

# Novel Representation Learning Technique using Graphs for Performance Analytics

Tarek Ramadan  
Oracle Corporation  
tarek.ramadan@oracle.com

Ankur Lahiry  
Texas State University  
vty8@txstate.edu

Tanzima Z. Islam  
Texas State University  
tanzima@txstate.edu

**Abstract**—The performance analytics domain in High Performance Computing (HPC) uses tabular data to solve regression problems, such as predicting the execution time. Existing Machine Learning (ML) techniques leverage the correlations among features given tabular datasets, not leveraging the relationships between samples directly. Moreover, since high-quality embeddings from raw features improve the fidelity of the downstream predictive models, existing methods rely on extensive feature engineering and pre-processing steps, costing time and manual effort. To fill these two gaps, we propose a novel idea of transforming tabular performance data into graphs to leverage the advancement of Graph Neural Network-based (GNN) techniques in capturing complex relationships between features and samples. In contrast to other ML application domains, such as social networks, the graph is not given; instead, we need to build it. To address this gap, we propose graph-building methods where nodes represent samples, and the edges are automatically inferred iteratively based on the similarity between the features in the samples. We evaluate the effectiveness of the generated embeddings from GNNs based on how well they make even a simple feed-forward neural network perform for regression tasks compared to other state-of-the-art representation learning techniques. Our evaluation demonstrates that even with up to 25% random missing values for each dataset, our method outperforms commonly used graph and Deep Neural Network (DNN)-based approaches and achieves up to 61.67% & 78.56% improvement in MSE loss over the DNN baseline respectively for HPC dataset and Machine Learning Datasets.

**Index Terms**—Graph Neural Network, High Performance Computing, Performance Analytics, Representation Learning

## I. INTRODUCTION

Simulations leverage HPC systems to evaluate numerous what-if scenarios of otherwise experimentally intractable phenomena. HPC configurations (e.g., number of threads, nodes, power cap, thread binding) significantly impact system utilization and execution time [1]–[4]. Estimating how long an application will run based on the configurations *before submitting the job* is a critical problem in HPC since scheduling systems use this estimation as a hard cut-off for aborting the application and scheduling a new job [5], [6]. Once an application gets evicted from an HPC system, the scheduler adds it at the back of the job queue. At a time, thousands of jobs can be waiting in the queue, meaning a user may need to wait for their job to be scheduled again, delaying the overall turnaround time of science. On the other hand, over-predicting the runtime causes the system to be idle once the application finishes leaving billion-dollar systems under-utilized. Hence, over and under

application	algorithm	bw_level	task_count	power_cap	thread_count	runtime
CoMD	pak	1	NaN	64	24	54.4034
CoMD	pak	1	4096	80	24	46.5075
CoMD	pak	1	4096	115	24	42.4216
mg	rand	4	2048	64	16	41.7099
amr	pak	1	2048	NaN	20	47.5113
CoMD	rand	1	NaN	64	24	53.0328
CoMD	NaN	NaN	2048	80	16	38.9581
CoMD	pak	1	2048	80	20	42.6380
CoMD	pak	1	2048	NaN	NaN	44.8226
CoMD	NaN	1	2048	115	NaN	37.1900

Fig. 1: A sample performance dataset. It contains numerical and categorical features in the first six columns and the target value in the last (Runtime). Note that values colored in red (NaN) are missing values. Streaming HPC performance monitoring systems often miss recording feature values or cannot measure the runtime until the application ends, resulting in missing feature values and unlabeled samples.

predictions of the runtime is undesirable.

The performance analytics research area in HPC leverages ML techniques to enable precise runtime prediction. Since the success of downstream ML tasks depends on data representation [7], [8], in this paper, we investigate the problem of extracting the most meaningful information from seemingly unrelated user inputs (i.e., configurations and algorithms) to preserve the correlations between input features and target performance metrics (e.g., runtime). Most HPC performance datasets are organized in tabular format. Figure 1 shows an example of a tabular dataset collected for HPC performance analytics, where the columns contain features including application, algorithm, power\_cap, and the rows contain samples. The target column is denoted as runtime, which describes how long an application takes when being run with a specific algorithm, bandwidth level, task count, power cap, and thread count. While the existing performance analytics research captures user inputs and the respective application performance in tabular format, this format only allows most downstream ML models to exploit relationships across the features within each sample but not across all samples directly.

Finding an effective representation learning technique for tabular data is an active research area. Tree-based models

such as XGBoost [9] have achieved encouraging enhancement in improving the fidelity of regression problems for real-world applications [10]–[12]. **Since missing measurements of hardware features is common in HPC (represented as NaN values for certain features), existing supervised algorithms perform poorly or become inapplicable [13] in predicting the execution time using NaN-filled fetures.** Moreover, performance samples in real-time monitoring systems arrive in a stream, which makes tabular methods rebuild models from scratch every time a new sample appears. To overcome these issues, existing literature has proposed an Attentive Interpretable Tabular Learning (AITL) technique such as the TabNet [14] model. Nonetheless, the TabNet model also uses deep learning techniques to leverage the correlations among features within each sample while leaving relationships among the samples largely unexploited. To overcome this shortcoming, the literature uses extensive feature engineering and pre-processing steps, which come at a high cost and requires a human in the loop.

To address the gaps mentioned above, **we propose to investigate a novel approach of transforming tabular data into a graph data structure to model the correlations among features and samples explicitly.** Specifically, *we hypothesize that using a graph structure to describe a performance dataset (spanning many samples) makes similarities between samples and features explicit, thus constructing the idea of neighborhoods.* The rationale for organizing samples based on similarity is that performance samples with similar feature values, e.g., number of threads = 4 and 6, are likely to result in similar runtimes. In contrast, samples with vastly different features, e.g., number of threads = 4 and 16, are likely to result in widely different runtimes.

However, unlike social networks or disease propagation domains, a performance graph is not given. We overcome this gap by proposing two different approaches for constructing a performance graph called Performance in a Graph (PinG) and evaluating them based on how they impact the effectiveness of the downstream analysis techniques. Another motivation of the PinG formulation is that for streaming performance analytics systems, where labeled data is scarce, and it is too expensive to rebuild models when new samples arrive, new unlabeled samples can be placed in the existing dataset based on their inherent neighborhood similarity. **To our knowledge, the idea of transforming tabular performance data into graphs to improve regression models’ effectiveness has never been explored in the HPC performance analytics domain.**

To summarize the contributions of this paper, we:

- Transform tabular performance data into a graph data structure (PinG) to capture relationships between samples and features.
- Build the PinG structure using two methods (since the performance graph is not given) and explore their effectiveness in improving the performance of downstream regression tasks.
- Develop a novel representation learning technique that can automatically refine the edges between samples

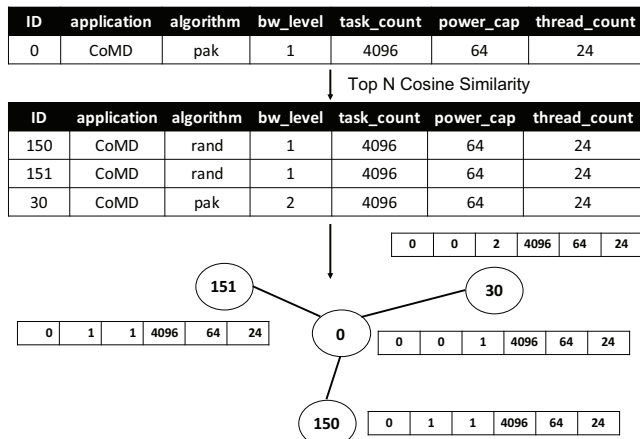


Fig. 2: Illustration of the Single-graph construction approach using  $N$  Cosine similarity. Here,  $N$  is a user-defined parameter that determines the density of the performance graph.

(nodes in PinG) based on feature and sample similarities using self-supervised Graph Neural Networks (GNNs).

- Develop an end-to-end framework that implements the proposed data transformation and representation learning techniques to build effective embeddings for explainable downstream models.

We evaluate the effectiveness of the generated embeddings from GNNs based on how well they make even a simple feed-forward neural network perform for regression tasks compared to other state-of-the-art representation learning techniques. **Since the characteristics of problems in HPC are unique, e.g., regression tasks, missing labels, missing measurements, streaming data, scarcity of training samples, this paper focuses solely on the applicability of our proposed approach in the context of the HPC domain.** Thus, we use ten HPC performance datasets collected on three supercomputing facilities. To study the applicability of our proposed design on ML datasets, we also present evaluation using three ML benchmarks. Our evaluation demonstrates that even with up to 25% random missing values for each dataset, our method outperforms commonly used graph and DNN-based approaches. Specifically, our method achieves up to 61.67% & 78.56% improvement in Mean Square Error (MSE) loss over DNN for the HPC and ML datasets, respectively.

## II. DESIGN

We hypothesize that holistically capturing the relationships across samples and parameters enabled by the graph formulation of tabular data can improve the effectiveness of the downstream ML tasks.

### A. Rationale for PinG

The rationale for organizing data as a graph compared to the state-of-the-art dictionary-based approach is that a graph structure can inherently describe relationships among features and samples, allowing downstream ML tasks to capture information from relevant neighbors by exploiting these




Fig. 3: Outcome of the batched-graph construction approach.

interrelations. However, unlike other research domains where the graph structure is explicitly provided, the graph from HPC data needs to be constructed. While data samples can directly map to nodes, edges between samples need to be defined explicitly or inferred automatically. On the other hand, the dictionary-based approach only leverages feature extraction to generate the representation vector, whereas the graph structure leverages both feature extraction and the relationship between the samples in generating the representation vector.

The graph design adopts the architecture of an undirected weighted graph  $G = (V, E, A)$  to represent a performance dataset, where each node  $V_i \in V$  denotes a sample, and  $edge \in E$  indicates a relationship with another measurement.  $A \in \mathbb{R}^{N \times N}$  is an adjacency matrix that specifies the weights on the edges, where  $A_{i,j}$  corresponds to the edge weight between nodes  $V_i$  and  $V_j$  that can be calculated based on their similarities. For data that belongs to the Euclidean space,  $A_{i,j}$  can be computed using a distance metric between  $V_i$  and  $V_j$ . However, for non-Euclidean or dense data such as a tree, building  $A_{i,j}$  using Euclidean distance is not meaningful. Hence, instead of using explicit distance measures to construct the final graph, we propose to *first*, construct a fully connected initial graph using measurements as node features and explicit distance measures such as cosine similarity to calculate the initial edge weights, and *then*, iteratively refine the node and edge features through an automated graph edge inference method using self-supervised learning.

### B. Initial PinG Construction

This section presents two approaches for generating the initial PinG graph from the tabular data.

**Single-Graph Construction Approach (SGC):** In this approach, we propose to build a single graph for the whole dataset where each sample represents a node. The initial edges in PinG are computed using the Cosine similarity algorithm by measuring the cosine angle between the current sample and all other samples. A high cosine similarity indicates that the angle between the two samples is small; therefore, the two samples are similar to each other. Using this technique, we calculate

the cosine similarity for each sample with all different samples in the dataset and then choose the top  $N$  similar samples. Here,  $N$  is a hyper-parameter that users can adjust for better model performance or specify based on memory constraints. Figure 2 demonstrates the cosine-similarity-based neighbor selection for a given node where  $N = 3$ . This step can be parallelized using a data-parallel ML technique that we will pursue in the future. After building the edges between similar samples, the edge weights are initialized to 1 and passed to a GNN model using self-supervised learning. The model assigns higher edge weight to similar data and lower edge weight to dissimilar data, thus explicitly encoding the similarity between samples and their features. Algorithm 1 shows the algorithmic steps of this process.

---

#### Algorithm 1 Single Graph Construction Approach

---

- 1: **Input:**  $x$  is a tabular dataset.,  $n$  minimum number of neighbours
  - 2: **Output:**  $graph$ , a NetworkX graph generated from the tabular dataset
  - 3:  $distance \leftarrow cosine\_distance(dataset)$
  - 4: declare empty arrays of  $sources$  and  $destinations$
  - 5: **for all**  $index, distance \in distances$  **do**
  - 6:    $source \leftarrow [1, N + 1]$
  - 7:   **append**  $source$  in  $sources$
  - 8:    $similars \leftarrow sort(distances)$
  - 9:    $destination \leftarrow similars.top(N + 1)$
  - 10:   **append**  $destination$  in  $destinations$
  - 11: **end for**
  - 12:  $edges \leftarrow zipped(sources, destinations)$
  - 13:  $graph \leftarrow$  an empty NetworkX graph
  - 14: create  $[1..N+1]$  nodes in the  $graph$
  - 15: add edges from the  $edges$  list
  - 16: add node features and edge features in the  $graph$
  - 17: remove self-loop from the  $graph$
  - 18: **return**  $graph$
- 

**Batched-Graph Construction Approach (BGC):** In this approach, we propose to generate clusters of samples using unsupervised learning to find groups within the data based on similarity to build  $M$  graphs, where  $M$  is the number of clusters. This approach is memory efficient since it uses batches of samples instead of all samples for building each sub-graph centered around each sample and leverage the averaging effect. The BGC method selects different cluster sizes to create different-sized batches, as illustrated in Figure 3, thus the cluster size is a hyperparameter to our algorithm. To further optimize this approach, we can reduce the number of edges by connecting the top  $N$  similar neighbors to each node, as explained in the first approach. We demonstrate the process in Algorithm 2. In this paper, we choose Hierarchical Density-based Spatial Clustering of Applications with Noise (HDBscan) [15] due to its parameter flexibility. HDBscan encodes the input samples into transformed space according to their density, then builds a minimum spanning tree (MST)

to find the cluster hierarchy, and finally extracts the generated cluster.

---

**Algorithm 2** Batched-Graph Construction Approach

---

```

1: Input:  $x$  is a tabular dataset.,  $N$  minimum number of
   neighbours
2: Output:  $graphs$ , batched generation of graphs
3:  $graphs \leftarrow$  empty array
4:  $clusters \leftarrow$  a group of clusters from the sample
5: for all  $group \in clusters$  do
6:    $distances \leftarrow cosine\_distance(group)$ 
7:   declare empty arrays of  $sources$  and  $destinations$ 
8:   for all  $index, distance \in distances$  do
9:      $source \leftarrow [1, N + 1]$ 
10:    append  $source$  in  $sources$ 
11:     $similars \leftarrow sort(distances)$ 
12:     $destination \leftarrow similars.top(N + 1)$ 
13:    append  $destination$  in  $destinations$ 
14:     $edges \leftarrow zipped(sources, destinations)$ 
15:     $graph \leftarrow$  an empty NetworkX graph
16:    create  $[1...N+1]$  nodes in the  $graph$ 
17:    add edges from the  $edges$  list
18:    add node features and edge features in the  $graph$ 
19:    remove self-loop from the graph
20:    append  $graph$  in  $graphs$ 
21:  end for
22: end for
23: return  $graphs$ 

```

---

### C. Self-Supervised Learning

Since a large number of streaming performance samples in HPC are unlabeled, supervised learning is not ideal. In contrast, Self-supervised learning (SSL) algorithms can recognize and understand patterns from unlabeled samples. The main idea of SSL is to automatically generate a representation of the input data and refine it using feedback from the model, showing promising results in Natural Language Processing (NLP) [16], [17]. Hence, in this paper, we propose leveraging self-supervised learning to develop an automated edge-inference-based method that optimizes the PinG edge weights.

### D. Putting It Altogether for Representation Learning

Graph Neural Network (GNN)-based representation learning techniques [18] can represent nodes, edges, and sub-graphs in low-dimensional vectors. These techniques capture complex and non-linear relationships among samples and features across the whole dataset to build the embeddings. Figure 4 presents our proposed end-to-end representation learning pipeline. After the initial graph is constructed using one of the SGC or BGC methods presented in Section II-B, the initial graph is passed to a convolution layer for creating self-supervision and then passed to a graph-based learning model for building embeddings. The embeddings are then extracted from the second-to-last layer and passed to a feed-forward neural network to calculate MSE. The error then propagates

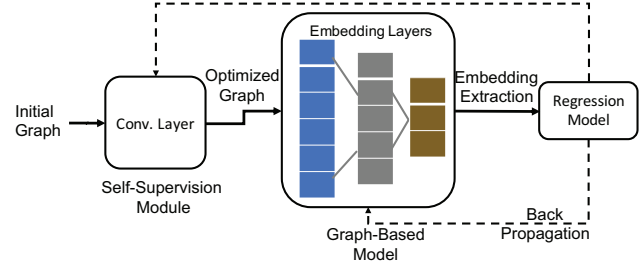


Fig. 4: The end-to-end pipeline of learning the optimized graph from initially constructed PinG formulation. Input graph and its edge weights are fed to a GNN model with self-supervision for constructing an embedding for the entire dataset. Here, solid lines represent forward propagation, and dotted lines represent backpropagation to enhance the edge weights.

back to the self-supervision and graph-based models to provide feedback and fine-tune.

Specifically, in this work, we propose to use GNN [19] to learn the embeddings from the constructed graphs. The main advantage of the GNN architecture is that the learning equation leverages edges connected between a node to its neighbors to compute that node’s vector representation. This operation is known as message passing. One of the most recent GNN architectures is GraphSAGE [20], which learns each node’s vector representation inductively by allowing complex aggregation functions for message passing. Thus, message passing makes GNN generalize to unseen samples during the deployment of the model by representing the unseen sample to the neighboring nodes. The GraphSAGE model generates a vector representation for each node and then applies semi-supervised node regression to produce a regression value for each node. We refer to the representation learning pipeline that uses SGC as the input as Self-supervised Graph Neural Network (SSGNN), and that using BGC as Self-supervised Batched Graph Neural Network (SSBGNN). In Section V, we compare the performance of using GraphSAGE as the main driver of SSGNN and SSBGNN with that of other methods such as Graph Convolutional Network (GCN) and Graph Attention Layer (GAT). **Our extensive experiments show that the choice of the driver method can impact accuracy of our proposed method and is dataset dependent.**

## III. IMPLEMENTATION

We implement the proposed representation learning pipeline using the Networkx (NX) [21] and Deep Graph Library (DGL) [22] in Python with the PyTorch back end.

### A. Pre-processing

Since the core of our approach is to find a better representation of the input data automatically, we used minimal processing, precisely two pre-processing techniques: standardization and imputation of missing values.

**Standardization:** To standardize the input data into a common scale, our pipeline transforms the data using the standard scaler algorithm provided by sklearn [23]. The standard

scaler algorithm removes the mean and scales the data based on the unit variance so that the scaling of each feature is independent of the other features; however, eventually, they are all scaled in the same way and have equal weights.

**Imputation:** To address the missing values in the input data, our pipeline uses the `Simpler Imputer` function provided by `sklearn`. The `Simpler Imputer` function finds the missing values and replaces all missing values in the column based on a selected strategy. We used the feature column’s mean to replace any missing values.

### B. Classifier Optimization

**Drop-out Layer:** We randomly drop some of the edges to the subsequent layers’ neurons during the training phase. The drop-out probability is a hyperparameter. This technique is only applied to the input and hidden layers of the model during training and does not affect the testing set.

**Optimizer:** The loss gets calculated based on the loss function in every epoch, and the optimizer tweaks the model parameters to minimize the loss. This work uses the Adam optimizer [24].

**Activation Function:** The pipeline uses an activation function in each neuron node in the neural network. Based on multiple experiments, we choose Rectified Linear Unit (ReLU) as the activation function.

**Removing self loops:** A node’s closest neighbor when computing cosine distance is the node itself, which introduces self-loops in the graph. So, while choosing the top  $N$  neighbors, we consider the top  $N + 1$  neighbors first instead (line 12 in Algorithm 2 & line 9 in Algorithm 1) and then remove those self-loops from the induced graph.

## IV. EXPERIMENTAL SETUP

### A. System

We run our experiments in parallel using the following systems: Penguin Computing On Demand (POD), Google Colab, Texas Advanced Computing Center (TACC). The POD system consists of 96 AMD EPYC CPUs. Each CPU has 8GB of RAM based on the 64-bit x86 architecture. We also run our evaluations on Google Colab using a machine with Intel(R) Xeon(R) CPU with 26.75GB RAM and one NVIDIA Tesla P100 GPU with 16GB memory. Since this work involves running various models simultaneously, we leverage supercomputers at the TACC center to run them in parallel.

### B. Dataset Description

To thoroughly evaluate the benefit of our proposed graph construction and representation learning technique, we use 10 HPC applications and three ML benchmarks in this work. The Vulcan and Cab datasets contain performance data collected from the XSBench [33] and the OpenMC [34] applications on an IBM BlueGene Q and an Intel Sandy Bridge systems, respectively. The Catalyst dataset combines the performance samples from MiniAMR [26], CG [27], FT [27], LU [27], MG [27], CoMD [28], Kripke [29] miniapps. Table I summarizes the size, the number of features, and the target variables for each dataset.

TABLE I: Regression dataset descriptions. “#S”, “#F”, and “Std.Dev.” refer to the number of samples, number of features, and standard deviation in the target values of the corresponding dataset.

Dataset	#S	#F	Domain	Target	Mean	Std.Dev
Catalyst [2]	3992	10	HPC	Runtime	36.88	17.38
Cab [25]	320	148	HPC	Runtime	21.04	15.65
Vulcan [25]	321	172	HPC	Runtime	23.51	29.43
MiniAMR [26]	301	9	HPC	Runtime	34.68	18.06
CG [27]	961	9	HPC	Runtime	31.48	13.75
FT [27]	554	9	HPC	Runtime	34.15	16.90
LU [27]	554	9	HPC	Runtime	43.10	22.09
MG [27]	513	9	HPC	Runtime	48.65	23.98
CoMD [28]	553	9	HPC	Runtime	41.61	4.71
Kripke [29]	554	9	HPC	Runtime	28.29	4.49
Airfoil Self-Noise [30]	1503	6	ML	SSPL	124.84	6.90
IoT Telemetry [31]	1000	6	ML	Temp	22.36	2.56
Beijing AQI [32]	1000	15	ML	PT08.S4 (NO2)	9.76	43.69

### C. Baselines

In this paper, we compare our proposed GNN-based method to two graph-based representation learning techniques—GCN [35] and GAT [36], and four non-graph-based methods commonly used for tabular data—DNN [37], XGBoost [9], LightGBM [38], and Denoising Autoencoder (DAE) [39]. The graph-based methods use the same graph construction approach as the SSGNN approach. Table II represents the hyperparameters for each model and the ranges of their values used during hyperparameter tuning. We leverage the Optuna framework [40] to find the best hyperparameters for each model automatically.

### D. Performance Metrics

Since all datasets in HPC performance analytics contain continuous target values, we evaluate the performance of regression tasks by comparing the predicted runtime with the actual one. We use the MSE loss that calculates the average squares of the error between the predicted and the actual labels. Using different baseline and graph-based methods, we generate embeddings and evaluate their effectiveness by assessing how well a simple linear regression model performs when fed with the embeddings. Since this experiment aims to attribute any benefit to having superior embeddings, we keep the prediction task’s model simple.

## V. RESULTS

In all experiments, we use 5-fold cross-validation by dividing the data into five train/test splits, generating five different results, and calculating the average. **We ensure that the**

TABLE II: Hyper-parameters for different models.

Model	Hyper-parameters	Search range
DNN, XGBoost, LightGBM, DAE, GAT, GCN, SSGNN, and SSBGNN	Hidden Layer Dimension	[25, 600]
	Learning Rate	$[1e^{-5}, 1e^{-1}]$
	Dropout	[0, 1]
	Optimizers	[Adam, Adagrad, Adadelta, SGD, RMSprop]
	Activation function	[relu, elu, leaky_relu]
GAT, GCN, SSGNN, and SSBGNN	Hidden Layer Aggregation	[pool, mean, gcn]
	Self-Supervision Conv. Layer activation function	[relu, elu, leaky_relu, sigmoid, tanh]
	Number of Attention Heads (GAT model only)	[1, 2]

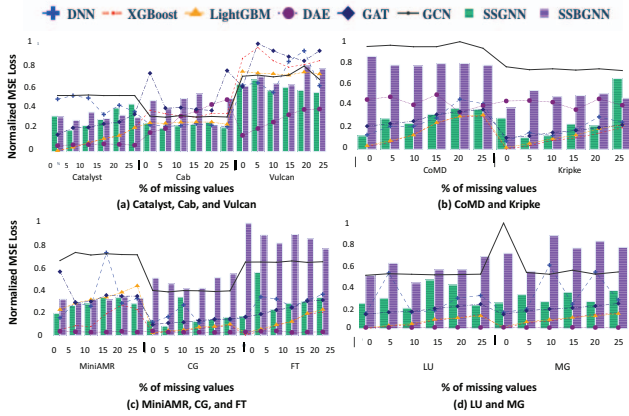


Fig. 5: Normalized MSE Loss for HPC dataset. The X-axis shows the % of missing feature values and the Y-axis shows the normalized MSE value for different methods.

**train and test sets are disjoint.** Table I shows all thirteen datasets from HPC and ML domains and Table II shows the hyperparameters for the models used.

#### A. Impact of Graph Representation

The experiment aims to test our initial hypothesis that using a graph structure to describe a performance dataset improves a predictive regression model’s effectiveness. Figure 5 summarizes the performance of HPC datasets.

From Figures 5 and 6, we can observe that: (1) The graph-based representation learning approaches outperform DNN. Although, DNN models can implicitly leverage higher-order correlations through connections across multiple layers; our experiments demonstrate that *explicitly modeling* the relations between data samples in PinG improves the effectiveness of a predictive regression model. (2) In Figures 5 and 6, we observe that one of our proposed methods (SSGNN or SSBGNN) outperforms GCN in all experiments across all datasets. This is because, our proposed method (GraphSAGE with self-supervision) captures information across the entire tabular dataset by allowing messages to pass across neighbors. In contrast, GCN aggregates information across neighbors of a fixed

TABLE III: Summary of improvements using our proposed methods compared to the baselines across datasets.

Type	# of Exps.	Methods	Best % Improvement	#Exps. where any of Our Methods Outperform
HPC	60	DNN	61.67	32
		XGBoost	38.19	12
		LightGBM	36.03	18
		DAE	76.04	14
		GAT	70.55	26
		GCN	86.96	60
ML	18	DNN	78.56	8
		XGBoost	56.80	6
		LightGBM	41.25	6
		DAE	20.38	2
		GAT	74.88	10
		GCN	99.99	18




Fig. 6: Normalized MSE Loss for ML datasets. The X-axis represents the % of missing feature values and the Y-axis shows the normalized MSE for different methods (the lower, the better).

distance, thus failing to capture long-distance relationships among the samples. (3) Our proposed method performs on par or better than GAT’s attention-based aggregation strategy in 36 out of 78 experiments. (4) Figure 5 shows that the graph-based methods outperform XGBoost and LightGBM where the target values have a wide range of distribution, such as in the case of the Vulcan dataset compared to the Catalyst dataset, as shown in Table I’s mean and standard deviation columns. Figure 6a shows that the SSGNN method outperforms the gradient- and neural-network-based methods in the IoT Telemetry dataset.

#### B. Impact of Missing Data

This experiment aims to evaluate the impact of the PinG formulation on the performance of regression tasks when missing data is present in the features. The rationale for conducting this experiment is that streaming HPC performance monitoring systems often miss recording feature values or cannot measure target values (unlabeled) until the application ends. To study the impact of missing data in the features on the effectiveness of graph-based representation learning techniques, we randomly inject missing values (NaN) in 5%, 10%, 15%, 20%, and 25% of the samples and compare the generated embeddings from all the methods listed in Table III. The GCN and GAT methods use the graph constructed using our proposed single graph construction method (Section II-B).

From Figures 5 and 6, we observe that although the decision-tree and gradient-based methods outperform graph-based methods, our method of using GNN with self-supervision improves the efficacy of neural network-based methods in building effective representations, which enhances

the performance of downstream regressions.

### C. Discussion

Table III summarizes the best %-improvement in MSE loss of our proposed methods (either SSGNN or SSBGNN) compared to all other baseline methods. The rationale for reporting the best of these two methods is that any of these methods being better supports our hypothesis that transforming performance data into a graph improves the performance of downstream analyses. The table also reports how frequently our proposed methods perform better than the baselines. From Table III, we can observe that even though GAT and GCN work on the transformed graph, our proposed embedding learning algorithm outperforms GCN and GAT in 100% and > 43% of the experiments, respectively, for the HPC datasets. Our proposed method outperforms DNN more than 50% of the times, and achieves 61.67% improvement in MSE across all HPC datasets. It is important to note that even though the tabular methods perform better for most of the HPC datasets, these methods require model rebuilding when new unlabeled performance samples arrive in a streaming performance monitoring system, which makes tabular methods computationally infeasible. On the other hand, finding a neighborhood for new samples based on similarity in a graph enables quick approximation of their labels, making a graph-based performance representation suitable for streaming scenarios.

Our experimental results show that the embedding learning pipeline with single-graph construction approach, SSGNN, performs better than the batched-graph construction approach in nearly all cases, making SSGNN a superior graph construction method for the PinG formulation.

## VI. RELATED WORK

### A. Deep Modeling Explainability using Graphs

Researchers have transformed tabular data into graphs to produce an explainable model called TableGraphNet [41]. TableGraphNet architecture builds one or more graphs for each data sample. The graph node comprises the feature attributes, and the graph edges contain the distance between the feature attributes. Using the generated graphs for the input samples, they extract node and attribute-centric features for every attribute. Their evaluations demonstrate that TableGraphNet architecture performs similarly to a regular Deep Neural Network model based on three classifications and eight regression datasets, whereas our architecture demonstrates significant improvement compared to DNN.

### B. Tabular Data Prediction using Multiplex Graphs

TabGNN [42] framework takes tabular data as input, constructs a directed multiplex graph based on the table columns referred to as features, and encodes each sample for an initial node embedding. It uses a graph neural network to produce latent feature embedding for each data sample. After getting the final representation embedding, TabGNN uses Auto Feature Engine (AutoFE) to choose significant features and then uses a Multi-Layer Perceptron to get the final prediction. Researchers

observe that the framework heavily relies on AutoFE and DeepFM [43], two popular DNN-based feature engineering methods for tabular data. Their experiments show that utilizing AutoFE+TabGNN and DeepFM+TabGNN outperform using AutoFE or TabGNN alone, but requires longer time to train. In contrast, our method leverages a similar graph-building method without the expensive DNN-based feature engineering step, thus requiring less training and testing time. As TabGNN uses DNN based approaches and our methods outperform DNN based approaches in most of the settings, we anticipate that our methods will either perform on-par or better than TabGNN.

### C. Knowledge Graph for Efficient Meta-learning

Recent work such as Automated Relational Meta-learning (ARML) [44] shows that graph structures improve the performance of Meta-Learning [45]. On a 2D toy application with random numbers, ARML leverages the meta-knowledge graph to obtain a more fine-grained structure than other gradient-based meta-learning implementations. The single limitation of this implementation is that it can only be used with meta-learning algorithms and has been evaluated on image classification datasets. The ARML architecture uses DNN to transform the input data into a feature representation. This architecture complements our approach, where our proposed method can convert tabular data to a graph formulation and then leverage ARML to produce a fine-grained representation.

## VII. CONCLUSIONS

This paper proposes a new data transformation approach that converts tabular data into graphs and develops a new representation learning pipeline using a self-supervision-based automated edge-inference-based technique. Graph representation learning leverages relationships between samples and features to create fine-grained embeddings, effectively improving the model's performance and outperforming the standard deep learning-based techniques. We evaluate our approach using different models on multiple HPC and ML datasets and find that relationships captured between samples and features in a graph formulation make regression tasks robust against missing values and perform better.

## VIII. ACKNOWLEDGEMENT

The work was supported in part by the grant DE-SC0022843 funded by the U.S. Department of Energy, Office of Science. The work was also supported in part by the Exascale Computing Project (ECP) funded by the U.S. Department of Energy, Office of Science. A part of the research was also supported by the Research Enhancement Program at Texas State University.

## REFERENCES

- [1] G. Panwar, D. Zhang, Y. Pang, M. Dahshan, N. DeBardleben, B. Rattindran, and X. Jian, "Quantifying memory underutilization in hpc systems and using it to improve performance via architecture support," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO '52. New York, NY, USA: Association for Computing Machinery, 2019, p. 821–835. [Online]. Available: <https://doi.org/10.1145/3352460.3358267>

- [2] T. Patki, J. J. Thiagarajan, A. Ayala, and T. Z. Islam, "Performance optimality or reproducibility: that is the question," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–30.
- [3] T. Z. Islam, J. J. Thiagarajan, A. Bhatele, M. Schulz, and T. Gamblin, "A machine learning framework for performance coverage analysis of proxy applications," in *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2016, pp. 538–549.
- [4] D. Beckingsale, O. Pearce, I. Laguna, and T. Gamblin, "Apollo: Reusable models for fast, dynamic tuning of input-dependent code," in *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2017, pp. 307–316.
- [5] W. Tang, N. Desai, D. Buettner, and Z. Lan, "Analyzing and Adjusting User Runtime Estimates to Improve Job Scheduling on the Blue Gene/P," in *Parallel Distributed Processing (IPDPS), 2010 IEEE International Symposium on*, 2010, pp. 1–11.
- [6] Y. Fan, P. Rich, W. E. Allcock, M. E. Papka, and Z. Lan, "Trade-off between prediction accuracy and underestimation rate in job runtime estimates," in *2017 IEEE International Conference on Cluster Computing (CLUSTER)*, Sep. 2017, pp. 530–540.
- [7] R. Jonschkowski and O. Brock, "State representation learning in robotics: Using prior knowledge about physical interaction," in *Robotics: Science and Systems*, 2014.
- [8] B. Yoshua, A. C. Courville, and P. Vincent, "Unsupervised feature learning and deep learning: A review and new perspectives," *CoRR*, vol. abs/1206.5538, 2012. [Online]. Available: <http://arxiv.org/abs/1206.5538>
- [9] T. Chen, T. He, M. Benesty, V. Khotilovich, Y. Tang, H. Cho, K. Chen *et al.*, "Xgboost: extreme gradient boosting," *R package version 0.4-2*, vol. 1, no. 4, pp. 1–4, 2015.
- [10] Z. Ding, H. Nguyen, X.-N. Bui, J. Zhou, and H. Moayed, "Computational intelligence model for estimating intensity of blast-induced ground vibration in a mine based on imperialist competitive and extreme gradient boosting algorithms," *Natural Resources Research*, vol. 29, no. 2, pp. 751–769, 2020.
- [11] A. I. A. Osman, A. N. Ahmed, M. F. Chow, Y. F. Huang, and A. El-Shafie, "Extreme gradient boosting (xgboost) model to predict the groundwater levels in selangor malaysia," *Ain Shams Engineering Journal*, vol. 12, no. 2, pp. 1545–1556, 2021.
- [12] S. K. Kiangala and Z. Wang, "An effective adaptive customization framework for small manufacturing plants using extreme gradient boosting-xgboost and random forest ensemble learning algorithms in an industry 4.0 environment," *Machine Learning with Applications*, vol. 4, p. 100024, 2021.
- [13] T. Emmanuel, T. Maupong, D. Mpoeleng, T. Semong, B. Mphago, and O. Tabona, "A survey on missing data in machine learning," *Journal of Big Data*, vol. 8, no. 1, pp. 1–37, 2021.
- [14] S. O. Arık and T. Pfister, "Tabnet: Attentive interpretable tabular learning," in *AAAI*, vol. 35, 2021, pp. 6679–6687.
- [15] L. McInnes, J. Healy, and S. Astels, "hdbscan: Hierarchical density based clustering," *J. Open Source Softw.*, vol. 2, no. 11, p. 205, 2017.
- [16] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [17] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [18] M. Wu, S. Pan, L. Du, and X. Zhu, "Learning graph neural networks with positive and unlabeled nodes," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 15, no. 6, pp. 1–25, 2021.
- [19] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE transactions on neural networks*, vol. 20, no. 1, pp. 61–80, 2008.
- [20] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," *Advances in neural information processing systems*, vol. 30, 2017.
- [21] A. Hagberg, P. Swart, and D. S. Chult, "Exploring network structure, dynamics, and function using networkx," Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 2008.
- [22] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai *et al.*, "Deep graph library: A graph-centric, highly-performant package for graph neural networks," *arXiv preprint arXiv:1909.01315*, 2019.
- [23] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in python," *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [25] Tanzima, "tzislam/2019-protocols-data: Data for protocols," Sep. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.3403038>
- [26] C. T. Vaughan and R. F. Barrett, "Enabling tractable exploration of the performance of adaptive mesh refinement," in *2015 IEEE International Conference on Cluster Computing*, Sep. 2015, pp. 746–752.
- [27] "The NAS Parallel Benchmarks." [Online]. Available: <http://lscience.nas.nasa.gov/Software/NPB/>
- [28] O. Pearce, H. Ahmed, R. W. Larsen, P. Pirkelbauer, and D. F. Richards, "Exploring dynamic load imbalance solutions with the comd proxy application," *Future Generation Computer Systems*, 2017. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167739X17300560>
- [29] A. J. Kunen, T. S. Bailey, and P. N. Brown, "KRIPKE - A Massively Parallel Transport Mini-App," in *American Nuclear Society M&C 2015*.
- [30] P. D. Brooks, Thomas and M. Marcolini, "Airfoil Self-Noise." UCI Machine Learning Repository, 2014, DOI: <https://doi.org/10.24432/C5VW2C>.
- [31] "Environmental Sensor Telemetry Data." [Online]. Available: <https://www.kaggle.com/datasets/garystafford/environmental-sensor-data-132k>
- [32] S. Chen, "Beijing Multi-Site Air-Quality Data." UCI Machine Learning Repository, 2019, DOI: <https://doi.org/10.24432/C5RK5G>.
- [33] J. Tramm, A. Siegel, T. Islam, and M. Schulz, "XSBench—the Development and Verification of a Performance Abstraction for Monte Carlo Reactor Analysis," *The Role of Reactor Physics toward a Sustainable Future (PHYSOR)*, 2014.
- [34] P. K. Romano and B. Forget, "The openmc monte carlo particle transport code," *Annals of Nuclear Energy*, vol. 51, pp. 274–281, 2013.
- [35] S. Haykin and N. Network, "A comprehensive foundation," *Neural networks*, vol. 2, no. 2004, p. 41, 2004.
- [36] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," *arXiv preprint arXiv:1710.10903*, 2017.
- [37] Schmidhuber, Jürgen, "Deep learning in neural networks: An overview," *Neural networks*, vol. 61, pp. 85–117, 2015.
- [38] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "Lightgbm: A highly efficient gradient boosting decision tree," *Advances in neural information processing systems*, vol. 30, 2017.
- [39] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, P.-A. Manzagol, and L. Bottou, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion." *Journal of machine learning research*, vol. 11, no. 12, 2010.
- [40] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A next-generation hyperparameter optimization framework," in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 2623–2631.
- [41] G. Terejanu, J. Chowdhury, R. Rashid, and A. Chowdhury, "Explainable deep modeling of tabular data using tablegraphnet," *arXiv preprint arXiv:2002.05205*, 2020.
- [42] X. Guo, Y. Quan, H. Zhao, Q. Yao, Y. Li, and W. Tu, "Tabgmn: Multiplex graph neural network for tabular data prediction," *arXiv preprint arXiv:2108.09127*, 2021.
- [43] H. Guo, R. Tang, Y. Ye, Z. Li, and X. He, "Deepfm: a factorization-machine based neural network for ctr prediction," *arXiv preprint arXiv:1703.04247*, 2017.
- [44] H. Yao, X. Wu, Z. Tao, Y. Li, B. Ding, R. Li, and Z. Li, "Automated relational meta-learning," *arXiv preprint arXiv:2001.00745*, 2020.
- [45] R. Vilalta and Y. Drissi, "A perspective view and survey of meta-learning," *Artificial intelligence review*, vol. 18, no. 2, pp. 77–95, 2002.