

VolcanoML: Speeding up End-to-End AutoML via Scalable Search Space Decomposition

Yang Li^{†‡}, Yu Shen[†], Wentao Zhang[†], Jiawei Jiang[‡], Bolin Ding[§], Yaliang Li[§]
Jingren Zhou[§], Zhi Yang[†], Wentao Wu^{*}, Ce Zhang[‡], Bin Cui^{†∇}

[†]Key Laboratory of High Confidence Software Technologies (MOE), School of EECS, Peking University

[‡]Department of Computer Science, ETH Zürich [§]Alibaba Group ^{*}Microsoft Research, Redmond

[∇]Institute of Computational Social Science, Peking University (Qingdao)

[†]{liyong.cs, shenyu, wentao.zhang, yangzhi, bin.cui}@pku.edu.cn [‡]{jiawei.jiang, ce.zhang}@inf.ethz.ch

[§]{bolin.ding, yaliang.li, jingren.zhou}@alibaba-inc.com ^{*}wentao.wu@microsoft.com

ABSTRACT

End-to-end AutoML has attracted intensive interests from both academia and industry, which automatically searches for ML pipelines in a space induced by feature engineering, algorithm-/model selection, and hyper-parameter tuning. Existing AutoML systems, however, suffer from scalability issues when applying to application domains with large, high-dimensional search spaces. We present VOLCANOML, a scalable and extensible framework that facilitates systematic exploration of large AutoML search spaces. VOLCANOML introduces and implements basic building blocks that decompose a large search space into smaller ones, and allows users to utilize these building blocks to compose an *execution plan* for the AutoML problem at hand. VOLCANOML further supports a Volcano-style *execution model* – akin to the one supported by modern database systems – to execute the plan constructed. Our evaluation demonstrates that, not only does VOLCANOML raise the level of expressiveness for search space decomposition in AutoML, it also leads to actual findings of decomposition strategies that are significantly more efficient than the ones employed by state-of-the-art AutoML systems such as `auto-sklearn`.

PVLDB Reference Format:

Yang Li, Yu Shen, Wentao Zhang, Jiawei Jiang, Bolin Ding, Yaliang Li, Jingren Zhou, Zhi Yang, Wentao Wu, Ce Zhang, and Bin Cui. VolcanoML: Speeding up End-to-End AutoML via Scalable Search Space Decomposition. PVLDB, 14(11): 2167 - 2176, 2021.
doi:10.14778/3476249.3476270

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/PKU-DAIR/mindware>.

1 INTRODUCTION

In recent years, researchers in the database community have been working on raising the level of abstractions of machine learning (ML) and integrating such functionality into today’s data management systems, e.g., SystemML [21], SystemDS [6],

Snorkel [62], ZeroER [78], TFX [3, 7], “Query 2.0” [79], Krypton [57], Cerebro [58], ModelDB [74], MLFlow [81], DeepDive [12], HoloClean [63], ActiveClean [43], and NorthStar [42]. End-to-end AutoML systems [29, 80, 84] have been an emerging type of systems that has significantly raised the level of abstractions of building ML applications. Given an input dataset and a user-defined utility metric (e.g., validation accuracy), these systems automate the search of an end-to-end ML pipeline, including *feature engineering*, *algorithm/model selection*, and *hyper-parameter tuning*. Open-source examples include `auto-sklearn` [18], TPOT [60], and `hyperopt-sklearn` [41], whereas most cloud service providers, e.g., Google, Microsoft, Amazon, Alibaba, etc., all provide their proprietary services on the cloud. As machine learning has become an increasingly indispensable functionality integrated in modern data (management) systems, an efficient and effective end-to-end AutoML component also becomes increasingly important.

End-to-end AutoML provides a powerful abstraction to automatically navigate and search in a given complex search space. However, in our experience of applying state-of-the-art AutoML systems in a range of real-world applications, we find that such a system running “fully automatically” is rarely enough – often, developing a successful ML application involves multiple iterations between a user and an AutoML system to iteratively improve the resulting ML artifact.

Motivating Practical Challenge One such type of interaction, which inspires this work, is the *enrichment of search space*. We observe that the default search space provided by state-of-the-art AutoML systems is often not enough in many applications. This was not obvious to us at all in the beginning and it is not until we finish building a range of real-world applications that we realize this via a set of concrete examples. For example, in one of our astronomy applications [65], the feature normalization function is domain-specific and not supported by most, if not all, AutoML systems. Similar examples can also be found when searching for suitable ML models via AutoML. In one of our meteorology applications, we need to extend the models with meteorology-specific loss functions. We saw similar problems when we tried to extend existing AutoML systems with pre-trained feature embeddings coming from TensorFlow Hub, to include newly arXiv’ed models to enrich the “Model Base” [48], or to support Cosine annealing as for tuning.

Technical Challenge: Scalability over the Search Space “Why is it hard to extend the search space, as a user, in an end-to-end

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 14, No. 11 ISSN 2150-8097.
doi:10.14778/3476249.3476270

AutoML system?” The answer to this question is a complex one that is *not* completely technical: some aspects are less technical such as engineering decisions and UX designs, however, there are also more *technically fundamental* aspects. An end-to-end AutoML system contains an optimization algorithm that navigates a joint search space induced by *feature engineering*, *algorithm selection*, and *hyper-parameter tuning*. Because of this joint nature, the search space of end-to-end AutoML is complex and huge while the enrichment is only going to make it even larger. As we will see, handling such a huge space is already challenging for existing systems, and further enriching it will make it even harder to scale.

Many existing systems such as `auto-sklearn` [18] and TPOT [60] deal with the entire composite search space *jointly*, which naturally leads to the scalability bottleneck. Decomposing a joint space has been explored for some subspaces (e.g., only algorithm and hyper-parameters as in [49, 54]), however, none of them has been applied to a search space as large as that of end-to-end AutoML. One challenge is that there exist many different ways to decompose the same space, but only *some* of them can perform well. *Without a structured, high-level abstraction for search space decomposition to explore different strategies, it is very hard to scale up an end-to-end AutoML system to accommodate the search space that will only get larger in the future.*

Summary of Technical Contributions In this paper, we focus on designing a system, VOLCANO_{ML}, which is *scalable to a large search space*. Our technical contributions are as follows.

C1. System Design: A Structured View on Decomposition. The main technical contribution of VOLCANO_{ML} is to provide a flexible and principled way of decomposing a large search space into multiple smaller ones. We propose a novel system abstraction: a set of *VOLCANO_{ML} building blocks* (Section 3), each of which takes charge of a smaller sub-search space whereas a *VOLCANO_{ML} execution plan* (Section 4) consists of a *tree* of such building blocks — the root node corresponds to the original search space and its child nodes correspond to different subspaces. Under this abstraction, optimizing in the joint space is conducted as optimization problems over different smaller subspaces. The execution model is similar to the classic “Volcano” query evaluation model in a relational database [20, 44] (thus the name VOLCANO_{ML}): The system asks the root node to take one iteration in the optimization process, which *recursively* invokes one of its child nodes to take one iteration on solving a smaller-scale optimization problem over its own subspace; this recursive invocation procedure will continue until a leaf node is reached. This flexible abstraction allows us to explore different ways that the same joint space can be decomposed. Together with a range of additional optimizations (Section 4), VOLCANO_{ML} can often support more scalable search process than the existing AutoML systems such as `auto-sklearn` and TPOT.

C2. Large-scale Empirical Evaluations. We conducted intensive empirical evaluations, comparing VOLCANO_{ML} with state-of-the-art AutoML systems including `auto-sklearn` and TPOT. We show that (1) under the same search space as `auto-sklearn`, VOLCANO_{ML} significantly outperforms `auto-sklearn` and TPOT — over 30 classification tasks and 20 regression tasks — VOLCANO_{ML} outperforms the *best* of `auto-sklearn` and TPOT on a majority of tasks; (2) using an enriched search space with additional feature engineering operators,

VOLCANO_{ML} performs significantly better than `auto-sklearn`; and (3) using an *enriched search space* with an additional data processing stage and functionalities beyond what `auto-sklearn` and TPOT currently support (i.e., an additional embedding selection stage using pre-trained models on TensorFlow Hub), VOLCANO_{ML} can deal with input types such as images efficiently.

Moving Forward The VOLCANO_{ML} abstraction enables a structured view of optimizing a black-box function via decomposition. This structured view itself opens up interesting future directions. For example, one may wish to *automatically* decompose a search space given a workload, just like what a classic query optimizer would do for relational queries. For constrained optimizations, we also imagine techniques similar to traditional “*push-down selection*” could be applied in a similar spirit. We explore the possibility of automatically searching for the best plan in Section 4 and discuss the limitations of this simple strategy and the exciting line of future work that could follow. While the full treatment of these aspects are beyond the scope of this paper, we hope the VOLCANO_{ML} abstraction can serve as a foundation for these future endeavors.

2 RELATED WORK

AutoML is a topic that has been intensively studied over the last decade. We briefly summarize related work in this section and readers can consult latest surveys [25, 29, 80, 84] for more details.

End-to-End AutoML. End-to-end AutoML aims to automate the development process of the end-to-end ML pipeline, including feature preprocessing, feature engineering, algorithm selection, hyper-parameter tuning, and model ensembling [9, 82, 83]. Often, this is modeled as a black-box optimization problem [30] and solved jointly [18, 60, 71]. Apart from grid search and random search [4], genetic programming [55, 60] and Bayesian optimization (BO) [5, 16, 28, 67, 68] has become prevailing frameworks for this problem. Several meta-learning approaches [11, 18, 27, 72] can guide ML practitioners to design better search spaces for AutoML tasks. Many end-to-end AutoML systems have raised the abstraction level of ML. `auto-weka` [71], `hyperopt-sklearn` [41] and `auto-sklearn` [18] are the main representatives of BO-based AutoML systems. `auto-sklearn` is one of the most popular open-source framework. TPOT [60] and ML-Plan [55] use genetic algorithm and hierarchical task networks planning respectively to optimize over the pipeline space, and require discretization of the hyper-parameter space. AlphaD3M [14] integrates reinforcement learning with Monte Carlo tree search (MCTS) to solve AutoML problems but without imposing efficient decomposition over hyper-parameters and algorithm selection. AutoStacker [10] focuses on ensembling and cascading to generate complex pipelines, and solves the CASH problem [18] via random search. Furthermore, a growing number of commercial enterprises also export their AutoML services to their users, e.g., H2O [45], Microsoft’s Azure Machine Learning [2], Google’s Prediction API [23], Amazon Machine Learning [53] and IBM’s Watson Studio AutoAI [31].

Automating Individual Components. Apart from end-to-end AutoML, many efforts have been devoted to studying sub-problems in AutoML: (1) feature engineering [36–39, 59], (2) algorithm selection [15, 18, 41, 49, 54, 71], and (3) hyper-parameter tuning [5, 17, 26, 28, 32, 34, 40, 47, 50, 61, 66, 68, 69, 77]. Meta-learning

methods [19, 22, 76] for hyper-parameter tuning can leverage auxiliary knowledge acquired from previous tasks to achieve faster optimization. Several systems offer a subset of functionalities in the end-to-end process. Microsoft’s NNI [64] helps users to automate feature engineering, hyper-parameter tuning, and model compression. Recent work [54] leverages the ADMM optimization framework to decompose the CASH problem [18], and solves two easier sub-problems. Berkeley’s Ray [56] and OpenBox [51] provide the `tune` module [50, 52] to support scalable hyper-parameter tuning tasks in a distributed environment. Featuretools [35] is a Python library for automatic feature engineering. Unlike these works, we focus on deriving an end-to-end solution to the AutoML problem, where the sub-problems are solved in a joint manner.

3 VOLCANOML AND BUILDING BLOCKS

The goal of VOLCANOML is to enable scalability with respect to the underlying AutoML search space. As a result, its design focuses on the *decomposition* of a given search space. In this section, we first introduce key building blocks in VOLCANOML, and in Section 4 we describe how multiple building blocks are put together to compose a VOLCANOML *execution plan* in a modular way. Later in Section 5, we introduce additional optimizations for these building blocks.

3.1 Search Space of End-to-End AutoML

We describe the search space of end-to-end AutoML following `auto-sklearn`. The input to the system is a dataset D , containing a set of training samples. The user also provides a pre-defined metric, e.g., validation accuracy or cross-validation accuracy, to measure the *utility* of a given ML pipeline. The output of an end-to-end AutoML system is an ML pipeline that achieves good utility.

To find such an ML pipeline, the system searches over a large search space of possible pipelines and picks one that maximizes the pre-defined utility. This search space is a composition of (1) feature engineering operators, (2) ML algorithms/models, and (3) hyper-parameters.

Feature Engineering. The feature engineering process takes as input a dataset D and outputs a new dataset D' . It achieves this by transforming the input dataset via a set of data transformations. In `auto-sklearn`, it further defines multiple *stages* of the feature engineering process: (1) *preprocessing*, (2) *rescaling*, (3) *balancing*, and (4) *feature transforming*. For each stage, the system chooses a single transformation to apply. For example, for *feature transforming*, the system can choose among `no_processing`, `kernel_pca`, `polynomial`, `select_percentile`, etc.

ML Algorithms. Given a transformed dataset D' , the system then picks an ML algorithm to train. Since different ML algorithms are suitable for different types of tasks, the system needs to consider a diverse range of possible ML algorithms. Taking `auto-sklearn` as an example, the search space for ML algorithms contains `Linear_Model`, `Support_Vector_Machine`, `Discriminant_Analysis`, `Random_Forest`, etc.

Hyper-parameters. Each ML algorithm has its own sub-search space for hyper-parameter tuning – if we choose to use a certain ML algorithm, we also have to specify the corresponding hyper-parameters. The hyper-parameters fall into three categories: continuous (e.g., `sub_sample_rate` for `Random_Forest`), discrete

(e.g., `maximal_depth` for `Decision_Tree`), and categorical (e.g., `kernel_type` for `Lib_SVM`).

If the system makes a concrete pick for each of the above decisions, then it can compose a concrete ML pipeline and evaluate its utility. This is often an expensive process since it involves training an ML model. To find the optimal ML pipeline, the system evaluates the utility of different ML pipelines in an iterative manner following a *search strategy*, and picks the one that maximizes the utility.

For example, `auto-sklearn` handles the above search space *jointly* and optimizes it with Bayesian optimization (BO) [67]. Given an initial set of function evaluations, BO proceeds by fitting a surrogate model to those observations, specifically a *probabilistic Random Forest* in `auto-sklearn`, and then chooses which ML pipeline to evaluate from the search space by optimizing an acquisition function that balances exploration and exploitation.

3.2 Building Blocks

Unlike `auto-sklearn`, VOLCANOML decomposes the above search space into smaller subspaces. One interesting design decision in VOLCANOML is to introduce a *structured abstraction* to express different *decomposition strategies*. A decomposition strategy is akin to an *execution plan* in relational database management systems, which is composed of *building blocks* akin to relational operators. A building block itself can be viewed as an *atomic* decomposition strategy. We next present the details of the building blocks implemented by VOLCANOML, and we will introduce how to use these blocks to compose VOLCANOML execution plans in Section 4.

Goal. The goal of VOLCANOML is to solve:

$$\min_{x_1, \dots, x_n} f(x_1, \dots, x_n; D),$$

where x_1, \dots, x_n is a set of n variables and each of them has domain \mathbb{D}_{x_i} for $i \in [n]$. Together, these n variables define a search space $(x_1, \dots, x_n) \in \prod_i \mathbb{D}_{x_i}$. D corresponds to the input dataset, which is a *set* of input samples. In our setting, $f(\cdot)$ is a black-box function that we can only evaluate (but not exploiting the derivative). Given constant c in the composite domain $c \in \prod_i \mathbb{D}_{x_i}$, we use the notation

$$f(\{(x_1, \dots, x_n) = c\}; D)$$

as the value of evaluating f by substituting (x_1, \dots, x_n) with c .

Subgoal. One key decision of VOLCANOML is to solve the optimization problem on a search space by decomposing it into multiple smaller subspaces, each of which will be solved by one *building block*. We define optimizing over each of these smaller subspaces as a *subgoal* of the original problem. Formally, a subgoal g is defined by two components: $\bar{x}_g \subseteq \{x_1, \dots, x_n\}$ as a subset of variables, and $\bar{c}_g \in \prod_{x_i \in \bar{x}_g} \mathbb{D}_{x_i}$ as an assignment in the domain of all variables in \bar{x}_g . Let $\bar{x}_{-g} = \{x_1, \dots, x_n\} - \bar{x}_g$ be all variables that are *not* in \bar{x}_g .

Each subgoal defines a function f_g over a smaller search space, which is constructed by *substituting* all variables in \bar{x}_g with \bar{c}_g :

$$f_g = f[\bar{x}_g/\bar{c}_g] : z \in \prod_{x_i \in \bar{x}_{-g}} \mathbb{D}_{x_i} \mapsto f(\{\bar{x}_g = \bar{c}_g; \bar{x}_{-g} = z\}; D).$$

Building Block. Each subgoal g corresponds to one building block $B_{g,D}$, whose goal is to solve

$$\min_{\bar{x}_{-g}} f_g(\bar{x}_{-g}; D).$$

A building block $B_{g,D}$ imposes several assumptions on g and D . First, given an assignment \bar{c}_{-g} to \bar{x}_{-g} , it is able to evaluate the value of the function $f_g(\bar{c}_{-g}, D)$. Note that such an evaluation can often be expensive and VOLCANOML tries to minimize the number of times that such a function is evaluated. Second, given a dataset D , a building block has the knowledge about how to subsample a smaller dataset $\tilde{D} \subseteq D$ and then conduct evaluations on such a subset $x \mapsto f_g(x; \tilde{D})$. Third, we assume that the building block has access to a cost model about the cost of an evaluation at x , $C_{g,D,x}$.

Interfaces. All implementations of a building block follow an interactive optimization process. A building block exposes several interfaces. First, one can initialize a building block via

$$B_{g,D} \leftarrow \text{init}(f, \bar{x}_g, \bar{c}_g, D),$$

which creates a building block. Second, one can query the current best solution found in $B_{g,D}$ by

$$\hat{x} \leftarrow \text{get_current_best}(B_{g,D}).$$

Furthermore, one can ask $B_{g,D}$ to iterate once via

$$\text{do_next!}(B_{g,D}),$$

where ‘!’ indicates potential change on the state of the input $B_{g,D}$.

Last but not least, one can query a building block about its *expected utility* (EU) if given K more budget units (e.g., seconds) via

$$[l, u] \leftarrow \text{get_eu}(B_{g,D}, K).$$

By adopting a similar design principle used in the existing AutoML systems [18, 54, 60], in VOLCANOML we estimate EU by *extrapolation* into the “future” with more available budget. Given the inherent uncertainty in our estimation method, rather than returning a single point estimate, we instead return a lower bound l and an upper bound u . We refer readers to [49] for the details of how the lower and upper bounds are established. Moreover, one can query a building block about its *expected utility improvement* (EUI) via

$$\delta \leftarrow \text{get_eui}(B_{g,D}).$$

Note that, different from EU, EUI is the *expected improvement* over the current observed utility if given K more budget units. In VOLCANOML, we estimate EUI by taking the mean of the observed improvements from history, following Levine et al [46].

3.3 Three Types of Building Blocks

Decomposition is the cornerstone of VOLCANOML’s design. Given a search space, apart from exploring it jointly, there are two classical ways of decomposition — to partition the search space via conditioning on different values of a certain variable (in a similar spirit of *variable elimination* [13]), or to decompose the problem into multiple smaller ones by introducing equality constraints (in a similar spirit of *dual decomposition* [8]). This inspires VOLCANOML’s design, which supports three types of building blocks: (1) *joint block* that simply optimizes the input subspace using Bayesian optimization; (2) *conditioning block* that further divides the input subspace into smaller ones by conditioning on one particular input variable; and (3) *alternating block* that partitions the input subspace into two and optimizes each one *alternately*. Note that both *conditioning block* and *alternating block* would generate new building blocks with smaller subgoals. We next present the implementation details for each type of building block.

Algorithm 1: The `do_next!` of conditioning block

Input: A conditioning block $B_{g,D}$, budget K .

- 1 Let B_1, \dots, B_m be all active (have not been eliminated) child blocks;
- 2 **for** $1 \leq i \leq L$ **do**
- 3 **for** $1 \leq j \leq m$ **do**
- 4 $\text{do_next!}(B_j)$;
- 5 **for** $1 \leq j \leq m$ **do**
- 6 $[l_j, u_j] \leftarrow \text{get_eu}(B_j, K)$;
- 7 Eliminate child blocks that are *dominated* by others, using $[l_j, u_j]$ for $1 \leq j \leq m$;

3.3.1 Joint Block. A joint block directly optimizes its subgoal via Bayesian optimization (BO) [67]. Specifically, BO based method - SMAC [28] has been used by many applications where evaluating the objective function is computationally expensive. It constructs a probabilistic surrogate model M to capture the relationship between the input variables \bar{x} and the objective function value ψ , and refines M iteratively using past observations (\bar{x}, ψ) s.

The implementation of `do_next!` for a joint block consists of the following three steps: 1) Use the surrogate model M to select \bar{x} that maximizes an acquisition function. In our implementation, we use *expected improvement* (EI) [33] as the acquisition function, which has been widely used in BO community. 2) Evaluate the selected \bar{x} and obtain its value about the objective function (i.e., the subgoal) $\psi = f_g(\bar{x}) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma^2)$, where \mathcal{N} is the normal distribution. 3) Refit the surrogate model M on the observed (\bar{x}, ψ) s.

Early-Stopping based Optimization. For large datasets, early-stopping based methods, e.g., Successive Halving [32], Hyperband [47], BOHB [17], MFES-HB [50], etc, can terminate the evaluations of poorly-performed configurations in advance, thus speeding up the evaluations. VOLCANOML supports MFES-HB [50], which combines the benefits of Hyperband and Multi-fidelity BO [70, 77], to optimize a joint block, in addition to vanilla BO.

3.3.2 Conditioning Block. A conditioning block decomposes its input \bar{x} into $\bar{x} = \{x_c\} \cup \bar{y}$, where x_c is a single variable with domain \mathbb{D}_{x_c} . It then creates one new building block for each possible value $d \in \mathbb{D}_{x_c}$ of x_c :

$$\min_{\bar{y}} g_d(\bar{y}; D) \equiv f(\{x_c = d, \bar{y}\}; D).$$

As a result, $|\mathbb{D}_{x_c}|$ new (child) building blocks are created.

The conditioning block aims to identify optimal value for x_c , and many previous AutoML researches have used Bandit algorithms for this purpose [32, 49, 50, 54]. In VOLCANOML, we follow these previous work and model it as a multi-armed bandit (MAB) problem, while our framework is flexible enough to incorporate other algorithms when they are available. There are $|\mathbb{D}_{x_c}|$ arms, where each arm corresponds to a child block. Playing an arm means invoking the `do_next!` primitive of the corresponding child block.

Algorithm 1 illustrates the implementation of `do_next!` for a conditioning block. It starts by playing each arm L times in a Round-Robin fashion (lines 2 to 4). Here, L is a user-specified configuration parameter of VOLCANOML. In our current implementation, we set $L = 5$. We then obtain the lower and upper bounds of the expected utility of each child block by invoking its `get_eu` primitive (lines 5 to 6), and eliminate child blocks that are dominated by others (line 7). The elimination works as follows. Consider two blocks B_i and B_j : if the upper bound u_i of B_i is less than the lower bound l_j of B_j , then the block B_i is eliminated. An eliminated arm/block will not be played in future invocations of `do_next!`.

Algorithm 2: The init of alternating block

Input: An alternating block $B_{g,D}$ with search space $\bar{x} = \bar{y} \cup \bar{z}$.

```
1 Initialize  $\bar{y}$  and  $\bar{z}$  with default values  $\bar{y}_0$  and  $\bar{z}_0$ ;  
2  $B_1 \leftarrow \text{init}(f, \bar{z}, \bar{z}_0, D)$ ;  
3  $B_2 \leftarrow \text{init}(f, \bar{y}, \bar{y}_0, D)$ ;  
4 for  $1 \leq i \leq L$  do  
5   do_next( $B_1$ );  
6    $\bar{y}_i \leftarrow \text{get\_current\_best}(B_1)$ ;  
7   set_var( $B_2, \bar{y}, \bar{y}_i$ );  
8   do_next( $B_2$ );  
9    $\bar{z}_i \leftarrow \text{get\_current\_best}(B_2)$ ;  
10  set_var( $B_1, \bar{z}, \bar{z}_i$ );
```

Algorithm 3: The do_next! of alternating block

Input: An alternating block $B_{g,D}$ with budget K .

```
1  $\delta_1 \leftarrow \text{get\_eui}(B_1)$ ;  
2  $\delta_2 \leftarrow \text{get\_eui}(B_2)$ ;  
3 if  $\delta_1 \geq \delta_2$  then  
4    $\bar{z}_{\text{best}} \leftarrow \text{get\_current\_best}(B_2)$ ;  
5   set_var( $B_1, \bar{z}, \bar{z}_{\text{best}}$ );  
6   do_next( $B_1$ );  
7 else  
8    $\bar{y}_{\text{best}} \leftarrow \text{get\_current\_best}(B_1)$ ;  
9   set_var( $B_2, \bar{y}, \bar{y}_{\text{best}}$ );  
10  do_next( $B_2$ );
```

Remark: We have simplified the above elimination criterion by using the lower and upper bounds calculated given K budget units for *each arm*. In fact, these K budget units are *shared* by all the arms, and as a result, each arm actually has fewer budget units than K . Our assumption is that, K is sufficiently large so that one can play *all arms* until (the observed distribution of rewards of) *each arm* converges. Otherwise, the lower and upper bounds obtained may be *over-optimistic*, and as a result, may lead to incorrect eliminations. Fortunately, our assumption usually holds in practice, where arms converge relatively fast.

3.3.3 Alternating Block. An alternating block decomposes its input search space into $\bar{x} = \bar{y} \cup \bar{z}$, and explores \bar{y} and \bar{z} in an *alternating* way. Similarly, we also model the optimization in alternating block as an MAB problem. Algorithm 2 illustrates how its `init` primitive works. It first creates two child blocks B_1 and B_2 , which will focus on optimizing for \bar{y} and \bar{z} respectively (lines 1 to 3). It then (again) views B_1 and B_2 as two arms and plays them using Round-Robin (lines 4 to 10). Note that, when B_1 optimizes \bar{y} (resp. when B_2 optimizes \bar{z}), it uses the current best \bar{z} found by B_2 (resp. the current best \bar{y} found by B_1). This is done by the `set_var` primitive (invoked at line 7 for B_2 and line 10 for B_1).

One problem of our alternating MAB formulation is that the utility improvements of the two building blocks often vary dramatically in practice. For example, some applications are very sensitive to the features being used (e.g., normalized vs. non-normalized features) while hyper-parameter tuning will offer little or even no improvement. In this case, we should spend more resources on looking for good features instead of tuning hyper-parameters. Our key observation is that, the *expected utility improvement* (EUI) decays as optimization proceeds. As a result, we propose to use EUI as an indicator that measures the *potential* of pulling an arm further. Algorithm 3 illustrates the details of this idea when used to implement the `do_next!` primitive.

Specifically, Algorithm 3 starts by polling the EUI of both child blocks (lines 1 and 2). Recall that the EUI is estimated by taking the




Figure 1: Two different execution plans for the same optimization problem. Each plan corresponds to a different way to decompose the same search space (x, y, z, w) .

mean of historic observations. It then compares the EUIs and picks the arm/block with larger EUI to play next (lines 3 to 10). Before pulling the winner arm, again it will use the current best settings found by the other arm/block (lines 4 to 6, lines 8 to 10).

3.3.4 Discussion: Pros and Cons of Building Blocks. While the joint block is the most straightforward way to solve the optimization problem associated, it is difficult to scale Bayesian optimization to a large search space [49, 75]. The alternating block addresses this scalability issue by decomposing the search space into two smaller subspaces, though with the assumption that the improvements of the two subspaces are *conditionally independent* of each other. As a result, the alternating block is a better choice when such an assumption approximately holds. The conditioning block is capable of pruning the search space *as optimization proceeds*, when bad arms are pulled less often or will not be played anymore, with the limitation that it can only work for conditional variables that are *categorical*. For non-categorical variables, one possible way to use conditioning blocks is to split the value range of variables. For example, given a numerical variable that ranges from 1 to 3, we split it into two ranges, which are [1, 2) and [2, 3). During the optimization iteration, we first choose one sub-range and then optimize the splitted space along with its corresponding subspace.

In addition, VOLCANOML uses bandit-based algorithms from the existing literature [46, 49] as default in both the alternating and conditioning block, and other bandit-based algorithms, such as successive halving [32], Hyperband [47], BOHB [17] and MFES-HB [50], can also be used in these blocks. In Section 4, we explore the possibility of automatically choosing building blocks to use by maximizing the empirical accuracy of different execution plans.

4 VOLCANOML EXECUTION PLAN

Given a pre-defined search space, the input of VOLCANOML is (1) a dataset D , (2) a utility metric (e.g. cross-validation accuracy) which defines the objective function f , and (3) a time budget. VOLCANOML then decomposes a large search space into an *execution plan*, following some specific *decomposition strategy*.

VOLCANOML Execution Plan. Due to space limitation, we omit the formal definition of a VOLCANOML execution plan. Intuitively, a VOLCANOML execution plan is a *tree* of building blocks. The root node corresponds to a building block solving the problem f with the entire search space, which can be further decomposed into multiple building blocks if necessary, as previously described. As an example, Figure 1 illustrates two possible execution plans for $f(x, y, z, w; D)$. **Plan 1** contains only a single root building block as a joint block, whereas **Plan 2** first introduces a conditioning block

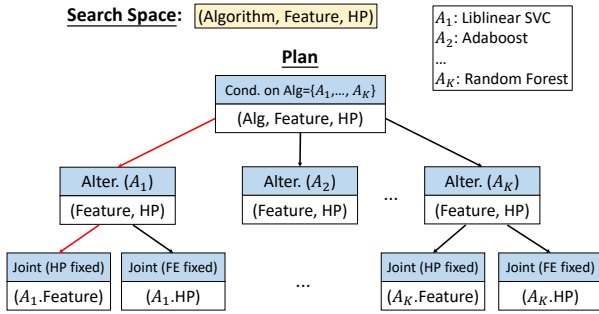


Figure 2: VOLCANOML’s execution plan for the same search space as explored by auto-sklearn. Here ‘Alg’ and ‘HP’ correspond to Algorithm and hyper-parameters respectively.

on x , and then creates one lower level of building blocks for each possible value of x (in Figure 1, we assume that $|\mathbb{D}_x| = 3$).

VOLCANOML Execution Model. To execute a VOLCANOML execution plan, we follow a Volcano-style execution that is similar to a relational database [24] — the system invokes the `do_next!` of the root node, which then invokes the `do_next!` of one of its child nodes, propagating until the leaf node. At any time, one can invoke the `get_current_best` of the root node, which returns the current best solution for the entire search space.

VOLCANOML Plan for auto-sklearn. Figure 2 presents a VOLCANOML execution plan for the same search space explored by auto-sklearn, which consists of the joint search of *algorithms*, *features*, and *hyper-parameters*. Instead of conducting the search process in a single joint block, as was done by auto-sklearn, VOLCANOML first decomposes the search space via a conditioning block on *algorithms* — this introduces a MAB problem in which each arm corresponds to one particular algorithm. It then further decomposes each of the conditioned subspaces via an alternating block between feature engineering and hyper-parameter tuning. The whole subspace of feature engineering (resp. that of hyper-parameter tuning) is optimized by a joint block.

Concretely, Figure 2 shows a search space for AutoML with K choices of ML algorithms. During each iteration, starting from the root node, VOLCANOML selects the child node to optimize until it reaches a leaf node, and then optimizes over the subspace in the leaf node. As shown by the red lines in Figure 2, in this iteration, VOLCANOML only tunes the feature engineering pipeline of algorithm A_1 while fixing its algorithm hyper-parameters.

Alternative Execution Plans. Note that the execution plan in Figure 2 is not the only possible one. Our flexible and scalable framework in VOLCANOML allows us to explore different execution plans before reaching the proposed one. We enumerate five possible plans in a coarse-grained level, and the results show that the proposed plan performs best. The reason why we choose this plan is due to the fundamental property of the AutoML search space — we observe that, the optimal choices of *features* are different across *algorithms*, which implies that we can first decompose the search space along *ML algorithms*. The improvements introduced by feature engineering and hyper-parameter tuning are largely complementary, and thus we can optimize them *alternately*. For feature engineering (resp. hyper-parameter tuning), the subspace is small enough to be handled by a single joint block efficiently.

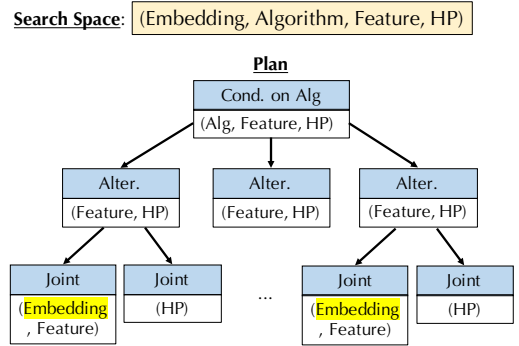


Figure 3: VOLCANOML’s execution plan for a larger search space enriched by an additional embedding selection stage.

VOLCANOML Plan for Enriched Search Space. We can easily extend VOLCANOML and enable functionalities that are not supported by most AutoML systems. For example, Figure 3 illustrates an execution plan for a search space with an additional stage — *embedding selection*. Given an input, e.g., image or text, we first choose embeddings based on a collection of TensorFlow Hub pre-trained models, and then conduct algorithm selection, feature engineering, and hyper-parameter tuning. We use an execution plan as illustrated in Figure 3, having the embedding selection step jointly optimized together with the feature engineering.

Discussion: Automatic Plan Generation. In principle, the design of VOLCANOML opens up the opportunity for “automatic plan generation” — given a collection of benchmark datasets, one could automatically search for the best decomposition strategy of the search space and come up with a physical plan automatically. While the full treatment of this problem is beyond the scope of this paper, we illustrate the possibility with a very simple strategy. We automatically enumerate all possible execution plans in a coarse-grained level, and find that our manually specified execution plan in Figure 2 outperforms the alternatives. There is still an open question that whether we can support finer-grained partition of the search space (e.g., different plans for different subspace of features), and moreover, whether we can conduct efficient automatic plan optimization without enumerating all possible plans. These are exciting future directions and we expect the endeavor to be non-trivial. We hope that this paper sets the ground for this line of research in the future (e.g., rule-based heuristics or reinforcement learning).

Further Optimization with Meta-learning. VOLCANOML supports meta-learning based techniques — given previous runs of the system over similar workloads, to transfer the knowledge and better help the workload at hand — to accelerate the optimization process of building blocks.

5 EXPERIMENTAL EVALUATION

We compare VOLCANOML with state-of-the-art AutoML systems and commercial AutoML platforms. In our evaluation, we focus on three perspectives: (1) the *performance* of VOLCANOML given the *same* search space explored by existing systems, (2) the *scalability* of VOLCANOML given larger search spaces, and (3) the *extensibility* of VOLCANOML to integrate new components into the search space of AutoML pipelines.

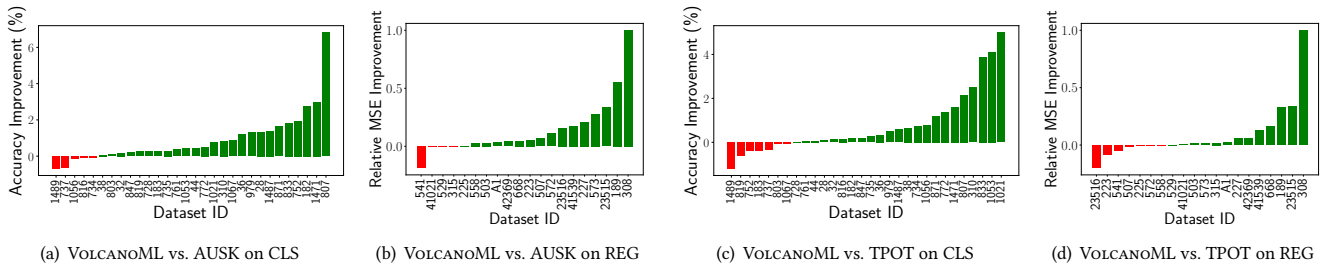


Figure 4: End-to-End results on 30 OpenML classification (CLS) datasets and 20 OpenML regression (REG) datasets.




Figure 5: Average test errors on four large datasets with different time budgets.

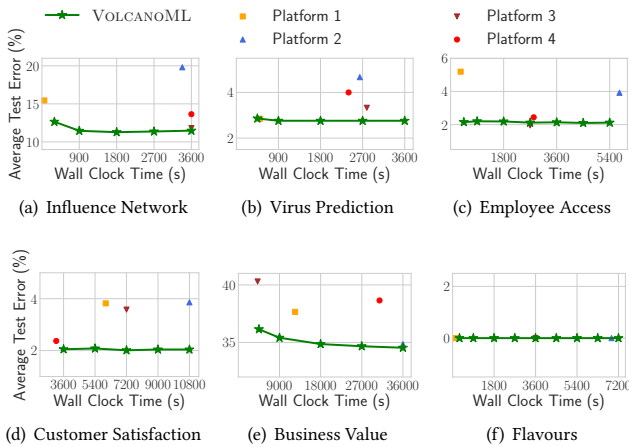


Figure 6: Test errors on 6 Kaggle competitions compared with four commercial platforms.

5.1 Experimental Setup

AutoML Systems. We evaluate VOLCANOML as well as two open source AutoML systems: `auto-sklearn` [18] and TPOT [60]. In addition, we also compare VOLCANOML with four commercial AutoML platforms from Google, Amazon AWS, Microsoft Azure, and Oracle. Both VOLCANOML and `auto-sklearn` support meta-learning, while TPOT does not. For fair comparison with TPOT, we also use VOLCANOML⁻ and AUSK⁻ to denote the versions of VOLCANOML and `auto-sklearn` when meta-learning is disabled. Our implementation of VOLCANOML is available at <https://github.com/VolcanoML>.

Datasets. We use 60 real-world ML datasets from the OpenML repository [73], including 40 for classification (CLS) tasks and 20

for regression (REG) tasks. 10 of the 40 classification datasets are relatively large, each with 20k to 110k data samples; the other 30 are of medium size, each with 1k to 12k samples. In addition, we also use datasets from six Kaggle competitions to compare VOLCANOML with four commercial platforms.

AutoML Tasks. We define three kinds of real-world AutoML tasks, including (1) a general classification task on 30 medium datasets, (2) a general regression task on 20 medium datasets, and (3) a large-scale classification task on 10 large datasets.

To test the scalability of the participating systems, we design three search spaces that include 20, 29, and 100 hyper-parameters, where the smaller search space is a subset of the larger one. We run VOLCANOML and the baseline AutoML systems against each of the three search spaces. The time budget is 900 seconds for the smallest search space and 1,800 seconds for the other two, when performing the general classification task (1); the time budget is increased to 5,400 and 86,400 seconds respectively, when performing the general regression task (2) and the large-scale classification task (3).

Utility Metrics. Following [18], we adopt the metric *balanced accuracy* for all classification tasks – compared with standard (classification) accuracy, it assigns equal weights to classes and takes the average of class-wise accuracy. For regression tasks, we use the *mean squared error* (MSE) as the metric.

In our evaluation, we repeat each experiment 10 times and report the average utility metric. In each experiment, we use four fifths of the data samples in each dataset to search for the best ML pipeline and report the utility metric on the remaining fifth.

Methodology for Comparing AutoML Systems. To compare the overall test result of each AutoML system on a wide range of datasets, we use the *average rank* as the metric following [1]. For each dataset, we rank all participant systems based on the result

Table 1: Average ranks on 30 classification (CLS) datasets and 20 regression (REG) datasets with three different search spaces. (The lower is the better)

Search Space - Task	TPOT	AUSK ⁻	AUSK	VOLCANOML ⁻	VOLCANOML
Small - CLS	3.09	3.07	3.01	2.94	2.89
Medium - CLS	3.2	3.32	3.27	2.78	2.43
Large - CLS	3.29	3.77	3.57	2.72	1.65
Small - REG	2.98	3.02	3.0	3.02	2.98
Medium - REG	2.95	3.3	3.12	2.75	2.88
Large - REG	3.1	3.85	3.82	2.15	2.08

of the best ML pipeline they have found so far; we then take the average of their ranks across different datasets.

5.2 End-to-End Comparison

We first evaluate the participant AutoML systems given the search space explored by `auto-sklearn`. Figure 4 presents the results of VOLCANOML compared to `auto-sklearn` (AUSK) and TPOT on the 30 classification tasks and the 20 regression tasks, respectively. For classification tasks, we plot the classification accuracy improvement (%); for regression tasks, we plot the *relative MSE improvement* Δ , which is defined as $\Delta(m_1, m_2) = \frac{s(m_2) - s(m_1)}{\max(s(m_2), s(m_1))}$, where $s(\cdot)$ is MSE on the test set. Overall, VOLCANOML outperforms `auto-sklearn` and TPOT on 25 and 23 of the 30 classification tasks, and on 17 and 15 of the 20 regression tasks, respectively.

We also conduct experiments to evaluate VOLCANOML with different *time budgets*. Figure 5 presents the results on four large classification datasets. VOLCANOML exhibits consistent performance over different time budgets. Notably, on `Higgs`, VOLCANOML achieves 27.2% test error within 4 hours, which is better than the performance of the other two systems given 24 hours.

We further study the *scalability* of the participant systems on the three aforementioned search spaces. Table 1 summarizes the results in terms of the average ranks. We have two observations: First, without meta-learning, VOLCANOML achieves the best average rank for both the classification and regression tasks – on the small search space (with 20 hyper-parameters), VOLCANOML performs slightly better than `auto-sklearn` and TPOT, and it performs significantly better on the medium (with 29 hyper-parameters) and large (with 100 hyper-parameters) search spaces. Second, with meta-learning, the average rank of VOLCANOML is dramatically improved compared with `auto-sklearn`. Overall, VOLCANOML with meta-learning achieves the best result over large search space.

5.3 Search Space Enrichment

We now focus on evaluating the *extensibility* of VOLCANOML.

Adding Data_Balancing Operator. In the first experiment, we implement “`smote_balancer`” – a new feature engineering operator, and incorporate it into the aforementioned *balancing* stage of feature engineering (FE) (Section 3.1). Note that `auto-sklearn` cannot support this fine-grained enrichment of the search space. Table 2 presents the results of `auto-sklearn`, VOLCANOML without enrichment, and VOLCANOML with enrichment, on five imbalanced datasets. We observe that enriching the search space brings further improvement, e.g., VOLCANOML with enrichment outperforms `auto-sklearn` by 3.57% (balanced accuracy) on the dataset `pc2`.

Supporting Embedding Selection. In the second experiment, we add a new stage “embedding selection” into the FE pipeline, with

Table 2: Test accuracy (%) of VOLCANOML with and without the enrichment of “`smote_balancer`” operator.

Dataset	AUSK	VOLCANOML ⁻	VOLCANOML
<code>sick</code>	97.29	97.31	97.34
<code>pc2</code>	86.70	86.91	90.27
<code>abalone</code>	66.86	65.97	67.32
<code>page-blocks(2)</code>	94.70	95.29	96.69
<code>hypothyroid(2)</code>	99.62	99.64	99.64

two candidate embedding-extraction operators (i.e., two pre-trained models). This allows VOLCANOML to deal with images, which are not easily supported by both `auto-sklearn` and TPOT. We implement two pre-trained models to generate embeddings for images, and we evaluate VOLCANOML with the enriched search space on the Kaggle dataset `dogs-vs-cats`. We observe that VOLCANOML achieves 96.5% test accuracy, which is significantly better than 69.7% obtained by `auto-sklearn` without considering embeddings.

5.4 Comparison with 4 Industrial Platforms

In addition, we run additional experiments on six Kaggle datasets to compare VOLCANOML with four commercial AutoML platforms: 1) Google Cloud AutoML, 2) Microsoft Azure Automated ML, 3) Oracle data science, and 4) Amazon AWS Sagemaker AutoPilot. Here, we anonymously refer to these platforms as Platform 1-4. Figure 6 show the results, and the Appendix contains the experiment details. We observe that, given the same time budget (i.e., fix the x-axis to some time budget), VOLCANOML is at least comparable with, often outperforms, the considered commercial platforms.

6 CONCLUSION

In this paper, we have presented VOLCANOML, a scalable and extensible framework that allows users to design decomposition strategies for large AutoML search spaces in an expressive and flexible manner. VOLCANOML introduces novel building blocks that enable expressing search space decomposition strategies in a *structured* fashion. Moreover, VOLCANOML introduces a Volcano-style execution model to execute the decomposition strategies it yields. Experimental evaluation demonstrates that VOLCANOML can generate more efficient decomposition strategies that also lead to performance-wise better ML pipelines, compared to state-of-the-art AutoML systems.

ACKNOWLEDGMENTS

This work is supported by the National Key Research and Development Program of China (No.2018YFB1004403), NSFC (No.61832001, U1936104), Beijing Academy of Artificial Intelligence (BAAI), and Kuaishou-PKU joint program. Bin Cui and Jiawei Jiang are the corresponding authors.

Ce Zhang and the DS3Lab gratefully acknowledge the support from the Swiss National Science Foundation (Project Number 200021_184628), Innosuisse/SNF BRIDGE Discovery (Project Number 40B2-0_187132), European Union Horizon 2020 Research and Innovation Programme (DAPHNE, 957407), Botnar Research Centre for Child Health, Swiss Data Science Center, Alibaba, Cisco, eBay, Google Focused Research Awards, Oracle Labs, Swisscom, Zurich Insurance, Chinese Scholarship Council, and the Department of Computer Science at ETH Zurich.

REFERENCES

- [1] Rémi Bardenet, Mátys Brendel, Balázs Kégl, and Michele Sebag. 2013. Collaborative hyperparameter tuning. In *International conference on machine learning*. PMLR, 199–207.
- [2] Jeff Barnes. 2015. Azure machine learning. *Microsoft Azure Essentials*. 1st ed, Microsoft (2015).
- [3] Denis Baylor, Eric Breck, Heng-Tze Cheng, Noah Fiedel, Chuan Yu Foo, Zakaria Haque, Salem Haykal, Mustafa Ispir, Vihan Jain, Levent Koc, et al. 2017. Tfx: A tensorflow-based production-scale machine learning platform. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1387–1395.
- [4] James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research* 13 (2012), 281–305.
- [5] James S Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. In *Advances in neural information processing systems*. 2546–2554.
- [6] Matthias Boehm, Iulian Antonov, Sebastian Baunsgaard, Mark Dokter, Robert Günthör, Kevin Innerebner, Florian Klezin, Stefanie Lindstaedt, Arnab Phani, Benjamin Rath, et al. 2019. SystemDS: A declarative machine learning system for the end-to-end data science lifecycle. *arXiv preprint arXiv:1909.02976* (2019).
- [7] Eric Breck, Neoklis Polyzotis, Sudip Roy, Steven Whang, and Martin Zinkevich. 2019. Data Validation for Machine Learning. In *MLSys*.
- [8] Claus C Carøe and Rüdiger Schultz. 1999. Dual decomposition in stochastic integer programming. *Operations Research Letters* 24, 1-2 (1999), 37–45.
- [9] Rich Caruana, Alexandru Niculescu-Mizil, Geoff Crew, and Alex Ksikes. 2004. Ensemble selection from libraries of models. In *Proceedings, Twenty-First International Conference on Machine Learning, ICML 2004*. <https://doi.org/10.1145/1015330.1015432>
- [10] Boyuan Chen, Harvey Wu, Warren Mo, Ishanu Chattopadhyay, and Hod Lipson. 2018. Autostacker: A compositional evolutionary learning system. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 402–409.
- [11] Alex G.C. de Sá, Walter José G.S. Pinto, Luiz Otavio V.B. Oliveira, and Gisele L. Pappa. 2017. RECIPE: A grammar-based framework for automatically evolving classification pipelines. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*.
- [12] Christopher De Sa, Alex Ratner, Christopher Ré, Jaeho Shin, Feiran Wang, Sen Wu, and Ce Zhang. 2016. Deepdive: Declarative knowledge base construction. *ACM SIGMOD Record* 45, 1 (2016), 60–67.
- [13] Rina Dechter. 1998. Bucket elimination: A unifying framework for probabilistic inference. In *Learning in graphical models*. Springer, 75–104.
- [14] Iddo Drori, Yamuna Krishnamurthy, Remi Rampin, Raoni De, Paula Lourenco, Jorge Piazentin Ono, Kyunghyun Cho, Claudio Silva, and Juliana Freire. 2018. AlphaD3M: Machine Learning Pipeline Synthesis. *AutoML Workshop at ICML* (2018).
- [15] Valeria Efimova, Andrey Filchenkov, and Viacheslav Shalakov. 2017. Fast Automated Selection of Learning Algorithm And its Hyperparameters by Reinforcement Learning. In *International Conference on Machine Learning AutoML Workshop*.
- [16] Katharina Eggensperger, Matthias Feurer, Frank Hutter, James Bergstra, Jasper Snoek, Holger Hoos, and Kevin Leyton-Brown. 2013. Towards an empirical foundation for assessing bayesian optimization of hyperparameters. In *NIPS workshop on Bayesian Optimization in Theory and Practice*, Vol. 10. 3.
- [17] Stefan Falkner, Aaron Klein, and Frank Hutter. 2018. BOHB: Robust and efficient hyperparameter optimization at scale. In *International Conference on Machine Learning*. PMLR, 1437–1446.
- [18] Matthias Feurer, Aaron Klein, Katharina Eggensperger, Jost Springenberg, Manuel Blum, and Frank Hutter. 2015. Efficient and robust automated machine learning. In *Advances in neural information processing systems*. 2962–2970.
- [19] Matthias Feurer, Benjamin Letham, and Eytan Bakshy. 2018. Scalable meta-learning for bayesian optimization using ranking-weighted gaussian process ensembles. In *AutoML Workshop at ICML*.
- [20] Hector Garcia-Molina, Jeffrey D. Ullman, and Jennifer Widom. 2008. *Database Systems: The Complete Book* (2 ed.). Prentice Hall Press, USA.
- [21] Amol Ghoting, Rajasekar Krishnamurthy, Edwin Pednault, Berthold Reinwald, Vikas Sindhwani, Shirish Tatikonda, Yuanyuan Tian, and Shivakumar Vaithyanathan. 2011. SystemML: Declarative machine learning on MapReduce. In *2011 IEEE 27th International Conference on Data Engineering*. IEEE, 231–242.
- [22] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and D Sculley. 2017. Google vizard: A service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1487–1495.
- [23] Google. 2020. Google Prediction API, <https://developers.google.com/prediction>.
- [24] Goetz Graefe. 1994. Volcano—An Extensible and Parallel Query Evaluation System. *IEEE Transactions on Knowledge and Data Engineering* (1994).
- [25] Xin He, Kaiyong Zhao, and Xiaowen Chu. 2021. AutoML: A Survey of the State-of-the-Art. *Knowledge-Based Systems* 212 (2021), 106622.
- [26] Yi-Qi Hu, Yang Yu, Wei-Wei Tu, Qiang Yang, Yuqiang Chen, and Wenyuan Dai. 2019. Multi-Fidelity Automatic Hyper-Parameter Tuning via Transfer Series Expansion. *AAAI* (2019).
- [27] Frank Hutter, Holger Hoos, and Kevin Leyton-Brown. 2014. An efficient approach for assessing hyperparameter importance. In *31st International Conference on Machine Learning, ICML 2014*.
- [28] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2011. Sequential model-based optimization for general algorithm configuration. In *International Conference on Learning and Intelligent Optimization*. Springer, 507–523.
- [29] Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren (Eds.). 2018. *Automated Machine Learning: Methods, Systems, Challenges*. Springer. In press, available at <http://automl.org/book>.
- [30] Frank Hutter, Jörg Lücke, and Lars Schmidt-Thieme. 2015. Beyond manual tuning of hyperparameters. *KI-Künstliche Intelligenz* 29, 4 (2015), 329–337.
- [31] IBM. 2020. IBM Watson Studio AutoAI. <https://www.ibm.com/cloud/watson-studio/autoai>.
- [32] Kevin Jamieson and Ameet Talwalkar. 2016. Non-stochastic best arm identification and hyperparameter optimization. In *Artificial Intelligence and Statistics*. 240–248.
- [33] Donald R Jones, Matthias Schonlau, and William J Welch. 1998. Efficient global optimization of expensive black-box functions. *Journal of Global optimization* 13, 4 (1998), 455–492.
- [34] Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and Barnabás Póczos. 2017. Multi-fidelity bayesian optimisation with continuous approximations. In *International Conference on Machine Learning*. PMLR, 1799–1808.
- [35] James Max Kanter and Kalyan Veeramachaneni. 2015. Deep feature synthesis: Towards automating data science endeavors. In *2015 IEEE International Conference on Data Science and Advanced Analytics, DSAA 2015, Paris, France, October 19-21, 2015*. IEEE, 1–10.
- [36] Gilad Katz, Eui Chul Richard Shin, and Dawn Song. 2016. Explorekit: Automatic feature generation and selection. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 979–984.
- [37] Ambika Kaul, Saket Maheshwary, and Vikram Pudi. 2017. Autolearn—Automated feature generation and selection. In *2017 IEEE International Conference on data mining (ICDM)*. IEEE, 217–226.
- [38] Udayan Khurana, Horst Samulowitz, and Deepak Turaga. 2018. Feature engineering for predictive modeling using reinforcement learning. In *32nd AAAI Conf. Artif. Intell. AAAI 2018*.
- [39] Udayan Khurana, Deepak Turaga, Horst Samulowitz, and Srinivasan Parthasarathy. 2016. Cognito: Automated feature engineering for supervised learning. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*. IEEE, 1304–1307.
- [40] Aaron Klein, Stefan Falkner, Simon Bartels, Philipp Hennig, and Frank Hutter. 2017. Fast Bayesian Optimization of Machine Learning Hyperparameters on Large Datasets. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. 528–536.
- [41] Brent Komer, James Bergstra, and Chris Eliasmith. 2014. Hyperopt-sklearn: automatic hyperparameter configuration for scikit-learn. In *ICML workshop on AutoML*, Vol. 9. Citeseer.
- [42] Tim Kraska. 2018. Northstar: An interactive data science system. *Proceedings of the VLDB Endowment* 11, 12 (2018), 2150–2164.
- [43] Sanjay Krishnan, Jiannan Wang, Eugene Wu, Michael J Franklin, and Ken Goldberg. 2016. Activeclean: Interactive data cleaning for statistical modeling. *Proceedings of the VLDB Endowment* 9, 12 (2016), 948–959.
- [44] Hai Lan, Zhifeng Bao, and Yuwei Peng. 2021. A survey on advancing the dbms query optimizer: Cardinality estimation, cost model, and plan enumeration. *Data Science and Engineering* 6, 1 (2021), 86–101.
- [45] Erin LeDell and S Poirier. 2020. H2o automl: Scalable automatic machine learning. In *Proceedings of the AutoML Workshop at ICML*, Vol. 2020.
- [46] Nir Levine, Koby Crammer, and Shie Mannor. 2017. Rotting bandits. In *Advances in NIPS*. 3074–3083.
- [47] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. 2018. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Proceedings of the International Conference on Learning Representations* (2018), 1–48.
- [48] Tian Li, Jie Zhong, Ji Liu, Wentao Wu, and Ce Zhang. 2018. Ease. ml: Towards multi-tenant resource sharing for machine learning workloads. *Proceedings of the VLDB Endowment* 11, 5 (2018), 607–620.
- [49] Yang Li, Jiawei Jiang, Jinyang Gao, Yingxia Shao, Ce Zhang, and Bin Cui. 2020. Efficient Automatic CASH via Rising Bandits. In *AAAI*. 4763–4771.
- [50] Yang Li, Yu Shen, Jiawei Jiang, Jinyang Gao, Ce Zhang, and Bin Cui. 2021. MFES-HB: Efficient Hyperband with Multi-Fidelity Quality Measurements. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 8491–8500.
- [51] Yang Li, Yu Shen, Wentao Zhang, Yuanwei Chen, Huaijun Jiang, Mingchao Liu, Jiawei Jiang, Jinyang Gao, Wentao Wu, Zhi Yang, et al. 2021. OpenBox: A Generalized Black-box Optimization Service. *arXiv preprint arXiv:2106.00421* (2021).
- [52] Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E Gonzalez, and Ion Stoica. 2018. Tune: A Research Platform for Distributed Model Selection and Training. *arXiv preprint arXiv:1807.05118* (2018).

- [53] Edo Liberty, Zohar Karmin, Bing Xiang, Laurence Rouesnel, Baris Coskun, Ramesh Nallapati, Julio Delgado, Amir Sadoughi, Yury Astashonok, Piali Das, et al. 2020. Elastic Machine Learning Algorithms in Amazon SageMaker. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 731–737.
- [54] Sijia Liu, Parikshit Ram, Djallel Bouneffouf, Gregory Bramble, Andrew R Conn, Horst Samulowitz, and Alexander G Gray. 2020. An ADMM Based Framework for AutoML Pipeline Configuration. (2020), 4892–4899.
- [55] Felix Mohr, Marcel Wever, and Eyke Hüllermeier. 2018. ML-Plan: Automated machine learning via hierarchical planning. *Machine Learning* 107, 8 (2018), 1495–1515.
- [56] Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. 2018. Ray: A distributed framework for emerging {AI} applications. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*. 561–577.
- [57] Supun Nakandala, Arun Kumar, and Yannis Papakonstantinou. 2019. Incremental and approximate inference for faster occlusion-based deep cnn explanations. In *Proceedings of the 2019 International Conference on Management of Data*. 1589–1606.
- [58] Supun Nakandala, Yuhao Zhang, and Arun Kumar. 2020. Cerebro: A data system for optimized deep learning model selection. *Proceedings of the VLDB Endowment* 13, 12 (2020), 2159–2173.
- [59] Fatemeh Nargesian, Horst Samulowitz, Udayan Khurana, Elias B Khalil, and Deepak S Turaga. 2017. Learning Feature Engineering for Classification.. In *Ijcai*. 2529–2535.
- [60] Randal S Olson and Jason H Moore. 2019. TPOT: A tree-based pipeline optimization tool for automating machine learning. In *Automated Machine Learning*. Springer, 151–160.
- [61] Matthias Poloczek, Jialei Wang, and Peter Frazier. 2017. Multi-information source optimization. In *Advances in Neural Information Processing Systems*. 4288–4298.
- [62] Alexander Ratner et al. 2017. Snorkel: Rapid Training Data Creation with Weak Supervision. *PVLDB* (2017).
- [63] Theodoros Rekatsinas, Xu Chu, Ihab F Ilyas, and Christopher Ré. 2017. HoloClean: Holistic Data Repairs with Probabilistic Inference. *Proceedings of the VLDB Endowment* 10, 11 (2017).
- [64] Microsoft Research. 2020. Microsoft NNI. <https://github.com/Microsoft/nni>.
- [65] Kevin Schawinski et al. 2017. Generative Adversarial Networks recover features in astrophysical images of galaxies beyond the deconvolution limit. *MNRAS Letters* (2017).
- [66] Rajat Sen, Kirthevasan Kandasamy, and Sanjay Shakkottai. 2018. Noisy blackbox optimization with multi-fidelity queries: A tree search approach. *arXiv preprint arXiv:1810.10482* (2018).
- [67] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* 104, 1 (2015), 148–175.
- [68] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*. 2951–2959.
- [69] Kevin Swersky, Jasper Snoek, and Ryan P Adams. 2013. Multi-task bayesian optimization. In *Advances in neural information processing systems*. 2004–2012.
- [70] Shion Takeno, Hitoshi Fukuoka, Yuhki Tsukada, Toshiyuki Koyama, Motoki Shiga, Ichiro Takeuchi, and Masayuki Karasuyama. 2020. Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization. In *International Conference on Machine Learning*. PMLR, 9334–9345.
- [71] Chris Thornton, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2013. Auto-WEKA: Combined selection and hyperparameter optimization of classification algorithms. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 847–855.
- [72] Jan N Van Rijn and Frank Hutter. 2018. Hyperparameter importance across datasets. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2367–2376.
- [73] Joaquin Vanschoren, Jan N Van Rijn, Bernd Bischl, and Luis Torgo. 2014. OpenML: networked science in machine learning. *ACM SIGKDD Explorations Newsletter* 15, 2 (2014), 49–60.
- [74] Manasi Vartak et al. 2016. ModelDB: A System for Machine Learning Model Management. In *HILDA*.
- [75] Ziyu Wang, Masrouz Zoghi, Frank Hutter, David Matheson, and Nando De Freitas. 2013. Bayesian optimization in high dimensions via random embeddings. In *Twenty-Third International Joint Conference on Artificial Intelligence*.
- [76] Martin Wistuba, Nicolas Schilling, and Lars Schmidt-Thieme. 2016. Two-stage transfer surrogate model for automatic hyperparameter optimization. In *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 199–214.
- [77] Jian Wu, Saul Toscano-Palmerin, Peter I Frazier, and Andrew Gordon Wilson. 2020. Practical multi-fidelity bayesian optimization for hyperparameter tuning. In *Uncertainty in Artificial Intelligence*. PMLR, 788–798.
- [78] Renzhi Wu, Sanya Chaba, Saurabh Sawlani, Xu Chu, and Saravanan Thirumuganathan. 2020. Zeroer: Entity resolution using zero labeled examples. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 1149–1164.
- [79] Weiyuan Wu, Lampros Flokas, Eugene Wu, and Jiannan Wang. 2020. Complaint-driven training data debugging for query 2.0. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 1317–1334.
- [80] Quanming Yao, Mengshuo Wang, Yuqiang Chen, Wenyuan Dai, Yu-Feng Li, Wei-Wei Tu, Qiang Yang, and Yang Yu. 2018. Taking human out of learning applications: A survey on automated machine learning. *arXiv preprint arXiv:1810.13306* (2018).
- [81] M. Zaharia et al. 2018. Accelerating the Machine Learning Lifecycle with MLflow. *IEEE Data Eng. Bull.* (2018).
- [82] Wentao Zhang, Jiawei Jiang, Yingxia Shao, and Bin Cui. 2020. Efficient diversity-driven ensemble for deep neural networks. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. IEEE, 73–84.
- [83] Wentao Zhang, Jiawei Jiang, Yingxia Shao, and Bin Cui. 2020. Snapshot boosting: a fast ensemble framework for deep neural networks. *Science China Information Sciences* 63, 1 (2020), 1–12.
- [84] Marc-André Zöllner and Marco F Huber. 2019. Survey on automated machine learning. *arXiv preprint arXiv:1904.12054* (2019).