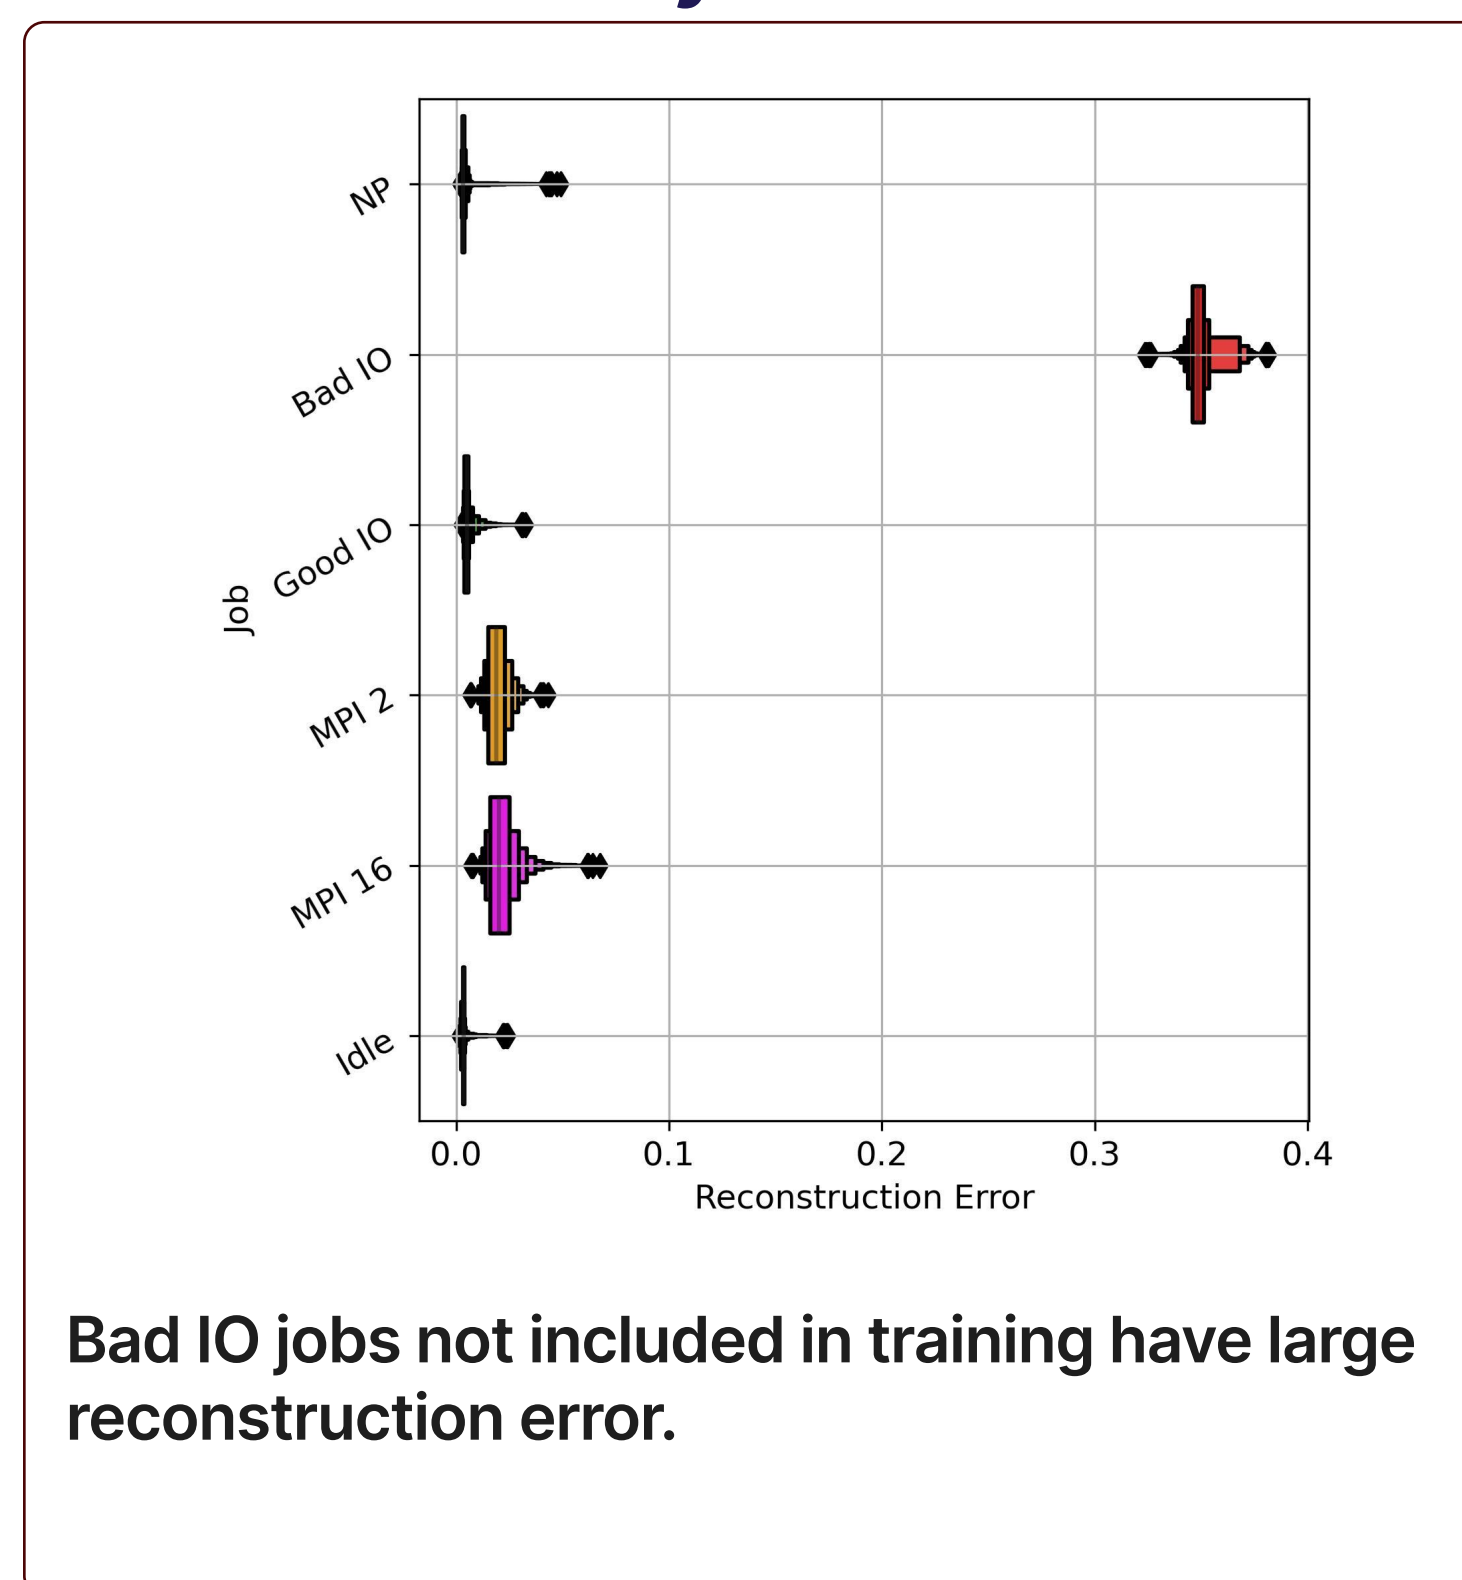
We aim to develop model(s) to capture the salient features of a high throughput computing cluster to be used for different downstream tasks such as anomaly detection and diagnosis.

## Results for Vanilla Autoencoder (AE)

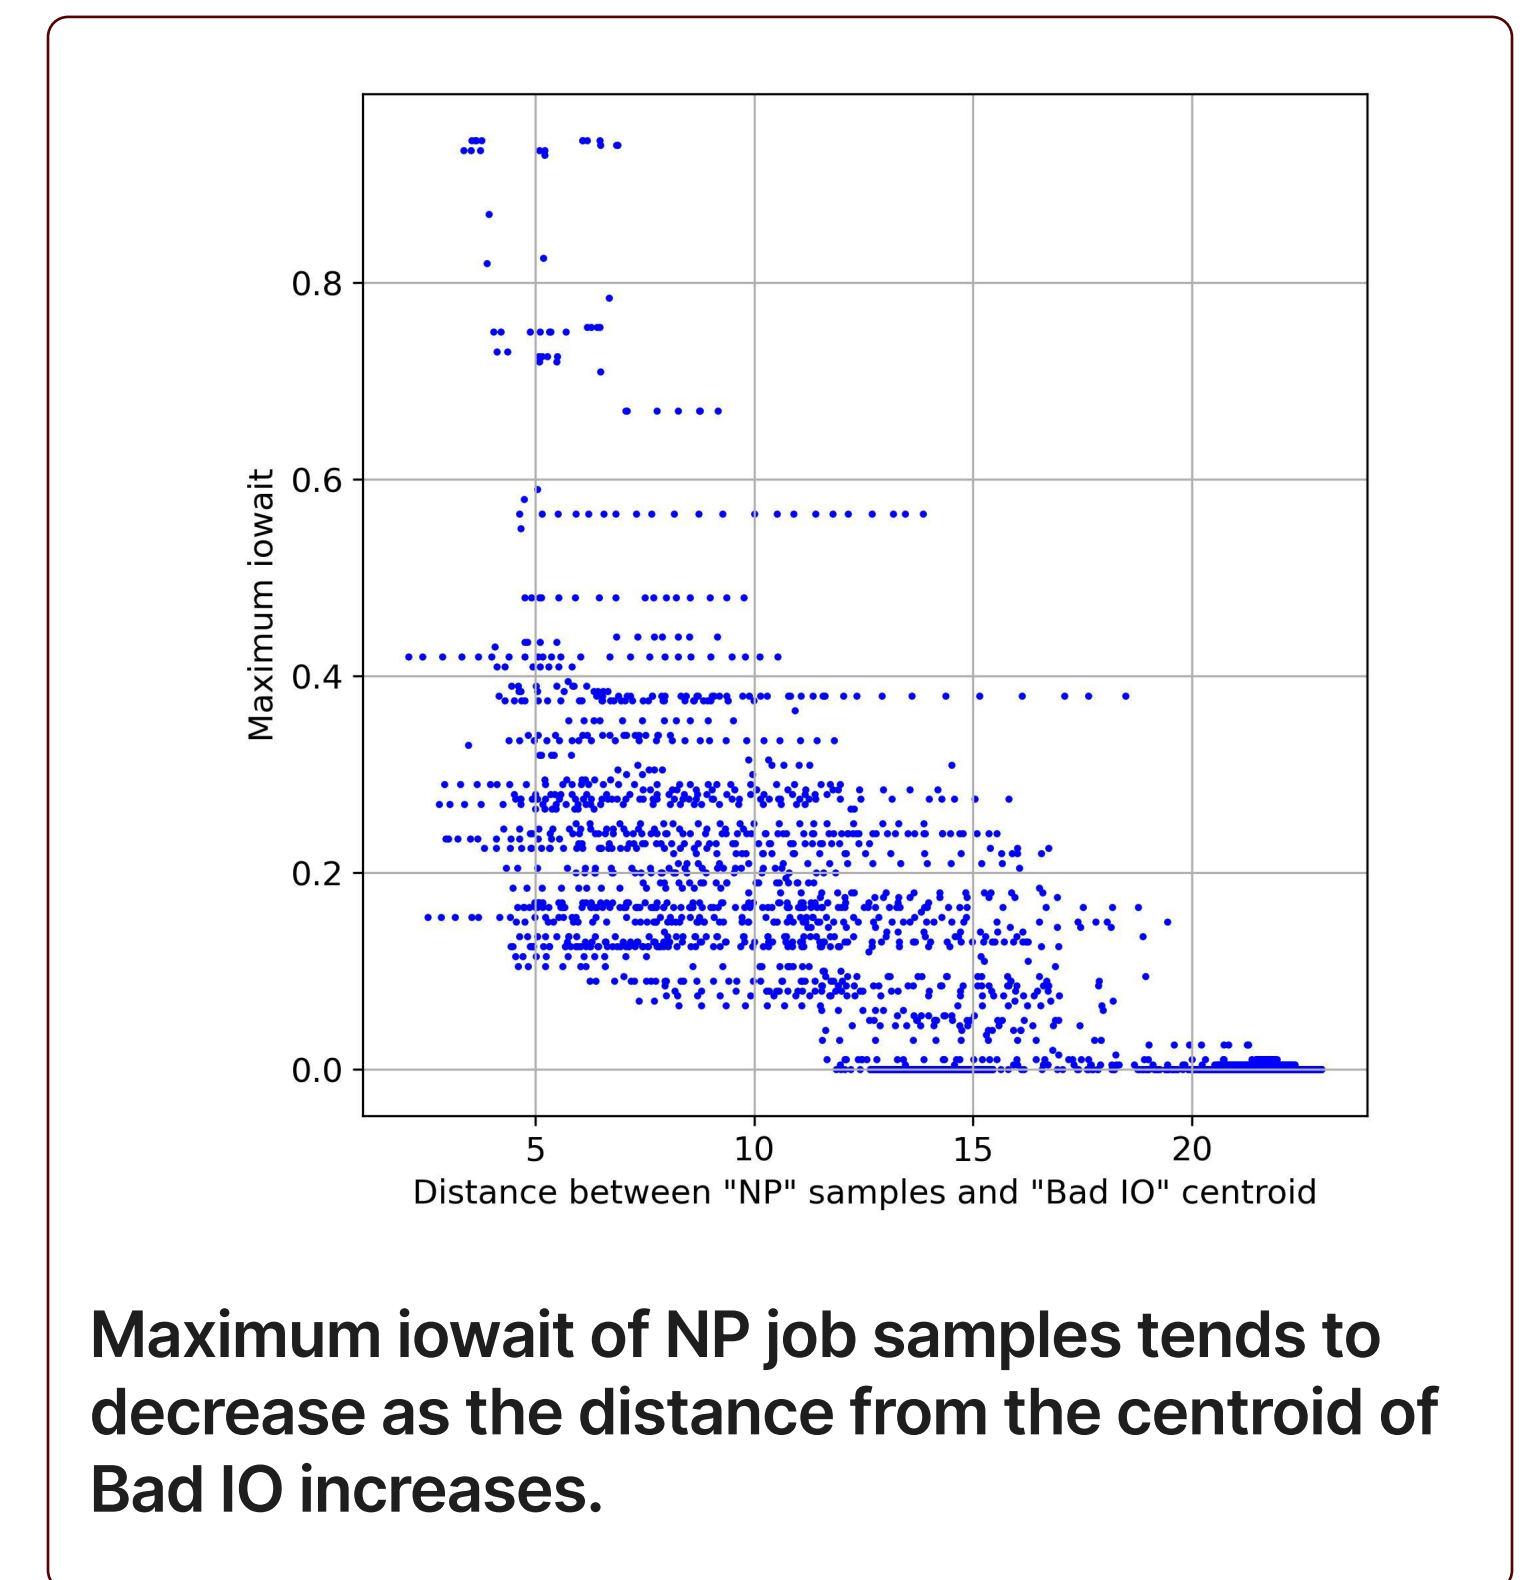
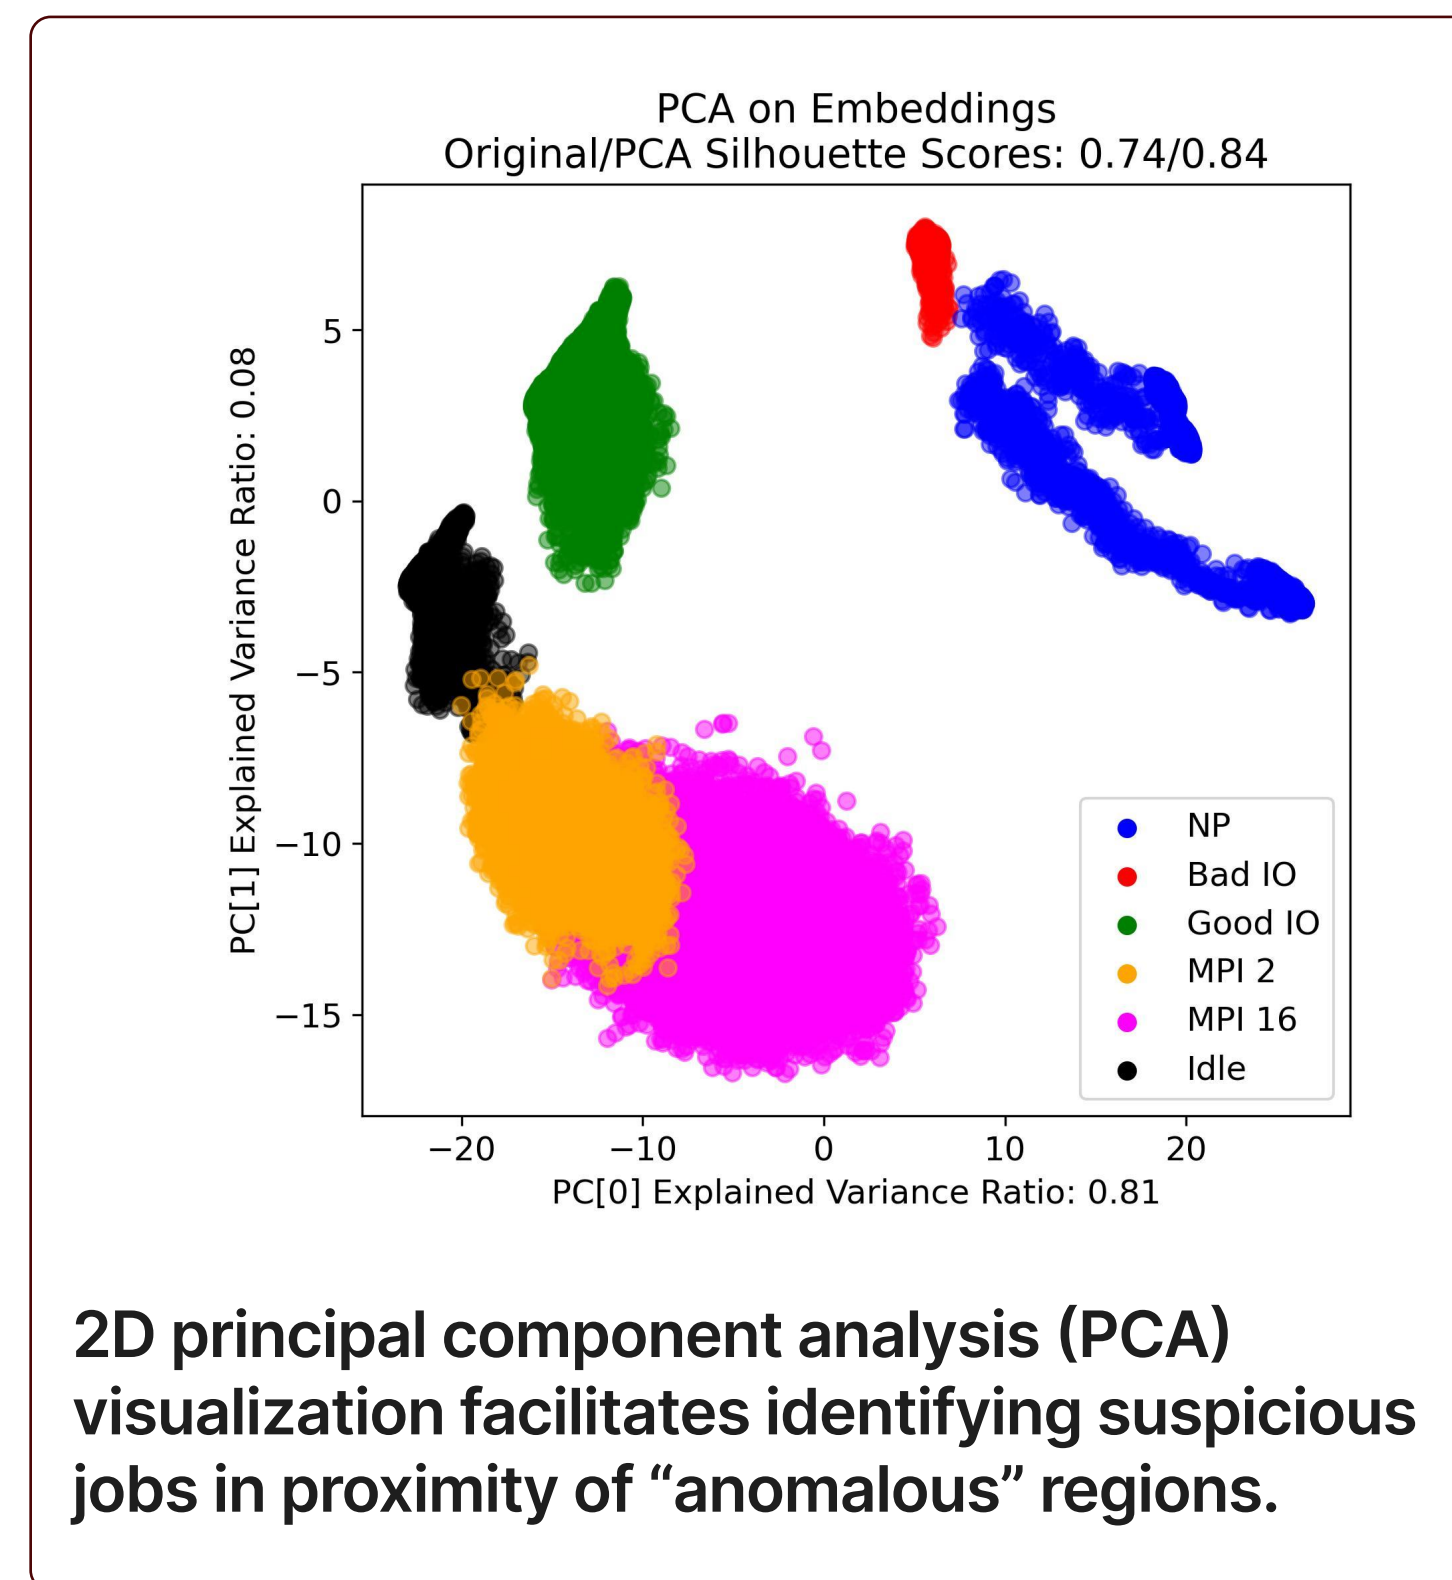
### Capturing salient features



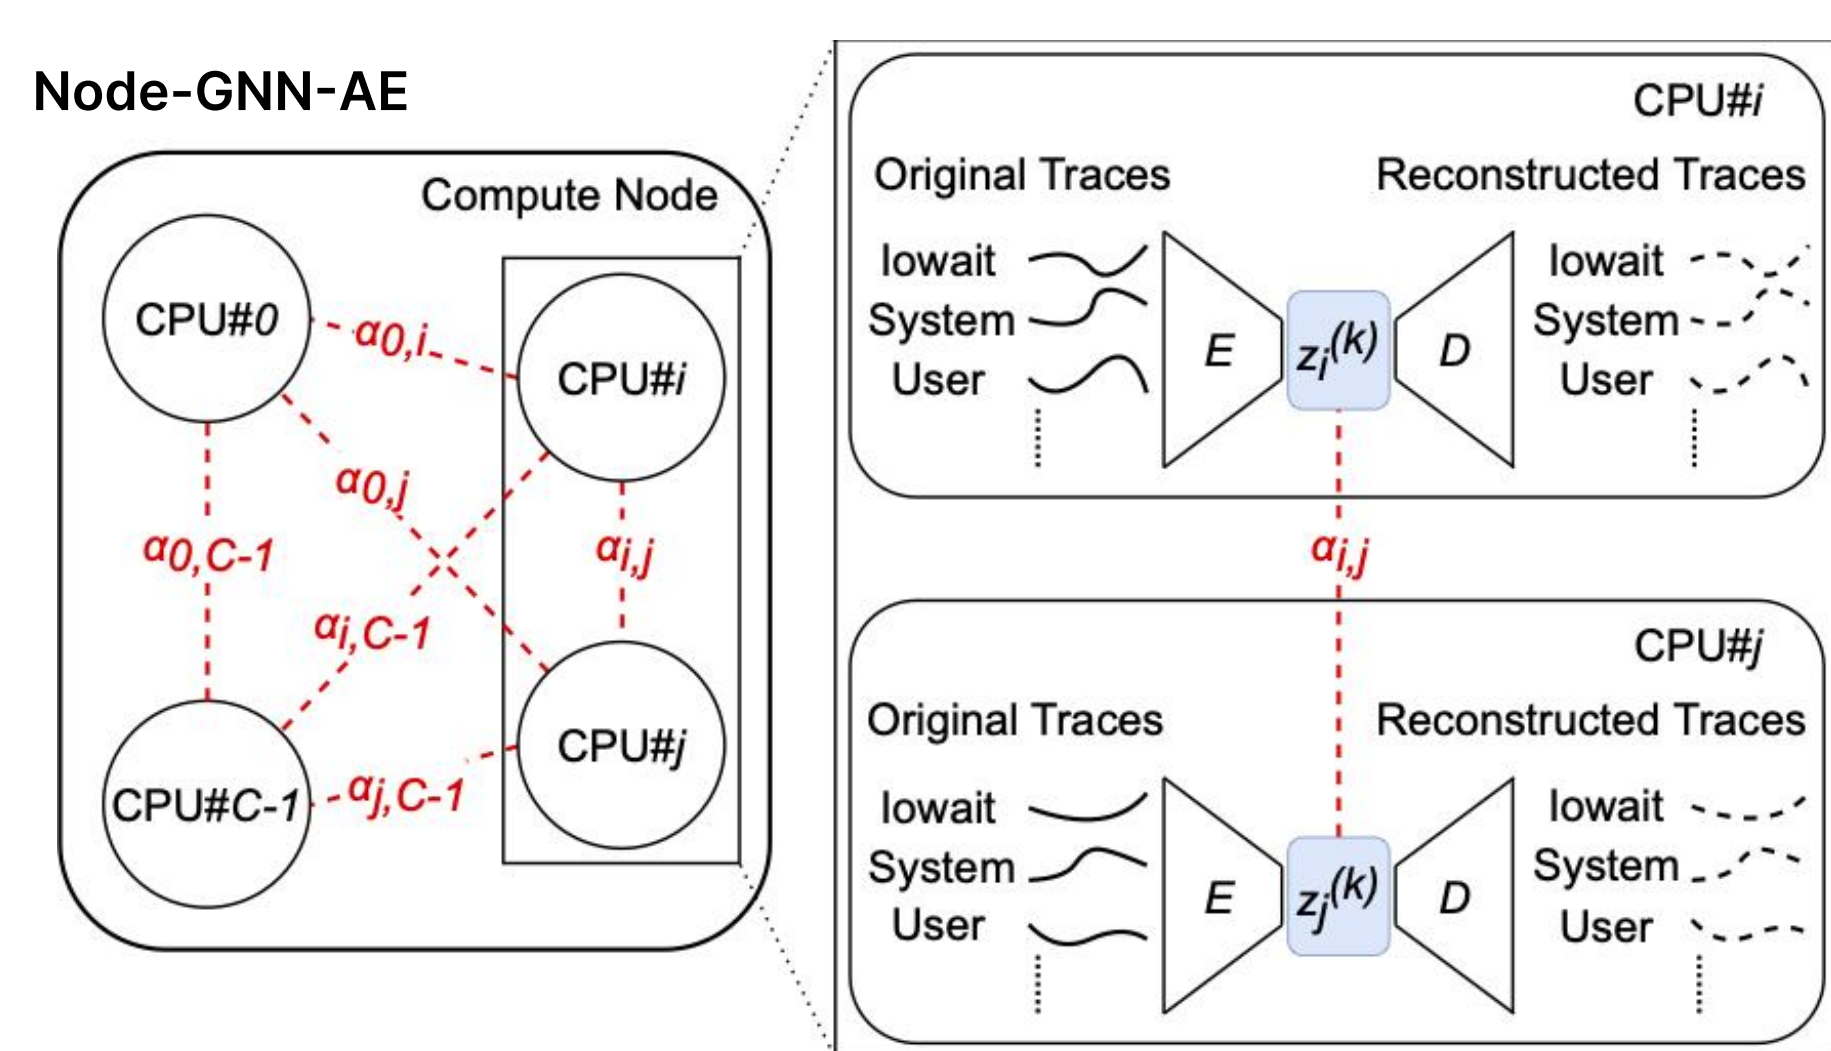
### Anomaly detection



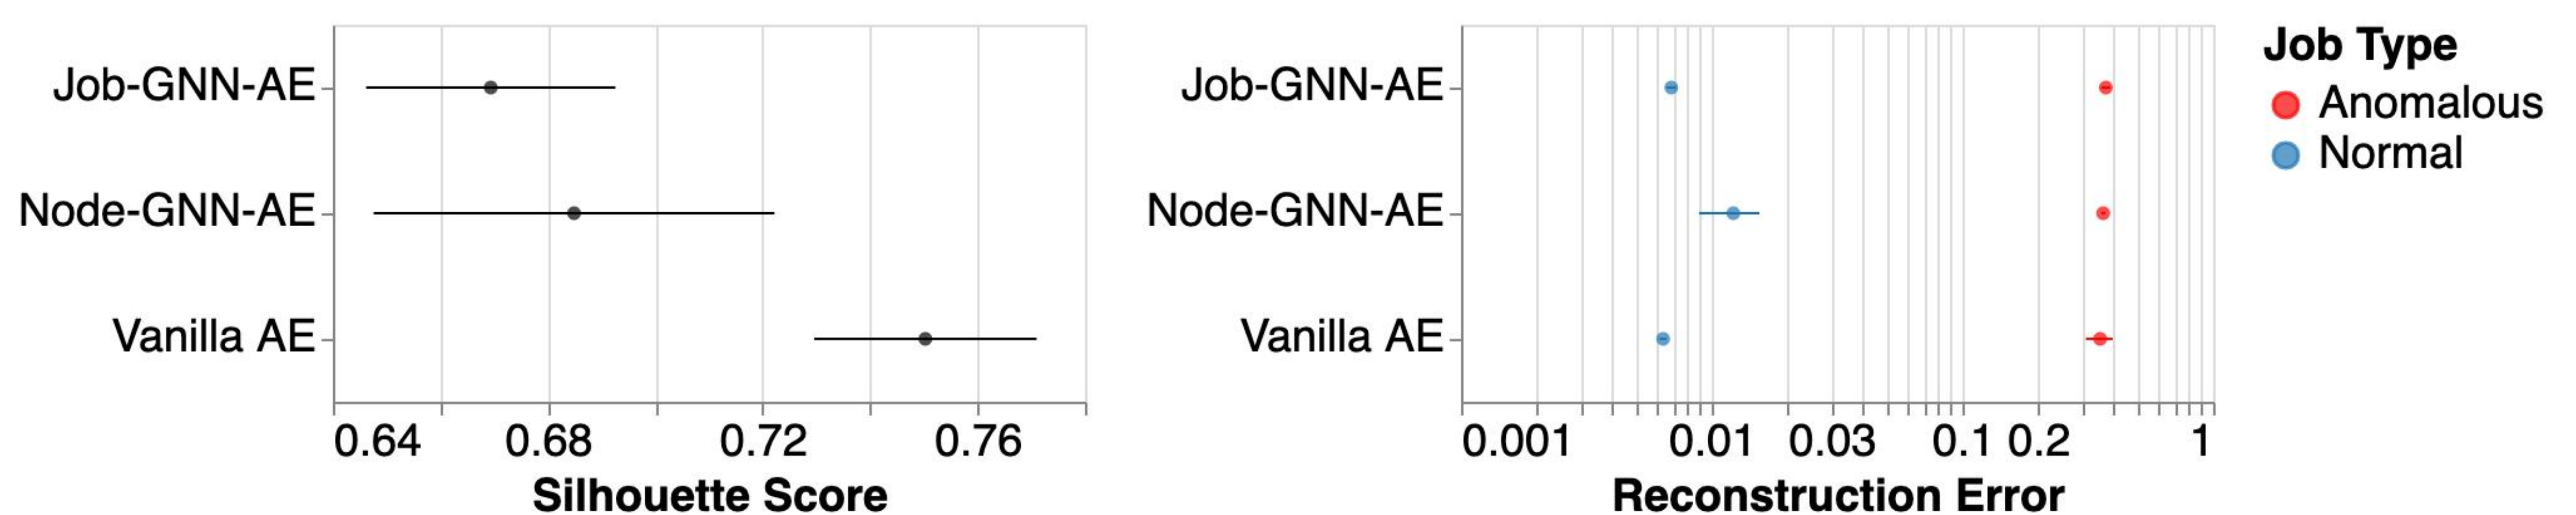
### Diagnosis



## Model architectures and performance



Reconstruction Error and Silhouette Score statistics on an ensemble of size 10 with random initializations for each of the model architectures



## Data Description

Prometheus with a custom CGroup-v2 exporter is used to collect Linux kernel CPU metrics with a flag to identify the job running on each thread.

Job	Description
Nuclear Physics (NP)	Likelihood fit of partial wave analysis [1]
Good IO (Bad IO)	Bad and Good IO generate several flat files for binning, with Bad IO specifically tuned to generate IO pressure through common bad coding practices.
MPI	CPU-bounded message passing interface (version: MPICH) job running on (2/16) threads for random amount of time before executing Allreduce
Idle	Refers to any thread not running any job

## Higher sampling improves differentiation

Silhouette scores quantify the separability between different job types [2].

