

# Lynceus: Cost-efficient Tuning and Provisioning of Data Analytic Jobs

Maria Casimiro<sup>1,4</sup>, Diego Didona<sup>2</sup>, Paolo Romano<sup>1</sup>, Luís Rodrigues<sup>1</sup>, Willy Zwaenepoel<sup>3</sup>, David Garlan<sup>4</sup>

<sup>1</sup>INESC-ID and Instituto Superior Técnico, Universidade de Lisboa

<sup>2</sup>IBM Research Zurich, <sup>3</sup>University of Sydney, <sup>4</sup>Carnegie Mellon University

**Abstract**—Modern data analytic and machine learning jobs find in the cloud a natural deployment platform to satisfy their notoriously large resource requirements. Yet, to achieve cost efficiency, it is crucial to identify a deployment configuration that satisfies user-defined QoS constraints (e.g., on execution time), while avoiding unnecessary over-provisioning.

This paper introduces Lynceus, a new approach for the optimization of cloud-based data analytic jobs that improves over state-of-the-art approaches by enabling significant cost savings both in terms of the final recommended configuration and of the optimization process used to recommend configurations.

Unlike existing solutions, Lynceus optimizes in a joint fashion both the cloud-related (i.e., which and how many machines to provision) and the application-level (e.g. the hyper-parameters of a machine learning algorithm) parameters. This allows for a reduction of the cost of recommended configurations by up to  $3.7\times$  at the 90-th percentile with respect to existing approaches, which treat the optimization of cloud-related and application-level parameters as two independent problems.

Further, Lynceus reduces the cost of the optimization process (i.e., the cloud cost incurred for testing configurations) by up to  $11\times$ . Such an improvement is achieved thanks to two mechanisms: i) a timeout approach which allows to abort the exploration of configurations that are deemed suboptimal, while still extracting useful information to guide future explorations and to improve its predictive model — differently from recent works, which either incur the full cost for testing suboptimal configurations or are unable to extract any knowledge from aborted runs; ii) a *long-sighted and budget-aware* technique that determines which configurations to test by predicting the long-term impact of each exploration — unlike state-of-the-art approaches for the optimization of cloud jobs, which adopt greedy optimization methods.

**Index Terms**—cloud computing, machine learning platforms, optimization, virtual machines, Bayesian optimization

## I. INTRODUCTION

Many enterprises run data analytic jobs in the cloud, such as training deep neural networks or building recommender systems. This sort of jobs is known to require a very large amount of computational resources and recent studies, e.g., [1], have shown that training large AI models can produce five times the lifetime emissions of the average American car (including the manufacturing of the car) and incur cloud costs of up to 3 million USDs. Further, data analytic jobs are often recurrent, i.e., they execute multiple times on similar datasets, with similar performances [2], [3].

As such, to reduce operational costs, it is crucial to ensure that jobs are deployed over the cheapest set of cloud resources

that complies with user specific constraints, e.g., on job execution time [4], [5], i.e., it is crucial to optimize the cloud provisioning process so as to avoid over-provisioning. Additionally, the efficiency of data analytic jobs can be substantially affected by the correct tuning of a multitude of application-level parameters — e.g., the hyper-parameters of a machine learning (ML) model can influence its training time [6]. With hundreds or thousands of possible combinations of cloud platform and job parameters, it is extremely challenging to identify the configuration that minimizes the provisioning cost and meets the target performance constraints.

**Existing solutions and their limitations.** State-of-the-art approaches to optimize the deployment of cloud jobs rely on either online or offline learning approaches to find (near) optimal cloud configurations.

Offline learning techniques require the availability of large training sets, collected by profiling different applications [7]–[9], or rely on *a priori* knowledge about the internal structure of the job [10]–[12]. These approaches either impose an expensive and time-consuming offline training phase, or require expert domain knowledge to model the performance of a job. We are instead interested in approaches that require no prior knowledge on the target job or other jobs. Therefore in this work we focus on online approaches.

Bayesian Optimization (BO) is a well-established online approach to tackle complex optimization problems [13], [14] and has recently emerged as a prominent solution to optimize the execution of data analytic jobs [2], [15], [16]. BO approaches profile the job on different configurations iteratively, building at each step a statistical performance model of the job. This model is then used to decide the next configuration to try, and ultimately to identify the best configuration for the job. Unfortunately, BO-based approaches targeting the optimization of cloud jobs suffer from several critical limitations.

1) *Disjoint optimization of cloud and application parameters.* Existing BO-based approaches [2], [15], [16] treat the optimization of cloud-related and application-level parameters as independent problems, thus neglecting the existence of important inter-dependencies between cloud and application configurations. Note that this limitation also affects existing offline approaches [7]–[9]. In Section III, we quantify the relevance of adopting a joint, cross-layer optimization approach by means of an experimental study based on three ML jobs, each deployed over 384 (cloud and application level) configurations.

Our study shows that approaches that optimize application and cloud parameters in a disjoint fashion are largely sub-optimal: they find the globally optimal configuration less than 50% of the times; further, the 90-th percentile of the cost of the recommend solutions is from  $1.2\times$  to  $3.7\times$  larger than the global optimum.

2) *Myopic optimization policy.* BO techniques employed by the state-of-the-art solutions [2] demonstrate significant limitations due to their *short-sighted* nature. In fact, at each step of the optimization process, existing solutions profile the configuration that is expected to maximize an immediate reward, such as the Expected Improvement (EI) or the model’s accuracy [17]–[19]. Such greedy approaches are likely to lead to a sub-optimal exploration of the configuration space and require testing a large number of configurations.

3) *Sampling sub-optimal configurations.* BO techniques can lead to exploring sub-optimal configurations, especially in the early stages of the optimization process when the model still has very limited information on the job. Existing BO-based approaches [20], [21] address this problem by cancelling the exploration of configurations that are detected to be of lower quality w.r.t. the best configuration identified so far and disregarding any performance metrics obtained during that exploration. However, such an approach suffers from a major drawback: simply cancelling the exploration translates into redundant computational time and money wasted, since the model will not learn from these explorations. This can lead to impoverishing the knowledge of the model on large portions of the configuration space, which has detrimental effects on its accuracy and, as such, on its ability to recommend high quality configurations.

**Lynceus.** This paper presents Lynceus (“Lynx-eyed”), an innovative tool to provision and tune data analytic jobs on the cloud. Lynceus addresses the challenges identified above by combining the following novel features.

First, Lynceus adopts a cross-layer, holistic approach that optimizes the parameters controlling the cloud deployment as well as the ones defining the application-level configuration *in a joint fashion* and eschews the need for any *a priori* knowledge about the target job.

Second, Lynceus introduces a novel *long-sighted* and *budget-aware* optimization method. Differently from existing *greedy* BO-based techniques that maximize a one-step reward, Lynceus plans which configurations to explore by simulating several *exploration paths*, i.e., sequences of configurations to sample sequentially. For each path, Lynceus estimates the path’s expected exploration cost and the advantages stemming from its exploration (i.e., improvement of model’s accuracy and/or of the currently known optimum). While simulating, Lynceus keeps into account predefined constraints both on the configuration’s performance and on the cumulative cost of the exploration. By simulating the outcomes of sampling a sequence of configurations, Lynceus determines more cost-effective ways to explore the configuration space. As we will show in Section VI-F, this technique allows Lynceus to the

reduce the cost incurred to optimize the configuration of ML jobs by up to  $6.3\times$ .

Third, Lynceus uses a new method to cope with the exploration of configurations that turn out to be sub-optimal. Unlike previous approaches [20], [21], Lynceus not only cancels the sampling of sub-optimal configurations to save money and time, but it also exploits the information on *when* such configurations exceeded the cost of the current optimum. Lynceus derives an updated prediction of the actual cost of the job with the sub-optimal configuration. This new prediction leverages the original model’s prediction of the execution cost for that configuration, and the time at which sampling was interrupted. This updated prediction is then fed back to the model, which allows for effectively enhancing its knowledge on the regions close to sub-optimal configurations. The use of this technique allows Lynceus to further enhance the cost effectiveness of its own optimization process by up to  $1.8\times$  w.r.t. existing approaches that discard information on sub-optimal configurations.

Overall, when combining its innovative features, Lynceus can reduce the 90-th percentile of the cost incurred to find a solution within 10% from the optimum by up to  $11\times$  when compared to state-of-the-art approaches. The experimental results reported in this work were obtained using both existing datasets [2], [16], as well as new datasets obtained by exhaustively deploying three TensorFlow jobs (distributed training of neural networks) over a large 5-dimensional configuration space encompassing 384 configurations.

**Contributions:** We make five main contributions:

- I) We propose Lynceus, a novel *long-sighted and budget-aware* approach to the tuning and provisioning of data analytic jobs, which we will make available as open source;
- II) We demonstrate the advantages of optimizing both cloud and application parameters jointly;
- III) We develop a new method for extracting useful information from the partial exploration of sub-optimal configurations;
- IV) We quantify the gains achievable by Lynceus via an extensive experimental study based on 26 diverse jobs;
- V) We make available to the systems’ community a dataset encompassing three Tensorflow jobs deployed on EC2, each including 384 configurations defined over 5 dimensions [22].

## II. RELATED WORK

This section discusses four main kinds of related work: *i)* systems to tune and provision data analytic jobs; *ii)* systems to optimize cloud applications; *iii)* BO approaches to tune generic applications; and *iv)* variants of the BO approach.

**Optimization of data analytic jobs.** Elastizer [10], ARIA [23] and MRTuner [11] model the internals of map-reduce jobs and obtain performance models to tune and provision them. Cumulon [24] targets matrix-based big data analysis jobs. Ernest [12] optimizes diverse job types but requires knowledge about the structure of the internal workflow of jobs, e.g., the communication pattern. Scout [16] exploits the availability of historical information on previous cloud jobs to enable transfer learning

and navigate through the search space more effectively. Unlike these approaches, Lynceus needs no *a priori* information about the target job or other jobs. CherryPick [2] and Arrow [15] rely on a greedy BO approach to select the best cloud infrastructure for a job. We discuss the limitations of such an approach in more detail in Section III and quantify them in Section VI. In contrast, Lynceus implements a novel *long-sighted* and *budget-aware* BO approach to achieve higher accuracy and better cost-efficiency. In addition, Lynceus tackles jointly the problems of selecting the best cloud infrastructure and optimizing the job’s tuning parameters.

**Optimization of cloud applications.** Paragon [25], Quasar [7], Selecta [8] and Paris [9] optimize the choice of the infrastructure for cloud applications. These systems employ black-box approaches to performance prediction that rely on the availability of abundant training data on different applications. Lynceus targets scenarios in which such data is not available, and requires running only the target job to infer its performance-cost function.

**BO approaches to tuning systems.** iTuned [20], Otter-tune [26], ProteusTM [27] and Metis [28] use BO approaches to optimize the tuning parameters of data platforms. These systems use the traditional BO approach, whose limitations we discuss in detail in Section III. Adding to these limitations, ProteusTM requires the availability of previous performance traces of other applications. Differently from the previous approaches, BOAT [29] extends BO to allow system experts to provide a probabilistic performance model of the target application, so as to speed up the optimization phase. This approach requires expert domain knowledge on the target application. Instead, Lynceus embraces a full black-box approach based on a novel long-sighted and budget-aware BO approach and does not require previous performance traces.

**BO with look-ahead.** The ML community has recently proposed non-greedy BO variants that rely on a *look-ahead* scheme that takes into account future steps in the exploration of the configuration space [17]–[19]. These approaches target the optimization of hyper-parameters of ML models, where testing any configuration has a unitary cost, and there is a fixed budget for exploring which is expressed in terms of the number of configurations that can be tested. Lynceus draws from these approaches, but augments them to capture specific idiosyncrasies of cloud environments, and hence to make them suitable in the context of job optimization in the cloud.

In particular, in cloud environments, testing different configurations results in different costs. Note that the cost of exploring a configuration depends on the duration of the job running in that configuration and is not known *a priori*. This is particularly relevant when there is a fixed budget for the optimization process, since it is not known how different explorations will affect the budget available for future explorations. Lynceus copes with this challenge by employing a black-box predictive model to estimate these costs. Further, Lynceus avoids wasting budget in testing suboptimal configurations. It achieves this goal via a novel technique to early stop the execution of a job

on suboptimal configurations, while still being able to leverage the knowledge attained from the partially completed job to increase the quality of the model.

### III. PROBLEM FORMALIZATION AND CHALLENGES

Lynceus seeks to find the optimal configuration to run a job, while meeting a target performance constraint. A configuration  $x$  is a tuple  $\langle N, H, P \rangle$ , where  $N$  is the number of VMs rented from the cloud provider,  $H$  encodes the hardware characteristics of the VMs (e.g., VCPUs and RAM), and  $P$  represents the settings of job-specific tuning parameters (e.g., hyper-parameters of a ML training process). We define the optimal configuration  $x^*$  as the one that minimizes the (monetary) cost of executing the job, and that is able to finish it in at most  $T_{max}$  time. The cost of executing the job with configuration  $x$ , noted  $C(x)$ , is given by the product of the time taken to run the job with  $x$ , noted  $T(x)$ , and the price per unit of time of renting the cloud configuration  $x$ , noted  $U(x)$ . We assume a pay-by-the-minute/second pricing scheme, which is typical nowadays in major PaaS infrastructures [30]–[32].

The optimization process relies on profiling the target job on a subset of configurations. We note such sub-set  $S$ , and we denote by  $C_S$  the cumulative cost of running the job on the configurations in  $S$ . Furthermore, we consider an additional constraint on the maximum cost of the profiling phase,  $C_S$ , which must not exceed a budget  $B$ . The problem can then be formalized as follows:


$$\begin{cases} \min C(x) \\ \text{s.t. } T(x) \leq T_{max} \\ \text{s.t. } C_S \leq B \end{cases}$$

Let us now discuss the key challenges to deriving an efficient and practical solution to this optimization problem.


**I) Lack of *a priori* information.** Given the heterogeneity and complexity of modern data analytic jobs, building white-box models capable of accurately predicting their performances, independently of their nature, is not a plausible solution in practice. Moreover, gathering data concerning previous optimizations of similar jobs can be too costly or impractical. In order to circumvent these issues, we advocate optimization methods that ensure two key properties.

**Black box approach.** *The optimization process should assume no knowledge about the target job, nor about the cloud infrastructure.* In fact, jobs can have very different structures (e.g., map-reduce vs parameter-server) [10], [12], and modeling the performance of cloud infrastructures is notoriously a complex task [33]. A black-box approach to the job optimization process reduces the modeling effort and is more flexible.

**No available data.** *The optimization process should not rely on a priori performance information about other jobs.* As noted before, some existing techniques to optimize application performance rely on the availability of huge training data [7], [8]. This information helps in bootstrapping and improving the optimization process. Unfortunately, obtaining large amounts of training data is very costly and time-consuming. Hence,



(a) Normalized cost of training 3 Neural Network models with Tensorflow, depending on the configuration (y axis in log scale).



(b) CDF of the normalized cost (w.r.t. the optimum) achievable by optimizing hyper-parameters and cloud configuration separately.

Figure 1: Challenges of job optimization. Costs are normalized w.r.t. the cost of the optimal configuration (CNO). a) Few configurations are close-to-optimal and many are highly sub-optimal. b) Disjoint optimization may lead to identifying sub-optimal configurations.

such approach is fit for large service providers (such as Amazon, Google and Azure), but constrains and is impractical for most cloud users (e.g., small and medium enterprises).

**II) Complexity of the optimization process.** The plethora of VMs offered by cloud providers, along with the multitude of tunable application-level parameters, generate a search space with hundreds of configurations, with largely different performances. Next, we present empirical data that demonstrates: i) the complexity of the problem at hand, and ii) the necessity for tuning application and cloud parameters in a joint fashion.

**Very few close-to-optimal configurations.** *The configuration space includes few close-to-optimal configurations and many highly sub-optimal ones.* To quantify the complexity of finding the optimal configuration of modern cloud jobs, we measured the performance of training three ML models (Multilayer, CNN and RNN) with Tensorflow on AWS, while varying the cloud infrastructure and job hyper-parameters. In total, we considered 384 configurations. More details about these experiments are provided in Section VI. Figure 1a shows the cost of running a job in each configuration, normalized w.r.t. the cost of the optimal configuration. Note that the cost of a “bad” configuration can be 3 orders of magnitude worse than the cost of the optimal one. In addition, depending on the job, only 5 to 20 configurations have a cost within a factor of two w.r.t. the optimal one. These configurations correspond to 1.5% and 5% of the size of the configuration space, respectively.

**The need for joint optimization.** *The cloud infrastructure and the hyper-parameters of the job must be optimized simultaneously.* An approach to simplify the optimization process could be optimizing these two aspects separately, as done by recent systems [7], [25]. This approach, that we call disjoint optimization, first finds the optimal hyper-parameters by profiling the job on a reference cloud infrastructure  $c^\dagger$ , and then finds the optimal cloud settings for the job running with these parameters. Disjoint optimization, however, implicitly assumes that the optimal hyper-parameters for  $c^\dagger$  are also optimal for other cloud settings. In reality, this is usually not the case. Therefore, disjoint optimization is prone to missing the best combination of hyper-parameters and cloud settings.

To illustrate this fact we apply disjoint optimization to our jobs using all possible configurations as  $c^\dagger$ , and we measure the cost of the configuration that is identified as optimal.

We note that in this experiment we assume that both (i) the hyper-parameter optimization on  $c^\dagger$  and (ii) the subsequent optimization of the cloud configuration are always able to identify the best solution. Hence, these results are an upper bound on the effectiveness of any practical solution using disjoint optimization. Figure 1b reports the CDF of the cost of the configuration identified via this ideal disjoint optimization (using different choices for the initial reference configuration  $c^\dagger$ ), normalized w.r.t. the cost of the actual optimal one. For all jobs, disjoint optimization finds the overall optimal configuration less than 50% of the times. The 50-th percentile of the normalized cost obtained ranges from 1.2 to 2, and the 90-th percentile from 1.2 to 3.7, depending on the job.

#### IV. BACKGROUND ON BAYESIAN OPTIMIZATION

Bayesian Optimization (BO) is a sequential strategy to find the optimum of a function  $f$  with an unknown closed form and whose evaluation is expensive [13], [14].

**BO operations.** BO builds a statistical model of  $f$  iteratively as follows: (i) evaluate  $f$  on a set of initial points  $x_1 \dots x_n$  and create a training set  $S$  with the pairs  $\langle x_i, f(x_i) \rangle$ ; (ii) build a model  $M$  over  $S$  with a regression algorithm; (iii) use an *acquisition function* to determine the next point  $x_m$  to evaluate; (iv) evaluate  $f(x_m)$ , and update  $S$  and  $M$ ; (v) repeat steps (ii) to (iv) until a stopping criterion is satisfied. In Lynceus, a point is a configuration, and the target function to minimize is the cost of running a job.

**Acquisition function.** Given the current model of  $f$ , the acquisition function determines which point to evaluate next, among the set of points that are not yet in  $S$ . The acquisition function used by Lynceus is based on the *constrained expected improvement* ( $EI_c$ ) [34]. The  $EI_c(x)$  for configuration  $x$  is computed as the product of the probability that  $x$  respects a given constraint, noted  $P_C(x)$ , and the Expected Improvement of  $x$ , noted  $EI(x)$ . As its name suggests,  $EI(x)$  estimates by how much configuration  $x$  is expected to improve over the currently known optimum. Such expectation is computed taking into account both the expected value of  $f(x)$  as predicted by the model, as well as the *uncertainty* of the model on this prediction.  $EI(x)$  can be computed in closed form, assuming that  $f$  follows a normal distribution [14]. Specifically,  $EI(x) = (y^* - \mu(x))\Phi(z) + \sigma(x)\phi(z)$ , where

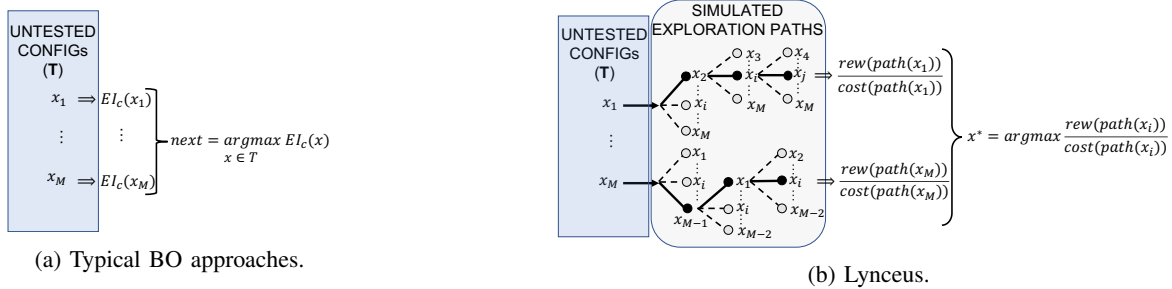


Figure 2: Selecting the next configuration in: (a) BO approaches; (b) Lynceus. BO approaches maximize a one-step acquisition function,  $EI_c$  in this case, that estimates the reward of sampling the *next* configuration. Lynceus, instead, speculates on different exploration paths, by *simulating* exploration paths. Dark/empty circles indicate the configurations selected/discarded at each step. For each configuration  $x$ , Lynceus computes the expected reward-to-cost ratio of the path that starts with  $x$ . Finally, it suggests the configuration  $x^*$  that maximizes the long-term expected reward-to-cost ratio.

$\mu(x)$ , resp.  $\sigma(x)$ , is the mean, resp. variance, of the prediction of the model of  $x$ ;  $\Phi$ , resp.,  $\phi$ , is the pdf, resp., CDF, of a standard normal distribution; and  $z = (y^* - \mu(x))/\sigma(x)$ .

In Lynceus,  $y^*$  is the cost of the cheapest configuration profiled so far such that running a job takes at most  $T_{max}$  time. If there is no such configuration,  $y^*$  is estimated as the cost of the most expensive configuration in  $S$  plus three times the maximum standard deviation over the predictions on the points not in  $S$  [18].

$P_C(x)$  can be computed by training a regression algorithm on the target constraint variable, whose value is known for each point in  $S$ . In Lynceus,  $P_C(x) = P(T(x) \leq T_{max})$ . Instead of training a separate model for  $T(x)$ , Lynceus reuses the model that it already builds for  $C(x)$ , by leveraging the fact that  $C(x) = T(x) \cdot U(x)$ , where  $U(x)$  is known. As such, rather than computing  $P(T(x) \leq T_{max})$ , Lynceus computes  $P(C(x) \leq T_{max} \cdot U(x))$ .

At each iteration, BO samples the configuration  $x \notin S$  that maximizes  $EI_c(x)$ .  $EI_c(x)$  has a high value not only if  $x$  is predicted –on average– to be a good point, but also if the uncertainty on  $y(x)$  is high. This allows for balancing *exploitation* (testing points that are considered good) and *exploration* (testing uncertain points) with the goal of improving the models’ quality.

**Regression model.** Computing  $EI_c$  in closed form requires a regression model that assigns to each point  $x$  a cdf  $p(x)$  that is normally distributed  $N(\mu(x), \sigma(x))$ . To meet this requirement, Lynceus uses a *bagging ensemble* [35] of decision trees, i.e., a set of decision trees, each trained over a uniform random sub-set of  $S^1$ . Then, Lynceus obtains  $\mu(x)$  and  $\sigma(x)$  based on the output of the individual predictors evaluated at  $x$ . Lynceus uses these values to compute the  $EI_c(x)$ , assuming that the  $p(x)$  associated with the ensemble of learners is normally distributed [36], [37].

**Stopping criterion.** Typical BO-based systems [2], [27], [38] stop the exploration phase once they detect that only marginal

improvements are predicted by the model, e.g., when the  $EI_c$  falls below 1% for all unexplored configurations. Lynceus supports this classic stopping criterion and complements it to keep into account user defined constraints on the maximum budget available for the exploration phase.

## V. LYNCEUS

Lynceus takes as input the budget  $B$ , the maximum job runtime  $T_{max}$  and a set  $H$  of possible configurations. Lynceus then proceeds in an iterative fashion, similarly to typical BO approaches. At each iteration, Lynceus indicates a new configuration on which to profile the job. Once the job completes, the corresponding cost and performance information are used to update the regression model. The budget is also reduced by the cost incurred to run the job on the configuration. Lynceus stops when there are no more configurations to try with the available budget, or if the  $EI_c$  for the unexplored configurations is marginal (below 1%). The configuration recommended in the end is the one, among those sampled by Lynceus, with the lowest cost and with runtime within  $T_{max}$ .

### A. Determining which configurations to sample

Figure 2 provides an overview of the selection process in typical BO approaches (Figure 2a) and in Lynceus (Figure 2b). At each iteration of the optimization process, Lynceus speculates about several *exploration paths*. Each path corresponds to a possible sequence of configurations to be explored. To select the best path, the outcomes of testing the configurations in each path are simulated using a black-box model. These simulation results are then used to compute the *reward* and the *cost* of each path.

The *reward* of a path corresponds, intuitively, to the aggregate reward resulting from exploring all the configurations in that path. The reward of a single configuration is given by its  $EI_c$ , that is, the cost improvement brought by that configuration, as predicted by the model, over the best configuration found so far. The *cost* of a path captures the predicted budget required to sample all configurations of the path. Finally, Lynceus explores the *first* configuration of the path with the

<sup>1</sup>Note that Lynceus can also operate using Gaussian Processes, as done by other BO approaches [13], [14]. We opted for a bagging ensemble of learners, since it offers more flexibility in the choice of the base learners to use.

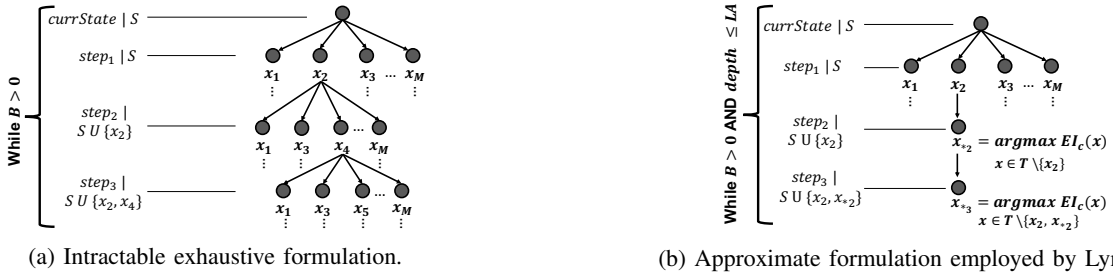


Figure 3: (a) exhaustive approach that enumerates the entire set of  $M!$  exploration paths, given a set of  $M$  unexplored configurations. (b) The search heuristics of Lynceus: (i) an in-breadth search, at step 1, spanning all untested configurations; (ii) an in-depth search, at step  $i > 1$ , which uses a BO-inspired principle and selects the configuration that maximizes the  $EI_c$ . The figure omits the technique used for incorporating the outcome of previous simulations in the model.

best reward/cost ratio. This approach renders the optimization process of Lynceus long-sighted and budget-aware.

*Long-sighted:* By analyzing in foresight a *sequence* of exploration steps (using a bounded look-ahead horizon) Lynceus defines effective exploration policies, which can intentionally sacrifice the immediate reward stemming from the next exploration in order to maximize the reward in the long term. This contrasts with existing BO approaches [17]–[19], which use a greedy policy that maximizes a one-step/myopic acquisition function (such as  $EI$ ).

*Budget-aware:* Lynceus dynamically adjusts its “explorative” nature depending on the budget currently available. Compared to conventional BO schemes, Lynceus tends to favor the exploration of uncertain configurations, provided that this does not compromise the budget available for future, less “risky” explorations. As a result, Lynceus adopts more explorative policies in the initial phase of the optimization process, when the model still has limited knowledge on the actual cost function and is, thus, more error prone. As the exploration progresses and the available budget diminishes, Lynceus tends to use a more risk-averse approach and to exploit the model’s knowledge to maximize shorter term rewards.

**Challenges and solutions.** Designing long-sighted optimization schemes, such as the one employed by Lynceus, requires tackling two main challenges. The first is related to the fact that the number of distinct exploration paths grows factorially with the unexplored configurations. As such, an idealized exhaustive approach, that analyzes all distinct exploration paths (illustrated in Figure 3a) would incur prohibitive computational costs in practical settings, forcing the use of approximations, i.e., search heuristics. The second challenge is tied to the simulation of the outcomes of exploration steps at depth  $i > 1$ . Such a simulation requires incorporating in the model used at step  $i$  the effects of performing all previous explorations at steps  $j < i$ . However, configuration  $x$  at step  $j$  was not actually tested, but only simulated via a (Gaussian) black-box model that associates a non-null probability to *any* possible cost value of  $x$ . To ensure that the effects of exploring configuration  $x$  at step  $j$  are taken into account in the model used at step  $j+1$ , it would be necessary to marginalize over all possible cost values, and corresponding probabilities, predicted

for  $x$  by the model at step  $j$ . Unfortunately, the closed form solution of such a nested marginalization problem implies prohibitive computational costs [39] even for two-steps look-ahead. Thus, approximation techniques are required to make the problem tractable.

Lynceus tackles the above challenges by means of three approximations, which ensure its scalability and viability.

1) The exploration paths considered by Lynceus are generated using a search heuristic that aims to balance the computational complexity of the optimization process and the effectiveness of the resulting exploration policy. This is achieved by using, in the first step, a breadth search policy that considers all untested configurations. At any subsequent step, instead, Lynceus employs a depth-first approach that selects the configuration that maximizes the  $EI_c$ , based on the current model’s state. This BO-inspired heuristic allows for pruning significantly the search space, as it avoids that a path branches to consider all cases corresponding to choosing each possible configuration for the next step (except in the first).

2) The in-depth simulation of a path is limited by a look-ahead window of size  $LA$ . Namely, the maximum length of an exploration path is limited to at most  $LA$  steps, in addition to the first one. A path can be shorter than  $LA$  steps in case the budget is depleted before reaching the  $LA$ -th step. If  $LA$  is 0, Lynceus collapses to the traditional BO approach, where a single-step reward is maximized. Figure 3b illustrates the combined use of these two heuristics.

3) To make the problem of simulating exploration paths mathematically tractable, Lynceus discretizes the cost distribution output by the black-box model using the Gaussian-Hermite (G-H) quadrature [40, Chapter 5.3]. The G-H quadrature is used to approximate the value of integrals of the form  $f(x)e^{-x^2}$  (such as the normal distribution that Lynceus associates with the outputs of its bagging regression model). The G-H quadrature produces  $K$  (value, weight) pairs associated with the approximated function. In Lynceus, each value is a cost, and each weight captures, roughly speaking, the likelihood of the corresponding cost.

With these approximations, Lynceus simulates only  $M$  paths ( $M$  being the number of unexplored configurations) of length  $LA$ . Thus, Lynceus’ complexity is  $\mathcal{O}(K^{LA})$  since the G-H

quadrature yields  $K$  sub-trees at each step, up to depth  $LA$ .

For space constraints we have to omit the pseudo-code detailing this part of Lynceus’s logic, which can, however, be found in the extended technical report [41].

### B. Sampling of sub-optimal configurations.

Inaccuracies in the model can lead to exploring sub-optimal configurations, whose sampling can take a significant amount of resources (both cost and time). Prior works in the literature on BO [20], [21] suggest coping with this issue by simply aborting the exploration of configurations that are found to be worse than the best configuration identified so far. While such a simplistic approach does allow for saving resources, it is easy to see that it also suffers from a major drawback: since the cost of the configuration is unknown (because the testing was terminated prematurely) no feedback is given to the model and, as such, the (economical and temporal) resources spent prior to canceling its sampling are wasted in vain.

**Challenges and solutions.** Addressing the previous limitation requires tackling two tightly intertwined challenges.

The first challenge is related to *how* to predict the cost of a configuration whose sampling has been timed out. Given the vast heterogeneity of modern cloud applications, we argue that, to maximize its interoperability, the technique used to perform this prediction should be fully generic and transparent, i.e., it should impose no requirements or make no assumption on the underlying application. This excludes, for instance, designs that require the application to externalize periodic information of its progress rate or of its expected termination time.

The second challenge is related to *when* to time out the sampling of a configuration: the later this is done, the more accurate the prediction can be (being fully accurate if, as an extreme, the sampling is not timed out at all), but also the smaller the gain in terms of money (and time) saved due to cancelling the configuration’s sampling. iTuned [20], for instance, takes a rather conservative approach and times out the sampling of a configuration once it has been in execution for twice as long as the fastest configuration found so far. Vizier [21], instead, adopts a more aggressive (and non-transparent) approach based on monitoring the progress rate of the application in the configuration under test, say  $x$ , and comparing it with prior runs in different configurations. If after  $t$  time units, the progress rate in  $x$  turns out to be slower than the median of the progress rate at time  $t$  for all the configurations tested so far,  $x$  is deemed sub-optimal and its exploration is cancelled. Note that, unlike Lynceus, neither iTuned nor Vizier provide any information on the quality of timed out configurations. Hence, the choice of when to time out the sampling of a configuration only affects the amount of resources that are spent. For Lynceus, the amount of information that the model gains about sampled configurations and its accuracy both depend on the timeout instant.

Compared to prior works, Lynceus addresses the latter problem by seeking a different trade-off regarding when to time out the sampling of a configuration  $x$ . Whenever the cost of  $x$  is found to exceed the cost, noted  $C(x^*)$ , of the cheapest

solution found so far (which must meet the user defined performance constraints), Lynceus times out the exploration. Since in the cloud the cost per unit of time of configuration  $x$ ,  $U(x)$ , is known *a priori*, it follows that the time out of the sampling of  $x$  will occur at time  $t = C(x^*)/U(x)$ . At this time, Lynceus knows that configuration  $x$  is not optimal.

At this point Lynceus relies on the black-box model it uses to estimate the cost of unknown configurations in order to predict the cost of the timed out configuration  $x$ . More precisely, the output distribution of the cost model, which follows a Gaussian distribution  $N(\mu(x), \sigma(x))$ , is conditioned to be strictly larger than  $C(x^*)$ . This yields a truncated Gaussian distribution [42], [43], whose expected value can be computed in closed form as:

$$E(C(x) \mid C(x) > C(x^*)) = \mu(x) + \sigma(x) \cdot \lambda(\alpha)$$

where  $\lambda(\alpha) = \frac{\phi(\alpha)}{1 - \Phi(\alpha)}$ ,  $\alpha = \frac{C(x^*) - \mu(x)}{\sigma(x)}$  and  $\phi/\Phi$  denote the pdf/CDF of a standard normal distribution, respectively. In Lynceus we use the expected value of this truncated Gaussian distribution, which is guaranteed to be larger than the cost of the current optimum,  $C(x^*)$ , to estimate the cost of  $x$  and feed this information back to the model. This allows Lynceus to, unlike previous work, leverage the knowledge attained from suboptimal explorations which are timed-out earlier to prevent resource exhaustion to update its knowledge base.

## VI. EVALUATION

Our evaluation addresses the following main questions:

- By how much can Lynceus reduce the cost of the optimization process w.r.t existing approaches (§ VI-C)?
- To what extent do the various features of Lynceus contribute to its effectiveness (§ VI-D)?
- How sensitive is Lynceus’ performance to the technique used to time out sub-optimal configurations and to the LA setting (§ VI-E and § VI-F)?
- What computational costs does Lynceus incur (§ VI-G)?

### A. Datasets

We consider two datasets of heterogeneous data analytic jobs. The first dataset is composed of three Tensorflow jobs, which are characterized by a large configuration space defined over 5 dimensions. The second dataset is composed of several Hadoop and Spark jobs that encompass smaller configuration spaces defined over 3 dimensions. These jobs have been used in the evaluation of the Scout [15] and CherryPick [2] systems.

**Tensorflow Dataset.** We consider the distributed training of three neural network models, i.e., CNN, RNN and Multilayer, over the MNIST dataset [44]. The networks were implemented with Tensorflow [45] using the parameter-server approach [46] and the ADAM optimizer [47]. The job terminates when the accuracy of the model reaches 0.85. We set a timeout of 10 minutes, after which a job is forcefully terminated.

Table I describes the tuning parameters that were considered. These include three hyper-parameters of the learning algorithm and two cloud-related parameters, yielding a total of 12 and 32 combinations of parameters, respectively. We

Parameter	Learning rate	Batch size	Training mode
Values	$\{10^{-3}, 10^{-4}, 10^{-5}\}$	{16, 256}	{sync, async}
VM type	VM characteristics	#VMs	
t2.small	{1 VCPU, 2 GB RAM}	{8, 16, 32, 48, 64, 80, 96, 112}	
t2.medium	{2 VCPU, 4 GB RAM}	{4, 8, 16, 24, 32, 40, 48, 56}	
t2.xlarge	{4 VCPU, 16 GB RAM}	{2, 4, 8, 12, 16, 20, 24, 28}	
t2.2xlarge	{8 VCPU, 32 GB RAM}	{1, 2, 4, 6, 8, 10, 12, 14}	

Table I: TensorFlow parameters and cloud configurations.

run our jobs on AWS EC2 and use 4 types of VMs. The VM clusters comprise 8 to 112 CPUs, for a total of 32 different cluster compositions. We did not use GPUs to train these models, since CPUs are known to be more cost-efficient than GPUs when training neural models with the MNIST dataset [48], [49]. Table I summarizes the cluster combinations that we use. Overall, the configuration space for these jobs is composed of a total of 384 configurations.

**Scout and CherryPick datasets.** The Scout dataset [16] is composed of 18 Hadoop and Spark jobs of the HiBench [50] and spark-ref [51] benchmarks. The CherryPick dataset [2] is composed of 5 jobs: TPC-H [52], TPC-DS [53], Terasort, Spark Kmeans [54], and Spark Regression [54]. These jobs stress CPU, network and memory resources differently, hence allowing us to evaluate Lynceus in heterogeneous use cases.

Both sets of jobs were run on AWS EC2, using different sets of VM types, but not varying any application-level parameter. Overall, the Scout dataset considers a total of 69 different configurations, whereas the configuration space for the CherryPick dataset ranges from 47 to 72 points. Additional details on the jobs can be found in the original papers.

## B. Methodology

**Compared systems.** We compare Lynceus with the traditional BO approach, used by state-of-the-art systems to optimize data analytic jobs, such as CherryPick [2] and Arrow [15]. We refer to this approach as BO. We also consider a simple random approach (RND) to establish a baseline on the complexity of the optimization task. We consider four values for the look-ahead parameter (LA) in Lynceus, i.e.,  $LA=\{0, 1, 2, 3\}$ .  $LA=0$  corresponds to a traditional BO approach, using as acquisition function the *EIc* per dollar [13], i.e., the ratio between the *EIc* and the expected cost of a configuration. Lynceus and BO use a bagging ensemble of 10 random trees to build the cost model of the job, as in recent BO systems [27], [36], [37].

**Experiments.** We perform our evaluation via a simulation approach, which uses the performance data of the Tensorflow, Scout and CherryPick datasets. In each experiment we run an optimizer 100 times against a target job. Each run uses a different set of initial configurations to bootstrap the model.

**Metrics.** We use the Cost Normalized w.r.t. the Optimum (CNO) to evaluate the quality of the configurations recommended by an optimizer. Noting  $x^*$  the optimal configuration, and  $x$  the configuration suggested by an optimizer, the CNO achieved by the optimizer is  $cost(x)/cost(x^*)$ . Hence, the lower the CNO, the better. The optimal value for CNO is 1. In

order to evaluate the cost-effectiveness of the optimizers we also measure the monetary cost consumed by each optimizer during the exploration phase (to deploy and sample a job in different cloud configurations).

**Budget.** To ensure a fair comparison with the baselines considered in this study, which do not consider any constraint on the exploration cost, we set the budget to infinity for Lynceus. By looking at how the budget is used over time it is possible to infer the behavior of Lynceus for different budget values.

**Default settings.** We set the initial number of samples,  $N$ , in a way that accounts for the size of the configuration space of each job. Specifically, noting with  $\mathcal{C}$  the job’s configuration space, we define  $N$  as the max of (i) 3% of the cardinality of  $\mathcal{C}$  (a percentage also used in previous works [27]) and (ii) the number of dimensions of  $\mathcal{C}$ . Unless stated otherwise, Lynceus uses  $LA=2$  and the Truncated Gaussian timeout policy. We evaluate different timeout policies in Section VI-E and lower values for  $LA$  in Section VI-F. We do not report results for larger  $LA$  values as, in our experiments, the gains deriving from setting  $LA = 3$  were marginal w.r.t.  $LA = 2$ . Finally, we set the time constraint for each job in such a way that it is satisfied by roughly half of the possible configurations.

## C. Cost of the optimization process

**Tensorflow jobs.** We evaluated the advantages of jointly optimizing cloud and application parameters in Section III. Thus, in the following, we assume, for fairness, that all the compared solutions are faced with the same configuration space that includes both cloud and application parameters.

Figure 4 reports the 90-th percentile of the CNO achieved by: Lynceus, BO with look-ahead (set to 2, as in Lynceus), BO without look-ahead and RND, as a function of the 90-th percentile of the exploration cost, for the three Tensorflow jobs. The costs corresponding to the initial sampling phase are not explicitly shown, as they are the same for all approaches, but are taken into account and added to the first cost represented in each plot. Lynceus consistently outperforms the baseline approaches by reaching the optimal configuration at a lower exploration cost. In particular, for Multilayer (Fig. 4c), while Lynceus reaches the optimal configuration after spending  $\sim \$19.5$ , the base BO technique requires spending  $\sim \$230$ , which corresponds to an improvement of more than  $11\times$ . For CNN (Fig. 4a) and RNN (Fig. 4b) the cost reduction to identify the optimum is lower, but still substantial, amounting to  $\sim 2\times$ .

Fig. 5 shows the CDF of the exploration cost to find a configuration that is either  $2\times$  (Fig. 5a) or 10% (Fig. 5b) away from the optimum, allowing us to assess the effectiveness of Lynceus in identifying configurations at different distances from the optimum. At the 90-th percentile, Lynceus spends  $\sim 2.7\times$  less than BO to identify configurations  $2\times$  away from the optimum (Fig. 5a), with gains up to  $\sim 6.2\times$  when considering the more challenging problem of identifying configurations that are only 10% worse than the optimum (Fig. 5b).

**Scout and CherryPick jobs.** Fig. 6 reports the CDFs of the exploration cost to identify configurations at 10% from the






Figure 4: 90-th percentile of the CNO achieved by Lynceus, BO with look-ahead 2, BO without look-ahead and RND for the Tensorflow jobs, as a function of the 90-th percentile of the exploration cost.

optimum, for the Scout and CherryPick datasets. At the 90-th percentile, the gains of Lynceus over BO remain remarkable but less pronounced than for the Tensorflow datasets – i.e., BO spends  $\sim 60\%$  and  $\sim 48\%$  more for the Scout and CherryPick datasets, respectively. This is due to the lower dimensionality of the search space (and thus to the lower complexity of the optimization problem), which decreases the benefits achievable by employing a more careful planning policy.

#### D. Breakdown of the improvements

Let us now quantify the benefits deriving from the two key novel features of Lynceus: i) the use of look-ahead to identify which configurations to sample, and ii) the timeout mechanism when sampling sub-optimal configurations. To this end, let us return to Figure 4. By comparing the base BO technique to the baseline using BO with LA=2, one can quantify the gains achieved using look-ahead. The benefits stemming from the timeout policy can instead be assessed by comparing Lynceus with the baseline of BO equipped with LA=2.

Overall, both mechanisms play an important role. While the use of look-ahead appears to be particularly relevant in the early stage of the optimization process, when the identified

configurations are still relatively distant from the optimum, the timeout tends to provide the largest gains the closer Lynceus is to the optimum. For instance when the CNO is equal to 2, the use of look-ahead allows for achieving a cost reduction ranging from 40% to 50%, across all networks. From that point on, the benefits of look-ahead, although still significant in RNN and CNN, tend to become less relevant when compared to the ones stemming from the use of timeout. This can be explained considering that, in the early stage of the optimization, planning ahead which configurations to sample allows for exploring the configuration space in a more cost-effective way. Once configurations closer to the optimum are found, the use of timeout becomes extremely effective by imposing a strict upper bound on the cost of future explorations.

#### E. Sensitivity to the timeout implementation

Figure 7 reports the 90-th percentile of the CNO, as a function of the 90-th percentile of the exploration cost when using various policies for estimating the full cost of running a job in a “timed out” configuration, namely: the Truncated Gaussian (TG) approach used by Lynceus; an approach (NO-INFO) similar to the one used in iTunes [20] and Vizier [21], where configurations are timed out if their cost exceeds by  $2\times$  the cost of the current optimum and where the model is provided with no information on timed out configurations; an Ideal approach that updates the model with the exact cost of fully executing the job; a Max Cost approach that feeds the model with the highest cost seen so far; a baseline that does not use the timeout feature (No); an approach (Linear) that assumes the availability of information on the job’s progress and uses a linear model to estimate its full cost. For Linear, as we consider jobs for training neural networks, we use as progress indicator the accuracy reached by the model upon timeout and linearly extend it until the desired target accuracy (85%). In order to evaluate the timeout feature in isolation, these results were obtained disabling the look-ahead.

The TG approach is consistently better than the others, getting very close to the Ideal timeout for the Multilayer and CNN networks (Figs. 7c and 7a). The Max cost and NO-INFO approaches are visibly worse than the others, showing that such simplistic approaches are clearly unfit and the relevance of feeding the model with accurate information when a configuration is timed out. The shortcomings of the Linear approach are also evident for Multilayer, where it

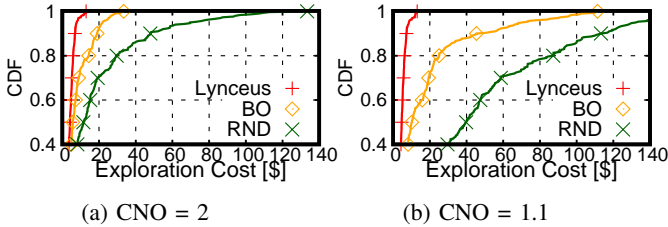


Figure 5: CDFs of the exploration cost for CNO=2 and CNO=1.1 (Tensorflow datasets).

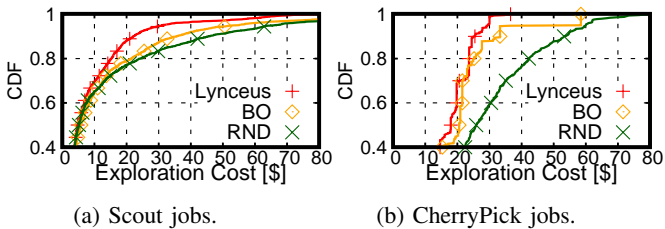


Figure 6: CDFs of the exploration cost for CNO = 1.1 (Scout and CherryPick datasets).




Figure 7: 90-th percentile of the CNO achieved by several versions of the timeout as a function of the 90-th percentile of the exploration cost. Lynceus’ method (TG) is the closest to an ideal approach that updates the model with exact values.




Figure 8: CDFs of the exploration cost achieved by Lynceus without the timeout feature and with  $LA=\{3, 2$  (default value),  $1, 0\}$ , for a CNO of 10%. The use of look-ahead provides clear benefits, but the gains diminish using horizons beyond 2.

clearly is outperformed by TG. We argue that this is due to the inadequacy of a linear model in predicting the time remaining to achieve the target accuracy for the job. While this issue might be tackled using more complex (non-linear) models, these data show that such a complexity is, in practice, unnecessary, since the proposed TG approach is very close to the ideal approach.

#### F. Sensitivity to the LA setting

Figure 8 compares the CDFs of the exploration cost achieved by Lynceus (which uses  $LA=2$ ) with the CDFs of three versions of Lynceus that use  $LA=\{3,1,0\}$ . The CDFs were obtained for a fixed CNO of 10% and without using the timeout mechanism. The plots allow us to draw two main conclusions. First, even small look-ahead horizons can significantly enhance the efficiency of the exploration process. For instance, at the 90-th percentile,  $LA=1$  allows for achieving a five-fold reduction of the exploration cost versus a greedy approach ( $LA=0$ ). Second, deeper look-ahead horizons tend to have diminishing gains, becoming marginal beyond  $LA=2$ . This is unsurprising, since the deeper the look-ahead horizon, the larger the probability that the model-based simulation of future explorations is stale, yielding limited benefits.

#### G. Prediction time

Table II reports the average time needed to predict the next configuration while varying the look-ahead’s depth. In particular, the table refers to RNN, but we have obtained similar values for CNN and Multilayer, since the cardinality of the search space is the same. We report results for the Tensorflow jobs, which have the largest configuration space among the jobs we consider and, as such, impose the largest computational costs. The simulations are run on machines with Intel Xeon Gold 6138 CPU with 20 physical cores and

Optimizer	BO, Lyn (LA=0)	Lyn (LA=1)	Lyn (LA=2)	Lyn (LA=3)
Avg seconds to next()	0.05	0.36	0.99	2.4

Table II: Average time needed to select the next config.

64 GB of main memory. As expected, Lynceus’ prediction time grows with the length of the look-ahead window. With  $LA=2$  (the default for Lynceus) the average computation time is around one second — a latency that we argue to be perfectly affordable in the context of data analytic jobs.

## VII. CONCLUSIONS AND FUTURE WORK

We presented Lynceus, a new tool to provision and tune data analytic jobs. Lynceus implements a novel approach that combines cross-layer optimization, budget awareness, long-sightedness, and the ability to cancel sub-optimal sampling while still improving the model. Lynceus consistently outperforms state-of-the-art approaches, identifying configurations that are up to  $3.7\times$  cheaper — thanks to the joint optimization of cloud and application parameters — and reducing the cost of the optimization process by up to  $11\times$  — thanks to its novel optimization method. As a final note, Lynceus can be extended to consider multiple constraints (e.g., one may want to enforce that the energy consumed to execute the job is also within a given threshold) and to take into account the costs associated with bootstrapping VMs during the exploration phase. An evaluation of these mechanisms is left for future work.

## REFERENCES

- [1] E. Strubell, A. Ganesh, and A. McCallum, “Energy and policy considerations for deep learning in NLP,” in *ACL*, 2019.
- [2] O. Alipourfard, H. H. Liu, J. Chen, S. Venkataraman, M. Yu, and M. Zhang, “Cherry-pick: Adaptively unearthing the best cloud configurations for big data analytics,” in *NSDI*, 2017.

- [3] Y. Zhang, G. Prekas, G. M. Fumarola, M. Fontoura, I. Goiri, and R. Bianchini, "History-based harvesting of spare cycles and storage in large-scale datacenters," in *OSDI*, 2016.
- [4] C. Curino, D. E. Difallah, C. Douglas, S. Krishnan, R. Ramakrishnan, and S. Rao, "Reservation-based scheduling: If you're late don't blame us!" in *SoCC*, 2014.
- [5] A. Tumanov, T. Zhu, J. W. Park, M. A. Kozuch, M. Harchol-Balter, and G. R. Ganger, "Tetrished: Global rescheduling with adaptive plan-ahead in dynamic heterogeneous clusters," in *EuroSys*, 2016.
- [6] S. Gupta, W. Zhang, and F. Wang, "Model accuracy and runtime tradeoff in distributed deep learning: A systematic study," in *IJCAI*. AAAI Press, 2017.
- [7] C. Delimitrou and C. Kozyrakis, "Quasar: Resource-efficient and qos-aware cluster management," in *ASPLOS*, 2014.
- [8] A. Klimovic, H. Litz, and C. Kozyrakis, "Selecta: Heterogeneous cloud storage configuration for data analytics," in *ATC*, 2018.
- [9] 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 *SoCC*, 2017.
- [10] H. Herodotou, F. Dong, and S. Babu, "No one (cluster) size fits all: Automatic cluster sizing for data-intensive analytics," in *SoCC*, 2011.
- [11] J. Shi, J. Zou, J. Lu, Z. Cao, S. Li, and C. Wang, "Mrtuner: A toolkit to enable holistic optimization for mapreduce jobs," *VLDB*, vol. 7, no. 13, 2014.
- [12] S. Venkataraman, Z. Yang, M. Franklin, B. Recht, and I. Stoica, "Ernest: Efficient performance prediction for large-scale advanced analytics," in *NSDI*, 2016.
- [13] E. Brochu, V. M. Cora, and N. de Freitas, "A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning," Tech. Rep. arXiv:1012.2599, 2010.
- [14] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *J. of Global Optimization*, vol. 13, no. 4, pp. 455–492, 1998.
- [15] C. Hsu, V. Nair, V. W. Freeh, and T. Menzies, "Arrow: Low-level augmented bayesian optimization for finding the best cloud VM," in *ICDCS*, 2018.
- [16] 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.
- [17] R. Lam, K. Willcox, and D. H. Wolpert, "Bayesian optimization with a finite budget: An approximate dynamic programming approach," in *NIPS*, 2016.
- [18] R. Lam and K. Willcox, "Lookahead bayesian optimization with inequality constraints," in *NIPS*, 2017.
- [19] J. Gonzalez, M. Osborne, and N. Lawrence, "Glasses: Relieving the myopia of bayesian optimisation," in *AIStats*, 2016.
- [20] S. Duan, V. Thummala, and S. Babu, "Tuning database configuration parameters with ituned," *VLDB*, vol. 2, no. 1, 2009.
- [21] D. Golovin, B. Solnik, S. Moitra, G. Kochanski, J. Karro, and D. Sculley, "Google vizier: A service for black-box optimization," in *SIGKDD*, 2017.
- [22] M. Casimiro, D. Didona, P. Romano, L. Rodrigues, W. Zwaenepoel and D. Garlan, "Tensorflow dataset," "link provided upon paper acceptance".
- [23] A. Verma, L. Cherkasova, and R. H. Campbell, "Aria: Automatic resource inference and allocation for mapreduce environments," in *ICAC*, 2011.
- [24] B. Huang, S. Babu, and J. Yang, "Cumulon: Optimizing statistical data analysis in the cloud," in *ACM SIGMOD*, 2013.
- [25] C. Delimitrou and C. Kozyrakis, "Paragon: Qos-aware scheduling for heterogeneous datacenters," in *ASPLOS*, 2013.
- [26] D. Van Aken, A. Pavlo, G. J. Gordon, and B. Zhang, "Automatic database management system tuning through large-scale machine learning," in *ACM SIGMOD/PODS*, 2017.
- [27] D. Didona, N. Diegues, A.-M. Kermarrec, R. Guerraoui, R. Neves, and P. Romano, "Proteustm: Abstraction meets performance in transactional memory," in *ASPLOS*, 2016.
- [28] Z. L. Li, C.-J. M. Liang, W. He, L. Zhu, W. Dai, J. Jiang, and G. Sun, "Metis: Robustly tuning tail latencies of cloud systems," in *ATC*, 2018.
- [29] V. Dalibard, M. Schaarschmidt, and E. Yoneki, "Boat: Building auto-tuners with structured bayesian optimization," in *WWW*, 2017.
- [30] Jeff Bar, AWS blog, "Per-Second Billing for EC2 Instances and EBS Volumes," <http://tinyurl.com/ycajp922>, 2017, accessed 4 Jan. 2020.
- [31] Microsoft Azure, "Azure Billing per minute," <https://azure.microsoft.com/en-us/pricing/details/cloud-services/>, 2018, accessed 4 Jan. 2020.
- [32] Paul Nash, "Per second billing in Google Cloud," <https://cloudplatform.googleblog.com/2017/09/extending-per-second-billing-in-google.html>, 2017, accessed 4 Jan. 2020.
- [33] H. Khazaei, J. Mistic, and V. B. Mistic, "A fine-grained performance model of cloud computing centers," *IEEE Transactions on Parallel and Distributed Systems*, vol. 24, no. 11, 2013.
- [34] J. R. Gardner, M. J. Kusner, Z. Xu, K. Q. Weinberger, and J. P. Cunningham, "Bayesian optimization with inequality constraints," in *ICML*, 2014.
- [35] L. Breiman, "Bagging predictors," in *Machine Learning*, vol. 24, no. 2. Springer, 1996, pp. 123–140.
- [36] F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Sequential model-based optimization for general algorithm configuration," in *LION*, 2011.
- [37] C. Thornton, F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Auto-weka: Combined selection and hyperparameter optimization of classification algorithms," in *SIGKDD*, 2013.
- [38] J. Zeng, P. Romano, J. Barreto, L. Rodrigues, and S. Haridi, "Online tuning of parallelism degree in parallel nesting transactional memory," in *IPDPS*, 2018.
- [39] M. A. Osborne, R. Garnett, and S. J. Roberts, "Gaussian processes for global optimization," in *LION*, 2009.
- [40] A. Gil, J. Segura, and N. M. Temme, *Numerical Methods for Special Functions*, 1st ed. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2007.
- [41] M. Casimiro, D. Didona, P. Romano, L. Rodrigues, W. Zwaenepoel, and D. Garlan, "Lynceus: Tuning and provisioning data analytic jobs on a budget," 2019. [Online]. Available: <http://arxiv.org/abs/1905.02119>
- [42] J. Burkardt, "The truncated normal distribution," *Department of Scientific Computing Website, Florida State University*, 2014.
- [43] D. R. Barr and E. T. Sherrill, "Mean and variance of truncated normal distributions," *The American Statistician*, vol. 53, no. 4, pp. 357–361, 1999.
- [44] L. Deng, "The mnist database of handwritten digit images for machine learning research [best of the web]," in *IEEE Signal Processing Magazine*, vol. 29, no. 6. IEEE, 2012.
- [45] M. Abadi *et al.*, "Tensorflow: A system for large-scale machine learning," in *OSDI*, 2016.
- [46] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, A. Ahmed, V. Josifovski, J. Long, E. J. Shekita, and B.-Y. Su, "Scaling distributed machine learning with the parameter server," in *OSDI*, 2014.
- [47] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [48] Z. Zhang, Y. Chen, and V. Saligrama, "Efficient training of very deep neural networks for supervised hashing," in *CVPR*, 2016, pp. 1487–1495.
- [49] Y. Wu, L. Liu, C. Pu, W. Cao, S. Sahin, W. Wei, and Q. Zhang, "A comparative measurement study of deep learning as a service framework," *IEEE Transactions on Services Computing*, 2019.
- [50] intel-hadoop, "HiBench Suite," <https://github.com/intel-hadoop/HiBench>, 2019, accessed 4 Jan. 2020.
- [51] Databricks, "Spark Performance Tests," <https://github.com/databricks/spark-perf>, 2019, accessed 4 Jan. 2020.
- [52] "TPC Benchmark H," <http://www.tpc.org/tpch/default.asp>, accessed 4 Jan. 2020.
- [53] "TPC Benchmark DS," <http://www.tpc.org/tpcds/default.asp>, accessed 4 Jan. 2020.
- [54] "Performance Tests for Spark," <https://github.com/databricks/spark-perf>, accessed 4 Jan. 2020.