Model-Agnostic Influence Analysis for Performance Data

Rahul Sridhar[1, 2], Rushil Anirudh[2], Jayaraman J. Thiagarajan[2], Nikhil Jain[2], Todd Gamblin[2]

[1]University of California, Irvine [2]Lawrence Livermore National Laboratory

Understanding influences of input parameters and different runs on performance are at the core of using machine learning in HPC for performance analysis. To circumvent the challenge of model bias introduced by existing approaches, we present a model-agnostic, graph based approach for influence analysis. Experiments with Kripke provide new insights into the typical predictive modeling pipeline.

1 Influence Analysis in HPC

- Performance measures such as execution time are sensitive to changes in parameter settings – number of threads, power limit, etc.
- **Goal:** Quantify the importance of parameters and samples in predicting variations in performance
  - ✓ Identify parameters significant for performance optimization
  - ✓ Identify sub-domains of interest in parameter space
- **Data:** Simulations with different parameter settings on Kripke (a Sn transport mini-app) and corresponding execution times

2 A Generic, Surrogate-Free Approach

- Existing approaches use surrogate models to measure parameter influence
- **Challenges:** Introduce unaccounted model bias and cannot generalize to sample influence analysis
- Proposed approach is model-agnostic and naturally applies to both cases
  - ✓ Utilizes tools from function analysis on graphs

3 Methodology

Step 1: Graph construction

- ✓ Sample runs → nodes (different parameter settings)
- ✓ Edge weights → Gower distance
- ✓ Signal $f: \mathcal{V} \rightarrow \mathbb{R}$ (execution time)

Signal on a graph

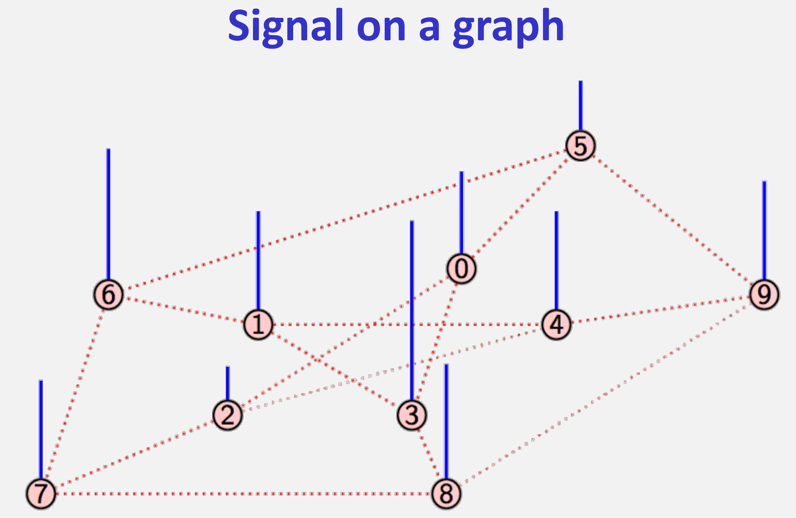

Gower distance

Variable type	Distance
Categorical	Match/No match
Continuous	Interval scaled
Ordinal	Converted to ranks and then interval scaled

Step 2: Spectral decomposition of a graph signal

- Decompose signal into its constituent Laplacian eigenfunctions
- Changes in graph connectivity → changes in the signal's spectrum

Step 3: Influence estimation

- **Parameter influence:** Use a leave-one-parameter-out strategy to study the changes in the spectrum at "high frequencies"
- **Sample influence:** Use a local filtering approach at each sample to study the changes in the spectrum at "low frequencies"

4 Identifying Influential Parameters

INPUT PARAMETER SPACE
1. Number of threads (OMP)
2. Power cap (PKG_LIMIT)
3. Loop nesting order
4. Other Kripke-specific parameters

➡ **Leave-One-Out Analysis** ➡ **Parameter Influences** ➡ ☺

Parameter influence metric: $E = |\sum_k^N \hat{x}_k^2 - \sum_k^N x_k^2|$ (energy in high frequencies)

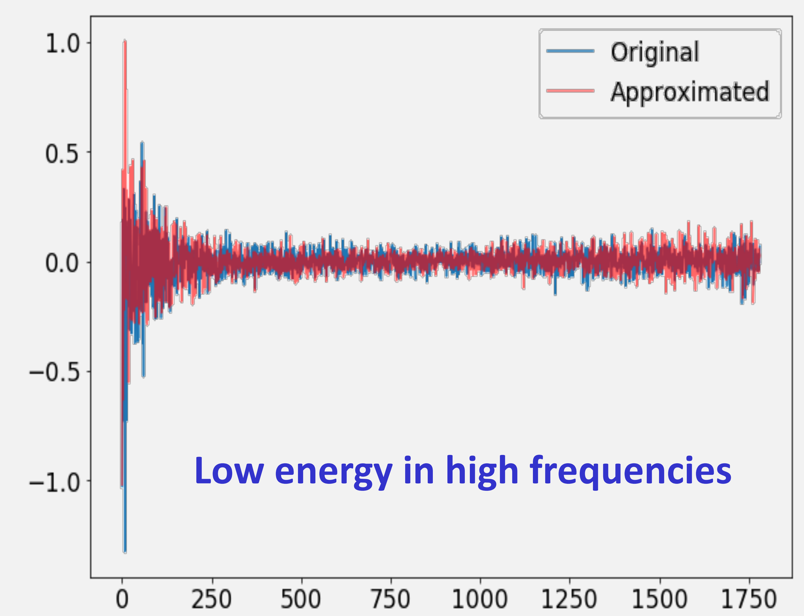

Removing a weak parameter

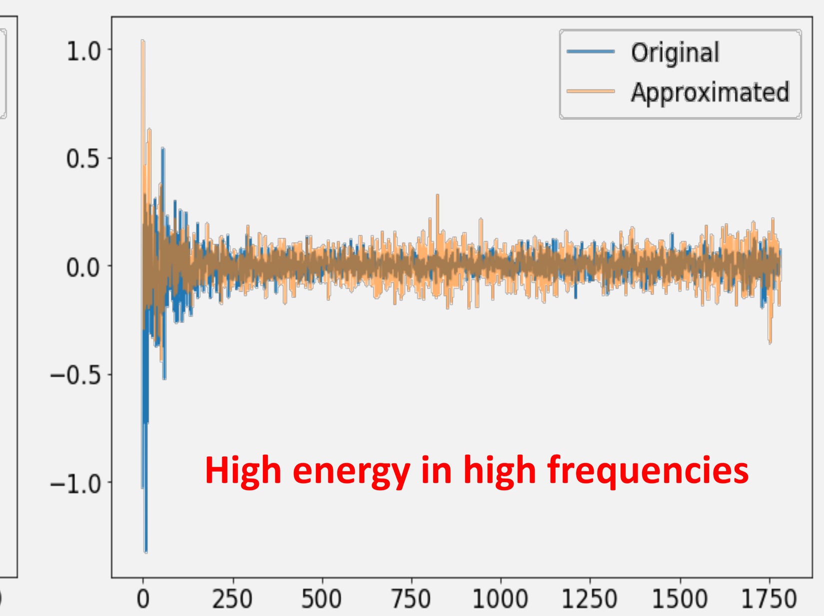

Removing a strong parameter

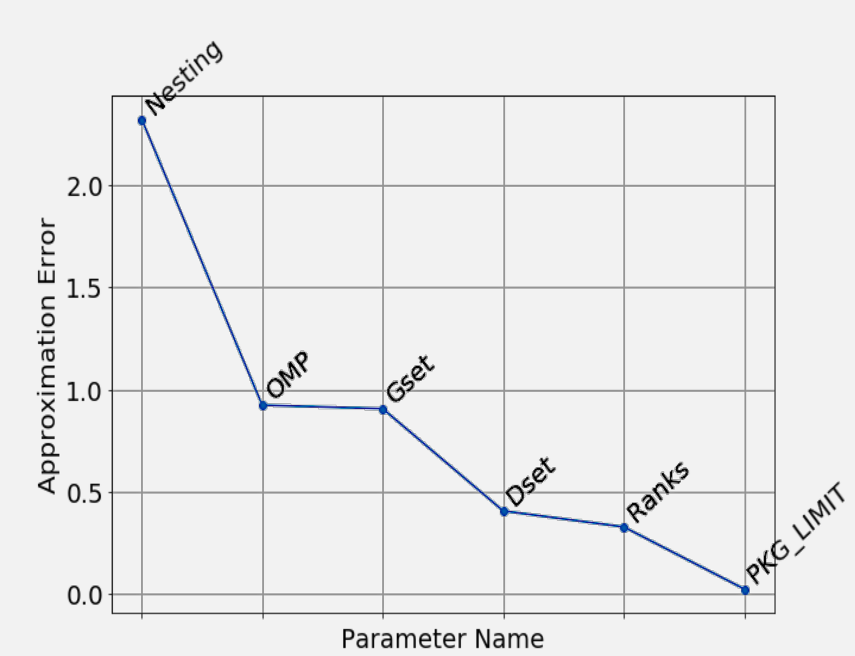

Parameter influence plot
Nesting
OMP, Gset
Dset, Ranks, PKG_LIMIT

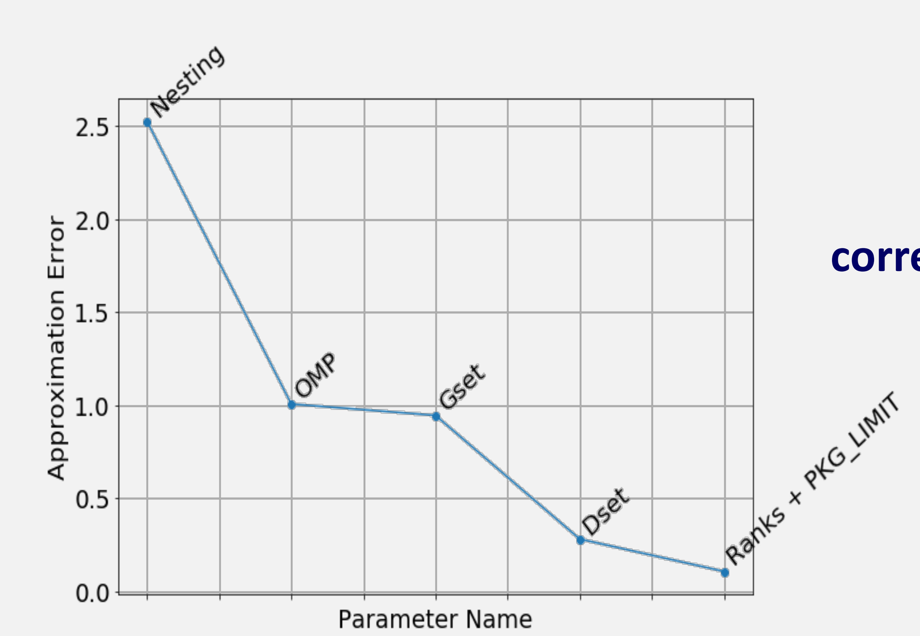

Are there higher order correlations between parameters?

Takeaway: Results align with intuition of domain experts and this can enable HPC analysts to tweak only the most influential parameters to control the performance variations

5 Estimating Sample Influence

Sample influence metric: $E = |\sum_1^k \hat{x}_k^2 - \sum_1^k x_k^2|$ (energy in low frequencies)

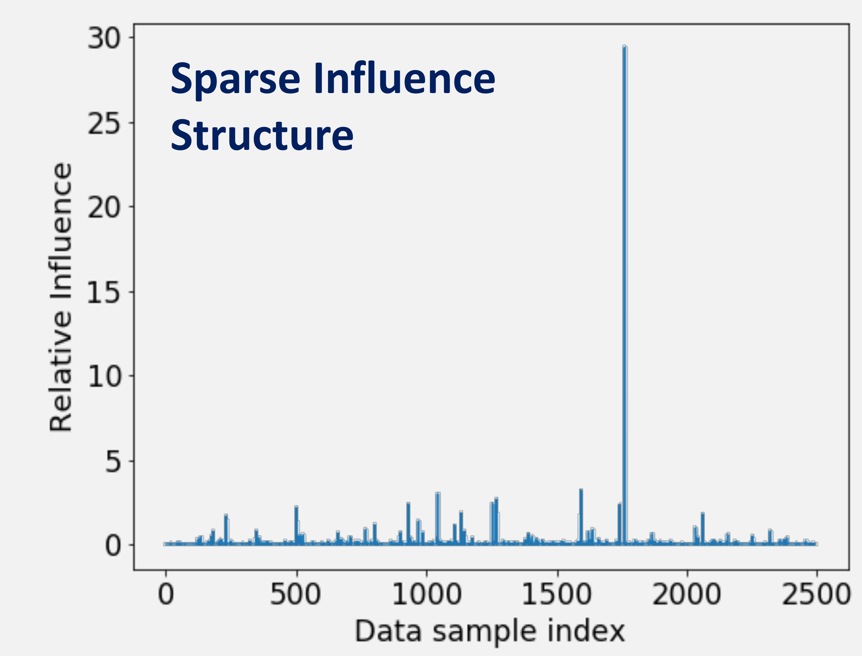

Influence of each data sample in the dataset

Signal is locally altered without changing the graph

Hypothesis: Removal of influential samples during training of a predictive model should affect its fidelity

Mean squared error

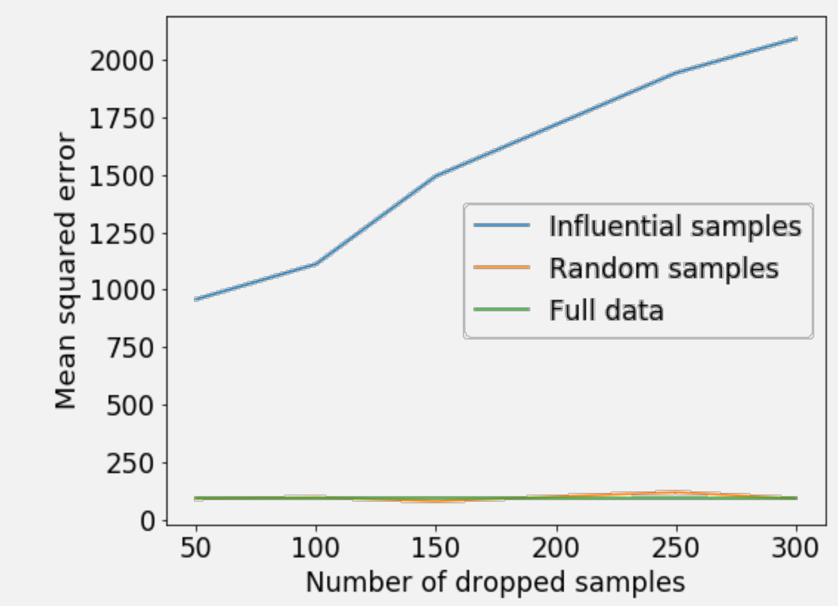

R^2-Statistic

Model – Gradient boosting regression

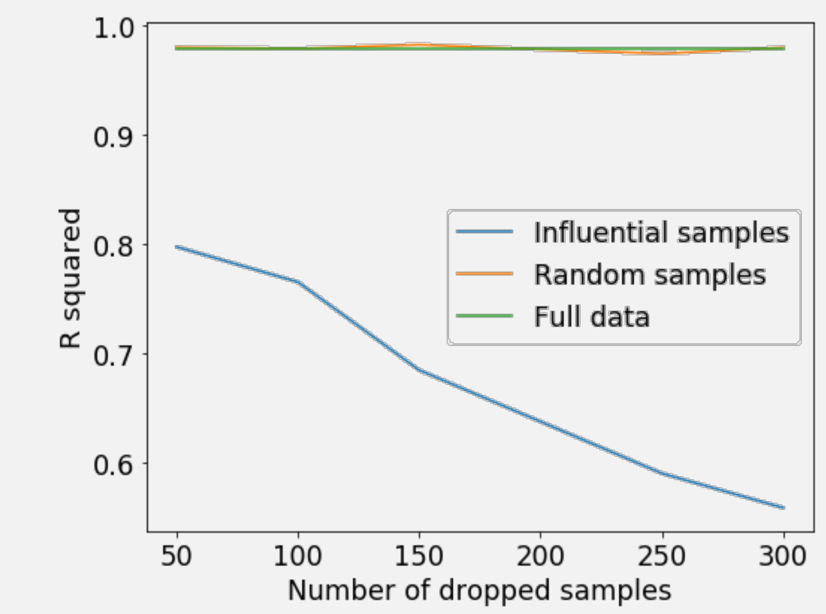

Takeaway: Analysis reveals a highly sparse influence structure and this will directly allow better exploration of high-dimensional parameter spaces to identify sub-domains of interest

References:

[1] A. J. Kunen et al. "KRIPKE - A Massively Parallel Transport Mini-App". American Nuclear Society M&C (2015)

[2] Shuman, David I., et al. "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains". IEEE Signal Processing Magazine 30.3 (2013): 83-98

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

LLNL-POST-735268