

LatentTune: Efficient Tuning of High Dimensional Database Parameters via Latent Representation Learning*

Sein Kwon	Youngwan Jo	Seungyeon Choi	Jieun Lee	Huijun Jin	Sanghyun Park
Computer Science	Computer Science	Computer Science	Computer Science	Computer Science	Computer Science
Yonsei University	Yonsei University	Yonsei University	Yonsei University	Yonsei University	Yonsei University
Seoul, Korea	Seoul, Korea	Seoul, Korea	Seoul, Korea	Seoul, Korea	Seoul, Korea
seinkwon97@yonsei.ac.kr	jyy1551@yonsei.ac.kr	tmddus1553@yonsei.ac.kr	jieun199624@yonsei.ac.kr	jinjuijun@yonsei.ac.kr	sanghyun@yonsei.ac.kr

Abstract—As data volumes continue to grow, optimizing database performance has become increasingly critical, making the implementation of effective tuning methods essential. Among various approaches, database parameter tuning has proven to be a highly effective means of enhancing performance. Recent studies have shown that machine learning techniques can successfully optimize database parameters, leading to significant performance improvements. However, existing methods still face several limitations. First, they require substantial time to generate large training datasets. Second, to cope with the challenges of high-dimensional optimization, they typically optimize only a subset of parameters rather than the full configuration space. Third, they often rely on information from similar workloads instead of directly leveraging information from the target workload. To address these limitations, we propose LatentTune, a novel approach that differs fundamentally from traditional methods. To reduce the time required for data generation, LatentTune incorporates a data augmentation strategy. Furthermore, it constructs a latent space that compresses information from all database parameters, enabling the optimization of the full configuration space. In addition, LatentTune integrates external metric information into the latent space, allowing for precise tuning tailored to the actual target workload. Experimental results demonstrate that LatentTune outperforms baseline models across four workloads on MySQL and RocksDB, achieving up to 1332% improvement for RocksDB and 11.82% throughput gain with 46.01% latency reduction for MySQL.

Index Terms—Database parameter tuning, Bayesian Optimization, Database Management Systems, Machine Learning

I. INTRODUCTION

Recently, the volume of data has grown significantly, and the necessity for high-performance database management systems (DBMS) to process large amounts of data efficiently has increased [1]. One of the approaches to improving DBMS performance is database parameter tuning [2], which refers to the process of improving the performance of the database by adjusting the values of the parameters. Conventional database parameter tuning requires Database Administrators (DBAs) to modify parameter values; however it is difficult to tune parameters efficiently due to the complex relationship between numerous parameters. Furthermore, different databases have distinct parameters, and DBAs must account for these varia-

tions along with the diversity of workloads. To overcome these limitations, recent research has been conducted on automatic database parameter tuning using machine learning (ML) techniques, achieving improved performance in various workloads and DBMSs [3]–[8]. These studies have employed various ML-based optimization algorithms, including Bayesian Optimization (BO) [9] and reinforcement learning (RL) [10] to guide the tuning process more effectively.

Limitations of Existing Methods. Machine learning models require large amounts of training data consisting of configurations and their corresponding performance metrics for effective training [11]. In conventional approaches, such datasets are typically generated by sampling configuration values and evaluating their performance using benchmarking tools [12]–[14].

However, this process is time-consuming, making it difficult to collect data efficiently. For example, when we limited the benchmarking duration to 300 seconds per configuration in MySQL [15], generating a dataset of 1,000 configurations required substantial amounts of time, highlighting practical constraints in creating sufficiently large training datasets.

Another challenge arises from the parameter selection process used to address high-dimensional optimization. To reduce the search space, existing BO-based methods [3]–[6] commonly adopt selection algorithms such as SHapley Additive exPlanations (SHAP) [16] and LASSO [17] to identify and optimize only the top- k parameters that most significantly affect performance. RL-based approaches [7], [8], on the other hand, utilize reinforcement learning algorithms such as Deep Deterministic Policy Gradient (DDPG) [18] to cope with high-dimensional continuous spaces by learning effective representations of actions and states.

However, the selected parameters can vary depending on the algorithm, even under the same workload. As shown in Table I, only three to five parameters overlap between the SHAP and LASSO algorithms for RocksDB [19], and the ranking order also differs. Figure 1 presents the performance impact of tuning with RGPE [20], one of our baseline models, using the top 5 and top 10 parameters. The results indicate that variations

TABLE I: Top-10 parameters using SHAP and LASSO for RocksDB R90W10 and R50W50 workloads. Overlapped parameters are highlighted in bold.

R90W10		R50W50	
SHAP (top-10)	Lasso (top-10)	SHAP (top-10)	Lasso (top-10)
target_file_size_base	compression_type	compression_ratio	compression_ratio
write_buffer_size	target_file_size_multiplier	write_buffer_size	max_background_flushes
level0_file_num_compaction_trigger	num_levels	max_write_buffer_number	num_levels
min_write_buffer_number_to_merge	target_file_size_base	min_write_buffer_number_to_merge	target_file_size_multiplier
compression_ratio	bloom_locality	lwb0_file_num_compaction_trigger	bloom_locality
block_size	max_background_compactions	max_bytes_for_level_base	memtable_bloom_size_ratio
max_background_compactions	max_background_flushes	compaction_pri	level0_stop_writes_trigger
num_levels	level0_stop_writes_trigger	target_file_size_multiplier	max_background_compactions
memtable_bloom_size_ratio	memtable_bloom_size_ratio	level0_stop_writes_trigger	level0_slowdown_writes_trigger
level0_slowdown_writes_trigger	level0_slowdown_writes_trigger	max_background_flushes	block_size




Fig. 1: Performance improvement with Top-5 and Top-10 Parameters on RocksDB.

in parameter selection can lead to substantial differences in performance. Moreover, focusing only on the top- k parameters may fail to capture critical interactions among the remaining parameters and overlook subtle but influential effects on database performance, ultimately reducing the stability of the tuning results.

In addition to these challenges, adapting to diverse workloads remains a significant difficulty in database parameter tuning. Some existing methods attempt to identify the most similar workload from a data repository by computing its similarity to the target workload. However, this approach does not directly optimize for the target workload itself, and inaccurate workload mapping may lead to suboptimal tuning results due to the mismatch between the assumed and actual workload characteristics. Even when two workloads appear similar based on internal metrics such as query latency or cache hit ratio, deeper structural differences, including variations in query patterns, transaction sizes, and read/write ratios, can substantially affect which parameters are most critical for optimization. Since many internal metrics primarily capture low-level system behaviors rather than high-level workload structures, relying solely on past workload data may lead to ineffective tuning, ultimately degrading database performance. This highlights the need for a more target specific optimization approach that directly leverages performance feedback from the actual workload.

Our Approach. To address these limitations, we propose **LatentTune**, a framework that efficiently tunes high-dimensional database parameters through latent representation learning. When the training dataset is limited, LatentTune employs data augmentation techniques using Latin Hypercube Sampling (LHS) [21] and TabNet [22]. Unlike conventional methods,

our approach avoids the need for benchmarking to construct configuration and metric data pairs. To handle the challenge of high-dimensional parameter spaces, LatentTune leverages an AutoEncoder (AE) [23] to compress all parameter configurations into a low-dimensional latent space, thereby enabling the joint optimization of all parameters. Furthermore, by injecting direct target-workload information into the latent space, LatentTune supports workload-specific tuning, removing the need to rely on similarity-based workload mapping. This integration leads to more accurate and effective tuning across diverse workload scenarios. Our key contributions are as follows:

- LatentTune reduces the time required for dataset generation by employing data augmentation techniques.
- We introduce a novel framework, LatentTune, which addresses the high-dimensionality challenge in database parameter tuning by projecting the full set of configuration parameters into a compact latent space. This enables full-parameter optimization, in contrast to conventional methods that rely on selecting only a subset of parameters.
- LatentTune incorporates target workload information directly into the latent space, enabling workload-specific tuning that adapts to diverse workload characteristics.
- The experimental results demonstrate that LatentTune achieves superior performance through comparison with various baseline models.

II. PRELIMINARIES

A. Automatic Database Parameter Tuning

Database parameter tuning is a method of optimizing database parameters (such as memory allocation, disk storage, cache size etc.) to enhance database performance. The configuration of the database refers to the settings and parameters that determine the operation of the database system. If the parameters are specified as P , the configuration can be defined as $\text{conf} = \{P_1, \dots, P_j\}$, where j is the number of parameters. Figure 2 depicts the basic pipeline of ML-based database parameter tuning. The overall flow of the tuning process can be described as follows:

- ① **Tuning Request.** When the controller receives information about the DBMS and the workload to be tuned, it initiates the tuning process.
- ② **Data Repository.** During the tuning process, the DBMS and workload information provided

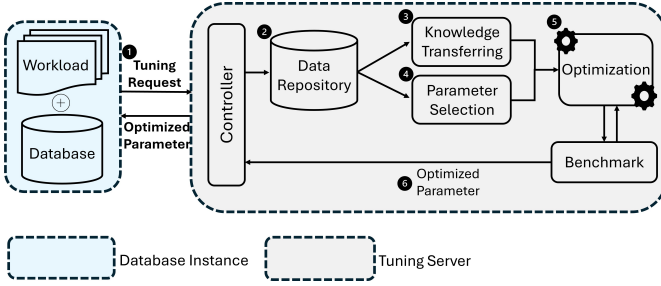


Fig. 2: Architecture of Database Parameter Tuning

in the tuning request are stored in the data repository. ③ **Knowledge Transferring.** To optimize the target workload, the process calculates the similarity between the target and stored workloads in the data repository using a distance metric such as Euclidean distance [24]. The most similar workload information is then utilized to guide the tuning process. The objective of this procedure is to facilitate efficient tuning across diverse workloads. ④ **Parameter Selection.** To address the challenge of optimization in high-dimensional search spaces, a parameter selection algorithm identifies the parameters that most strongly influence database performance. ⑤ **Optimization.** The optimization algorithm, such as BO or RL, optimizes the top- k parameters that have a significant impact on database performance (④) and information about the target workload (③). ⑥ **Optimized Parameter.** The optimized parameters are passed to the controller, which then applies these parameters in the actual database.

B. Related Work

Automatic database parameter tuning techniques are divided into two categories: search-based methods [25] and ML-based methods [5]–[8]. Search-based methods allow tuning without detailed information about the internals of the database system, and ML-based methods tune the parameters using ML models.

Search-based method. Search-based tuning is a technique for optimizing database performance by finding database parameters according to predefined rules. BestConfig [25] uses Divide and Diverge Sampling (DDS) for parameter sampling in high-dimensional parameter spaces and Recursive Bound and Search (RBS) to find optimal configurations in the bounded space around it. However, search-based methods can become unstable with low-quality data and rely on heuristic techniques that are often time-consuming.

ML-based method. ML-based methods use ML models to optimize parameters. For instance, OtterTune [5] selects parameters that are important for performance to optimize in a high-dimensional search space and uses BO with a Gaussian Process (GP) [26] to enhance performance. However, since the set of parameters optimized by OtterTune varies depending on the parameter selection algorithm used, it may result in unstable performance. CDBTune [7] uses Reinforcement Learning (RL), specifically the DDPG [18], which integrates the DQN [27] and actor-critic [28] methodologies to optimize in high-dimensional spaces. However, CDBTune requires extensive training time and benchmarking, resulting

in long tuning durations. Ranking weighted Gaussian process ensemble (RGPE) [20] is an ensemble model for BO-based optimization, which combines information about the workload and historical knowledge to optimize the process. However, RGPE faces challenges when optimizing high-dimensional data, presenting a similar problem to OtterTune, where the selection of parameters for optimization does not guarantee stable performance. Furthermore, RGPE experiences limitations in setting appropriate hyperparameter values for the model.

III. METHOD

Our proposed framework consists of several modules, as illustrated in Figure 3. To address the challenge of collecting large amounts of training data for database configuration tuning, our approach begins with data augmentation applied to a small set of measured configurations and their corresponding performance metrics. By leveraging both the original and augmented samples, we train a low-dimensional representation space that captures the essential relationships between configuration parameters and performance outcomes. The encoder compresses high-dimensional configurations into a compact latent space, while the decoder reconstructs them back into the original configuration space. Finally, Bayesian Optimization (BO) is performed in the latent space to efficiently search for high-performing configurations, and the optimal latent vectors are decoded into complete configuration sets for deployment. Further details of each module are provided in the following subsections.

A. Data Augmentation

ML-based methods require large amounts of training data to improve model accuracy. However, in the field of database tuning research, there are no available public datasets due to the significant influence of various hardware conditions and the specific workloads (sets of tasks or queries) each database system processes. Consequently, the training data generation must be conducted individually, however it requires considerable resources and time. To address this limitation, we employ Sample Augmentation process.

The Sample Augmentation module augments the configuration and metric data pairs. We define the original dataset as $\mathcal{D} = \{(\text{conf}_i, y_i)\}_{i=1}^N$, where each configuration $\text{conf}_i \in \mathbb{R}^d$ is a d -dimensional space of database parameters, and $y_i \in \mathbb{R}^e$ is an e -dimensional vector of external performance metrics. Similarly, the augmented configuration and metric pairs are denoted by $\text{conf}'_i \in \mathbb{R}^d$ and $y'_i \in \mathbb{R}^e$. The resulting augmented dataset is represented as $\hat{\mathcal{D}} = \{(\text{conf}'_i, y'_i)\}_{i=1}^M$.

In the configuration sampling phase, we employ LHS [21] to generate configuration data conf' . LHS is a stratified sampling technique that enables uniform and efficient exploration of multi-dimensional parameter spaces. It ensures that samples are evenly distributed within each dimension, thereby achieving a well-balanced coverage of the overall input space. This systematic approach allows for a more accurate reflection of the individual and combined impact of parameters on

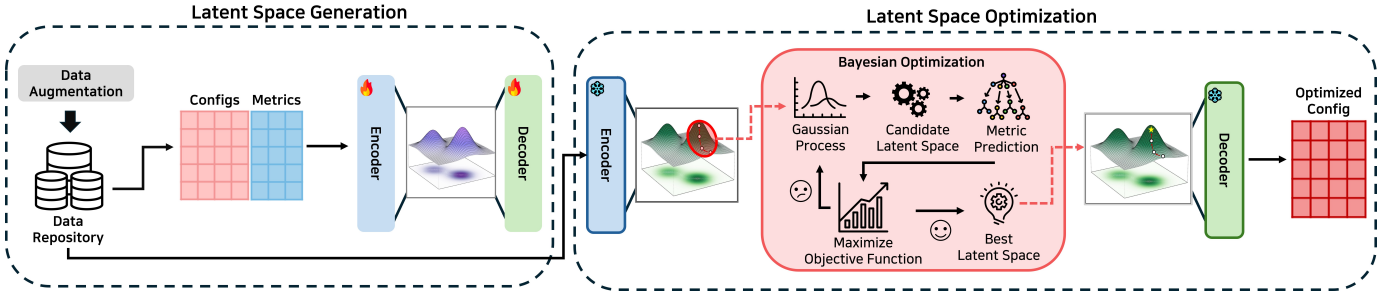


Fig. 3: Overview Architecture of LatentTune

performance metrics [20], [29]–[31]. In contrast, conventional methods such as Simple Random Sampling [32] and Gaussian Sampling [33] often suffer from poor coverage in high-dimensional spaces, particularly when parameter ranges are wide. These methods tend to produce clustered or sparse samples, leading to biased or inefficient exploration of the configuration space.

Given these limitations, LHS enables more uniform and comprehensive coverage of the search space, making it particularly suitable for database parameter tuning scenarios. Therefore, we adopt LHS for generating configuration samples in our framework.

The process of measuring metric values using benchmarking tools is time-consuming and highly resource-intensive. Thus, we utilized a prediction model to predict metric values for $conf'$. Specifically, we employ TabNet [22], which is suitable for tabular data and efficient for multi-target label prediction. TabNet employs a mask-based attention mechanism to effectively identify and prioritize important features among numerous columns, thereby enhancing predictive performance by focusing on highly influential attributes while suppressing less informative ones. For example, database configuration data usually contain numerous parameters and corresponding values organized in a tabular format. TabNet systematically identifies and learns the importance of each parameter, enabling accurate predictions of external performance metrics.

TabNet is trained with the original dataset \mathcal{D} and used to predict metric values y' for the augmented configuration data $conf'$. The augmented dataset $\hat{\mathcal{D}}$ is then constructed by combining the LHS sampled configurations $conf'$ with the corresponding predicted values y' . The prediction accuracy of TabNet is evaluated in the Section IV-C.

B. Latent Space Generation

We leverage an AE [23] to construct a latent space that captures the underlying relationships between database configurations and external metrics. During the latent space optimization (LSO) phase, the decoder of the trained AE is employed to reconstruct the optimized latent vectors into the original high-dimensional configuration space, enabling direct deployment of the resulting configurations.

Training. To generate an informative latent space, LatentTune trains the AE on both the original dataset \mathcal{D} and the augmented

dataset $\hat{\mathcal{D}}$. Each training sample is formed by concatenating a configuration and its corresponding performance metric, resulting in an input data $I \in \mathbb{R}^t$, where $t = d + e$, with d and e denoting the dimensions of configuration and metric, respectively. The reason for concatenating the configuration and metric is to inject workload information into the latent space. Since configurations are composed of arbitrary numerical values, it is difficult to extract meaningful semantic representations in the latent space using configuration data solely. Therefore, by concatenating the configuration and its corresponding performance metric as input, the latent space is able to capture the dependencies between parameters and performance. This approach also allows the latent space to implicitly encode information about diverse target workloads.

The AE is trained to minimize the reconstruction loss between the input and its decoded output. Specifically, it minimizes the following objective:

$$\mathcal{L}_{AE} = \frac{1}{N + M} \sum_{i=1}^{N+M} \|x_i - \mathcal{D}_{\theta}(\mathcal{E}_{\phi}(x_i))\|^2 \quad (1)$$

Here, x_i represents the input data (concatenated configuration and metric), $\mathcal{E}(\cdot)$ and $\mathcal{D}(\cdot)$ are denote the encoder and decoder, respectively. Once training is complete, the encoder compresses the high-dimensional input data into a lower-dimensional latent space, while the decoder reconstructs it back to the original input data space. The LSO process then performs optimization in this latent space, and the optimized latent vectors are decoded to generate final configuration candidates.

Encoder. $\mathcal{E} : \mathbb{R}^t \rightarrow \mathbb{R}^l$ maps the high-dimensional input data I to a low-dimensional latent representation $z \in \mathbb{R}^l$, where $l \ll t$, preserving essential features while reducing dimensionality.

Decoder. $\mathcal{D} : \mathbb{R}^l \rightarrow \mathbb{R}^t$ reconstructs the latent space back into the original high-dimensional space. Specifically, the optimized latent vector z^* obtained from the LSO step is reconstructed through the trained decoder to produce the final output, consisting of a configuration and its corresponding metric.

C. Latent Space Optimization.

The latent space $z \in [0, 1]$ extracted via the LSG step is optimized by BO [9]. BO is a probabilistic optimization approach designed to efficiently explore and exploit the search

space of black-box functions, especially in scenarios where function evaluations are computationally expensive and no explicit functional form is available. In each iteration, BO aims to maximize the value of a black-box objective function by iteratively constructing a probabilistic surrogate model and selecting candidate points through an acquisition function that balances exploration and exploitation.

In this process, we employ a Gaussian Process (GP) [26] as a surrogate model in the latent space to capture the posterior distribution of the objective. Given a candidate latent vectors z , we use the Expected Improvement (EI) acquisition function, which quantifies the expected improvement over the best observed performance. Formally,

$$EI(z) = (\mu(z) - f(z^*) - \xi) \cdot \Phi(\gamma(z)) + \sigma(z) \cdot \phi(\gamma(z)) \quad (2)$$

where $\mu(z)$ and $\sigma(z)$ are the GP predicted mean and standard deviation, respectively, and $f(z^*)$ is the best observed performance so far. The normalized improvement $\gamma(z)$ is computed as:

$$\gamma(z) = \frac{\mu(z) - f(z^*) - \xi}{\sigma(z)} \quad (3)$$

where ξ is a small positive constant that balances exploration and exploitation.

The objective function is typically evaluated as performance (i.e., metrics) by a benchmarking tool. However, since benchmarking is a time-consuming process, we implemented a performance prediction model based on TabNet. It is important to note that this model does not share weights with the TabNet model used in the sample augmentation step.

TabNet is trained on the latent space z , generate through the LSG process, along with the corresponding metrics y and y' . In other words, the input to the prediction model is the latent space, and the output is the metric value. As TabNet's predictions are multi-label, a score function is defined to convert these predict value into a scalar function, facilitating smoother comparisons. The objective of the optimization process is to maximize this score function.

For MySQL, the primary goal is to maximize Throughput while minimizing Latency. To reflect this objective, we define the score function as follows:

$$f_{\text{MySQL}}(x) = \frac{\text{Throughput}(x)}{\text{Latency}(x)} \quad (4)$$

For RocksDB, multiple performance metrics must be optimized simultaneously. Specifically, given four performance metrics (TIME, RATE, WAF, and SAF), the details of each metric are provided in Section IV-A, Evaluation Metrics. we define the score function as:

$$f_{\text{RocksDB}}(x) = -w_0 \cdot \text{TIME} + w_1 \cdot \text{RATE} - w_2 \cdot \text{WAF} - w_3 \cdot \text{SAF} \quad (5)$$

where, TIME, WAF, and SAF are metrics that should be minimized, and RATE is a metric that should be maximized. w_0, w_1, w_2 , and w_3 represent the weights assigned to each of the four metrics.

TABLE II: MySQL Workload Information

MySQL							
Workload Index	Scale Factor	Data Size	Read	Insert	Scan	Update	Read Modify Write
A	12000	15GB	50%	-	-	50%	-
B			95%	-	-	5%	-
E			-	5%	95%	-	-
F			50%	-	-	-	50%

IV. EXPERIMENT

A. Experiment Setup

Hardware. All MySQL performance measurements were performed on a server running the Ubuntu 20.04.6 operating system with an Intel® Core™ i7-11700 @ 2.52GHZ, 32 GB of RAM, and a 256 GB disk. All RocksDB performance measurements were performed with a server running the CentOS 7.9 operating system on an Intel® Core™ i5-8500 CPU @ 3.00 GHz, 32 GB of RAM, and a 256 GB disk.

Tuning Setting. We evaluated performance using MySQL v5.7.37 [15] and Rocks-DB v6.25.0 [19]. We select 138 tunable parameters in MySQL via OtterTune without knobs related to debugging, security, or path settings. The initial dataset consists of 1,000 samples, and it is augmented to 5,000 samples through the data augmentation process. A discussion of the optimal number of augmented samples is presented in 4.3 section. The experiment was repeated five times with different random seed values for each workload, and the average values were calculated. In the baseline experiment, the top- k selected parameters for tuning were set to 5, 10, and 15, respectively, and the best-performing values were used for comparison.

Workload. We conducted experiments on four different workloads for both MySQL and RocksDB. The details of all workloads are shown in Table II and III. **YCSB**(Yahoo! Cloud Serving Bench-mark) [34] is a benchmark tool offers several workloads, each defining the rate and pattern of various operations (load, insert, read, update, scan, etc.) on the database. It can be used to compare and evaluate the performance of various database systems, including MySQL. For RocksDB workloads, we experimented with four workloads with different percentages for read, write, and update.

Evaluation Metrics. The metrics we used to compare the performance improvement between the baseline and Latent-Tune are as follows. **MySQL** uses two evaluation metrics: Throughput, which represents the amount of query processing that can be completed within a given time, and Latency, which refers to the delay time required to process the queries.

Higher throughput and lower latency indicate better performance, the latency being measured at the 99th percentile. **RocksDB** performance metrics include TIME, which measures

TABLE III: RocksDB Workload Information

RocksDB					
Workload Index	Value Size	# of Entry	READ	WRITE	UPDATE
R90W10	16384	65472	90%	10%	X
R50W50			50%	50%	
R10W90			10%	90%	
UPDATE			-	-	O

the time taken to complete database operations for a given workload; RATE, which measures the number of operations processed per second; Write Amplification Factor (WAF) which indicates the write amplification ratio of the workload for a specific configuration; and Space Amplification Factor (SAF), which represents the space amplification ratio during the data storage process. Good performance is defined by lower TIME, WAF, and SAF values, and a higher RATE value. To compare the performance improvement rates of the four metrics in RocksDB, we calculated the following score:

$$SCORE = \log\left(\frac{TIME_{default}}{TIME}\right) + \log\left(\frac{RATE}{RATE_{default}}\right) + \log\left(\frac{WAF_{default}}{WAF}\right) + \log\left(\frac{SAF_{default}}{SAF}\right) \quad (6)$$

This score represents the performance improvement of the optimized parameters compared to the default settings, with the logarithms of individual improvement rates added for easier comparison.

Implementation Details. Our model projected to 32 dimensions for MySQL, and 16 for RocksDB. We set the same 300 iterations for all BO processes in our experiments.

B. Efficiency Comparison

We compared LatentTune with OtterTune, CDBTune, and RGPE on four workloads for MySQL and RocksDB. Each of the three baselines tuned a selection of the top parameters in terms of impact on database performance, and each model included a workload mapping process.

Figure 4 presents the experimental results on MySQL. The results demonstrate that LatentTune achieved the highest improvement rates in both throughput and latency across all workloads. For workload B, only slight performance improvements are observed across both LatentTune and the baseline methods. This is because workload B is a read-heavy workload composed of lightweight operations, which inherently place minimal stress on the system. As a result, changes in database configurations have limited impact on throughput. However, latency can still be significantly affected by parameters related to concurrency control and I/O scheduling, leading to noticeable latency improvements with LatentTune.

Figure 5 shows the experimental results on RocksDB. To facilitate a clear comparison of performance improvements, the enhanced performance over the default configuration is

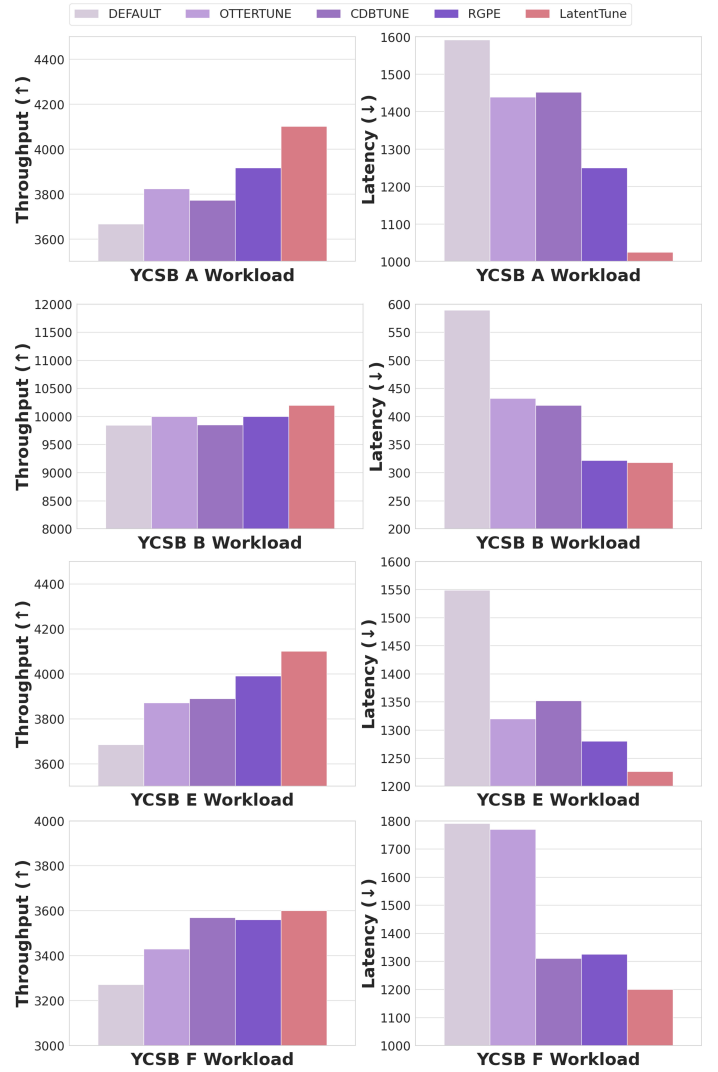


Fig. 4: Performance improvement for MySQL YCSB A,B,E,F workload

represented using the score function. For all four workloads, LatentTune demonstrates improved performance with a maximum score of 9.19 and this result indicates that the LatentTune outperforms the baseline model.

Furthermore, we observed performance improvements across all four metrics compared to the baselines, enabling balanced optimization across multiple aspects of the database. The experimental results indicate that optimizing with information from all parameters is more effective than tuning only the most important ones, as the latter approach ignores the impact of the unselected parameters. Moreover, LatentTune demonstrates the ability to tune various workloads without a workload mapping process, unlike the three baselines. This is because the target workload information is injected into the low-dimensional space.




Fig. 5: Performance improvement for RocksDB workloads

C. Analysis of Sample Augmentation

1) Performance comparison based on number of data

To demonstrate the efficiency of data augmentation in the ML-based tuning process, we compared the AE's reconstruction loss, and LatentTune's performance improvement rate based on the amount of data. Figure 6 shows the AE's reconstruction loss based on the number of training data. Reconstruction loss refers to the deviation between the AE's input and output data by the decoder. In other words, a lower loss indicates higher-quality reconstruction performance. The experimental results show that the reconstruction loss of the AE gradually decreases as the number of training data increases, and the lowest loss is observed when the model is trained with 5,000 datasets. As the number of data points exceeds 5,500, increased reconstruction loss is observed because the information in the data is limited, and the AE's ability to learn is also constrained. This means that even with data augmentation, there is an upper limit to the amount of relevant information that the model can learn, and AE performance can be expected to decrease if that threshold is exceeded. Moreover, when a large amount of data is augmented, the relative proportion of the original data decreases. As a result, the accuracy of the information in the latent space compressed by the AE diminishes, leading to a subsequent increase in reconstruction loss. This consequently leads to a decline in the performance enhancement achieved by LatentTune.

Figure 7 shows LatentTune's throughput improvement rate based on the number of training data. Similar to the lowest reconstruction loss, the highest improvement rate is observed when the number of training data is 5,000.

In conclusion, the AE exhibited optimal performance when the data were augmented for a total of 5,000 datasets. Furthermore, the latent space generated from a well-trained AE successfully compresses information regarding the database parameters and target workload. These results also confirm that LatentTune maximizes database performance by optimizing the latent space. Accordingly, we set the number of augmented data samples to 5,000 when configuring our model experiments.

2) Prediction models for sample augmentation

We compared the accuracy of the prediction model used in benchmarking stage during the sample augmentation process. We selected five models for comparison: basic multilayer



Fig. 6: Reconstruction Loss with number of data.




Fig. 7: Performance Improvement with number of data.

perceptron (MLP) [35], two tree-based models [36], [37], and two models with boosting techniques [38], [39].

As shown in Table IV, Random Forest or XGBoost achieved the highest accuracy in MySQL with fewer prediction labels. However, as shown on the right side of Table IV, the results for RocksDB demonstrate that TabNet achieved superior accuracy with four labels. Although random forest showed good accuracy for MySQL, it failed to achieve consistent prediction accuracy across all metrics when predicting performance for RocksDB. The reason for this is that TabNet dynamically allocates important features for each label when making predictions for multi-label tasks. Therefore, we selected TabNet for sample augmentation due to its consistently high prediction accuracy across all labels, even with multi-label data.

D. Analysis of Latent Space

1) Efficiency of Latent Space Optimization

To evaluate the effectiveness of the proposed latent space optimization method, we analyzed the latent vectors for both optimized and non-optimized configurations in MySQL under two different workloads, YCSB A and YCSB F.

As shown in Figure 8, the contour density plots illustrate that the latent vectors of non-optimized configurations (red) and optimized configurations (blue) form clearly distinct clusters in both workloads. This separation demonstrates that latent space optimization effectively captures and preserves meaningful structures from high-dimensional parameter spaces, ensuring that optimized configurations are well-represented within specific regions of the latent space. Furthermore, the clear distinction between the two clusters indicates the efficiency of our latent space optimization approach. To further validate this, we reconstructed the latent space representations into the original high-dimensional space using the decoder of a trained AE and evaluated their actual performance. Figure 9 shows the contour density plots of the throughput-latency relationship for both workloads. The optimized configurations (blue) exhibit high throughput and low latency, aligning with optimal performance regions. In contrast, the non-optimized configurations (red) are concentrated in suboptimal regions characterized by lower throughput and higher latency. The clear distinction in contour density confirms that our method successfully optimizes configurations, and this pattern is consistently observed across different workloads.

Figure 10 presents the inter-quartile range (IQR) of throughput and latency, respectively, as a detailed analysis of the per-

TABLE IV: Metric Prediction Model Performance for MySQL (YCSB E, YCSB F) and RocksDB (R10W90, UPDATE) in Sample Augmentation.

YCSB E					R10W90							
Model	Throughput		Latency		TIME		RATE		WAF		SAF	
	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)
MLP	0.08	0.810	0.09	0.780	0.45	0.0060	0.31	0.00900	0.40	0.0030	0.45	0.06000
Decision Tree	0.84	0.140	0.83	0.144	0.83	0.0010	0.87	0.00100	0.50	0.0020	0.99	0.00005
Random Forest	0.99	0.006	0.96	0.026	0.85	0.0007	0.93	0.00090	0.76	0.0010	0.99	0.00002
XGBoost	0.99	0.001	0.97	0.022	0.88	0.0005	0.93	0.00090	0.86	0.0007	0.99	0.00003
CatBoost	0.98	0.016	0.95	0.033	0.89	0.0005	0.92	0.00080	0.87	0.0007	0.99	0.00007
TabNet	0.98	0.010	0.95	0.035	0.90	0.0001	0.95	0.00003	0.87	0.0003	0.99	0.00001

YCSB F					UPDATE							
Model	Throughput		Latency		TIME		RATE		WAF		SAF	
	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)	R2(\uparrow)	MSE(\downarrow)
MLP	0.17	0.871	0.16	0.900	0.46	0.0030	0.35	0.007	0.51	0.030	0.46	0.0500
Decision Tree	0.98	0.018	0.95	0.045	0.80	0.0010	0.76	0.002	0.29	0.020	0.98	0.0004
Random Forest	0.98	0.012	0.97	0.022	0.90	0.0005	0.87	0.001	0.68	0.003	0.99	0.0005
XGBoost	0.98	0.017	0.97	0.023	0.92	0.0004	0.87	0.001	0.75	0.003	0.99	0.0006
CatBoost	0.97	0.021	0.96	0.037	0.92	0.0005	0.85	0.005	0.80	0.002	0.99	0.0005
TabNet	0.96	0.043	0.93	0.070	0.94	0.0002	0.87	0.001	0.83	0.002	0.99	0.0004




Fig. 8: Latent vectors density on YCSB A and YCSB F workload.

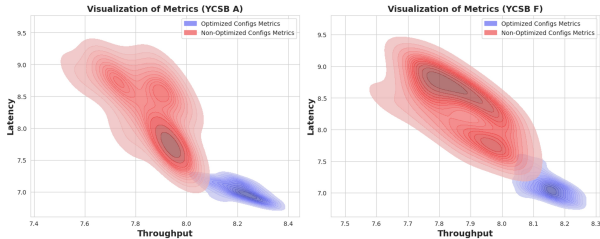


Fig. 9: Metrics density on YCSB A and YCSB F workload.

formance metrics shown in Figure 9. The results further confirm that optimized configurations achieve significantly higher throughput and lower latency compared to non-optimized configurations in both workloads. In particular, optimized configurations (blue) exhibit lower variance and no outliers, demonstrating that the optimization process improves both performance and stability.

2) Prediction models for objective function in LSO

In the LSO stage, instead of performing actual benchmarking in the objective function, we replaced the metric prediction model of the optimized configuration. We compare the accuracy of the prediction model for the optimized latent space, with the same settings as the experiment in Table IV.

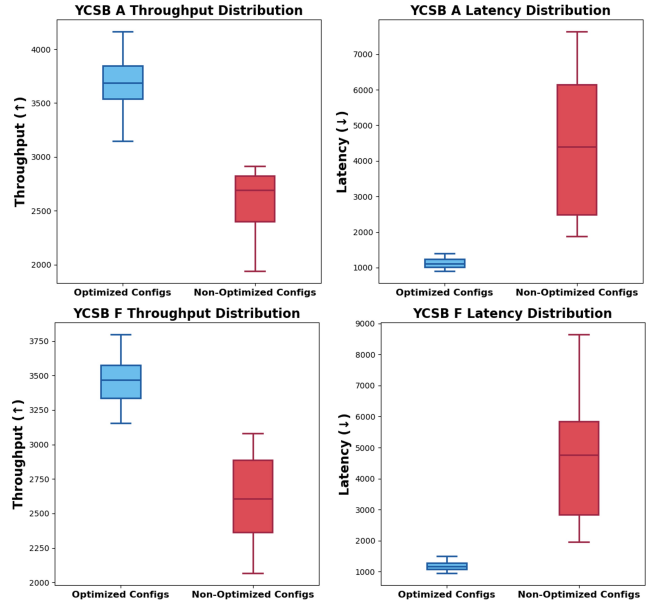


Fig. 10: Metrics distribution on YCSB A (top) and YCSB F (bottom) workloads.

As shown in Table V, the other prediction models either demonstrated poor performance or failed to predict accurately across all labels. However, TabNet achieved the highest accuracy in predicting metrics for both MySQL and RocksDB, and demonstrated consistent accuracy across all labels. Input data of the prediction model is in the form of tabular data, where the rows contain information about configurations and metrics, and the columns represent compressed key-features of the original data. As TabNet is well-suited for tabular data prediction, it outperforms the five comparison models and

TABLE V: Comparison of Metric Prediction Model Performance for MySQL (YCSB E, YCSB F) and RocksDB (R90W10, R10W90) in BO.

YCSB E					R90W10							
Model	Throughput		Latency		TIME		RATE		WAF		SAF	
	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)
MLP	0.28	0.850	0.24	0.600	0.41	0.010	0.48	0.010	0.48	0.006	0.78	0.020
Decision Tree	-0.47	0.080	-0.51	0.040	-0.30	0.040	-0.26	0.030	-0.29	0.170	0.22	0.080
Random Forest	0.18	0.040	0.16	0.020	0.37	0.010	0.43	0.010	0.41	0.007	0.68	0.003
XGBoost	0.16	0.050	0.14	0.020	0.44	0.010	0.50	0.010	0.59	0.005	0.74	0.020
CatBoost	0.45	0.030	0.42	0.010	0.49	0.010	0.54	0.010	0.61	0.005	0.79	0.020
TabNet	0.85	0.008	0.79	0.005	0.80	0.004	0.72	0.005	0.75	0.003	0.96	0.003

YCSB F					R10W90							
Model	Throughput		Latency		TIME		RATE		WAF		SAF	
	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)	R2(↑)	MSE(↓)
MLP	0.024	0.070	0.023	0.030	0.65	0.002	0.61	0.0070	0.18	0.002	0.82	0.010
Decision Tree	-0.34	0.120	-0.39	0.400	0.09	0.005	0.28	0.0080	-0.83	0.006	0.34	0.070
Random Forest	0.21	0.040	0.20	0.020	0.58	0.002	0.69	0.0030	0.13	0.003	0.74	0.020
XGBoost	0.29	0.040	0.23	0.020	0.62	0.002	0.77	0.0080	0.16	0.002	0.78	0.002
CatBoost	0.42	0.030	0.40	0.010	0.47	0.001	0.81	0.0020	0.22	0.002	0.73	0.020
TabNet	0.92	0.004	0.86	0.004	0.72	0.001	0.93	0.0007	0.78	0.001	0.79	0.002

demonstrates equivalent performance across all labels.

E. Analysis of Tuning Overhead and Efficiency

Table VI compares the tuning time across two MySQL workloads. Although RelTune did not achieve the shortest tuning time, it achieved the best overall optimization performance. The relatively higher time cost of RelTune arises from the additional training required to construct the latent space via the autoencoder. However, this cost remains moderate since the objective function in Bayesian Optimization (BO) is replaced with a lightweight prediction model. This result highlights a trade-off between tuning efficiency and optimization performance—RelTune slightly increases the preprocessing time but significantly accelerates convergence toward optimal configurations. Furthermore, RelTune exhibits superior data efficiency, achieving better results with fewer benchmark samples. In practical settings, such a modest overhead is acceptable, as the gain in tuning stability and reproducibility outweighs the additional computation.

F. Ablation Study

We compared the effect of each module of our design on different workloads. Our experiments are conducted using various combinations of Sample Augmentation and LSG. All combinations are optimized using BO, with the same number of iterations set for the optimization process.

Table VII shows the results of the ablation experiments conducted on MySQL. Performance improvement for each combination demonstrates that the design of LatentTune with its modules is effective across all workloads. LatentTune leads to best performance improvement. Overall, combinations without sample augmentation and LSG showed very low performance, as BO in high-dimensional space leads to the curse of dimensionality.

TABLE VI: Time cost (minutes) to complete tuning across workloads. Lower is better.

Method	YCSB A	YCSB B	YCSB E	YCSB F
OtterTune	48.52	50.16	49.70	48.24
CDBTune	362.21	370.15	381.54	421.85
RGPE	62.13	68.46	67.51	64.24
LatentTune	57.23	60.43	59.1	58.53

This curse of dimensionality hinders the learning of the surrogate model in BO and significantly reduces the efficiency of exploration. For this reason, the first combination shows performance that is inferior to the default in all workloads.

When only the augmentation module is applied, the overall increase in the number of data achieve the surrogate model train, resulting in slightly improved performance compared to the first combination. However, due to the curse of dimensionality, the performance remains similar to the default.

When only the LSG module is applied, the reduction in dimensionality addresses the cure of dimensionality. However, since training an autoencoder also requires a large amount of data, the latent space is not well-compressed with the limited data without augmentation. In other words, it fails to effectively compress important information from the data. This, in turn, degrades the performance of the prediction model used in BO to replace the objective function, leading to the poorest performance of the optimized configuration among all combinations.

V. CONCLUSIONS

In this paper, we addressed three limitations of ML-based database parameter tuning: the need for large training datasets,

TABLE VII: Ablation Study on MySQL with YCSB A,B and E,F. "Aug" refers to data augmentation, and "LSG" denotes latent space generation.

(a) YCSB A and YCSB B					
Module		YCSB A		YCSB B	
Aug	LSG	Throughput	Latency	Throughput	Latency
X	X	3651.22	1550	7448.72	714
O	X	3689.79	1281	9952.72	611
X	O	3314.65	1443	6570.68	903
O	O	4100.34	1024	10000.00	318
Default		3666.90	1591	9845.56	589

(b) YCSB E and YCSB F					
Module		YCSB E		YCSB F	
Aug	LSG	Throughput	Latency	Throughput	Latency
X	X	3570.29	1588	3049.66	1220
O	X	3996.52	1261	2934.99	1476
X	O	3439.19	1782	2196.67	7001
O	O	4100.45	1226	3600.76	1200
Default		3685.17	1549	3271.47	1791

the tendency to tune only a subset of parameters due to high dimensionality, and reliance on similar workloads rather than the actual target workload. To overcome these challenges, we proposed LatentTune, a latent space-based tuning framework that dimensionality reduction with data augmentation. This allows for effective training with small datasets and enables full-parameter optimization. By injecting workload-specific information into the latent space, LatentTune also facilitates better adaptation to diverse workloads.

Experimental results show that LatentTune significantly outperforms existing baselines, achieving up to 1332% improvement in RocksDB and 11.82% throughput 46.01% latency gains in MySQL. In future work, we plan to extend LatentTune to support real-time tuning in dynamic environments, moving closer to practical deployment in production database systems.

ACKNOWLEDGMENT

This research was supported by the National Research Foundation (NRF) funded by the Korean government (MSIT) (No. RS-2023-00229822).

REFERENCES

- [1] X. Zhou, C. Chai, G. Li, and J. Sun, "Database meets artificial intelligence: A survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 3, pp. 1096–1116, 2020.
- [2] X. Zhao, X. Zhou, and G. Li, "Automatic database knob tuning: a survey," *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [3] S. Duan, V. Thummala, and S. Babu, "Tuning database configuration parameters with ituned," *Proceedings of the VLDB Endowment*, vol. 2, no. 1, pp. 1246–1257, 2009.
- [4] A. Fekry, L. Carata, T. Pasquier, A. Rice, and A. Hopper, "Tuneful: An online significance-aware configuration tuner for big data analytics," *arXiv preprint arXiv:2001.08002*, 2020.
- [5] D. Van Aken, A. Pavlo, G. J. Gordon, and B. Zhang, "Automatic database management system tuning through large-scale machine learning," in *Proceedings of the 2017 ACM international conference on management of data*, 2017, pp. 1009–1024.
- [6] X. Zhang, H. Wu, Z. Chang, S. Jin, J. Tan, F. Li, T. Zhang, and B. Cui, "Restune: Resource oriented tuning boosted by meta-learning for cloud databases," in *Proceedings of the 2021 international conference on management of data*, 2021, pp. 2102–2114.

- [7] J. Zhang, Y. Liu, K. Zhou, G. Li, Z. Xiao, B. Cheng, J. Xing, Y. Wang, T. Cheng, L. Liu *et al.*, "An end-to-end automatic cloud database tuning system using deep reinforcement learning," in *Proceedings of the 2019 international conference on management of data*, 2019, pp. 415–432.
- [8] G. Li, X. Zhou, S. Li, and B. Gao, "Qtune: A query-aware database tuning system with deep reinforcement learning," *Proceedings of the VLDB Endowment*, vol. 12, no. 12, pp. 2118–2130, 2019.
- [9] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," *Advances in neural information processing systems*, vol. 25, 2012.
- [10] M. A. Wiering and M. Van Otterlo, "Reinforcement learning," *Adaptation, learning, and optimization*, vol. 12, no. 3, p. 729, 2012.
- [11] M. I. Jordan and T. M. Mitchell, "Machine learning: Trends, perspectives, and prospects," *Science*, vol. 349, no. 6245, pp. 255–260, 2015.
- [12] oltpbenchmark, "OLTPBenchmark.com," <http://oltpbenchmark.com>, 2012.
- [13] sysbench, "SysBench: A System Performance Benchmark," <http://sysbench.sourceforge.net>, 2024.
- [14] AuctionMark, "V. Angkanawaraphan and A. Pavlo. AuctionMark: A Benchmark for High-Performance OLTP Systems." <http://hstore.cs.brown.edu/projects/auctionmark>, 2022.
- [15] MySQL, "MySQL v5.7 Documentation," <https://dev.mysql.com/doc/refman/5.7/en/>, 2024.
- [16] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in neural information processing systems*, vol. 30, 2017.
- [17] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 58, no. 1, pp. 267–288, 1996.
- [18] T. Lillicrap, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [19] RocksDB, "RocksDB," <https://rocksdb.org/>, 2012.
- [20] X. Zhang, Z. Chang, Y. Li, H. Wu, J. Tan, F. Li, and B. Cui, "Facilitating database tuning with hyper-parameter optimization: a comprehensive experimental evaluation," *arXiv preprint arXiv:2110.12654*, 2021.
- [21] M. D. Shields and J. Zhang, "The generalization of latin hypercube sampling," *Reliability Engineering & System Safety*, vol. 148, pp. 96–108, 2016.
- [22] S. Ö. Arik and T. Pfister, "Tabnet: Attentive interpretable tabular learning," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 8, 2021, pp. 6679–6687.
- [23] P. Baldi, "Autoencoders, unsupervised learning, and deep architectures," in *Proceedings of ICML workshop on unsupervised and transfer learning*. JMLR Workshop and Conference Proceedings, 2012, pp. 37–49.
- [24] P.-E. Danielsson, "Euclidean distance mapping," *Computer Graphics and image processing*, vol. 14, no. 3, pp. 227–248, 1980.
- [25] Y. Zhu, J. Liu, M. Guo, Y. Bao, W. Ma, Z. Liu, K. Song, and Y. Yang, "Bestconfig: tapping the performance potential of systems via automatic configuration tuning," in *Proceedings of the 2017 symposium on cloud computing*, 2017, pp. 338–350.
- [26] C. Williams and C. Rasmussen, "Gaussian processes for regression," *Advances in neural information processing systems*, vol. 8, 1995.
- [27] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [28] V. Konda and J. Tsitsiklis, "Actor-critic algorithms," *Advances in neural information processing systems*, vol. 12, 1999.
- [29] M. M. Khan and W. Yu, "Robotune: high-dimensional configuration tuning for cluster-based data analytics," in *Proceedings of the 50th international conference on parallel processing*, 2021, pp. 1–10.
- [30] B. Cai, Y. Liu, C. Zhang, G. Zhang, K. Zhou, L. Liu, C. Li, B. Cheng, J. Yang, and J. Xing, "Hunter: an online cloud database hybrid tuning system for personalized requirements," in *Proceedings of the 2022 International Conference on Management of Data*, 2022, pp. 646–659.
- [31] J. Xin, K. Hwang, and Z. Yu, "Locat: Low-overhead online configuration auto-tuning of spark sql applications," in *Proceedings of the 2022 International Conference on Management of Data*, 2022, pp. 674–684.
- [32] P. West, "Simple random sampling of individual items in the absence of a sampling frame that lists the individuals," *New Zealand journal of forestry science*, vol. 46, pp. 1–7, 2016.
- [33] D. Micciancio and M. Walter, "Gaussian sampling over the integers: Efficient, generic, constant-time," in *Advances in Cryptology—CRYPTO 2017: 37th Annual International Cryptology Conference, Santa Barbara*,

CA, USA, August 20–24, 2017, *Proceedings, Part II* 37. Springer, 2017, pp. 455–485.

- [34] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears, “Benchmarking cloud serving systems with ycsb,” in *Proceedings of the 1st ACM symposium on Cloud computing*, 2010, pp. 143–154.
- [35] M.-C. Popescu, V. E. Balas, L. Perescu-Popescu, and N. Mastorakis, “Multilayer perceptron and neural networks,” *WSEAS Transactions on Circuits and Systems*, vol. 8, no. 7, pp. 579–588, 2009.
- [36] J. R. Quinlan, “Induction of decision trees,” *Machine learning*, vol. 1, pp. 81–106, 1986.
- [37] L. Breiman, “Random forests,” *Machine learning*, vol. 45, pp. 5–32, 2001.
- [38] T. Chen and C. Guestrin, “Xgboost: A scalable tree boosting system,” in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.
- [39] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, “Catboost: unbiased boosting with categorical features,” *Advances in neural information processing systems*, vol. 31, 2018.