Accelerating the Configuration Tuning of Big Data Analytics with Similarity-aware Multitask Bayesian Optimization

Ayat Fekry*, Lucian Carata*, Thomas Pasquier† and Andrew Rice*
*Computer Laboratory, University of Cambridge, UK
†Department of Computer Science, University of Bristol, UK

Abstract—One of the key challenges for data analytics deployment is configuration tuning. The existing approaches for configuration tuning are expensive and overlook the dynamic characteristics of the analytics environment (i.e. frequent changes in workload due to receiving evolving input sizes or change in the underlying cluster environment). Such workload/environment changes can cause significant performance degradation, with retuning the configuration to accommodate those changes can yield up to 85% potential execution time saving.

We propose SimTune, an approach that accommodates such changes through efficient configuration tuning. SimTune combines workload characterization and Multitask Bayesian optimization to identify similarity across workloads and accelerate finding near-optimal configurations. Our experimental results show that SimTune reduces the search time for finding close-to-optimal configurations by 56-73% (at the median) when compared to existing state-of-the-art techniques. This means that the amortization of the tuning cost happens significantly faster, enabling practical tuning in the rapidly changing environment of distributed analytics.

I. Introduction

With the ever growing data volume [1], the timely extraction of useful insights from data remains a challenge. The need to analyse large datasets has led to the wide adoption of Data Intensive Scalable Computing (DISC) Systems such as Hadoop [2], Spark [3] and Flink [4]. These DISC systems enable the manipulation and analysis of large amounts of data by distributing work over a cluster of machines. They are used to help organizations make better and faster decisions, as well as in research areas that range from biology [5] and physics [6] to astronomy [7].

One of the challenges in setting up DISC systems is to identify the right configuration in order to optimize processing speed. Misconfiguration can lead to either resource contention/exhaustion or under-utilization, with the former potentially triggering errors hours after the start of execution. Consequently, developers spend significant time and resources identifying the appropriate configuration for their workload. This laborious effort motivated work towards the automation of configuration tuning [8]–[13].

Generally, such explorations follow one of two strategies: either search based or model-based, wherein the latter a performance prediction model is built to guide the tuning. Both strategies require hundreds of executions to perform high dimensional configuration tuning [14], [15], with the cost of performing the optimisation itself starting to amortize as a function of both how fast the algorithm converges to very good configurations and how many bad (slow) configurations are explored in order to get there.

Virtually all of the existing work for DISC systems configuration tuning assume a *static* environment, such as a stable workload running in a fixed underlying cluster. Therefore, it is assumed that changes in the workload/environment are infrequent enough that re-tuning from scratch is cost effective. Re-tuning happens either by re-learning the models or restarting the search algorithm, which is slow and prohibitively expensive. Ignoring any prior knowledge acquired during workload tuning poses poor adaptivity to such changes and slows down the tuning process, incurring high re-tuning costs.

In a rapidly changing environment such as big data systems, it is crucial to perform tuning/retuning efficiently using a small number of executions, which eventually will accelerate the amortization of tuning costs. We propose solving this problem using a similarity-aware multitasked tuning approach (SimTune). We leverage neural encoding of workload execution metrics to detect workload similarities, then share the tuning knowledge gained previously across the similar workloads using multitask Bayesian Optimization (BO). Our main focus is on *how* to leverage workload similarity to accelerate the configuration tuning. We consider an existing *base* tuner and address accelerating this base tuner using our proposed tuning approach. The key contributions of this work are:

- Detecting workload similarity using a novel representation learning approach that transforms workload execution metrics into a low dimensional space using nonlinear neural encoding and minimizes the information loss compared to earlier work.
- Tuning the configuration parameters incrementally using a small number of workload executions, by employing multitask BO to share the gained tuning knowledge across the similar workloads.
- Obtaining comparable configurations to prior work but converging to those significantly faster, enabling a quicker amortization of the tuning costs.

SimTune is the first work to study *how* to share the tuning knowledge across similar workloads to enable *data-efficient* [16] tuning of DISC systems configurations. The

source code and experiment dataset is available at [17].

II. BACKGROUND

BO [19] is a method for minimizing blackbox functions f iteratively, using a limited number of samples. This is useful when it is expensive to evaluate f at a given point (such as running a big-data workload with a given configuration). BO is characterized by its *prior model* and *acquisition function*: the prior model represents a space of possible target functions f, and the acquisition function guides the selection of the next evaluation point based on the prior knowledge.

One of the widely accepted prior models for BO is Gaussian Process (GP). It represents a distribution over functions (a sample drawn from this process is a function) with given mean and covariance. Here, the mean function describes expected values at each point and the covariance function defines the smoothness of the functions which can be drawn as samples, encoding prior assumptions about the data that we want to model [20]. The GP maintains a probabilistic belief about what functions f are possible, given known characteristics and already seen data. This belief is updated by using an acquisition function, which determines the best point of f to sample next. After sampling, the prior belief about possible functions f is updated and a new sampling decision can be made, iteratively. At each step, the posterior distribution has filtered-out functions not consistent with the sampled data and will ideally have a narrower candidate function space. The GP acquisition function represents the metric by which the GP picks the next input sample to improve the probabilistic model function. It is typically a function that is cheap to evaluate at a given point x and its value is proportional to how useful evaluating $f(x)$ would be for the optimisation problem. Various acquisition functions have been proposed to define the way the GP samples the input space, e.g., random, sequential, Probability of Improvement (PI), and Expected Improvement (EI) [21]–[23].

Multitask Bayesian Optimization (MTBO) is an extension of the standard BO to enable modelling and minimizing different tasks. This modelling principally relies on defining a covariance function between input-task pairs. While the covariance function of the standard GP defines the relationship between inputs, the covariance function of Multitask GP (MTGP) defines the relationship between inputs and tasks. It learns the degree of correlation between tasks and utilizes this information to guide the search for the best point to sample. Figure 1 shows an example MTGP with three tasks, 1(a) shows the actual functions of the three tasks. Task 2 and 3 are strongly correlated, 1 and 3 are anti-correlated, and 1 and 2 are not correlated. The goal is to model task 3, the dots represent observations and the dashed lines represent the predictive mean function. The curve at the bottom represents the acquisition function(e.g., Expected Improvement(EI)) for each input location on Task 3, the next point to sample is the one with the maximum acquisition (maximum EI). The confidence interval indicates the range that the actual function should fall in with a high probability, the narrower the better.

Whereas the independent GP in 1(b) has a wide confidence interval and optimizing EI leads to a false minimum as the next point to sample, the multitask GP in 1(c) leverages the correlation across the tasks, thus the confidence interval is improved (narrowed) and the maximal EI point corresponds to the true minimum. We discuss MTGP as it applies to SimTune in § III.

III. SIMILARITY-AWARE MULTITASK TUNING

To accommodate the need for efficient configuration tuning, we propose a similarity-aware tuning approach that leverages the tuning knowledge between similar workloads and effectively accelerate the tuning process.

Assumptions: We assume the existence of readily available workloads tuned using a *base* tuner.

We are building upon our prior work on Tuneful [24], which we use as our *base* tuner, since it finds configurations comparable to the state-of-the-art using significantly fewer workload executions. Moreover, it tunes the configurations using single tasked GP (STGP), which has been widely adopted in several configuration tuning systems ([12], [25]).

While tuning, Tuneful also determines workload-specific significant configuration parameters that we can leverage in our similarity-aware tuning.

Our work addresses *how* to define similarity across workloads and *efficiently* leverage this similarity in twofold: 1) tune a new similar workload, 2) retune already seen workload to accommodate the ever growing input sizes. We start with learning how to characterize workloads to detect workload similarity, then employ MTBO to accelerate the tuning of these similar workloads.

A. Workload Characterization

As a first step, we identify the main components distinguishing each workload so that we detect the similarity across workloads. To achieve this, we monitor the workload's execution metrics then learn the main aspects that represent each workload.

1) Workload Monitoring:: We capture metrics that involve CPU time, number of tasks per stage, number of stages per job, input and output size, Garbage Collection (GC) time, execution time, data serialization/deserialization time, the size of shuffled data, memory spilled data and disk spilled data. We collect all the numeric metrics that relate to those workload execution features (a total of 24 metrics). The overhead of collecting these metrics is minimal as we leverage the existing Spark logs.

To build representative statistics of each metric, for each stage in the analytics workload, we calculate the 95% quantile of the metric across the tasks of the stage. Then, we average the captured metric across all the stages and represent them in terms of ratios. For example, we consider the amount of GC time relative to the total CPU time and the amount of shuffled data with respect to the total input data, rather than the quantitative values of each metric. This is useful since we ultimately use the characterization to detect similarity between
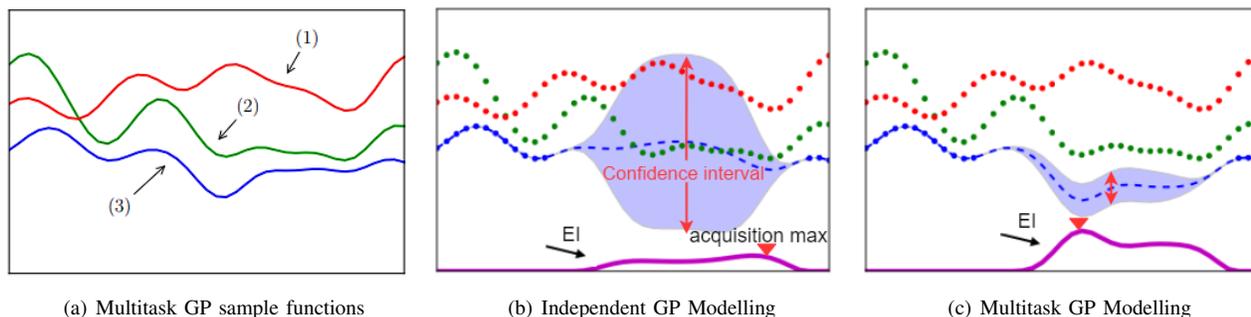
(a) Multitask GP sample functions   (b) Independent GP Modelling   (c) Multitask GP Modelling

Fig. 1: (a)The actual functions of a Multitask GP with three tasks. Task 2 and 3 are strongly correlated, 1 and 3 are anti-correlated, and 1 and 2 are not correlated. (b) Independent single tasked GP modelling for Task 3. (c) Multitask GP modelling for Task 3, utilizing the other tasks (figure source is [18] with minor edits applied for more clarity).

workloads, and in our context two workloads are similar if they inherit similarity in terms of their relative resource usage- not the quantitative resources usage. Intuitively and supported by our experiments, the workloads that are similar in their relative resource usage are expected to also have similar significant configuration parameters.

2) Workload Representation Learning:: The existing work on detecting similarities relies on first projecting the workload execution features in a low dimensional space then calculating the distance within this space. This is due to the fact that distance metrics such as L_p norm become less informative in the high dimensional space [26]. We embrace the same approach but differentiate in the way we represent the workload, by leveraging neural encoding to encode workload execution features in a low dimensional space. Previous work employs PCA for this purpose. However, by projecting the data in a linear space some of the nonlinear dependencies in the latent space are lost, leading to a higher data reconstruction loss. We instead use a nonlinear neural autoencoder to lower this loss. More details on evaluating workload representation are in § IV-B.

An autoencoder (AE) is a neural network that learns how to represent data in a low dimensional space. It has been widely adopted to encode data in different domains such as image processing, natural language processing and signal compression. It consists of two components: encoder and decoder, the encoder is a neural layer(s) learns $f(x)$, which transforms input $x \in \mathbb{R}^d$ into $h \in \mathbb{R}^n$ latent space, such that $n < d$. The decoder learns $g(h)$, which reconstructs the latent low dimensional h back to $\hat{x}$ of the same dimensionality as the original input space. The AE compares the decoder's output (reconstructed input) $\hat{x}$ to the encoder's original input x and update the neurons' weights accordingly to minimize the loss of the re-constructed input L as shown in Equation 1, where L is a function that penalizes $g(f(x))$ for not being similar to x (e.g., the mean square error).

$$L(x, \hat{x}) = ||x - g(f(x))||^2 \qquad (1)$$

As illustrated in Figure 2, our AE learns the workload representation as an offline phase. We use a one layer nonlinear

AE to encode the execution features into a low dimensional latent space. We made this decision as it minimizes the reconstruction error compared to linear AE and PCA, moreover it has a smaller number of model parameters compared to multilayer AE.

3) Similarity Analysis:: The aim of the similarity Analysis is to find a *source* – already tuned – workload and share its significant configuration parameters with the *target* workload (the one under tuning). The source workload is the most similar one based on the Manhattan distance between the encoded execution features h. We use Manhattan distance as it provides more contrast compared to other L_p norm distance metrics [26].

The distance is calculated between workloads running under the *same* configuration. We use a *fixed* representative configuration to run the workload once and extract workload execution metrics using the learnt workload representation (this works as workload fingerprint). We then match the workload fingerprint to one of the seen workloads. We experimented with varying the number of workload executions used for workload fingerprinting and how this influences the accuracy of workload matching. Using a single representative configuration was enough to provide the same matching as multiple executions, so we opted to use a single execution to limit the cost of workload fingerprinting. Two workloads are considered similar if they have the smallest distance in terms of their relative resource usage, since this implies that they have similar influential parameters. Figure 2 shows how the similarity analysis happens online, upon receiving a new *target* workload. As a first step, the Tuning Manager checks if this workload has a known similar workload or not. If not, it suggests the fingerprinting configuration to the Spark driver, which then forwards it to the Spark workers for execution. After the workload execution, the Tuning Manager provides the collected metrics to the Similarity Analyzer, which encodes the execution metrics of this workload, leveraging the learned workload representation from the offline phase. The workload with the smallest distance is selected as the source workload and we start tuning its significant parameters for the target
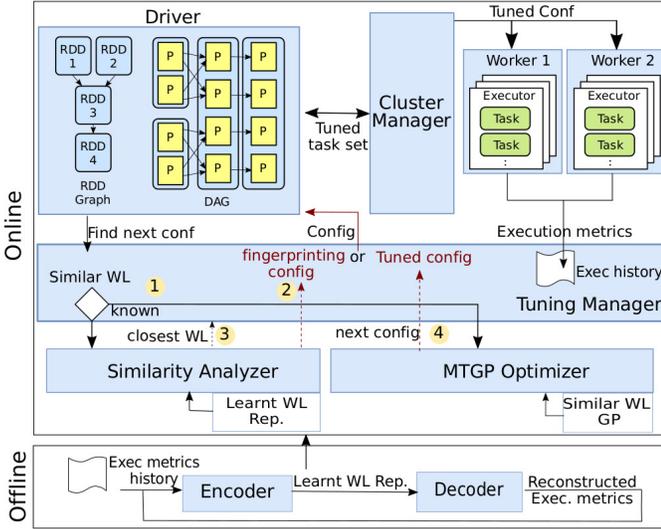
Fig. 2: SimTune in the offline and online phases to tune Spark workload.

B. Multitask Configuration Tuning

After finding a similar *source* workload – tuned using a *base* tuner – we address accelerating it using SimTune. We transfer the knowledge gained during tuning the *source* workload to the *target* workload. This knowledge includes: 1) the significant configuration parameters already determined for the *source* workload. 2) the execution samples evaluated during tuning this source workload. We assume that two workloads identified as similar will share the same significant parameters, but that the actual values they take in an optimal configuration of the *target* workload might be different.

The Tuning Manager passes the tuning knowledge of the source workload to the MTGP Optimizer to reuse the significant configuration parameters and leverages the previously seen execution samples. This knowledge guides the MTGP Optimizer to suggest the next configuration sample, with the highest potential to maximize the EI. In the beginning, it suggests the configuration samples that help learning the covariance across the different workloads (modelled as tasks). Then it decides the next sample to try based on this covariance (i.e. if the covariance is strong, it picks the sample that maximizes the EI of the source task, otherwise it picks the sample that maximizes the EI of the target task). More details on how MTGP works is in [18].

In Alg. 1, we show the steps of our multitask configuration tuning, its input arguments are as follows:

- w target workload;
- $W_{seen} = \{w_1, w_2, ..\}$ the matrix of the encoded execution features for the seen workloads;
- α the acquisition function;
- $GP = \{GP_1, GP_2, ..\}$ the Gaussian process models of the seen workloads;

Algorithm 1: Similarity-aware Multitasked configuration tuning

Input : α, w, W_{seen}, GP
1 **forall** $x_i \in W_{seen}$ **do**
2   calculate distance d between $w_{metrics}$ and x_i;
3 Find workload s with the smallest d ;
4 Transfer s significant parameters to w ;
5 Add new task t_j for workload w to GP_s ;
6 Find config: $x_i \leftarrow \arg\max_{x} \alpha(GP_s(x, t_j))$;
7 Evaluate config: $y_i \leftarrow C_w(x_i)$;
8 Update: $GP_s \leftarrow GP_s | (t_j, x_i, y_i)$;

Upon receiving a new *target* workload w and following fingerprinting it to find its execution metrics $w_{metrics}$ (as described in § III-A3), the algorithm works as follows: We start at line 2 with calculating the distance between the execution metrics of the target workload $w_{metrics}$ and the auxiliary workloads W_{seen}, the distance is calculated based on the learnt low dimensional encoding of workload execution metrics. We then select workload s with the smallest distance to use as the *source* task for the multitask tuning at line 3. The significant parameters of s is reused to tune w and a new task t_j is added to the GP of s for w (line 5). At line 6, we leverage the GP of s to select the next configuration to execute w, the next configuration is selected so that it maximizes α acquisition function over task t_j. Lastly, workload w is executed using the selected configuration (line 7) and the execution cost C is used to update the GP model of task t_j(line 8).

Figure 2 shows how the tuning takes place online during the normal workload executions. The MTGP Optimizer suggests the next configuration to explore, then the Tuning Manager feeds the execution cost[1] back to the MTGP optimizer to update the MTGP model. The MTGP optimizer detects if it is time to stop the optimization. It stops modelling a workload after suggesting a minimum of n samples (e.g 10 samples) and then after the expected improvement (EI) drops below 10%. A decision made to make sure that we balance between the exploration of the tuning space and exploitation of the best configuration found.

Resilience to workload mismatching: we can detect inaccuracies in workload matching through monitoring the gap between the predicted execution time by the MTGP model and the actual execution time. If a continuous degradation takes place, then we trigger workload matching again (given that we have seen sufficient workloads, otherwise we perform a standalone tuning for this workload using the base tuner). Even if a workload mismatching happens due to the limited number of seen workloads, SimTune guarantees that the found configuration is not worse than the reused state-of-the-art's configuration, which implies reusing the configuration sug-

---

[1]We focus on execution time through the paper, but the algorithm works similarly for any defined cost function.

Application	Abbr.	Input data sizes (DS)
Pagerank	PR	5, 10, 15, 20, 25 (million pages)
TPC-H benchmark	TPCH	20, 40, 60, 80, 100 (compressed GB)
Terasort	TR	20, 40, 60, 80, 100 (GB)
Bayes Classifier	Bayes	5, 10, 30, 40, 50 (million pages)
Wordcount	WC	32, 50, 80, 100, 160 (GB)

TABLE I: The set of applications and input sizes used to learn workload representation.

gested by the state-of-the-art for a similar workload (we refer to this as Direct Transfer and illustrate that more in § IV).

SimTune is implemented on top of Spearmint [27], a Python BO framework that implements MTBO based on [18].

IV. EVALUATION

A. Experimental setup

Cluster and configuration specification: We use a cluster of 4 AWS *h1.4xlarge* instances with 16 vCPUs, 64 GB memory, and 2TB storage each. We use HDFS [28] version 2.7 for accessing the shared data and Spark version 2.2.1 as the system under tuning. Our base tuner tunes 30 configuration parameters, with approximatively $2 \cdot 10^{40}$ configurations possible in total (this represents the size of the search space), a list describing each configuration parameter and its range can be found in the project repository [29]. We use the same ranges when evaluating the other tuning approaches.

We start with evaluating the workload representation learning then evaluate SimTune effectiveness and efficiency.

B. Workload Representation Learning

Environment setup: We used Keras v2.3.0 to build the AE, with epoch (the number of complete passes over the dataset to update the internal model parameters) of 300 and a sigmoid activation function. We tried the other possible nonlinear activation functions and selected the one with the minimal information loss. For setting the epoch, we searched till 1000 and the information loss converged at 300, we use the same number of epoch for the linear AE.

Dataset: We built a dataset of execution metrics using two well known big data benchmarks (Hibench [30] and TPC-H [31]). We selected five applications of heterogeneous characteristics from these benchmarks. For each application we experimented with five different input sizes, we captured the execution metrics of each application and input size pair under 100 random configurations for the 30 configuration parameters in [29], with a total of 2200 application executions that took 2188 compute hours to execute. We do not consider this time as an overhead of our tuning approach as this data will be readily available over time in a deployed system. Table I lists the applications and input sizes, the dataset is publicly available at the project repository [17].

To evaluate the accuracy of learning how to represent workload's execution features in a low dimensional space, we compare the loss of our representation using nonlinear AE against PCA and linear AE. Figure 3 shows the reconstruction
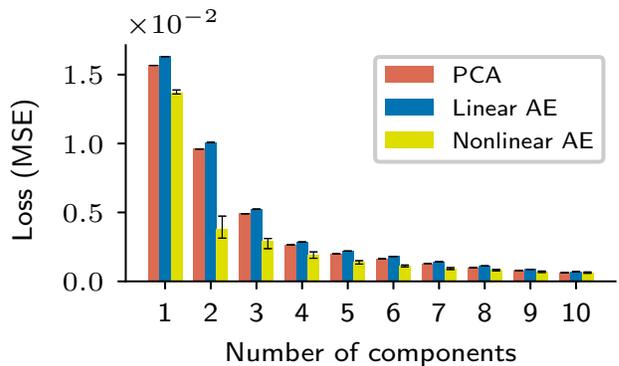

Fig. 3: Reconstruction loss over the number of encoded low dimensional components, showing median, 90th and 10th percentile of 10 experiments.
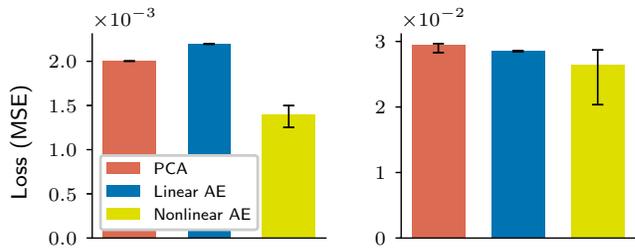

Fig. 4: Reconstruction loss for 20% holdout validation dataset (left) and test dataset of 500 workloads execution metrics running on a different cluster (right).

loss over the number of components/encoded dimensions during learning the representation using PCA, linear AE and nonlinear AE at the median, 90th and 10th percentile. The nonlinear encoder can represent the execution metrics using a smaller number of components/dimensions while maintaining less reconstruction loss compared to PCA and the linear AE. For example, at the 90th percentile the nonlinear AE can encode the execution metrics using 5 dimensions with a loss less than PCA and linear AE using 6 dimensions, since they restrict the encoding of data into a linear space that does not capture the nonlinear relationships across the execution metrics. We chose to represent the workload execution metrics using 5 dimensions, representing a good compromise between the reconstruction loss and the number of components for all the three approaches (e.g., PCA represents 92% of the data variance using 5 components).

We excluded 20% of the dataset as a hold out validation dataset to evaluate the learnt representation. Further, to evaluate the generalization of the learnt representation in a different environment, we built a test dataset that includes execution metrics of 5 applications running on a cluster of a different characteristics: a 20 Google Compute Engine [32] instances (1 driver + 19 workers), with the driver being an *n1-highmem-8* instance and the 19 workers being *n1-standard-*
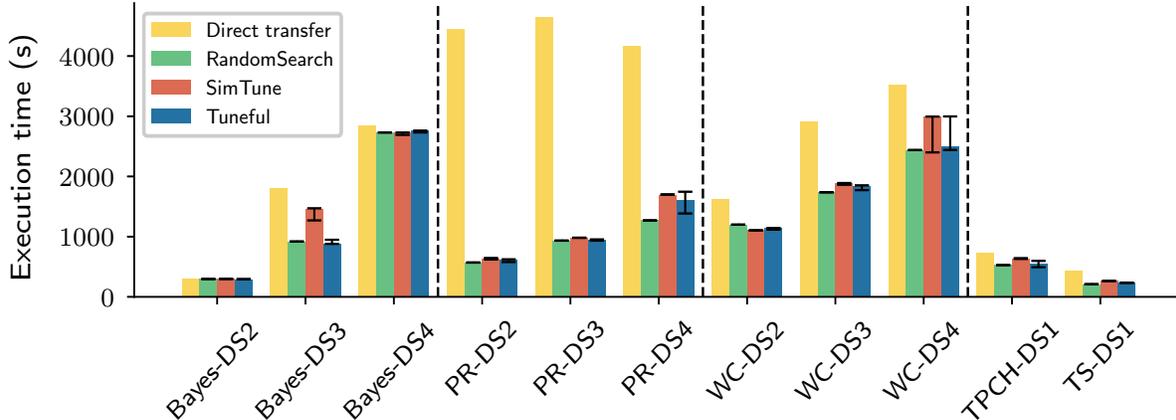
Fig. 5: Execution time of the direct transferred configuration, the tuned configurations using Tuneful, Random search and SimTune.

Workload	Bayes-DS2	Bayes-DS3	Bayes-DS4	PR-DS2	PR-DS3	PR-DS4
RandomSearch	12	1915	769	330	3147	4265
Tuneful	12	1168	684	416	1261	2452
SimTune	5	511	46	229	258	739

Workload	WC-DS2	WC-DS3	WC-DS4	TPCH-DS1	TS-DS1
RandomSearch	2351	452	1161	151	423
Tuneful	719	452	952	54	180
SimTune	111	612	1247	208	117

TABLE II: The median search time (m) till finding 10% of the optimal configuration using Tuneful, Random search and SimTune. The green values represent the total search time when a configuration within 10% of the optimal is not found.

16 instances. The dataset contains the execution metrics of 500 executions of TPC-H, PR, bayes, WC, TR and kmeans workloads. Figure 4 shows the reconstruction loss for the validation and test dataset. The nonlinear AE maintains the lowest loss for the validation dataset at the 90th percentile and generalizes better than the other approaches for the test dataset, with a smaller reconstruction loss at the median.

C. Tuning Effectiveness and Efficiency

We compare the configurations picked by our approach with the configuration tuned by: 1)*Direct transfer*, which implies transferring the tuned configuration from the source workload directly to the target workload. We compare against this to show the importance of retuning the configuration. 2)*Tuneful*, our base tuner and the state-of-the-art GP-based DISC system tuning approach, it starts with detecting workload-specific significant parameters then tunes them using single tasked GP [33]. We chose Tuneful to assess the accuracy of reusing the significant parameters of a similar workload against performing an independent detection of the significant parameters and tuning. 3) *Transfer learning and single tasked GP tuning* (TL+STGP), which implies transferring the significant parameters from a similar workload then performing configuration

tuning using Single Tasked Gaussian Process (STGP). We compare against this approach to evaluate the benefits of using MTGP against STGP. 4) For the sake of completeness, we also compare the tuned configuration against *RandomSearch*, with a budget of 100 execution samples and the random configurations are generated using low-discrepancy sequences [34], since they cover the search space quicker and more evenly than the standard random numbers.

Metrics: We use three metrics to evaluate our proposed approach: 1) The execution time of the tuned configuration, akin to the running cost of the tuned configurations. The target here is to obtain tuned configurations similar to what state-of-the-art achieve. 2) Search Cost, the amount of time and actual cost (in $) required by each approach to find good configurations while repeatedly running workloads in a cloud environment. The target is to get close-to-optimal configurations (within 10% of the estimated best configuration) significantly faster than the state-of-the-art. 3) *Amortization speed*, the number of needed workload executions to amortize the tuning cost. The target is to amortize the tuning cost after a small number of workload executions.

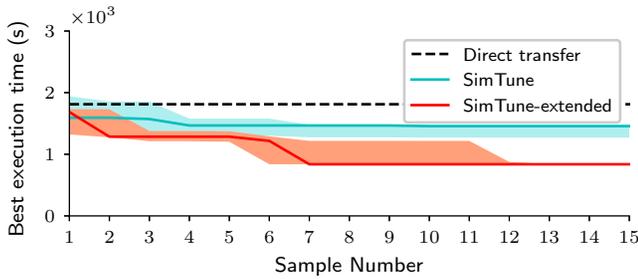The results are always presented as the median of 10

Fig. 6: Convergence speed of the Bayes workload when the significant parameters are transferred from Bayes-DS2.

experiments, bars represent the 10th and 90th percentile.

Methodology: To mimic the big data analytics environment when they receive similar workloads or growing input sizes, we consider a set of *Auxiliary workloads* that represents the seen workloads, then evaluate tuning a set of similar workloads with evolving input sizes.

Auxiliary workload set: consists of three workloads from three applications (bayes-DS1, PR-DS1 and WC-DS1), each tuned using Tuneful [33] (our base tuner), which performs an independent significant parameter detection and tuning for each workload. We excluded two applications (TPC-H and TS) from the auxiliary set workload to evaluate the effectiveness of tuning unseen applications.

To evaluate the effectiveness of SimTune in finding good configurations, for each application in the auxiliary workload set, we tune three workloads of evolving input sizes using SimTune and compare against direct transfer, Tuneful, and RandomSearch. We also consider tuning two workloads of two applications that were not included in the auxiliary workload set (TPC-H-DS1 and TS-DS1).

SimTune finds configurations comparable to the state-of-the-art outperforming the direct transfer by 32% at median and 79% at the 90th percentile: Figure 5 shows the execution time of the configurations found by the different tuning approaches, the direct transfer of the configurations does not guarantee near-optimal performance, with a potential performance degradation of up to 85% compared to retuning the configuration. Across all workloads, SimTune finds on average a configuration with an execution time within 8% of Tuneful's and 12% of RandomSearch's. This average is computed over the difference in median execution times for each workload that are shown in Figure 5.

However, for some workloads (e.g., Bayes-DS3 and WC-DS4) Tuneful found configurations outperform the ones found by SimTune by up to 36%. This is due to the limited set of auxiliary workloads. The next section describes how this can be mitigated through extending the set of auxiliary workloads.

To evaluate differences in time taken to find configurations close-to-optimal, we allow each approach to execute workloads until it finds the first configuration resulting in a runtime within 10% of the one produced by the estimated optimal configuration. This is defined as the best configuration ever found across all our tests for each workload, irrespective of the tuning approach or experiment. Table II shows the search time of each approach. Since for some workloads SimTune and Tuneful do not find a configuration within 10% of the estimated optimal configuration, we color these workloads in green in the table, the value for those workloads represent the total search time to find configuration has the execution time shown in Figure 5.

The median search time for the-state-of-the-art is 2.3-3.7X longer when compared to SimTune: Table II shows that overall and given this limited auxiliary workloads, SimTune significantly accelerate the search time of 6 workloads while finding a configuration within 10% of the estimated optimal. For the remaining workloads, the total search time of SimTune is still considerably smaller than the other approaches, while not only notably outperforming direct transfer, but also finding a configuration comparable to Tuneful for most of the workloads and within 17%-37% of RandomSearch.

This suggests that a significantly shorter search phase is possible even with a limited number of source workloads to transfer tuning data from. However, in around 45% of the cases this also means that the reached configuration is slightly further away from the optimum. This can be mitigated when a larger set of "auxiliary workloads" exists as shown in Figure 6. **Extending the tuning knowledge for SimTune enables a faster and more effective tuning:** Table II shows that Tuneful takes 2.5X longer time than SimTune for Bayes-DS3, while finding configuration outperforms direct transfer by 54% and saves execution time by 36% compared to the one found by SimTune. This is due to the limited set of auxiliary workloads, resulting SimTune to find configuration not within 10% of the optimal and the time in the table represents the total search time for this workload. However, Figure 6 shows the convergence speed of SimTune when extending the auxiliary workload set to tune the Bayes-DS3 workload. The area between the 90th and 10th percentile of execution time is shaded, and the line shows the median of 10 experiments. We plot the minimum execution time found by the tuning algorithm until each execution sample on the x axis. Extending the auxiliary workload set to include Bayes-DS2 workload allows SimTune to select Bayes-DS2 as the source workload, this enables the tuning to happen using 38% less search time compared to using a smaller auxiliary workload set, while finding a configuration comparable to Tuneful's outperforming the direct transfer by 54%. In comparison, the one found using a smaller auxiliary workload set outperforms direct transfer by 19%.

To evaluate the amount of cost saving, we estimate the search cost based on AWS [35] per-second pricing. SimTune saves the median tuning cost by 3.3X and 4.7X compared to Tuneful and RandomSearch.
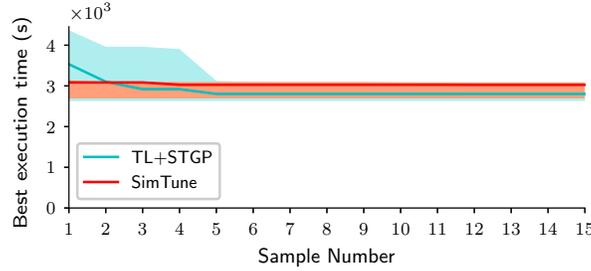
Finally, the algorithm overhead is negligible: a few seconds to select the most similar workload, transfer the significant parameters and pick the next configuration during tuning. This is due to the small number of samples that we use during the MTGP optimization.
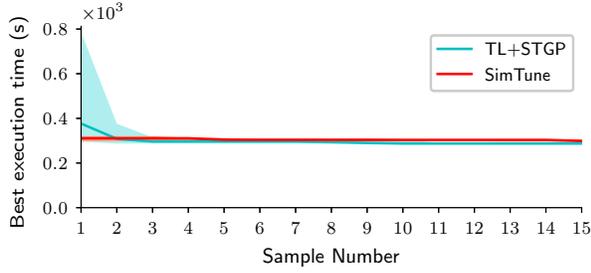
(a) Bayes-DS2
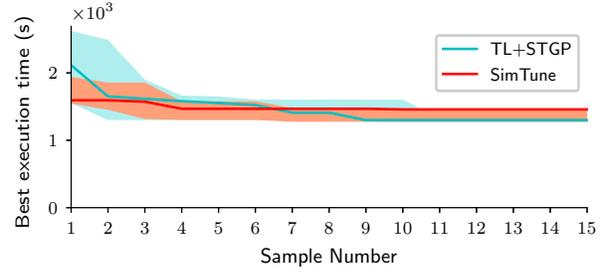
(b) Bayes-DS3

(c) Bayes-DS4

Fig. 7: Convergence speed of transfer learning and STGP against SimTune. The shaded area represent the space between the 90th and 10th percentile of execution time.

Single tasked versus multitasked tuning: In order to evaluate the benefits of MTGP over STGP, we experiment with TL+STGP for three evolving input sizes of the Bayes workloads. This is to show the benefits of sharing the tuning experience (execution samples tried during tuning) of a similar workload in limiting the exploration time. By leveraging similarity-aware MTGP optimization, as shown in Figure 7, at the 90th percentile TL+STGP tries configurations that are more costly with a wider search interval. SimTune finds configuration close to the ones found by TL+STGP (within 10%) while not only avoiding the cold start problem of TL+STGP but also bounding the search interval overall. It is worth noting that the closer the source workload is from the target workload (e.g., Bayes-DS2), the smaller the interpercentile range of SimTune and the faster the convergence to near-optimal configuration. This suggests that having more seen workloads in the auxiliary set has the potential to save more search time.

Our similarity-aware tuning accelerates the amortization of tuning costs: the previous experiments do not show the full story on how the different approaches compare in behaviour as they perform incremental tuning from one execution to the next. For that, it is useful to have a timeline view. Figure 8 shows the cumulative execution time of running Bayes-DS3 workload over multiple configurations, as determined iteratively by the tuning algorithms considered. Here, we compare against using the directly transferred configuration from the source workload, since it represents the static tuning approach
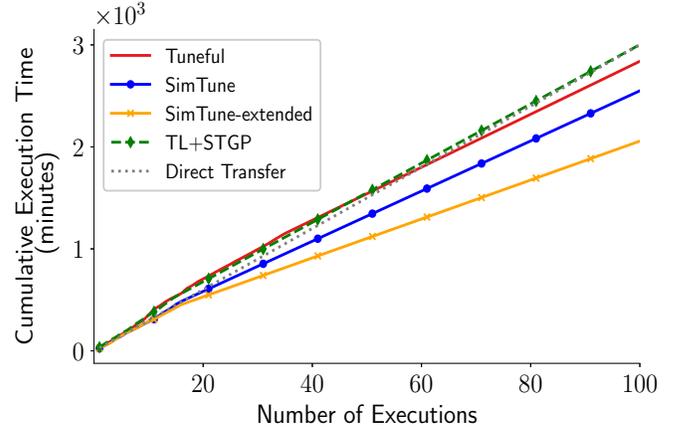

Fig. 8: Amortization speed of the different tuning approaches (Bayes-DS3 workload).

followed by existing tuning approaches, in which the configuration is tuned once and reused- ignoring the need for configuration retuning. The dotted line shows cumulative execution time for this configuration *without any tuning*. The dynamic tuning "pays off" only after the lines intersect the dotted line, even if good configurations were found much earlier (finding them required exploring some worse configurations). Better configurations are shown as lines with shallower slopes (e.g. SimTune-extended, while equivalent configurations appear as parallel lines (e.g. SimTune and Tuneful). Our proposed tuning

approach runs a fingerprinting execution sample once and tunes the configurations incrementally within 15 executions, then picks the best configuration it found and continues only using that.

As shown in Figure 8, it takes less than 5 executions to amortize the tuning cost using SimTune. This enables a quicker adoption of the need for workload retuning in a rapidly changing environment such as the big data environment. On the other hand, other approaches have higher tuning cost and need higher number of executions to amortize this cost (e.g. 60 executions).

This suggests that cloud providers, if able to observe executions similarity across clients, would be in the ideal position to offer tuned-configuration-as-a-service in a way that minimizes costs, even for workloads that are repeated just a couple of times.

V. Related Work

Some work has been proposed to characterize workloads, aiming to either detect similarity across workloads [36], come up with representative benchmarks of workloads [37], or reveal existing performance issues [38]. Recently, a few work has been proposed to characterize and cluster Hadoop and Spark analytics workloads, with the majority of work employing PCA to identify the most important metrics to lower the dimensionality of the captured execution metrics [37]. Besides finding the important execution metrics, some work has been proposed to cluster the workloads using KMeans, grouping the various workloads based on their execution patterns [36], [37]. Jia et al. [39] characterize data analytics workloads (most of them are Hadoop based workloads) in data centers based on their microarchitectural characteristics (e.g., L1 instruction cache misses per thousand instructions). They concluded that data analytics workloads have inherent different characteristics from traditional workloads (e.g., HPC workloads). Similarly, some work has been proposed characterizing in-memory analytics workloads [40], [41]. Jiang et al. [40] characterize Spark workloads and compare them against Hadoop and traditional HPC workloads. Their work shed light on the significant differences in Spark in terms of memory utilization, memory access and disk I/O frequency. Awan et al. [41] highlight thread scalability issues in Spark (i.e., Spark does not scale linearly when more than twelve threads are running). Chiba et al. [38] characterize the memory, network, JVM, and GC usage to tune the performance TPC-H workloads on Spark.

On the configuration tuning front, several solutions have been proposed for tuning the configurations of Spark and Hadoop workloads. These solutions are either search-based leveraging techniques such as hill climbing [10] and genetic algorithm [8], or model-based, wherein a performance prediction model is built to guide the tuning using SVM [11], regression trees [42], or hierarchical models [15]. Some configuration tuning systems have leveraged similarity to accelerate the tuning process: AROMA [11] clusters the Hadoop workloads then builds a performance model that guides the tuning of each workload cluster. Scout [43] and Vanir [44] exploit workload similarity to explore the search space more effectively for tuning cloud configurations (number of instances and their resource allocation). Ultimately, a joint solution for optimizing cloud instance and DISC framework configurations will be needed. Lastly, OtterTune [45] utilizes similarity to guide the tuning of similar workloads in DBMS.

None of the existing work studied the need for configuration retuning in a dynamic rapidly changing environment and how to address that efficiently in the high-dimensional configuration space of in-memory data analytics such as Spark workloads. We differ from this earlier work since we go a step further than characterizing the big data analytics workloads or tuning the configurations. Our work focuses on *how* to utilize similarity to accelerate the configuration tuning/retuning of in-memory data analytics workloads. SimTune shares a *wider* tuning knowledge across similar workloads and compares against existing similarity-aware solutions that are either based on static tuning (direct transfer) or transfer learning (TL+STGP).

VI. Conclusion

We presented SimTune, an approach for similarity-aware configuration tuning through leveraging workload characterization and MTBO. We showed that SimTune significantly reduces the exploration cost and accelerates the amortization of tuning costs, while finding configurations that are comparable the ones from the-state-of-the-art. We illustrated how SimTune works both when limited tuning knowledge is available and after extending the tuning knowledge, enabling the configuration tuning to happen faster while finding configuration closer to the optimal.

VII. Acknowledgement

References

[1] "Growth forecast for the worldwide big data and business analytics market through 2020," 2015, https://www.idc.com/getdoc.jsp?containerId=prUS41826116.

[2] "Apache Hadoop," 2016, http://hadoop.apache.org/.

[3] "Apache Spark: fast and general engine for large-scale data processing," 2015, https://spark.apache.org/.

[4] "Apache Flink," 2016, http://flink.apache.org/.

[5] R. Guo, Y. Zhao, Q. Zou, X. Fang, and S. Peng, "Bioinformatics applications on Apache Spark," *GigaScience*, vol. 7, no. 8, p. giy098, 2018. [Online]. Available: http://dx.doi.org/10.1093/gigascience/giy098

[6] K. Jacobs and K. Surdy, "Apache Flink: Distributed Stream Data Processing," Tech. Rep., 2016.

[7] Z. Zhang, K. Barbary, F. A. Nothaft, E. Sparks, O. Zahn, M. J. Franklin, D. A. Patterson, and S. Perlmutter, "Scientific computing meets big data technology: An astronomy use case," in *2015 IEEE International Conference on Big Data (Big Data)*. IEEE, 2015, pp. 918–927.

[8] G. Liao, K. Datta, and T. L. Willke, "Gunther: Search-based autotuning of MapReduce," in *European Conference on Parallel Processing*. Springer, 2013, pp. 406–419.

[9] H. Herodotou, H. Lim, G. Luo, N. Borisov, L. Dong, F. B. Cetin, and S. Babu, "Starfish: A self-tuning system for big data analytics." in *CIDR*, vol. 11, no. 2011, 2011, pp. 261–272.

[10] M. Li, L. Zeng, S. Meng, J. Tan, L. Zhang, A. R. Butt, and N. Fuller, "Mronline: Mapreduce online performance tuning," in *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing*. ACM, 2014, pp. 165–176.

[11] P. Lama and X. Zhou, "Aroma: Automated resource allocation and configuration of MapReduce environment in the cloud," in *Proceedings of the 9th international conference on Autonomic computing*. ACM, 2012, pp. 63–72.

[12] O. Alipourfard, H. H. Liu, J. Chen, S. Venkataraman, M. Yu, and M. Zhang, "Cherrypick: Adaptively unearthing the best cloud configurations for big data analytics." in *NSDI*, vol. 2, 2017, pp. 4–2.

[13] N. J. Yadwadkar, B. Hariharan, J. E. Gonzalez, B. Smith, and R. H. Katz, "Selecting the best VM across multiple public clouds: A data-driven performance modeling approach," in *Proceedings of the 2017 Symposium on Cloud Computing*. ACM, 2017, pp. 452–465.

[14] 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*. ACM, 2017, pp. 338–350.

[15] Z. Yu, Z. Bei, and X. Qian, "Datasize-aware high dimensional configurations auto-tuning of in-memory cluster computing," in *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2018, pp. 564–577.

[16] M. P. Deisenroth, D. Fox, and C. E. Rasmussen, "Gaussian processes for data-efficient learning in robotics and control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 2, pp. 408–423, 2015.

[17] "SimTune source code and dataset," 2020, https://github.com/ayat-khairy/simtune.

[18] K. Swersky, J. Snoek, and R. P. Adams, "Multi-task bayesian optimization," in *Advances in neural information processing systems*, 2013, pp. 2004–2012.

[19] J. Mockus, *Bayesian approach to global optimization: theory and applications*. Springer Science & Business Media, 2012, vol. 37.

[20] C. E. Rasmussen, "Gaussian processes in machine learning," in *Advanced lectures on machine learning*. Springer, 2004, pp. 63–71.

[21] J. Villemonteix, E. Vazquez, and E. Walter, "An informational approach to the global optimization of expensive-to-evaluate functions," *Journal of Global Optimization*, vol. 44, no. 4, p. 509, 2009.

[22] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," *arXiv preprint arXiv:0912.3995*, 2009.

[23] M. Hoffman, B. Shahriari, and N. Freitas, "On correlation and budget constraints in model-based bandit optimization with application to automatic machine learning," in *Artificial Intelligence and Statistics*, 2014, pp. 365–374.

[24] A. Fekry, L. Carata, T. Pasquier, A. Rice, and A. Hopper, "To tune or not to tune? in search of optimal configurations for data analytics," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ser. KDD '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 2494–2504. [Online]. Available: https://doi.org/10.1145/3394486.3403299

[25] 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.

[26] C. C. Aggarwal, A. Hinneburg, and D. A. Keim, "On the surprising behavior of distance metrics in high dimensional space," in *International conference on database theory*. Springer, 2001, pp. 420–434.

[27] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Advances in neural information processing systems*, 2012, pp. 2951–2959.

[28] "Hadoop distributed file system," 2018, https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html.

[29] "SimTune: configuration parameters under tuning," 2019, https://drive.google.com/file/d/1uHYGQjriw4IQyYTmOQukkWBYtIDD-4Tc/view?usp=sharing.

[30] S. Huang, J. Huang, J. Dai, T. Xie, and B. Huang, "The HiBench benchmark suite: Characterization of the MapReduce-based data analysis," in *Data Engineering Workshops (ICDEW), 2010 IEEE 26th International Conference on*. IEEE, 2010, pp. 41–51.

[31] "TPC-H SQL benchmark," 2014, http://www.tpc.org/tpch/.

[32] S. Krishnan and J. L. U. Gonzalez, "Google compute engine," in *Building Your Next Big Thing with Google Cloud Platform*. Springer, 2015, pp. 53–81.

[33] A. Fekry, L. Carata, T. Pasquier, A. Rice, and A. Hopper, "Tuneful: An online significance-aware configuration tuner for big data analytics," 2020, https://arxiv.org/pdf/2001.08002.pdf.

[34] I. M. Sobol, "On quasi-monte carlo integrations," *Mathematics and computers in simulation*, vol. 47, no. 2-5, pp. 103–112, 1998.

[35] "Amazon EC2 instance Pricing," 2018, https://aws.amazon.com/ec2/pricing/on-demand/.

[36] S. Aggarwal, S. Phadke, and M. Bhandarkar, "Characterization of hadoop jobs using unsupervised learning," in *2010 IEEE Second International Conference on Cloud Computing Technology and Science*. IEEE, 2010, pp. 748–753.

[37] Z. Jia, J. Zhan, L. Wang, R. Han, S. A. McKee, Q. Yang, C. Luo, and J. Li, "Characterizing and subsetting big data workloads," in *2014 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 2014, pp. 191–201.

[38] T. Chiba and T. Onodera, "Workload characterization and optimization of TPC-H queries on Apache Spark," in *2016 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 2016, pp. 112–121.

[39] Z. Jia, L. Wang, J. Zhan, L. Zhang, and C. Luo, "Characterizing data analysis workloads in data centers," in *2013 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 2013, pp. 66–76.

[40] T. Jiang, Q. Zhang, R. Hou, L. Chai, S. A. Mckee, Z. Jia, and N. Sun, "Understanding the behavior of in-memory computing workloads," in *2014 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 2014, pp. 22–30.

[41] A. J. Awan, M. Brorsson, V. Vlassov, and E. Ayguade, "Performance characterization of in-memory data analytics on a modern cloud server," in *2015 IEEE Fifth International Conference on Big Data and Cloud Computing*. IEEE, 2015, pp. 1–8.

[42] G. Wang, J. Xu, and B. He, "A novel method for tuning configuration parameters of Spark based on machine learning," in *High Performance Computing and Communications; IEEE 14th International Conference on Smart City; IEEE 2nd International Conference on Data Science and Systems (HPCC/SmartCity/DSS), 2016 IEEE 18th International Conference on*. IEEE, 2016, pp. 586–593.

[43] C.-J. Hsu, V. Nair, T. Menzies, and V. W. Freeh, "Scout: An experienced guide to find the best cloud configuration," *arXiv preprint arXiv:1803.01296*, 2018.

[44] M. Bilal, M. Canini, and R. Rodrigues, "Finding the right cloud configuration for analytics clusters," in *Proceedings of the 11th ACM Symposium on Cloud Computing*, ser. SoCC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 208–222. [Online]. Available: https://doi.org/10.1145/3419111.3421305

[45] 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*. ACM, 2017, pp. 1009–1024.