The Machine Learning Bazaar: Harnessing the ML Ecosystem for Effective System Development

Micah J. Smith
MIT
micahs@mit.edu

Carles Sala
MIT
csala@csail.mit.edu

James Max Kanter
Feature Labs
max.kanter@featurelabs.com

Kalyan
Veeramachaneni
MIT
kalyanv@mit.edu

ABSTRACT

As machine learning is applied more widely, data scientists often struggle to find or create end-to-end machine learning systems for specific tasks. The proliferation of libraries and frameworks and the complexity of the tasks have led to the emergence of “pipeline jungles” — brittle, ad hoc ML systems. To address these problems, we introduce the *Machine Learning Bazaar*, a new framework for developing machine learning and automated machine learning software systems. First, we introduce ML primitives, a unified API and specification for data processing and ML components from different software libraries. Next, we compose primitives into usable ML pipelines, abstracting away glue code, data flow, and data storage. We further pair these pipelines with a hierarchy of AutoML strategies — Bayesian optimization and bandit learning. We use these components to create a general-purpose, multi-task, end-to-end AutoML system that provides solutions to a variety of data modalities (image, text, graph, tabular, relational, etc.) and problem types (classification, regression, anomaly detection, graph matching, etc.). We demonstrate 5 real-world use cases and 2 case studies of our approach. Finally, we present an evaluation suite of 456 real-world ML tasks and describe the characteristics of 2.5 million pipelines searched over this task suite.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Software and its engineering** → *Abstraction, modeling and modularity*; *Software development techniques*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMOD'20, June 14–19, 2020, Portland, OR, USA
© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6735-6/20/06...\$15.00
<https://doi.org/10.1145/3318464.3386146>

KEYWORDS

machine learning; AutoML; software development; ML pipelines; ML primitives

ACM Reference Format:

Micah J. Smith, Carles Sala, James Max Kanter, and Kalyan Veeramachaneni. 2020. The Machine Learning Bazaar: Harnessing the ML Ecosystem for Effective System Development. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD'20)*, June 14–19, 2020, Portland, OR, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3318464.3386146>

1 INTRODUCTION

Once limited to conventional commercial applications, machine learning (ML) is now being widely applied in physical and social sciences, in policy and government, and in a variety of industries. This diversification has led to difficulties in actually creating and deploying real-world systems, as key functionality becomes fragmented across ML-specific or domain-specific software libraries created by independent communities. In addition, the process of building problem-specific end-to-end systems continues to be marked by ML and data management challenges, such as formulating achievable learning problems [29], managing and cleaning data and metadata [6, 38, 49], scaling tuning procedures [17, 34], and deploying models and serving predictions [3, 12]. In practice, engineers and data scientists often spend significant effort developing ad hoc programs for new problems: writing “glue code” to connect components from different software libraries, processing different forms of raw input, and interfacing with external systems. These steps are tedious and error-prone and lead to the emergence of brittle “pipeline jungles” [43].

These points raise the question, “How can we make building ML systems easier in practical settings?” A new approach is needed to designing and developing software systems that solve specific ML tasks. Such an approach should address a wide variety of input data modalities, such as images, text, audio, signals, tables, graphs, and learning problem types, such as regression, classification, clustering, anomaly detection, community detection, graph matching; it should

cover the intermediate stages involved, such as data preprocessing, munging, featurization, modeling, and evaluation; and it should enable AutoML functionality to fine-tune solutions, such as hyperparameter tuning and algorithm selection. Moreover, it should offer coherent APIs, fast iteration on ideas, and easy integration of new ML innovations. In sum, this ambitious goal would allow almost all end-to-end learning problems to be solved or built using a single framework.

To address these challenges, we present the *Machine Learning Bazaar*,¹ a framework for designing and developing ML and AutoML systems. We organize the ML ecosystem into composable software components, ranging from basic building blocks like individual classifiers to full AutoML systems. With our design, a user specifies a task, provides a raw dataset, and either composes an end-to-end pipeline out of pre-existing, annotated, ML primitives or requests a curated pipeline for their task (Figure 1). The resulting pipelines can be easily evaluated and deployed across a variety of software and hardware settings and tuned using a hierarchy of AutoML approaches. Using our own framework, we have created an AutoML system which we have entered in participation in DARPA’s Data-Driven Discovery of Models (D3M) program; ours is the first end-to-end, modular, publicly released system designed to meet the program’s goal.

To preview the potential of development using our framework, we highlight the Orion project within MIT for ML-based anomaly detection in satellite telemetry (Figure 2), as one of several successful real-world applications that uses *ML Bazaar* for effective ML system development (Section 4.1). The Orion pipeline processes a telemetry signal using several custom preprocessors, an LSTM predictor, and a dynamic thresholding postprocessor to identify anomalies. The entire pipeline can be represented in a short Python snippet, custom processing steps are easily implemented as modular components, two external libraries are integrated without glue code, and the pipeline can be tuned using composable AutoML functionality.

Our contributions in this paper include:

A composable framework for representing and developing ML and AutoML systems: Our framework enables users to specify a pipeline for any ML task, ranging from image classification to graph matching, through a unified API (Sections 2 and 3).

The first general-purpose automated machine learning system: Our system, AutoBazaar, is, to the best of our knowledge, the first open-source, publicly-available, system

with the ability to reliably compose end-to-end, automatically-tuned, solutions for 15 data modalities and problem types (Section 3.3).

Successful applications: We describe 5 successful application of our framework on real-world problems (Section 4).

A comprehensive evaluation: We evaluated our AutoML system against a suite of 456 ML tasks/datasets covering 15 ML task types, analyzing 2.5 million scored pipelines (Section 5).

Open-source libraries: Components of our framework have been released as the open-source libraries *MLPrimitives*, *MLBlocks*, *BTB*, *piex*, and *AutoBazaar*.

2 THE MACHINE LEARNING BAZAAR

The *ML Bazaar* is a composable framework for developing ML and AutoML systems based on a hierarchical organization of and unified API for the ecosystem of ML software and algorithms. One can use curated or custom software components for every aspect of the practical ML process, from featurizers for relational datasets to signal processing transformers to neural networks to pre-trained embeddings. From these *primitives*, data scientists can easily and efficiently construct ML solutions for a variety of ML task types, and ultimately, automate much of the work of tuning these models.

2.1 ML Primitives

A *primitive* is a reusable, self-contained, software component for ML paired with the structured *annotation* of its metadata. It has a well-defined *fit/produce* interface wherein it receives input data in one of several formats or types, performs computations, and returns the data in another format or type. With this categorization and abstraction, widely varying functionality required to construct ML pipelines can be collected in a single location. Primitives can be re-used in chained computations while minimizing glue code written by callers. An example primitive annotation is shown in Listing 1.

Primitives encapsulate different types of functionality. Many have a learning component, such as a random forest classifier. Many primitives, categorized as transformers, may have no learning component and only have a *produce* method, but are very important nonetheless. For example, the Hilbert and Hadamard transforms from signal processing would be important primitives to include when building an ML system to solve a task in Internet-of-Things.

Some primitives do not change the values in the data, but simply prepare or reshape the data. These *glue primitives* are intended to reduce glue code required to connect primitives

¹Just as one open-source community was described as “a great babbling bazaar of different agendas and approaches” [41], our framework is characterized by the availability of many compatible alternatives, a wide variety of libraries and custom solutions, a space for new contributions, and more.

Figure 1: Various ML task types that can be solved in *ML Bazaar* using composition of ML primitives (abbreviated here from fully-qualified names). Primitives are categorized into preprocessors, feature processors, estimators, and postprocessors and are drawn from many different ML libraries, such as scikit-learn, Keras, OpenCV, and NetworkX, as well as custom implementations. Many additional primitives and pipelines are available in our curated catalog.

into a full system. An example of this type of primitive is `pandas.DataFrame.unstack`.

2.1.1 Annotations. Each primitive is annotated with machine-readable metadata that enables its usage and automatic integration within an execution engine. Annotations allow us to unify a variety of primitives from disparate libraries, reduce the need for glue code, and provide information about the tunable hyperparameters. This full annotation² is provided in a single JSON file and has three major sections:

- **Meta-information.** This section has the name of the primitive, the fully-qualified name of the underlying implementation as a Python object, and other detailed metadata, such as the author, description, documentation URL, categorization, and the data modalities it is most used for. This information enables searching and indexing primitives.
- **Information required for execution.** This section has the names of the methods pertaining to `fit`/`produce` in the underlying implementation as well as the data types of the primitive's inputs and outputs. When applicable, for each primitive, we annotate the *ML data types* of declared inputs and outputs, i.e., recurring objects in ML that have a well-defined semantic meaning, such as a feature matrix X , a target vector y , or a space of class

```
{
  "name": "cv2.GaussianBlur",
  "contributors": [
    "Carles Sala <csala@csail.mit.edu>"
  ],
  "description": "Blur an image using a Gaussian filter.",
  "classifiers": {
    "type": "preprocessor", "subtype": "transformer"
  },
  "modalities": ["image"],
  "primitive": "mlprimitives.adapters.cv2.GaussianBlur",
  "produce": {
    "args": [{"name": "X", "type": "ndarray"}],
    "output": [{"name": "X", "type": "ndarray"}]
  },
  "hyperparameters": {
    "fixed": {
      "ksize_width": {"type": "int", "default": 3},
      "ksize_height": {"type": "int", "default": 3},
      "sigma_x": {"type": "float", "default": 0},
      "sigma_y": {"type": "float", "default": 0}
    },
    "tunable": {}
  },
  "...": "..."
}
```

Listing 1: JSON representation of the GaussianBlur transformer primitive using `MLPrimitives`. (Some fields are abbreviated or elided.) This primitive does not annotate any *tunable hyperparameters* but such a section marks hyperparameter types, defaults, and feasible values.

²The primitive annotation specification is described and documented in full in the associated `MLPrimitives` library.

labels classes. We provide a mapping between ML data types and synonyms used by specific libraries as

Figure 2: Representation and usage of the Orion pipeline for anomaly detection using the *ML Bazaar* framework. ML system developers or researchers describe the pipeline in a short Python snippet by a sequence of primitives annotated from several libraries (and optional additional parameters). Our framework compiles this into a graph representation (Section 2.2.2) by consulting meta-information associated with the underlying primitives (Section 2.1.1). Developers can then use our Python SDK to train the pipeline on “normal” signals and then identify anomalies in test signals. The *MLPipeline* provides a familiar interface but enables more general data engineering and ML processing. It also can expose the entire underlying hyperparameter configuration space for tuning by our AutoML libraries or others (Section 3).

necessary. This logical structure will help dramatically decrease the amount of glue code developers must write (Section 2.2.1).

- *Information about hyperparameters.* The third section details all the hyperparameters of the primitive — their names, descriptions, data types, ranges, and whether they are *fixed* or *tunable*. It also captures any conditional dependencies between the hyperparameters.

We have developed the open-source *MLPrimitives*³ library which contains a number of primitives adapted from different libraries (Table 1). For libraries that already provide a fit/produce interface or similar (e.g. scikit-learn), a primitive developer has to write the JSON specification and point to the underlying estimator class.

To support integration of primitives from libraries that need significant adaptation to the fit/produce interface,

MLPrimitives also provides a powerful set of adapter modules that assist in wrapping common patterns. These adapter modules then allow us to integrate many functionalities as primitives from the library without having to write a separate object for each — thus requiring us to write only an annotation file for each primitive. Keras is an example of such a library.

For developers, domain experts, and researchers, *MLPrimitives* enables easy *contribution* of new primitives in several ways by providing primitive templates, example annotations, and detailed tutorials and documentation. We also provide procedures to validate proposed primitives against the formal specification and a unit test suite. Finally, contributors can also write custom primitives.

Currently, *MLPrimitives* maintains a curated *catalog* of high-quality, useful primitives from 12 libraries,⁴ as well as

³<https://github.com/HDI-Project/MLPrimitives>

⁴As of *MLPrimitives* v0.2.4.

Source	Count	Source	Count
scikit-learn	39	XGBoost	2
MLPrimitives (custom)	27	LightFM	1
Keras	25	OpenCV	1
pandas	16	python-louvain	1
Featuretools	4	scikit-image	1
NumPy	3	statsmodels	1
NetworkX	2		

Table 1: Primitives in the curated catalog of MLPrimitives, by library source. Catalogs maintained by individual projects may contain more primitives.

custom primitives that we have created (Table 1). Each primitive is identified by a fully-qualified name to differentiate primitives across catalogs. The JSON annotations can then be mined for additional insights.

2.1.2 Designing for contributions. We considered multiple alternatives to the primitives API, such as representing all of them as Python data structures or classes, regardless of their type (i.e. transformers or estimators). One disadvantage of these alternatives is that it makes it more difficult for domain experts to contribute primitives. We have found that domain experts, such as engineers and scientists in the satellite industry, prefer writing functions rather than other constructs such as classes, and many domain-specific processing methods are simply transformers without a learning component.

2.1.3 Lightweight integration. Another option we considered was to enforce that every primitive—whether brought over from a library with a compatible API or otherwise—be integrated via a Python class with wrapper methods. We opted against this approach as it led to excessive wrapper code and created redundancy, which made it more difficult to write primitives. Instead, for libraries that are compatible, our design requires that we only create the annotation file.

2.1.4 Language independence. In this work, we focus on the wealth of ML functionality that exists in the Python ecosystem. Through *ML Bazaar*'s careful design, we could also support other common languages in data science like R, MATLAB, and Julia and enable multi-language pipelines. Starting from our JSON primitive annotation format, a multi-language pipeline execution backend would be built that uses language-specific kernels or containers and relies on an interoperable data format such as Apache Arrow. A language-independent format like JSON provides several additional advantages. It is both machine- and human- readable and writeable. It is also a natural format for storage and querying in NoSQL document stores, allowing developers to easily

query a knowledge base of primitives for the subset appropriate for a specific ML task type, for example.

2.2 Building ML pipelines

To solve practical learning problems, we must be able to instantiate and compose primitives into usable programs. These programs must be easy to specify with a natural interface, such that developers can easily compose primitives without sacrificing flexibility. We aim to support both end-users trying to build an ML solution for their specific problem who may not be savvy about software engineering, as well as system developers wrapping individual ML solutions in AutoML components. In addition, we provide an abstracted execution layer, such that learning, data flow, data storage, and deployment are handled automatically by various configurable and pluggable backends. As one realization of these ideas, we have implemented MLBlocks,⁵ a library for composing, training, and deploying end-to-end ML pipelines.

```
from mlblocks import MLPipeline
from mlblocks.datasets import load_umls
from mlblocks.discovery import load_pipeline

dataset = load_umls()
X_train, X_test, y_train, y_test = dataset.get_splits(1)
graph = dataset.graph
node_columns = ['source', 'target']

ppl = MLPipeline(load_pipeline('graph.link_prediction.nx.xgb'))
ppl.fit(X_train, y_train, graph=graph, node_columns=node_columns)
y_pred = ppl.predict(X_test, graph=graph,
                    node_columns=node_columns)
```

Listing 2: Usage of the MLBlocks library for a graph link prediction task. Curated pipelines in the MLPrimitives library can be easily loaded. Pipelines provide a familiar API but enable more general data engineering and ML.

2.2.1 Steps and Pipelines. We introduce *ML pipelines*, which collect multiple primitives into a single computational graph. Each primitive in the graph is instantiated in a *pipeline step*, which loads and interprets the underlying primitive and provides a common interface to run a step in a larger program.

We define a *pipeline* as a directed acyclic multigraph $L = \langle V, E, \lambda \rangle$, where V is a collection of pipeline steps, E are the directed edges between steps representing data flow, and λ is a joint hyperparameter vector for the underlying primitives. A valid pipeline — and its generalizations (Section 3.1) — must also satisfy acceptability constraints that require the inputs to each step to be satisfied by the outputs of another step connected by a directed edge.

⁵<https://github.com/HDI-Project/MLBlocks>

The term “pipeline” is used in the literature to refer to a ML-specific sequence of operations, and sometimes abused (as we do here) to refer to a more general computational graph or analysis. In our conception, we bring foundational data processing operations of raw inputs into this scope, like featurization of graphs, multi-table relational data, time series, text, and images, as well as simple data transforms, like encoding integer or string targets. This gives our pipelines a greatly expanded role, providing solutions to any ML task type and spanning the entire ML process beginning with the raw dataset.

Figure 3: Recovery of ML computational graph from pipeline description for a text classification pipeline. The ML data types that enable extraction of the graph, and stand for data flow, are labeled along edges.

2.2.2 Pipeline description interface. Large graph-structured workloads can be difficult to specify for end-users due to the complexity of the data structure and such workloads are an active area of research in data management. In *ML Bazaar*, we consider three aspects of pipeline representation: ease of composition, readability, and computational issues. First, we prioritize easily composing complex ML pipelines by providing a *pipeline description interface* (PDI) in which developers specify only the topological ordering of all pipeline steps in the pipeline without requiring any explicit dependency declarations. These steps can be passed to our libraries as Python data structures or loaded from JSON files. Full training-time (fit) and inference-time (produce) computational graphs can then be recovered (Algorithm 1). This is made possible by the meta-information provided in the primitive annotations, in particular, the ML data types of the primitive inputs and outputs. We leverage the observation that steps that modify the same ML data type can be grouped into the same subpath. In cases where this information does not uniquely identify a graph, the user can additionally provide an *input-output map* which serves to explicitly add edges to the graph, as well as other parameters to customize the pipeline.

Though it may be more difficult to read and understand these pipelines from the PDI alone as the edges are not shown

nor labeled, it is easy to accompany them with the recovered graph representation (Figures 2 and 3).

Input: pipeline description $S = (v_1, \dots, v_n)$, source node v_0 , sink node v_{n+1}
Output: directed acyclic multigraph $\langle V, E \rangle$

```

begin
   $S \leftarrow v_0 \cup S \cup v_{n+1}$ 
   $V \leftarrow \emptyset, E \leftarrow \emptyset$ 
   $U \leftarrow \emptyset$  // unsatisfied inputs
  while  $S \neq \emptyset$  do
     $v \leftarrow \text{popright}(S)$  // last pipeline step remaining
     $M \leftarrow \text{popmatches}(U, \text{outputs}(v))$ 
    if  $M \neq \emptyset$  then
       $V \leftarrow V \cup \{v\}$ 
      for  $(v', \sigma) \in M$  do
         $E \leftarrow E \cup \{(v, v', \sigma)\}$ 
        for  $\sigma \in \text{inputs}(v)$  do // unsatisfied inputs of v
           $U \leftarrow U \cup \{(v, \sigma)\}$ 
      else // isolated node
        return INVALID
    if  $U \neq \emptyset$  then // unsatisfied inputs remain
      return INVALID
  return  $\langle V, E \rangle$ 

```

Algorithm 1: Pipeline-Graph Recovery. Pipeline steps are added to the graph in reverse order and edges are iteratively added when the step under consideration produces an output that is required by an existing step. Exactly one graph is recovered if a valid graph exists. In cases where multiple graphs have the same topological ordering, the user can additionally provide an *input-output map* (which modifies the result of $\text{inputs}(v)/\text{outputs}(v)$ above) to explicitly add edges and thereby select from among several possible graphs.

The resulting graphs describe abstract computational workloads, but we must be able to actually execute them for purposes of learning and inference. From the recovered graphs, we could re-purpose many existing data engineering systems as backends for scheduling and executing the workloads [42, 45, 53]. In our MLBlocks execution engine, a collection of objects and a metadata tracker in a key-value store are iteratively transformed through sequential processing of pipeline steps. The Orion pipeline would be executed using MLBlocks as shown in Figure 2c.

2.3 Discussion

Why not scikit-learn? Several alternatives exist to our new ML pipeline abstraction (Section 2.2), such as scikit-learn’s Pipeline [11]. Ultimately, while our pipeline is inspired by these alternatives, it aims to provide more general data engineering and ML functionality. While the scikit-learn pipeline

sequentially applies a list of transformers to X and y only before outputting a prediction, our pipeline supports general computational graphs, accepts multiple data modalities as input simultaneously, produces multiple outputs, manages evolving metadata, and can use software from outside the scikit-learn ecosystem/design paradigm. For example, we can use our pipeline to construct entity sets [28] from multi-table relational data for input to other pipeline steps. We can also support pipelines in an unsupervised learning paradigm, such as in Orion, where we create the target y “on-the-fly” (Figure 3).

Where’d the glue go? To connect learning components from different libraries with incompatible APIs, data scientists end up writing “glue code”. Typically, this glue code is written within pipeline bodies. In *ML Bazaar*, we mitigate the need for this glue by pushing the need of API adaptation down to the level of primitive annotations, which are written once and reside in central locations, amortizing the adaptation cost. Moreover, the need for glue code arises in creating intermediate outputs and shaping of the data. We created a number of primitives that support these common programming patterns and miscellaneous needs in development of a ML pipeline. These are, for example, data reshaping primitives like `pandas.DataFrame.unstack`, data preparation primitives like `pad_sequences` required for Keras-based LSTMs, and utilities like `UniqueCounter` that count the number of unique classes.

Interactive development. Interactivity is an important aspect of data science development for beginners and experts alike, as they build understanding of the data and iterate on different modeling ideas. In *ML Bazaar*, the level of interactivity possible depends on the specific runtime library. For example, our *MLBlocks* library supports interactive development in a shell or notebook environment by allowing the inspection of intermediate pipeline outputs and by allowing pipelines to be iteratively expanded starting from a loaded pipeline description. Alternatively, ML primitives could be used as a backend pipeline representation for software that provides more advanced interactivity such as drag-and-drop. For interfaces that require low latency pipeline scoring to provide user feedback such as [13], *ML Bazaar*’s performance depends mainly on the underlying primitive implementations (Section 5).

Supporting new task types. While *ML Bazaar* handles 15 ML task types (Table 4), there are many more task types for which we do not currently provide pipelines in our default catalog (Section 5.5). To extend our approach to support new task types, it is generally sufficient to write several new primitive annotations for pre-processing input and post-processing output — no changes are needed to the core ML

Bazaar software libraries such as *MLPrimitives* and *MLBlocks*. For example, for the anomaly detection task type from the Orion project, several new simple primitives were implemented: `rolling_window_sequences`, `regression_errors`, and `find_anomalies`. Indeed, support for a certain task type is predicated on the availability of a pipeline for that task type rather than any characteristics of our software libraries.

Primitive versioning. The default catalog of primitives from the *MLPrimitives* library is versioned together, and library conflicts are resolved manually by maintainers through carefully specifying minimum and maximum dependencies. This strategy ensures that the default catalog can always be used, even if there are incompatible updates to the underlying libraries. Automated tools can be integrated to aid both end-users and maintainers in understanding potential conflicts and safely bumping library-wide versions.

3 AUTOML SYSTEM DESIGN AND ARCHITECTURE

From the components of the *ML Bazaar*, data scientists can easily and effectively build ML pipelines with fixed hyperparameters for their specific problems. To improve the performance of these solutions, we introduce the more general *pipeline templates* and *pipeline hypertemplates* and then present the design and implementation of AutoML primitives which facilitate hyperparameter tuning and model selection, either using our own library for Bayesian optimization or external AutoML libraries. Finally, we describe AutoBazaar, one specific AutoML system we have built on top of these components.

3.1 Pipeline templates and hypertemplates

Frequently, pipelines require hyperparameters to be specified at several places. Unless these values are fixed at annotation-time, hyperparameters must be exposed in a machine-friendly interface. This motivates pipeline templates and pipeline hypertemplates, which generalize pipelines by allowing a hierarchical tunable hyperparameter configuration space and provide first-class tuning support.

We define a *pipeline template* as a directed acyclic multigraph $T = \langle V, E, \Lambda \rangle$, where Λ is the joint hyperparameter configuration space for the underlying primitives. By providing values $\lambda \in \Lambda$ for the unset hyperparameters of a pipeline template, a specific pipeline is created.

In some cases, certain values of hyperparameters can affect the domains of other hyperparameters. For example, the type of kernel for a support vector machine results in different kernel hyperparameters, and preprocessors used to adjust for class imbalance can affect the training procedure of a downstream classifier. We call these *conditional hyperparameters*, and accommodate them with pipeline hypertemplates.

We define a *pipeline hypertemplate* as a directed acyclic multigraph $H = \langle V, E, \bigcup_j \Lambda_j \rangle$, where V is a collection of pipeline steps, E are directed edges between steps, and Λ_j is the hyperparameter configuration space for pipeline template T_j . A number of pipeline templates can be derived from one pipeline hypertemplate by fixing the conditional hyperparameters.

3.2 AutoML Primitives

Just as primitives units components of ML computation, AutoML primitives represent components of an AutoML system. We separate AutoML primitives into *tuners* and *selectors*. In our extensible AutoML library for developing AutoML systems, BTB,⁶ we provide various instances of these AutoML primitives.

3.2.1 Tuners. Given a pipeline template, an AutoML system must find a specific pipeline with fully-specified hyperparameter values to maximize some utility. Given pipeline template T and a function f that assigns a performance score to pipeline L_λ with hyperparameters $\lambda \in \Lambda$, the tuning problem is defined as $\lambda^* = \arg \max_{\lambda \in \Lambda} f(L_\lambda)$. We introduce *tuners*, AutoML primitives which provide a record/propose interface in which evaluation results are recorded to the tuner by the user or by an AutoML controller and new hyperparameters are proposed in return.

Hyperparameter tuning is widely studied and its effective use is instrumental to maximizing the performance of ML systems [4, 5, 18, 46]. One widely used approach to hyperparameter tuning is Bayesian optimization (BO), a black-box optimization technique in which expensive evaluations of f are kept to a minimum by forming and updating a meta-model for f . At each iteration, the next hyperparameter configuration to try is chosen according to an acquisition function. We structure these meta-models and acquisition functions as separate BO-specific AutoML primitives that can be combined together to form a tuner. Researchers have argued for different formulations of meta-models and acquisition functions [39, 46, 51]. In our BTB library for AutoML, we implement the GP-EI tuner, which uses a Gaussian Process meta-model primitive and an Expected Improvement (EI) acquisition function primitive, among several other tuners. Many other tuning paradigms exist, such as those based on evolutionary strategies [37, 40], adaptive execution [27, 33], meta-learning [21], or reinforcement learning [16]. Though we have not provided implementations of these in BTB, one could do so using our common API.

3.2.2 Selectors. For many ML task types, there may be multiple pipeline templates or pipeline hypertemplates available, each with their own tunable hyperparameters. The aim is

to balance the exploration-exploitation tradeoff while selecting promising pipeline templates to tune. For a set of pipeline templates $\mathcal{T}$, we define the selection problem as $T^* = \arg \max_{T \in \mathcal{T}} \max_{\lambda_T \in \Lambda_T} f(L_{\lambda_T})$. We introduce *selectors*, AutoML primitives which provide a `compute_rewards/select` API.

Algorithm selection is often treated as a multi-armed bandit problem where the score returned from a selected template can be assumed to come from an unknown underlying probability distribution. In BTB, we implement the UCB1 selector, which uses the upper confidence bound method [2], among several other selectors. Users or AutoML controllers can use selectors and tuners together to perform joint algorithm selection and hyperparameter tuning.

3.3 Building an AutoML system

Using the ML Bazaar framework, we have built AutoBazaar,⁷ an open-source, end-to-end, general-purpose, multi-task, automated machine learning system. It consists of several components: an AutoML controller; a pipeline execution engine; data stores for metadata and pipeline evaluation results; loaders and configuration for ML tasks, primitives, etc.; a Python language client; and a command-line interface. AutoBazaar is an open-source variant of the AutoML system we have developed for the DARPA D3M program.

We focus here on the core pipeline search and evaluation algorithms (Algorithm 2). The input to the search is a computational budget and an ML task, which consists of the raw data and task and dataset metadata — dataset resources, problem type, dataset partition specifications, and an evaluation procedure for scoring. Based on these inputs, AutoBazaar searches through its catalog of primitives and pipeline templates for the most suitable pipeline that it can build. First, the controller loads the train and test dataset partitions, $\mathcal{D}^{(train)}$ and $\mathcal{D}^{(test)}$, following the metadata specifications. Next, it loads from its default catalog and the user's custom catalog a collection of candidate pipeline templates suitable for the ML task type. Using the BTB library, it initializes a UCB1 selector and a collection of GP-EI tuners for joint algorithm selection and hyperparameter tuning. The search process begins and continues for as long as the computation budget has not been exhausted. In each iteration, the selector is queried to select a template, the corresponding tuner is queried to propose a hyperparameter configuration, a pipeline is generated and scored using cross validation over $\mathcal{D}^{(train)}$, and the score is reported back to the selector and tuner. The best overall pipeline found during the search, L^* , is re-fit on $\mathcal{D}^{(train)}$ and scored over $\mathcal{D}^{(test)}$. Its specification is returned to the user alongside the score obtained, s^* .

⁶<https://github.com/HDI-Project/BTB>

⁷<https://github.com/HDI-Project/AutoBazaar>


```

Input: task  $t = (M, f, \mathcal{D}^{(train)}, \mathcal{D}^{(test)})$ , budget  $B$ 
Output: best pipeline  $L^*$ , best score  $s^*$ 

begin
   $\mathcal{T} \leftarrow \text{load\_available\_templates}(M)$ 
   $A \leftarrow \text{init\_automl}(\mathcal{T})$  // bookkeeping
   $s^* \leftarrow +\infty, L^* \leftarrow \emptyset$ 
  while  $B > 0$  do
     $T \leftarrow \text{select}(A)$  // uses selector.select
     $\lambda \leftarrow \text{propose}(A, T)$  // uses T's tuner.propose
     $L \leftarrow (T, \lambda)$ 
     $s \leftarrow \text{cross\_validate\_score}(f, L, \mathcal{D}^{(train)})$ 
    record( $A, L, s$ ) // update selector and tuners
    if  $s < s^*$  then
      |  $s^* \leftarrow s, L^* \leftarrow L$ 
    decrease( $B$ )
   $s^* \leftarrow \text{fit\_and\_score}(f, L^*, \mathcal{D}^{(train)}, \mathcal{D}^{(test)})$ 
  return  $L^*, s^*$ 

```

Algorithm 2: Search and evaluation of pipelines in AutoBazaar. Detailed task metadata M is used by the system to load relevant pipeline templates and scorer function f is used to score pipelines.

4 APPLICATIONS

In this paper, we claim that *ML Bazaar* makes it easier to develop ML systems. We provide evidence for this claim in this section by describing 5 real-world use cases in which *ML Bazaar* is currently used to create both ML and AutoML systems. Through these industrial applications we examine the following questions: Does *ML Bazaar* support the needs of ML system developers? If not, how easy was it to extend?

4.1 Use cases

Anomaly detection for satellite telemetry. *ML Bazaar* is used by a communications satellite operator which provides video and data connectivity globally. This company wanted to monitor more than 10,000 telemetry signals from their satellites and identify anomalies, which might indicate a looming failure severely affecting the satellite's coverage. This time series/anomaly detection task was not initially supported by any of the pipelines in our curated catalog. Our collaborators were able to easily implement a recently developed end-to-end anomaly detection method [26] using pre-existing transformation primitives in *ML Bazaar* and by adding several new primitives: a primitive for the specific LSTM architecture used in the paper and new time series anomaly detection postprocessing primitives, which take as input a time series and time series forecast, and produce as output a list of anomalies, identified by intervals $\{[t_i, t_{i+1}]\}$. This design enabled rapid experimentation through substituting different time series forecasting primitives and comparing the results. In

current work, they apply ML pipelines to 82 publicly available satellite telemetry signals from NASA and evaluate the anomaly detections against 105 known anomalies. The work has been released as the open-source Orion project⁸ and is currently under active development.

Predicting clinical outcomes from electronic health records. Cardea is an open-source, automated framework for predictive modeling in health care on electronic health records following the FHIR schema. Its developers formulated a number of prediction problems including predicting length of hospital stay, missed appointments, and hospital readmission. All tasks in Cardea are multitable regression or classification. From *ML Bazaar*, Cardea uses the `featuretools.dfs` primitive to automatically engineer features for this highly-relational data and multiple other primitives for classification and regression. The framework also presents examples on a publicly available patient no-show prediction problem. The framework has been released as an open-source project.⁹

Failure prediction in wind turbines. *ML Bazaar* is also used by a multinational energy utility to predict critical failures and stoppages in their wind turbines. Most prediction problems here pertain to the time series classification ML task type. *ML Bazaar* has several time series classification pipelines available in its catalog and they enable usage of time series from 140 turbines to develop multiple pipelines, tune them, and produce prediction results. Multiple outcomes are predicted, ranging from stoppage and pitch failure to less common issues, such as gearbox failure. This library is released as the open-source GreenGuard project.¹⁰

Leaks and crack detection in water distribution systems. A global water technology provider uses *ML Bazaar* for a variety of ML needs, ranging from image classification for detecting leaks from images, to crack detection from time series data, to demand forecasting using water meter data. *ML Bazaar* provides a unified framework for these disparate needs. The team also builds custom primitives internally and uses them directly with the MLBlocks backend.

4.2 DARPA D3M program

DARPA's Data-Driven Discovery of Models (D3M) program, of which we are participants, aims to spur development of automated systems for model discovery for use by non-experts. Among other goals, participants aim to design and implement AutoML systems that can produce solutions to arbitrary ML tasks without any human involvement. We used *ML Bazaar* to create an AutoML system to be evaluated against

⁸<https://github.com/D3-AI/Orion>

⁹<https://github.com/D3-AI/Cardea>

¹⁰<https://github.com/D3-AI/GreenGuard>

System	Top pipeline	Beats Expert 1	Beats Expert 2	Rank
System 1	29	57	31	1
<i>ML Bazaar</i>	18	56	28	2
System 3	15	47	22	3
System 4	14	46	21	4
System 5	10	42	14	5
System 6	8	43	15	6
System 7	8	33	12	7
System 8	6	24	11	8
System 9	4	25	13	9
System 10	2	27	12	10

Table 2: Results from the DARPA D3M Summer 2019 evaluation (most recent conducted). Entries represent the number of ML tasks. “Top pipeline” is the number of tasks for which a system created a winning pipeline. “Beats Expert 1” and “Beats Expert 2” are the number of tasks for which a system beat the two expert team baselines. We highlight Systems 6 and 7 as they belong to the same teams as [44] and [16], respectively. (We are unable to comment on other systems as they have not yet provided public reports.) Rank is given based on number of top pipeline lines produced. The top 4 teams are consistent in their ranking even if a different column is chosen.

other teams from US academic institutions. Participants include ourselves (MIT), CMU, UC Berkeley, Brown, Stanford, TAMU, and others. Our system relies on AutoML primitives (Section 3) and other features of our framework, but does *not* use our primitive and pipeline implementations (neither MLPrimitives nor MLBlocks).

We present results comparing our system against other teams in the program. DARPA organizes an evaluation every 6 months (Winter and Summer). During evaluation, AutoML systems submitted by participants are run by DARPA on 95 tasks spanning several task types for three hours per task. At the end of the run, the best pipeline identified by the AutoML system is evaluated on held-out test data. Results are also compared against two independently-developed expert baselines (MIT Lincoln Laboratory and Exline).

Results from one such evaluation from Spring 2018 were presented by [44]. We make comparisons from the Summer 2019 evaluation, the results of which were released in August 2019 — the most recent evaluation as of this writing. Table 2 compares our AutoML system against 9 other teams. Given the same tasks and same machine learning primitives, this comparison highlights the efficacy of the AutoML primitives (BTB) in *ML Bazaar* only — it does not provide any evaluation of our other libraries. In its implementation, our system uses a GP-MAX tuner and a UCB1 selector. Across all metrics, our system places 2nd out of the 10 teams.

4.3 Discussion

Through these applications using the components of the *ML Bazaar*, several advantages surfaced.

4.3.1 Composability. One important aspect of *ML Bazaar* is that it does not restrict the user to use a single monolithic system, rather users can pick and choose parts of the framework they want to use. For example, Orion uses only MLPrimitives/MLBlocks, Cardea uses MLPrimitives but integrates the hyperopt library for hyperparameter tuning, our D3M AutoML system submission mainly uses AutoML primitives and BTB, and AutoBazaar uses every component.

4.3.2 Focus on infrastructure. The ease of developing ML systems for the task at hand freed up time for teams to think through and design a comprehensive ML infrastructure. In the case of Orion and GreenGuard, this led to the development of a database that catalogues the metadata from every ML experiment run using *ML Bazaar*. This had several positive effects: it allowed for easy sharing between team members, and it allowed the company to transfer the knowledge of what worked from one system to another system. For example, the satellite company plans to use the pipelines that worked on a previous generation of the satellites on the newer ones from the beginning. With multiple entities finding uses for such a database, creation of such infrastructure could be templated.

4.3.3 Multiple use cases. Our framework allowed the water technology company to solve many different ML task types using the same framework and API.

4.3.4 Fast experimentation. Once a baseline pipeline has been designed to solve a problem, we notice that users can quickly shift focus to developing and improving primitives that are responsible for learning.

4.3.5 Production ready. A fitted pipeline maintains all the learned parameters as well as all the data manipulations. A user is able to serialize the pipeline and load it into production. This reduces the development-to-production lifecycle.

5 EXPERIMENTAL EVALUATION

In this section, we experimentally evaluate *ML Bazaar* along several dimension. We also leverage our evaluation results to perform several case studies in which we show how a general-purpose evaluation setting can be used to assess the value of specific ML and AutoML primitives.

5.1 ML task suite

The *ML Bazaar* Task Suite is a comprehensive corpus of tasks and datasets to be used for evaluation, experimentation, and diagnostics. It consists of 456 ML tasks spanning

	min	p25	p50	p75	max
Number of examples	7	202	599	3,634	6,095,521
Number of classes [†]	2	2	3	6	115
Columns of X	1	3	9	22	10,937
Size (compressed)	3KiB	21KiB	145KiB	2MiB	36GiB
Size (inflated)	22KiB	117KiB	643KiB	7MiB	42GiB

Table 3: Summary of tasks in ML Bazaar Task Suite (n=456). [†] for classification tasks

15 task types. Tasks, which encompass raw datasets and annotated task descriptions, are assembled from a variety of sources, including MIT Lincoln Laboratory, Kaggle, OpenML, Quandl, and Crowdfunder. We created train/test splits and organized the folder structure. Other than this, we do not do any preprocessing (sampling, outlier detection, imputation, featurization, scaling, encoding, etc.), presenting data in its raw form as it would be ingested by end-to-end pipelines. Our publicly-available task suite can be browsed online¹¹ or through `piex`,¹² our library for exploration and meta-analysis of ML tasks and pipelines. The covered task types are shown in Table 4 and a summary of the tasks is shown in Table 3.

We made every effort to curate a corpus that was evenly balanced across ML task types. Unfortunately, in practice, available datasets are heavily skewed to traditional ML problems of single-table classification and our task suite reflects this deficiency (though 49% are not single-table classification). Indeed, among other evaluation suites, the OpenML 100 and the AutoML Benchmark [8, 19] are both exclusively comprised of single-table classification problems. Similarly, evaluation approaches for AutoML methods usually target the black-box optimization aspect in isolation [15, 20, 23] without considering the larger context of an end-to-end pipeline.

5.2 Pipeline search

We run the search process for all tasks in parallel on a heterogeneous cluster of 400 AWS EC2 nodes. Each ML task is solved independently on a node of its own over a 2-hour time limit. Metadata and fine-grained details about every pipeline evaluated are stored in a MongoDB document store. The best pipelines for each task, after checkpoints at 10, 30, 60, and 120 minutes of search, are selected by considering the cross-validation score on the training set and are then re-scored on the held-out test set.¹³

¹¹<https://mlbazaar.github.io>

¹²<https://github.com/HDI-Project/piex>

¹³Exact replication files and detailed instructions for the experiments in this section are included here: <https://github.com/micahjsmith/ml-bazaar-2019> and can be further analyzed using our `piex` library.

Figure 4: Execution time of AutoBazaar pipeline search attributable to different libraries/components. The box plot shows quartiles of the distribution, 1.5× IQR, and outliers. MLB Ext and BTB Ext refer to calls to external libraries providing underlying implementations, like the scikit-learn GaussianProcessRegressor used in the GP-EI tuner. The vast majority of execution time is attributed to the underlying primitives implemented in external libraries.

5.3 Computational bottlenecks

We first evaluate the computational bottlenecks of the AutoBazaar system. To assess this, we instrument AutoBazaar and our framework libraries (MLBlocks, MLPrimitives, BTB) to determine what portion of overall execution time for pipeline search is due to our runtime libraries vs. other factors such as I/O and underlying component implementation. The results are shown in Figure 4. Overall, the vast majority of execution time is due to execution of the underlying primitives (p25=90.2%, p50=96.2%, p75=98.3%). A smaller portion is due to the AutoBazaar runtime (p50=3.1%) and a negligible (p50<0.1%) portion of execution time is due to our other framework libraries and I/O. Thus, performance of pipeline execution/search is largely limited by the performance of the underlying physical implementation from the external library.

5.4 AutoML performance

One important attribute of AutoBazaar is the ability to improve pipelines for different tasks through tuning and selection. We measure the improvement in the best pipeline per task, finding that the average task improves its best score by 1.06 standard deviations over the course of tuning, and that 31.7% of tasks improve by more than 1 standard deviation (Figure 5). This demonstrates the AutoBazaar pipeline search effectiveness that a user may expect to obtain. However, as we describe in Section 3, there are many possible AutoML primitives that can be implemented using our tuner/selector APIs; a comprehensive comparison is beyond the scope of our work.

Data Modality	Problem Type	Tasks	Pipeline Template
graph	community detection	2	CommunityBestPartition
	graph matching	9	link_prediction_feat_extr graph_feat_extr CategoricalEncoder SimpleImputer StandardScaler XGBClassifier
	link prediction	1	link_prediction_feature_extraction CategoricalEncoder SimpleImputer StandardScaler XGBClassifier
	vertex nomination	1	graph_feature_extraction categorical_encoder SimpleImputer StandardScaler XGBClassifier
image	classification	5	ClassEncoder preprocess_input MobileNet XGBClassifier ClassDecoder
	regression	1	preprocess_input MobileNet XGBRegressor
multi table	classification	6	ClassEncoder dfs SimpleImputer StandardScaler XGBClassifier ClassDecoder
	regression	7	dfs SimpleImputer StandardScaler XGBRegressor
single table	classification	234	ClassEncoder dfs SimpleImputer StandardScaler XGBClassifier ClassDecoder
	collaborative filtering	4	dfs LightFM
	regression	87	dfs SimpleImputer StandardScaler XGBRegressor
	timeseries forecasting	35	dfs SimpleImputer StandardScaler XGBRegressor
text	classification	18	UniqueCounter TextCleaner VocabularyCounter Tokenizer pad_sequences LSTMTextClassifier
	regression	9	StringVectorizer SimpleImputer XGBRegressor
timeseries	classification	37	ClassEncoder dfs StandardImputer StandardScaler XGBClassifier ClassDecoder

Table 4: ML task types (data modality and problem type pairs) and associated ML tasks counts in the *ML Bazaar* Task Suite, along with default templates from AutoBazaar (i.e., where we have curated appropriate pipeline templates to solve a task).

Figure 5: Distribution of task performance improvement from *ML Bazaar* AutoML. Improvement for each task is measured as the score of the best pipeline less the score of the initial default pipeline, in standard deviations of all pipelines evaluated for that task.

5.5 Expressiveness of *ML Bazaar*

To further examine the expressiveness of *ML Bazaar* to solve a wide variety of tasks, we randomly selected 23 Kaggle competitions from 2018, comprising tasks ranging from image and time series classification to object detection and multi-table regression. For each task, we attempted to develop a solution using existing primitives and atalogs.

Overall, we were able to immediately solve 11 tasks. We did not currently support 4 task types: image matching (2 tasks), object detection within images (4 tasks), multi-label classification (1 task), and video classification (1 task). We could readily support these within our framework by developing new primitives and pipelines. In the remaining tasks,

multiple data modalities were provided to participants (i.e. some combination of image, text, and tabular data). To support these tasks, we would need to develop a new “glue” primitive for concatenating separately-featurized data from each resources to create a single feature matrix. Though our evaluation suite contains many examples of tasks with multiple data resources of different modalities, we had written pipelines customized to operate on certain common subsets (i.e. tabular + graph). We can never expect to have *already* implemented pipelines for the innumerable diversity of ML task types, but we can still write new primitives and pipelines using our framework to solve these problems.

5.6 Case study: evaluating ML primitives

When new primitives are contributed by the ML community, they can be incorporated into pipeline templates and pipeline hypertemplates, either to replace similar pipeline steps or to form the basis of new topologies. By running the end-to-end system on our evaluation suite, we can assess the impact of the primitive on general-purpose ML workloads (rather than overfit baselines).

In this first case study, we compare two similar primitives: annotations for the XGBoost (XGB) and random forest (RF) classifiers. We ran two experiments, one in which RF is used in pipeline templates and one in which XGB is substituted instead.

We consider 1.86×10^6 relevant pipelines to determine the best scores produced for 367 tasks. We find that the XGB pipelines substantially outperformed the RF pipelines, winning 64.9% of the comparisons. This confirms the experience of practitioners, who widely report that XGBoost is one of the most powerful ML methods for classification and regression.

5.7 Case study: evaluating AutoML primitives

The design of the *ML Bazaar* AutoML system and our extensive evaluation corpus allows us to easily swap in new AutoML primitives (Section 3.2) to see to what extent changes in components like tuners and selectors can improve performance in general settings.

In this case study, we revisit [46], which was influential for bringing about widespread use of Bayesian optimization for tuning ML models in practice. Their contributions include: (C1) proposing the usage of the Matérn 5/2 kernel for tuner meta-model (see Section 3.2.1), (C2) describing an integrated acquisition function that integrates over uncertainty in the GP hyperparameters, (C3) incorporating a cost model into an EI per second acquisition function, and (C4) explicitly modeling pending parallel trials. How important was each of these contributions to a resulting tuner?

Using *ML Bazaar*, we show how a more thorough *ablation study* [36], not present in [46], would be conducted to address these questions, by assessing the performance of our general-purpose AutoML system using different combinations of these four contributions. Here we focus on contribution C1. We run experiments using a baseline tuner with a squared exponential kernel (GP-SE-EI) and compare it with a tuner using the Matérn 5/2 kernel (GP-Matérn52-EI). In both cases, kernel hyperparameters are set by optimizing the marginal likelihood. Thus we can isolate the contributions of the proposed kernel in the context of general-purpose ML workloads.

In total, 4.31×10^5 pipelines were evaluated to find the best pipelines for a subset of 414 tasks. We find that there is no improvement from using the Matérn 5/2 kernel over the SE kernel — in fact, the GP-SE-EI tuner outperforms, winning 60.1% of the comparisons. One possible explanation for this negative result is that the Matérn kernel is sensitive to hyperparameters which are set more effectively by optimization of the integrated acquisition function. This is supported by the over-performance of the tuner using the integrated acquisition function in the original work; however, the integrated acquisition function is not tested with the baseline SE kernel, and more study is needed.

6 RELATED WORK

Researchers have developed numerous algorithmic and software innovations to make it possible to create ML and AutoML systems in the first place.

ML libraries. High-quality ML libraries have originated over a period of decades. For general ML applications, scikit-learn implements many different algorithms using a common API centered on the influential fit/predict paradigm [11]. For specialized analysis, libraries have been developed in separate academic communities, often with different and incompatible APIs [1, 7, 10, 24, 28, 32]. In *ML Bazaar*, we connect and link components of these libraries, only creating missing functionality ourselves.

ML systems. Prior work has provided several approaches for making it easier to develop ML systems. For example, caret [31] standardizes interfaces and provides utilities for the R ecosystem, but without enabling more complex pipelines. Recent systems have attempted to provide graphical interfaces, like [22] and Azure Machine Learning Studio. Development of ML systems is closely tied to the execution environments needed to train, deploy, and update the resulting models. In SystemML [9] and Weld [45], implementations of specific ML algorithms are optimized for specific runtimes. Velox [12] is an analytics stack component that efficiently serves predictions and manages model updates.

AutoML libraries. AutoML research has often been limited to solving sub-problems of an end-to-end ML workflow, such as data cleaning [14], feature engineering [28, 30], hyperparameter tuning [18, 21, 25, 27, 33, 40, 46, 48], or algorithm selection [25, 50]. Thus AutoML solutions are often not widely applicable or deployed in practice without human support. In contrast, *ML Bazaar* integrates many of these existing approaches and designs one coherent and configurable structure for joint tuning and selection of end-to-end pipelines.

AutoML systems. These AutoML libraries, if deployed, are typically one component within a larger system that aims to manage several practical aspects such as parallel and distributed training, tuning, and model storage, and even serving, deployment, and graphical interfaces for model building. These include ATM [47], Vizier [20], and Rafiki [52], as well as commercial platforms like Google AutoML, DataRobot, and Azure Machine Learning Studio. While these systems provide many benefits, they have several limitations. First, they each focus on a specific subset of ML use cases, such as computer vision, NLP, forecasting, or hyperparameter tuning. Second, these systems are designed as proprietary applications and do not support community-driven integration of new innovations. *ML Bazaar* provides a new approach to

developing such systems in the first place: it supports a wide variety of ML task types, and builds on top of a community-driven ecosystem of ML innovations. Indeed, it could serve as the backend for such ML services or platforms.

The DARPA D3M program [35], of which we are participants, aims to spur development of automated systems for model discovery for use by non-experts. Several differing approaches are being developed within this context. For example, Alpine Meadow [44] focuses on efficient search for producing interpretable ML pipelines with low latencies for interactive usage. It combines existing techniques from query optimization, Bayesian optimization, and multi-armed bandits to efficiently search for pipelines. AlphaD3M [16] formulates a pipeline synthesis problem and uses reinforcement learning to construct pipelines. In contrast, *ML Bazaar* is a framework to develop ML or AutoML systems in the first place. While we present our open-source AutoBazaar system, it is not the primary focus of our work and represents a single point in the design space of AutoML systems using our framework libraries. Indeed, one could use specific AutoML approaches like the ones described by Alpine Meadow or AlphaD3M for pipeline search within our own framework.

7 CONCLUSION

Throughout this paper, we have built up abstractions, interfaces, and software components for data scientists, data engineers, and other practitioners to effectively develop machine learning systems. Developers can use *ML Bazaar* to compose one-off pipelines, tunable pipeline templates, or full-fledged AutoML systems. Researchers can contribute individual ML or AutoML primitives and make them easily accessible to a broad base for inclusion in end-to-end solutions.

We have applied this approach to several real-world ML problems and entered our AutoML system in an automated modeling program. As we collect more and more scored pipelines, we expect opportunities will emerge for meta-learning and debugging on ML tasks and pipelines, as well as the ability to track progress and transfer knowledge within data science organizations. We will focus on several complementary extensions in future work. These include continuing to improve our AutoML system and making it more robust for everyday use by a diverse user base, and studying how to best support users of different backgrounds in using and interacting with ML and AutoML systems.

ACKNOWLEDGMENTS

The authors would like to thank Plamen Valentinov Kolev for contributions running experiments and testing datasets. The

authors also acknowledge the contributions of the following people: Laura Gustafson, William Xue, Akshay Ravikumar, Ihssan Tinawi, Alexander Geiger, Sarah Alnegheimish, Saman Amarasinghe, Stefanie Jegelka, Zi Wang, Benjamin Schreck, Seth Rothschild, Manuel Alvarez Campo, Sebastian Mir Peral, Peter Fontana, and Brian Sandberg. The authors are part of the DARPA Data-Driven Discovery of Models (D3M) program, and would like to thank the D3M community for the discussions around the design. This material is based on research sponsored by DARPA and Air Force Research Laboratory (AFRL) under agreement number FA8750-17-2-0126. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. The view and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies of endorsements, either expressed or implied, of DARPA and Air Force Research Laboratory (AFRL) or the U.S. Government. Authors MJS, CS, and KV acknowledge support and user feedback from Iberdrola S.A. and SES S.A.

REFERENCES

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> Software available from tensorflow.org.
- [2] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47, 2-3 (2002), 235–256.
- [3] Denis Baylor, Eric Breck, Heng-Tze Cheng, Noah Fiedel, Chuan Yu Foo, Zakaria Haque, Salem Haykal, Mustafa Ispir, Vihan Jain, Levent Koc, Chiu Yuen Koo, Lukasz Lew, Clemens Mewald, Akshay Naresh Modi, Neoklis Polyzotis, Sukriti Ramesh, Sudip Roy, Steven Euijong Whang, Martin Wicke, Jarek Wilkiewicz, Xin Zhang, and Martin Zinkevich. 2017. TFX: A TensorFlow-Based Production-Scale Machine Learning Platform. In *KDD*.
- [4] James Bergstra and Yoshua Bengio. 2012. Random Search for Hyper-Parameter Optimization. *JMLR* 13 (2012), 281–305.
- [5] James S Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. In *NIPS*. 2546–2554.
- [6] Anant Bhardwaj, Souvik Bhattacharjee, Amit Chavan, Amol Deshpande, Aaron J Elmore, Samuel Madden, and Aditya Parameswaran. 2015. DataHub: Collaborative Data Science & Dataset Version Management at Scale. In *CIDR*.
- [7] Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. O'Reilly Media, Inc.

- [8] Bernd Bischl, Giuseppe Casalicchio, Matthias Feurer, Frank Hutter, Michel Lang, Rafael G Mantovani, Jan N van Rijn, and Joaquin Vanschoren. 2017. OpenML benchmarking suites and the OpenML100. (2017). arXiv:arXiv:1708.03731
- [9] Matthias Boehm, Michael W Dusenberry, Deron Eriksson, Alexandre V Evfimievski, Faraz Makari Manshadi, Niketan Pansare, Berthold Reinwald, Frederick R Reiss, Prithviraj Sen, Arvind C Surve, and Shirish Tatikonda. 2016. Systemml: Declarative machine learning on spark. *PVLDB* 9, 13 (2016), 1425–1436.
- [10] G. Bradski. 2000. The OpenCV Library. *Dr. Dobb's Journal of Software Tools* (2000).
- [11] Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, Robert Layton, Jake VanderPlas, Arnaud Joly, Brian Holt, and Gaël Varoquaux. 2013. API design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*. 108–122.
- [12] Daniel Crankshaw, Peter Bailis, Joseph E. Gonzalez, Haoyuan Li, Zhao Zhang, Michael J. Franklin, Ali Ghodsi, and Michael I. Jordan. 2015. The Missing Piece in Complex Analytics: Low Latency, Scalable Model Management and Serving with Velox. In *CIDR*.
- [13] Andrew Crotty, Alex Galakatos, Emanuel Zgraggen, Carsten Binnig, and Tim Kraska. 2015. Vizdom: Interactive Analytics Through Pen and Touch. *Proc. VLDB Endow.* 8, 12 (Aug 2015), 2024–2027.
- [14] Dong Deng, Raul Castro, Fernandez Ziawasch, Abedjan Sibo, Ahmed Elmagarmid, Ihab F Ilyas, Samuel Madden, Mourad Ouzzani, and Nan Tang. 2017. The Data Civilizer System. *CIDR* (2017).
- [15] Ian Dewancker, Michael McCourt, Scott Clark, Patrick Hayes, Alexandra Johnson, and George Ke. 2016. A Strategy for Ranking Optimization Methods using Multiple Criteria. *ICML AutoML workshop* 2011 (2016), 1–12.
- [16] 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. *ICML 2018 AutoML Workshop*, 1–8.
- [17] Stefan Falkner, Aaron Klein, and Frank Hutter. 2018. BOHB: Robust and Efficient Hyperparameter Optimization at Scale. In *ICML*. 1436–1445.
- [18] Matthias Feurer, Aaron Klein, Katharina Eggensperger, Jost Springenberg, et al. 2015. Efficient and robust automated machine learning. In *NIPS*. 2962–2970.
- [19] Pieter Gijsbers, Erin Ledell, Janek Thomas, Sébastien Poirier, Bernd Bischl, and Joaquin Vanschoren. 2019. An Open Source AutoML Benchmark. *6th ICML Workshop on Automated Machine Learning* (2019), 1–8.
- [20] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and D Sculley. 2017. Google Vizier: A Service for Black-Box Optimization. In *KDD*.
- [21] Taciana A.F. Gomes, Ricardo B.C. Prudêncio, Carlos Soares, André L.D. Rossi, and André Carvalho. 2012. Combining meta-learning and search techniques to select parameters for support vector machines. *Neuro-computing* 75, 1 (2012), 3–13.
- [22] Ming Gong, Linjun Shou, Wutao Lin, Zhijie Sang, Quanjia Yan, Ze Yang, and Daxin Jiang. 2019. NeuronBlocks – Building Your NLP DNN Models Like Playing Lego. (2019). arXiv:arXiv:1904.09535
- [23] Isabelle Guyon, Kristin Bennett, Gavin Cawley, Hugo Jair Escalante, Sergio Escalera, Tin Kam Ho, Núria Macià, Bisakha Ray, Mehreen Saeed, Alexander Statnikov, and Evelyn Viegas. 2015. Design of the 2015 ChaLearn AutoML challenge. In *IJCNN*.
- [24] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. 2008. Exploring network structure, dynamics, and function using NetworkX. In *SciPy*, Gaël Varoquaux, Travis Vaught, and Jarrod Millman (Eds.). 11–15.
- [25] Martin Hirzel, Kiran Kate, Avraham Shinnar, Subhrajit Roy, and Parikshit Ram. 2019. Type-Driven Automated Learning with Lale. (May 2019). arXiv:arXiv:1906.03957
- [26] Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, and Tom Soderstrom. 2018. Detecting Spacecraft Anomalies Using LSTMs and Nonparametric Dynamic Thresholding. In *KDD*.
- [27] Kevin Jamieson and Amee Talwalkar. 2016. Non-stochastic Best Arm Identification and Hyperparameter Optimization. In *AISTATS*. 240–248.
- [28] James Max Kanter. 2015. Deep Feature Synthesis: Towards Automating Data Science Endeavors. In *DSAA*. 1–10.
- [29] James Max Kanter, Owen Gillespie, and Kalyan Veeramachaneni. 2016. Label, segment, featurize: A cross domain framework for prediction engineering. In *DSAA*. 430–439.
- [30] Udayan Khurana, Deepak Turaga, Horst Samulowitz, and Srinivasan Parthasarathy. 2016. Cognito: Automated feature engineering for supervised learning. In *ICDMW*. 1304–1307.
- [31] Max Kuhn. 2008. Building Predictive Models in R Using the caret Package. *Journal of Statistical Software* 28, 5 (2008), 159–160.
- [32] Maciej Kula. 2015. Metadata Embeddings for User and Item Cold-start Recommendations. In *Proceedings of the 2nd Workshop on New Trends on Content-Based Recommender Systems*, Vol. 1448. 14–21.
- [33] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Amee Talwalkar. 2017. Hyperband: a novel bandit-based approach to hyperparameter optimization. *JMLR* 18, 1 (2017), 6765–6816.
- [34] Liam Li, Kevin Jamieson, Afshin Rostamizadeh, Ekaterina Gonina, Moritz Hardt, Benjamin Recht, and Amee Talwalkar. 2018. Massively parallel hyperparameter tuning. (2018). arXiv:arXiv:1810.05934
- [35] Richard Lippmann, William Campbell, and Joseph Campbell. 2016. An Overview of the DARPA Data Driven Discovery of Models (D3M) Program. In *NIPS Workshop on Artificial Intelligence for Data Science*.
- [36] Zachary C Lipton and Jacob Steinhardt. 2018. Troubling trends in machine learning scholarship. (2018). arXiv:arXiv:1807.03341
- [37] Ilya Loshchilov and Frank Hutter. 2016. CMA-ES for Hyperparameter Optimization of Deep Neural Networks. (2016). arXiv:arXiv:1604.07269
- [38] Hui Miao, Ang Li, Larry S Davis, and Amol Deshpande. 2017. Modelhub: Deep learning lifecycle management. In *ICDE*. 1393–1394.
- [39] ChangYong Oh, Efstratios Gavves, and Max Welling. 2018. BOCK: Bayesian Optimization with Cylindrical Kernels. (2018). arXiv:arXiv:1806.01619
- [40] Randal S. Olson, Nathan Bartley, Ryan J. Urbanowicz, and Jason H. Moore. 2016. Evaluation of a Tree-based Pipeline Optimization Tool for Automating Data Science. In *GECCO*.
- [41] Eric Raymond. 1999. The cathedral and the bazaar. *Knowledge, Technology & Policy* 12, 3 (1999), 23–49.
- [42] Matthew Rocklin. 2015. Dask: Parallel Computation with Blocked algorithms and Task Scheduling. In *SciPy*, Kathryn Huff and James Bergstra (Eds.). 130–136.
- [43] D Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-Francois Crespo, and Dan Dennison. 2015. Hidden technical debt in machine learning systems. In *NIPS*. 2503–2511.
- [44] Zeyuan Shang, Emanuel Zgraggen, Benedetto Buratti, Ferdinand Kossmann, Philipp Eichmann, Yeunoh Chung, Carsten Binnig, Eli Upfal, and Tim Kraska. 2019. Democratizing data science through interactive curation of ml pipelines. In *SIGMOD*. 1171–1188.
- [45] Palkar Shoumik, James Thomas, Deepak Naryanan, Pratishtha Thaker, Rahul Palamuttam, Parimajan Negi, Anil Shanbhag, Malte Schwarzkopf, Holger Pirk, Saman Amarasinghe, Samuel Madden, and Matei Zaharia. 2018. Evaluating end-to-end optimization for data analytics applications in weld. *PVLDB* 11, 9 (2018), 1002–1015.

- [46] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. In *NIPS*. 2951–2959.
- [47] Thomas Swearingen, Will Drevo, Bennett Cyphers, et al. 2017. ATM: A distributed, collaborative, scalable system for automated machine learning. In *BigData*.
- [48] Chris Thornton, Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. 2013. Auto-WEKA: combined selection and hyperparameter optimization of classification algorithms. In *KDD*.
- [49] Tom van der Weide, Dimitris Papadopoulos, Oleg Smirnov, Michal Zielinski, and Tim van Kasteren. 2017. Versioning for End-to-End Machine Learning Pipelines. In *DEEM*. 1–9.
- [50] Jan N. van Rijn, Salisu Mamman Abdulrahman, Pavel Brazdil, and Joaquin Vanschoren. 2015. Fast Algorithm Selection Using Learning Curves. In *IDA*. 298–309.
- [51] Hao Wang, Bas van Stein, Michael Emmerich, and Thomas Back. 2017. A new acquisition function for Bayesian optimization based on the moment-generating function. In *SMC*. 507–512.
- [52] Wei Wang, Jinyang Gao, Meihui Zhang, Sheng Wang, Gang Chen, Teck Khim Ng, Beng Chin Ooi, Jie Shao, and Moaz Reyad. 2018. Rafiki: machine learning as an analytics service system. *PVLDB* 12, 2 (2018), 128–140.
- [53] Matei Zaharia, Reynold S. Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J. Franklin, Ali Ghodsi, Joseph Gonzalez, Scott Shenker, and Ion Stoica. 2016. Apache Spark: A Unified Engine for Big Data Processing. *Commun. ACM* 59, 11 (Oct 2016), 56–65.