

Machine Learning at Microsoft with ML.NET

Zeeshan Ahmed¹, Saeed Amizadeh¹, Mikhail Bilenko², Rogan Carr¹, Wei-Sheng Chin¹, Yael Dekel¹, Xavier Dupre¹, Vadim Eksarevskiy¹, Senja Filipi¹, Tom Finley¹, Abhishek Goswami¹, Monte Hoover¹, Scott Inglis¹, Matteo Interlandi¹, Najeeb Kazmi¹, Gleb Krivosheev¹, Pete Luferenko¹, Ivan Matantsev¹, Sergiy Matusevych¹, Shahab Moradi¹, Gani Nazirov¹, Justin Ormont¹, Gal Oshri¹, Artidoro Pagnoni¹, Jignesh Parmar¹, Prabhat Roy¹, Mohammad Zeeshan Siddiqui¹, Markus Weimer¹, Shauheen Zahirazami¹, Yiwen Zhu^{1*}

¹Microsoft, ²Yandex

ABSTRACT

Machine Learning is transitioning from an art and science into a technology available to every developer. In the near future, every application on every platform will incorporate trained models to encode data-based decisions that would be impossible for developers to author. This presents a significant engineering challenge, since currently data science and modeling are largely decoupled from standard software development processes. This separation makes incorporating machine learning capabilities inside applications unnecessarily costly and difficult, and furthermore discourage developers from embracing ML in first place.

In this paper we present ML.NET, a framework developed at Microsoft over the last decade in response to the challenge of making it easy to ship machine learning models in large software applications. We present its architecture, and illuminate the application demands that shaped it. Specifically, we introduce DataView, the core data abstraction of ML.NET which allows it to capture full predictive pipelines efficiently and consistently across training and inference lifecycles. We close the paper with a surprisingly favorable performance study of ML.NET compared to more recent entrants, and a discussion of some lessons learned.

ACM Reference Format:

Zeeshan Ahmed, Saeed Amizadeh, Mikhail Bilenko, Rogan Carr, Wei-Sheng Chin, Yael Dekel, Xavier Dupre, Vadim Eksarevskiy, Senja Filipi, Tom Finley, Abhishek Goswami, Monte Hoover, Scott Inglis, Matteo Interlandi, Najeeb Kazmi, Gleb Krivosheev, Pete Luferenko, Ivan Matantsev, Sergiy Matusevych, Shahab Moradi, Gani Nazirov, Justin Ormont, Gal Oshri, Artidoro Pagnoni, Jignesh Parmar, Prabhat Roy, Zeeshan Siddiqui, Markus Weimer, Shauheen Zahirazami, Yiwen Zhu. 2019. Machine Learning at Microsoft with ML.NET. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3292500.3330667>

1 INTRODUCTION

We are witnessing an explosion of new frameworks for building Machine Learning (ML) models [8, 9, 13, 21, 27–29, 33]. This profusion

*Pete Luferenko is now at Facebook. Gal Oshri is now at Google.

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.

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330667>

is motivated by the transition from machine learning as an art and science into a set of technologies readily available to every developer. An outcome of this transition is the abundance of applications that rely on trained models for functionalities that evade traditional programming due to their complex statistical nature. Speech recognition and image classification are only the most prominent such cases. This unfolding future, where most applications make use of at least one model, profoundly differs from the current practice in which data science and software engineering are performed in separate and different processes and sometimes even organizations. Furthermore, in current practice, models are routinely deployed and managed in completely distinct ways from other software artifacts. While typical software packages are seamlessly compiled and run on a myriad of heterogeneous devices, machine learning models are often relegated to be run as web services in relatively inefficient containers [6, 7, 18]. This pattern not only severely limits the kinds of applications one can build with machine learning capabilities [17], but also discourages developers from embracing ML as a core component of applications.

At Microsoft, we have encountered this phenomenon across a wide spectrum of applications and devices, ranging from services and server software to mobile and desktop applications running on PCs, Servers, Data Centers, Phones, Game Consoles and IOT devices. A machine learning toolkit for such diverse use cases, frequently deeply embedded in applications, must satisfy additional constraints compared to the recent cohort of toolkits. For example, it has to limit library dependencies that are uncommon for applications; it must cope with datasets too large to fit in RAM; it has to scale to many or few cores and nodes; it has to be portable across many target platforms; it has to be model class agnostic, as different ML problems lend themselves to different model classes; and, most importantly, it has to capture the full prediction pipeline that takes a test example from a given domain (e.g., an email with headers and body) and produces a prediction that can often be structured and domain-specific (e.g., a collection of likely short responses). The requirement to encapsulate predictive pipelines is of paramount importance because it allows for effectively decoupling application logic from model development. Carrying the complete train-time pipeline into production provides a dependable way for building efficient, reproducible, production-ready models [38].

The need for ML pipelines has been recognized previously. Python libraries such as Scikit-learn [27] provide the ability to author complex machine learning cascades. Python has become the most popular language for data science thanks to its simplicity, interactive nature (e.g., notebooks [15, 37]) and breadth of libraries (e.g.,

numpy [34], pandas [20], matplotlib [19]). However, Python-based libraries inherit many syntactic idiosyncrasies and language constraints (e.g., interpreted execution, dynamic typing, global interpreter locks that restrict parallelization), making them suboptimal for high-performance applications targeting a myriad of devices.

In this paper we introduce ML.NET: a machine learning framework allowing developers to author and deploy in their applications complex ML pipelines composed of data featurizers and state of the art machine learning models. Pipelines implemented and trained using ML.NET can be seamlessly surfaced for prediction without any modification: training and prediction, in fact, share the same code paths, and adding a model into an application is as easy as importing ML.NET runtime and binding the inputs/output data sources. ML.NET's ability to capture full, end-to-end pipelines has been demonstrated by the fact that 1,000s of Microsoft's data scientists and developers have been using ML.NET over the past decade, infusing 100s of products and services with machine learning models used by hundreds of millions of users worldwide.

ML.NET supports large scale machine learning thanks to an internal design borrowing ideas from relational database management systems and embodied in its main abstraction: *DataView*. *DataView* provides *compositional processing of schematized data* while being able to *gracefully and efficiently handle high dimensional data in datasets larger than main memory*. Like views in relational databases, a *DataView* is the result of computations over one or more base tables or views, is immutable and lazily evaluated (unless forced to be materialized, e.g., when multiple passes over the data are requested). Under the hood, *DataView* provides streaming access to data so that working sets can exceed main memory.

We run an experimental evaluation comparing ML.NET with Scikit-learn and H2O [13]. To examine runtime, accuracy and scalability performance, we set up several experiments over three different datasets and utilizing different data sample rates. Our experiments show that ML.NET outperforms both Sklearn and H2O in speed and accuracy, in most cases by a large margin.

Summarizing our contributions are:

- The introduction of ML.NET: a machine learning framework for authoring production-grade machine learning pipelines, which can then be easily integrated into applications running on heterogeneous devices;
- Discussion on the motivations pushing Microsoft to develop ML.NET, and on the lessons we learned from helping thousands of developers in building and deploying ML pipelines at enterprise and cloud scale;
- Introduction of the *DataView* abstraction, and how it translates into efficient executions through streaming data access, immutability and lazy evaluation;
- A set of experiments comparing ML.NET against the well known Scikit-learn and the more recent H2O, and proving ML.NET's state-of-the-art performance.

The remainder of the paper is organized as follows: Section 2 introduces the main motivations behind the development of ML.NET. Sections 3 introduces ML.NET's design and the *DataView* abstraction, while Section 4 drills into the details of ML.NET implementation. Section 5 contains our experimental evaluation. Lessons learned are introduced in Section 6. The paper ends with related works and conclusions, respectively in Sections 7 and 8.

2 MOTIVATIONS AND OVERVIEW

The goal of the ML.NET project is improving the development life-cycle of ML pipelines. While pipelines proceed from the initial experimentation stages to engineering for scaling up and eventual deployment into applications, they traditionally require significant (and sometimes complete) rewriting at significant cost. Aiming to reduce such costs by simplifying and unifying the underlying ML frameworks, we observed three interesting patterns:

Pattern 1: Many data scientists within Microsoft were not following the interactive pattern made popular by notebooks and Python ML libraries such as Sklearn. This was due to two key factors: the sheer size of production-grade datasets and accuracy of the final models. While an "accurate enough" model can be easily developed by iteratively refining the ML pipeline on data samples, finding the best model often requires experimenting over the entire dataset. In this context, interactive exploration is of limited applicability inasmuch as most of the datasets were large enough to not fit into main memory on a single machine. Working on large datasets has led to many of Microsoft's data scientists working in batch mode: running multiple scripts concurrently sweeping over models and parameters, and not performing much exploratory data analysis.

Pattern 2: In order to obtain the best possible model, data scientists focused their experiments on several state of the art algorithms from different model families. These methods were typically developed by different groups, often coming directly from researchers within our company. As a result, each model was originally implemented without a standard API, input format, or hyperparameter notation. Data scientists were therefore spending considerable effort on implementing glue code and ad-hoc wrappers around different algorithms and data formats to employ them in their pipelines. In general, getting all of the steps right required multiple iterations and significant time costs. Because ad-hoc pipelines are constructed to run in batch mode, there is no static, compile-time checking to detect any inconsistencies in dataflow across the pipelines. As a result, pipeline debugging was performed via log parsing and errors thrown at runtime, sometimes after multiple hours of training.

Pattern 3: Building production-grade applications making use of machine learning pipelines is a laborious task. As a first step, one needs to formalize the prediction task, and choose the components: feature construction and transformations, the training algorithm, hyper parameters and their tuning. Once the pipeline is developed and successfully trained, it must be integrated into the application and shipped in production. This process is usually performed by a different engineer (or team) than the one building the model, and a significant rewriting is often required because of various runtime constraints (e.g., a different hardware or software platform, constraints on pipeline size or prediction latency/throughput). Such rewriting is often done for particular applications, resulting in custom solutions: a process not sustainable at Microsoft scale.

Solving the issues revealed by the above patterns requires rethinking the ML framework for pipeline composition. Key requirements for it can be summarized as follows:

- (1) *Unification:* ML.NET must act as a unifying framework that can host a variety of models and components (with related idiosyncrasies). Once ML pipelines are trained, the same

pipeline must be deployable into any production environment (from data centers to IoT devices) with close to zero engineering cost. In the last decade 100s of products and services have employed ML.NET, validating its success as a unifying platform.

- (2) *Extensibility*: Data scientists are interested in experimenting with different models and features with the goal of obtaining the best accuracy. Therefore, it should be possible to add new components and algorithms with minimal reasonable effort via a general API that supports a variety of data types and formats. Since its inception, ML.NET has been extended with many components. In fact, a large fraction of the now more than 120 built-in operators started life as extensions shared between data scientists.
- (3) *Scalability and Performance*: ML.NET must be scalable and allow maximum hardware utilization—i.e., be fast and provide high throughput. Because production-grade datasets are often very large and do not fit in RAM, scalability implies the ability to run pipelines in out-of-memory mode, with data paged in and processed incrementally. As we show in the experiment section, ML.NET achieves good scalability and performance (up to several orders-of-magnitude) when compared to other publicly available toolkits.

ML.NET: Overview. ML.NET is a .NET machine learning library that allows developers to build complex machine learning pipelines, evaluate them, and then utilize them directly for prediction. Pipelines are often composed of multiple transformation steps that featurize and transform the raw input data, followed by one or more ML models that can be stacked or form ensembles. Note that “pipeline” is a bit of a misnomer, as they, in fact, are Direct Acyclic Graphs (DAGs) of operators. We next illustrate how these tasks can be accomplished in ML.NET on a short example¹; we will also exploit this example to introduce the main concepts in ML.NET.

```
1 var ml = new MLContext();
2 var data = ml.Data.LoadFromTextFile<SentimentData>(trainingDatasetPath);
3 var pipeline = ml.Transforms.Text.FeaturizeText("Features", "Text")
4   .Append(ml.BinaryClassification.Trainers.FastTree());
```

Figure 1: A pipeline for text analysis whereby input sentences are classified according to the expressed sentiment.

Figure 1 introduces a Sentiment Analysis pipeline (SA). The first item required for building a pipeline is the *MLContext* (line 1): the entry-point for accessing ML.NET features. In line 2, a *loader* is used to indicate how to read the input training data. In the example pipeline, the input schema (*SentimentData*) is specified explicitly, but in other situations (e.g., CSV files with headers) schemas can be automatically inferred by the loader. Loaders generate a *DataView* object, which is the core data abstraction of ML.NET. *DataView* provides a fully schematized non-materialized view of the data, and gets subsequently transformed by pipeline components.

The second step is feature extraction from the input *Text* column (line 3). To achieve this, we use the *FeaturizeText transform*. Transforms are the main ML.NET operators for manipulating data. Transforms accept a *DataView* as input and produce another

¹This and other examples can be accessed at <https://github.com/dotnet/machinelearning-samples>.

DataView. *FeaturizeText* is actually a complex transform built off a composition of nine base transforms that perform common tasks for feature extraction from natural text. Specifically, the input text is first normalized and tokenized. For each token, both char- and word-based ngrams are extracted and translated into vectors of numerical values. These vectors are subsequently normalized and concatenated to form the final *Features* column. Some of the above transforms (e.g., normalizer) are *trainable*: i.e., before producing an output *DataView* they are required to scan the whole dataset to determine internal parameters (e.g., scalars).

Subsequently, in line 4 we apply a *learner* (i.e., a trainable model) to the pipeline—in this case, a binary classifier called *FastTree*: an implementation of the MART gradient boosting algorithm [11]. Once the pipeline is assembled, we can train it by calling the *Fit* method on the pipeline object with the expected output prediction type (Figure 2). ML.NET evaluation is *lazy*: no computation is actually run until the *Fit* method (or other methods triggering pipeline execution) is called. This allows ML.NET to (1) properly validate that the pipeline is well-formed before computation; and (2) deliver state of the art performance by devising efficient execution plans.

```
var model = pipeline.Fit(data);
```

Figure 2: Training of the sentiment analysis pipeline. Up to this point no execution is actually triggered.

Once a pipeline is trained, a model object containing all training information is created. The model can be saved to a file (in this case, the information of all trained operators as well as the pipeline structure are serialized into a compressed file), or evaluated against a test dataset (Figure 3) or directly used for prediction serving (Figure 4). To evaluate model performance, ML.NET provides specific components called *evaluators*. Evaluators accept as input a *DataView* upon which a model has been previously applied, and produce a set of metrics. In the specific case of the evaluator used in Figure 3, relevant metrics are those used for binary classifiers, such as accuracy, Area Under the Curve (AUC), log-loss, etc.

```
var output = model.Transform(testData);
var metrics = mlContext.BinaryClassification.Evaluate(output);
```

Figure 3: Evaluating model accuracy using a test dataset.

Finally, serving the model for prediction is achieved by first creating a *PredictionEngine*, and then calling the *Predict* method with a list of *SentimentData* objects. Predictions can be served natively in any OS (e.g., Linux, Windows, Android, macOS) or device supported by the .NET Core framework.

```
var engine = ml.Model
    .CreatePredictionEngine<SentimentData, SentimentPrediction>(model);
var predictions = engine.Predict(PredictionData);
```

Figure 4: Serving predictions using the trained model.

3 SYSTEM DESIGN AND ABSTRACTIONS

In order to address the requirements listed in Section 2, ML.NET borrows ideas from the database community. ML.NET’s main abstraction is called *DataView* (Section 3.1). Similarly to (intensional)

database relations, the DataView abstraction provides compositional processing of schematized data, but specializes it for machine learning pipelines. The DataView abstraction is generic and supports both primitive operators as well as the composition of multiple operators to achieve higher-level semantics such as the `FeaturizeText` transform of Figure 1 (Section 3.2). Under the hood, operators implementing the DataView interface are able to gracefully and efficiently handle high-dimensional and large datasets thanks to *cursoring* (Section 3.3) which resembles the well-known iterator model of databases [12].

3.1 The DataView Abstraction

In relational databases, the term *view* typically indicates the result of a query on one or more tables (base relations) or views, and is generally immutable. Views (and tables) are defined over a *schema* which expresses a sequence of *columns names* with related *types*. The semantics of the schema is such that each data row outputs of a view must conform to its schema. Views have interesting properties which differentiate them from tables and make them appropriate abstractions for machine learning: (1) views are *composable*—new views are formed by applying transformations (queries) over other views; (2) views are *virtual*, i.e., they can be lazily computed on demand from other views or tables without having to materialize any partial results; and (3) since a view does not contain values, but merely computes values from its source views, it is *immutable* and *deterministic*: the same exact computation applied over the same input data always produces the same result. Immutability and deterministic computation (note that several other data processing systems such as Apache Spark [36] employ the same assumptions) enable transparent data caching (for speeding up iterative computations such as ML algorithms) and safe parallel execution. DataView inherits the aforementioned database view properties, namely: schematization, composability, lazy evaluation, immutability, and deterministic execution.

Schema with Hidden Columns. Each DataView carries schema information specifying the name and type of each view’s column. DataView schemas are ordered and, by design, multiple columns can share the same name, in which case, one of the columns *hides* the others: referencing a column by name always maps to the latest column with that name. Hidden columns exist because of immutability and can be used for debugging purposes: having all partial computations stored as hidden columns allows the inspection of the provenance of each data transformation. Indeed, hidden columns are never fully materialized in memory (unless explicitly required) therefore their resource cost is minimal.

High Dimensional Data Support with Vector Types. While the DataView schema system supports an arbitrary number of columns, like most schematized data systems, it is designed for a modest number of columns, typically, limited to a few hundred. Machine learning and advanced analytics applications often involve high-dimensional data. For example, common techniques for learning from text uses bag-of-words (e.g., `FeaturizeText`), one-hot encoding or hashing variations to represent non-numerical data. These techniques typically generate an enormous number of features. Representing each feature as an individual column is far from ideal, both from the perspective of how the user interacts with the information and how the information is managed in the

schematized system. The DataView solution is to represent each set of features as a single *vector* column. A vector type specifies an item type and optional dimensionality information. The item type must be a primitive, non-vector, type. The optional dimensionality information specifies the number of items in the corresponding vector values. When the size is unspecified, the vector type is variable-length. For example, the `TextTokenizer` transform (contained in `FeaturizeText`) maps a text value to a sequence of individual terms. This transformation naturally produces variable-length vectors of text. Conversely, fixed-size vector columns are used, for example, to represent a range of column from an input dataset.

3.2 Composing Computations using DataView

ML.NET includes several standard operators and the ability to compose them using the DataView abstraction to produce efficient machine learning pipelines. *Transform* is the main operator class: transforms are applied to a DataView to produce a derived DataView and are used to prepare data for training, testing, or prediction serving. *Learners* are machine learning algorithms that are trained on data (eventually coming from some transform) and produce predictive models. *Evaluators* take scored test datasets and produced metrics such as precision, recall, F1, AUC, etc. Finally, *Loaders* are used to represent data sources as a DataView, while *Savers* serialize DataViews to a form that can be read by a loader. We now details some of the above concepts.

Transforms. Transforms take a DataView as input and produce a DataView as output. Many transforms simply “add” one or more computed columns to their input schema. More precisely, their output schema includes all the columns of the input schema, plus some additional columns, whose values are computed starting from some of the input columns. It is common for an added column to have the same name as an input column, in which case, the added column hides the input column, as we have previously described. Multiple primitive transforms may be applied to achieve higher-level semantics: for example, the `FeaturizeText` transform of Figure 1 is the composition of 9 primitive transforms.

Trainable Transforms. While many transforms simply map input data values to output by applying some pre-defined computation logic (e.g., `Concat`), other transforms require “training”, i.e., their precise behavior is determined automatically from the input training data. For example, normalizers and dictionary-based mappers translating input values into numerical values (used in `FeaturizeText`) build their state from training data. Given a pipeline, a call to `Train` triggers the execution of all trainable transforms (as well as learners) in topological order. When a transform (learner) is trained, it produces a DataView representing the computation up to that point in the pipeline: the DataView can then be used by downstream operators. Once trained and later saved, the state of a trained transform is serialized such that, once loaded back the transform is not retrained.

Learners. Similarly to trainable transforms, learners are machine learning algorithms that take DataView as input and produce “models”: transforms that can be applied over input DataViews and produce predictions. ML.NET supports learners for *binary classification*, *regression*, *multi-class classification*, *ranking*, *clustering*, *anomaly detection*, *recommendation* and *sequence prediction* tasks.

3.3 Cursoring over Data

ML.NET uses *DataView* as a representation of a computation over data. Access to the actual data is provided through the concept of *row cursor*. While in databases queries are compiled into a chain of operators, each of them implementing an iterator-based interface, in ML.NET, ML pipelines are compiled into chains of *DataViews* where data is accessed through cursoring. A row cursor is a movable window over a sequence of data rows coming either from the input dataset or from the result of the computation represented by another *DataView*. The row cursor provides the column values for the current row, and, as iterators, can only be advanced forward (no backtracking is allowed).

Columnar Computation. In data processing systems, it is common for a down-stream operator to only require a small subset of the information produced by the upstream pipeline. For example, databases have columnar storage layouts to avoid access to unnecessary columns [31]. This is even more so in machine learning pipelines where featurizers and ML models often work on one column at a time. For instance, *FeaturizeText* needs to build a dictionary of all terms used in a text column, while it does not need to iterate over any other columns. ML.NET provides columnar-style computation model through the notion of *active columns* in row cursors. Active columns are set when a cursor is initialized: the cursor then enforces the contract that only the computation or data movement necessary to provide the values for the active columns are performed.

Pull-based Model, Streaming Data. ML.NET runtime performance are proportional to data movements and computations required to scan the data rows. As iterators in database, cursors are pull-based: after an initial setup phase (where for example active columns are specified) cursors do not access any data, unless explicitly asked to. This strategy allows ML.NET to perform at each time only the computation and data movements needed to materialize the requested rows (and column values within a row). For large data scenarios, this is of paramount importance because it allows efficient streaming of data directly from disk, without having to rely on the assumption that working sets fit into main memory. Indeed, when the data is known to fit in memory, caching provides better performance for iterative computations.

4 SYSTEM IMPLEMENTATION

ML.NET is the solution Microsoft developed for the problem of empowering developers with a machine learning framework to author, test and deploy ML pipelines. As introduced in Section 2, ML.NET is implemented with the goal of providing a tool that is easy to use, scalable over large datasets while providing good performance, and able to unify under a single API data transformations, featurizers, and state of the art machine learning models. In its current implementation, ML.NET comprises 2773K lines of C# code, and about 74K lines of C++ code, the latter used mostly for high-performance linear algebra operations employing SIMD instructions. ML.NET supports more than 80 featurizers and 40 machine learning models.

4.1 Writing Pipelines in ML.NET

ML.NET comes with several APIs, all covering the different use cases we observed during the years at Microsoft. All APIs eventually are compiled into the typed learning pipeline API with generics

shown in the SA example of Section 2. Beyond the typed API with generics, ML.NET supports a (1) command line / scripting API enabling data scientists to easily experiment with several pipelines; (2) a Graphical User Interface for users less familiar with coding; and (3) an *Entry Point* (EP) API allowing to execute and code-generate APIs in different languages (e.g., Scala and Python). Due to space constraint, next we only detail the EP API and one of its applications, namely *NimbusML* [24]: a Python API mirroring Scikit-learn pipeline API.

Entry Points and Graph Runner. The recommended way of interacting with ML.NET through other, non-.NET, programming languages is by composing, and exchanging *entry point graphs*. An EP is a JSON representation of a ML.NET operator. EPs descriptions are grouped into a manifest file: a JSON object that documents and defines the structure of any available EP. The operator manifest is code-generated by scanning the ML.NET assemblies through reflection and searching for specific types and class annotations of operators. Using the EP API, pipelines are represented as graphs of EP operators and input/output relationships which are all serialized into a JSON file. EP graphs are parsed in ML.NET by the *graph runner* component which generates and directly executes the related pipeline. Non-.NET APIs can be automatically generated starting from the manifest file so that there is no need to write and maintain additional APIs.

NimbusML and Interoperability with Scikit-learn. *NimbusML* is ML.NET's Python API mirroring Scikit-learn interface (*NimbusML* operators are actually subclasses of Scikit-learn components) and taking advantage of the EP API functionalities. Furthermore, data scientists can start with a Scikit-learn pipeline and swap Scikit-learn transformations or algorithms with ML.NET's ones to achieve better scalability and accuracy. To achieve this level of interoperability, however, data residing in Scikit-learn needs to be accessed from ML.NET (and vice-versa), but the EP API does not provide such functionality. To obtain such behavior in an efficient and scalable way, when *NimbusML* is imported into a Scikit-learn project, a .NET Core runtime as well as an instance of ML.NET are spawn within the same Python process. When a user triggers the execution of a pipeline, the call is intercepted on the Python side and an EP graph is generated, and submitted to the graph runner component of the ML.NET instance. If the data resides in memory in Scikit-learn as a Pandas data frame or a numpy array, the C++ reference of the data is passed to the graph runner through C#/C++ interop and wrapped around a *DataView*, which is then used as input for the pipeline. We used *Boost.Python* [2] as helper library to access the references of data residing in Scikit-learn. If the data resides on disk, a *DataView* is instead directly used. Figure 5 depicts the architecture of *NimbusML*.

4.2 Pipeline Execution

Independently from which API is used by the developer or the task to execute (training, testing or prediction), ML.NET eventually runs a learning pipeline. As we will describe shortly, thanks to lazy evaluation, immutability and the Just In Time (JIT) compiler provided by the .NET runtime, ML.NET is able to generate highly efficient computations. Internally, transforms consume *DataView* columns as input and produce one (or more) *DataView* columns as output. Columns are immutable whereby multiple downstream

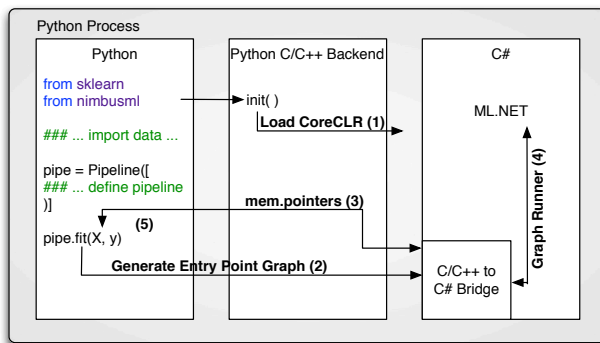


Figure 5: Architecture and execution flow of NimbusML. When NimbusML is imported, a CoreCLR instance with ML.NET is initialized (1). When a pipeline is submitted for execution, the entry point graph of NimbusML operators is extracted and submitted for execution (2) together with pointers to the datasets (3) if residing in memory. The graph runner executes the graph (4) and returns the result to the Python program as a data frame object (5).

operators can safely consume the same input without triggering any re-execution. Trainable transforms and learners, instead, need to be trained before generating the related output DataView column(s). Therefore, when a pipeline is submitted for execution (e.g., by calling Train), each trainable transform / learner is trained in topological order. For each of them, a one-time initialization cost is paid to analyze the cursors in the pipeline, e.g., each cursor checks the active columns and the expected input type(s).

CPU Efficiency. Output of the initialization process at each DataView’s cursor is a lambda function, named getter, condensing the logic of the operator into one single call. Each getter in turn triggers the generation of the getter function of the upstream cursor until a data source is found (e.g., a cached DataView or input data). When all getters are initialized, each upstream getter function is used in the downstream one, so that, from the outer cursor perspective, computation is represented as a chain of lambda function calls. Once the initialization process is complete, the cursor iterates over the input data and executes the training (or prediction) logic by calling its getter. At execution time, the chain of getter are JIT-compiled by the .NET runtime to form a unique, highly efficient function executing the whole pipeline (up to that point) on a single call. The process is repeated until no trainable operator is left in the pipeline.

Memory Efficiency. Cursors are inherently efficient from a memory allocation perspective. Advancing the cursor to the next row requires no memory allocation. Retrieving primitive column values from a cursor also requires no memory allocation. To retrieve vector column values from a cursor, the caller to the getter can optionally provide buffers into which the values should be copied. When the provided buffers are sufficiently large, no additional memory allocation is required. When the buffers are not provided or are too small, the cursor allocates buffers of sufficient size to hold the values. This cooperative buffer sharing protocol eliminates the need to allocate separate buffers for each row.

Parallel Computation. ML.NET provides 2 possibilities for improving performance through parallel processing: (1) from directly

inside the algorithm; and (2) using parallel cursoring. The former case is strictly related to the algorithm implementation. In the latter case, a transform requires a *cursor set* from its input DataView. Cursors sets are propagated upstream until a data source is found: at this point cursor set are mapped into available threads, and data is collaboratively scanned. From a callers perspective, cursor sets return a consolidated, unique, cursor, although, from an execution perspective, cursor’s data scan is split into concurrent threads.

4.3 Learning Tasks in ML.NET

We have surveyed the top learners by number of unique users within Microsoft. We will here subdivide the usage by what appears to be most popular.²

Gradient Boosting Trees. The most popular *single* learner is FastTree: a gradient boosting algorithm over trees. FastTree uses an algorithm that was originally engineered for web-page ranking, but was later adapted to other tasks—and in fact the ranking task, while still having thousands of unique users within Microsoft, is comparatively much less popular than the more classical tasks of binary classification and regression. This learner requires a representation of the dataset in memory to function. Interestingly, a random-forest algorithm based on the same underlying code sees only a small fraction of the usage of the boosting-based interface. As point of interest, a faster implementation of the same basic algorithm called LightGBM [16] was introduced few years ago, and is gaining in popularity. However, usage of LightGBM still remains a fraction of the original algorithm, possibly for reasons of inertia.

Linear Learners. While the most popular learner is based on boosted decision trees, one could argue that collectively linear learners see more use. Linear learners in contrast to the tree based algorithm do work well over streaming data. The most popular linear learners are basic implementations of such familiar algorithms like *averaged perceptron*, *online gradient descent*, *stochastic gradient descent*. These scale well and are quite simple to use, though they lack the sophistication of other methods. Following this “basic set” in popularity are a set of linear learners based on OWL-QN [1], a variant of L-BFGS capable of L1 regularization, and thus learning sparse linear prediction rules. These algorithms have some advantages, but because these algorithms on the whole require more passes over the dataset to converge to a good answer compared to the earlier stochastic methods, they are less popular. Even less popular still is an SDCA-based [32] algorithm, as well as a linear Pegasos SVM trainer [30]. Each still has in excess of a thousand users, but this is still considerably less than the other algorithms.

Other Learners. Compared to these supervised tasks, unsupervised tasks like clustering and anomaly detection are definitely part of the long tail, each having perhaps only hundreds of unique users. Even more obscure tasks like sequence classification and recommendation, despite supporting quite important products, seem to have only a few unique users. Readers may note neural networks and deep learning as a very conspicuous omission. We do internally have a neural network that sees considerable use, but in the open source version we instead provide an interface to other, already available, neural network frameworks.

² Note that using unique users to assess popularity is indeed wrong: just because a learner is not popular does not mean that its support is not strategically important.

5 EXPERIMENTAL EVALUATION

In this section we compare ML.NET against Sklearn and H2O. For ML.NET we employ its regular learning pipeline API (ML.NET) and the Python bindings through NimbusML. For NimbusML we test both reading from Pandas' Data Frames (NimbusML-DF) and the streaming API (NimbusML-DV) of DataView which allows to directly stream data from disk. We tried as much as we could to use the same data transforms and ML models across different toolkits in order to measure the performance of the frameworks and not of the data scientist. For the same reason, for each pipeline we use the default parameters. We report the total runtime (training plus testing), AUC for the classification problems, and Root Mean Square (RMS) for the regression one. To examine the scale-out performance of the frameworks, in the first set of experiments we train the pipelines over 0.1%, 1%, 10% and 100% of samples of the training data over all accessible cores. Finally, Section 5.4 contains a scale-up experiment where we measure how the performance change as we change the number of used cores. For this final set of experiments we only compare ML.NET/NimbusML and H2O because Scikit-learn is only able to use a single core.

All the experiments are run three times and the minimum runtime and the average accuracy are reported. Further information about the experiments are reported into the Reproducibility Section attached to the end of the paper.

Configuration. All the experiments in the paper are carried out on Azure on a D13v2 VM with 56GB of RAM, 112 GB of Local SSD and a single Intel(R) Xeon(R) @ 2.40GHz processor. We used ML.NET version 0.1, NimbusML version 0.6, Scikit-learn version 0.19.1 and H2O version 3.20.0.7.

Scenarios. In our evaluation we train models for four different scenarios. In the first scenario we aim at predicting the click through rate for an online advertisement. In the second scenario we train a model to predict the sentiment class for e-commerce customer reviews. In the third scenario we predict the delay for scheduled flights according to historical records. (Note that the first two are classification problems, while the third one is a regression problem.) In these first three scenario we allow the systems to uses all the available processing resources. Conversely, in the last scenario we report the performance as we scale-up the number of available cores. For this scenario we chose two different test situations: one where the dataset is a small sample (1%), and one where the dataset is large. In this way we can evaluate the difference scale-up performance.

Datasets. For each scenario we use a different dataset. For the first scenario we use the Criteo dataset [4]. The full training dataset includes around 45 million records, and the size of the training file is around 10GB. Among the 39 features, 14 are numeric while the remaining are categorical. In the set of experiments for the second scenario we employed the Amazon Review dataset [14]. In this case the full training dataset includes around 17 million records, and the size of the training file is around 9GB. For this scenario we only use one text column as the input feature. Finally, for the third scenario we use the Flight Delay dataset [26]. The training dataset includes around 1 million records, the size of the training file is around 1GB and each record contains 631 columns.

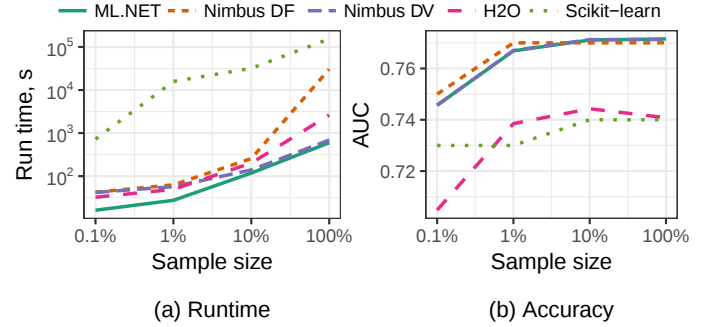


Figure 6: Experimental results for the Criteo dataset.

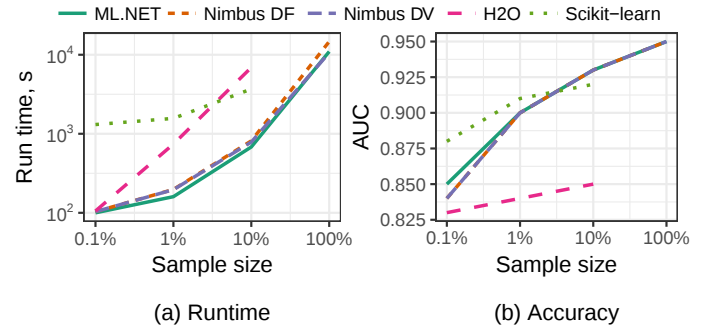


Figure 7: Experimental results for the Amazon dataset.

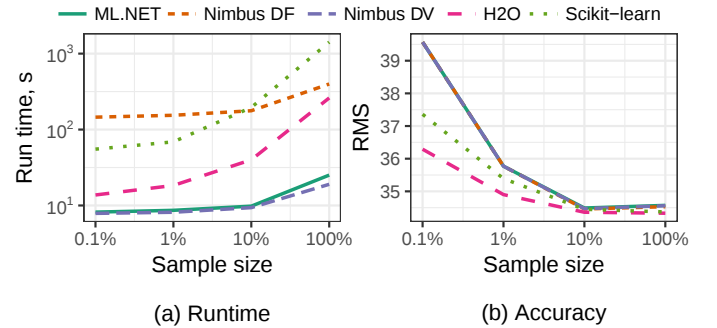


Figure 8: Experimental results for the Flight Delay dataset.

5.1 Criteo

In the first scenario we build a pipeline which (1) fills in the missing values in the numerical columns of the dataset; (2) encodes a categorical columns into a numeric matrix using a hash function; and (3) applies a gradient boosting classifier (LightGBM for ML.NET). Figure 6a shows the total runtime (including training and testing), while Figure 6b depicts the AUC on the test dataset. As we can see, ML.NET has the best performance, while NimbusML-DV ranks second. Both NimbusML-DF and H2O show good runtime performance, especially for smaller datasets. In this experiment Scikit-learn has the worst running time: with the full training dataset, ML.NET and NimbusML-DV train in around 10 minutes while Scikit-learn takes more than 2 days. Regarding the accuracy, we can notice that the results from NimbusML-DV/NimbusML-DF and ML.NET are very similar and all of them dominate Scikit-learn/H2O by a large margin. This is mainly due to the superiority of LightGBM versus the gradient boosting algorithm used in the latters.

5.2 Amazon

In this scenario the pipeline first featurizes the text column of the dataset and then applies a linear classifier. For the featurization part, we used the `FeaturizeText` transform in ML.NET and the `TfidfVectorizer` in Scikit-learn. The `FeaturizeText/TfidfVectorizer` extracts numeric features from the input corpus by producing a matrix of token ngrams counts. H2O implements the Skip-Gram word2vec model [22] as the only text featurizer. In both the classical approach based on ngrams and the neural network approach, text featurization is a heavy operation. In fact, both Sklearn and H2O throw overflow/memory errors when training with the full dataset because of the large vocabulary sizes. Therefore, no results are reported for Sklearn and H2O, trained with the full dataset. As we can see from Figure 7a, ML.NET is able to complete all experiments, and all versions (ML.NET, NimbusML-DF, NimbusML-DV) show similar runtime performance. Interestingly enough, NimbusML-DF is able to complete over the full dataset, while Scikit-learn is not. This is because, under the hood, NimbusML-DF uses DataView to execute the pipeline in a streaming fashion, whereas Scikit-learn materializes partial results in data frames. Regarding the measured accuracy (Figure 7b), Sklearn shows the highest AUC with 0.1% of the dataset, likely due to the different initial settings for the algorithm, but for the remaining data points ML.NET performs better.

5.3 Flight Delay

For this dataset we pre-process all the feature columns into numeric values and we compare the performance of a single operator in ML.NET/NimbusML-DV/NimbusML-DF versus Sklearn/H2O without a pipeline (LightGBM vs gradient boosting). The results for speed and accuracy are reported in Figure 8. Since this is a regression problem, we report the RMS on the test set. Interestingly, for this dataset NimbusML-DF runtime is considerably worst than ML.NET and NimbusML-DV, and even worst than Scikit-learn for small samples. We found that this is due to the fact that since we are applying the algorithm directly over the input data frame, and the algorithm requires several passes over the data, the overhead of accessing the data residing in C/C++ dominates the runtime. Additionally we found that the accuracy of ML.NET, especially for smaller samples, is worst than both H2O and Scikit-learn. By examining the execution, we found that with a boosting tree trained with small subset of the data, a simpler tree from Sklearn predicts better than LightGBM in ML.NET. However, with 0.1% of the dataset, those models are trained with 1000 samples and over 600 features. Therefore models can be easily overfitted. Models from ML.NET converge faster (with much smaller error for the training set, i.e. RMS = 18 for ML.NET and 28 for Sklearn) and are overfitted. In this specific case, Sklearn/H2O models have better performance over the test set as they are less overfitted. For large training sets, all systems converge to approximately the same accuracy, although ML.NET is more than 10× faster than Scikit-learn and H2O.

5.4 Scale-up Experiments

In Figure 9 we show how the performance of the systems change as we increase the number of cores. For these experiments we use both a small sample (Amazon 1%, depicted on the left-hand side

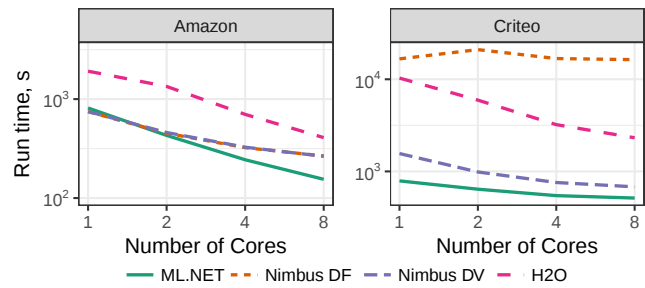


Figure 9: Scale-up experiments.

of Figure 9) and a full dataset (Criteo, right-hand side of Figure 9) to compare how the systems perform under different stress situations. For the Amazon sample, ML.NET and H2O scale linearly, while Nimbus scalability decreases due to the overheads between components. On the full Criteo experiments, we see that ML.NET and Nimbus-DV scalability is less compared to H2O, however the latter system is about 10× slower than the formers. Nimbus-DF performance do not increase as we increase the number of cores because of the overhead of reading/writing Data Frames. Recall that Scikit-learn is not reported for these experiments because parallel training is not supported and therefore performance do not change.

6 LESSONS LEARNED

The current version of ML.NET is the result of almost a decade of design choices. Originally there was no DataView concept nor pipeline but, instead, all featurizers and models computations were applied over an enumerable of *instances*: a vector of floats (sparse or dense) paired with a label (which was itself a float). With this design, certain types of model classes such as neural networks, recommender systems or time series were difficult to express. Similarly, there was no notion of intermediate computation nor any abstraction allowing to compose arbitrary sequences of operations. Because of these limitations, DataView was introduced. With DataView, developers can easily stick together different operations, even if internally each operation can be arbitrarily complex and produce an arbitrarily complex output.

In early versions of DataView, we explored using .NET's `IEnumerable`. At the beginning this was the most natural choice, but we soon found that such approach was generating too many memory allocations. This led to the cursor approach with cooperative buffer sharing. In the first version of DataView with cursoring, the implementation did not have any getter but rather a method `GetColumnValue<T>(int col, ref T val)`. However, the approach had the following problems: (1) every call had to verify that the column was active; (2) every call had to verify that T was of the right type; and (3) when this is part of a transform in a pipelines (as they often do) each access would be then be accompanied by a virtual method call to the upstream cursor's `GetColumnValue`. In contrast, consider the situation with the lambda functions provided by getters: (1) the verification of whether the column is active happens exactly once; (2) the verification of types happens exactly once; and (3) rather than every access being passed up through a chain of virtual function calls, only a getter function is used from the cursor, and every data access is executed directly and JIT-ed. The practical result of this is that, for some workloads, the "getter" method became an order of magnitude faster.

7 RELATED WORK

Scikit-learn has been developed as a machine learning tool for Python and, as such, it mainly targets interactive use cases running over datasets fitting in main memory. Given its characteristic, Scikit-learn has several limitations when it comes to experimenting over Big Data: runtime performance are often inadequate; large datasets and feature sets are not supported; datasets cannot be streamed but instead they can only be accessed in batch from main memory. Finally, multi-core processing is not natively supported because of Python's global interpreter lock (although some work exists [5] trying to solve some of these issues for embarrassingly parallel computations such as cross validation or tree ensemble models). ML.NET solves the aforementioned problems thanks to the DataView abstraction (Section 3.1) and several other techniques inspired by database systems. Nonetheless, ML.NET provides Sklearn-like Python bindings through NimbusML (Section 4.1) such that users already familiar with the former can easily switch to ML.NET.

MMLib [10], Uber's Michelangelo [21], H2O [13] and Salesforce's TransmogriAI [33] are machine learning systems built off Scikit-learn limitations. Differently than Scikit-learn, but similarly to ML.NET, MMLib, Michelangelo, H2O and TransmogriAI are not "data science languages" but enterprise-level environments for building machine learning models for applications. These systems are all JVM-based and they all provide performance for large dataset mainly through in-memory distributed computation (based on Apache Spark [36]). Conversely, ML.NET main focus is efficient single machine computation.

In Section 5 we compared against H2O because we deem this framework as the closest to ML.NET. While ML.NET uses DataView, H2O employs H2O Data Frames as abstraction of data. Differently than DataView however, H2O Data Frames are not immutable but "fluid", i.e., columns can be added, updated and removed by modifying the base data frame. Fluid vectors are compressed so that larger than RAM working sets can be used. H2O provides several interfaces (R, Python, Scala) and large variety of algorithms.

Other popular machine learning frameworks are TensorFlow [8], PyTorch [28], CNTK [29], MXNet [9], Caffe2 [3]. These systems however mostly focus on Deep Neural Network models (DNNs). If we look both internally at Microsoft, and at external surveys [25] we find that DNNs are only part of the story, whereas the great majority of models used in practice by data scientists are still generic machine learning models. We are however studying how to merge the two worlds [35].

8 CONCLUSIONS

Machine learning is rapidly transitioning from a niche field to a core element of modern application development. This raises a number of challenges Microsoft faced early on. ML.NET addresses a core set of them: it brings machine learning onto the same technology stack as application development, delivers the scalability needed to work on datasets large and small across a myriad of devices and environments, and, most importantly, allows for complete pipelines to be authored and shared in an efficient manner. These attributes of ML.NET are not an accident: they have been developed in response to requests and insights from thousands of data scientists at Microsoft who used it to create hundreds of services and products used by hundreds of millions of people worldwide every day.

ML.NET [23] and NimbusML [24] are open source and publicly available under the MIT license.

ACKNOWLEDGMENTS

We would like to thank all ML.NET contributors, and in particular Eric Erhardt, Costin Eseanu, Shon Katzenberger, and Sarthak Shah.

REFERENCES

- [1] Galen Andrew and Jianfeng Gao. 2007. Scalable training of L1-regularized log-linear models. In *ICML 2007*.
- [2] Boost.Python. 2019. <http://wiki.python.org/moi/boost.python>. (2019).
- [3] Caffe2. 2018. <http://caffe2.ai/>. (2018).
- [4] Criteo. 2014. Kaggle Challenge. (2014). <http://labs.criteo.com/2014/02/kaggle-display-advertising-challenge-dataset/>
- [5] Joblib Documentation. 2018. <http://media.readthedocs.org/pdf/joblib/latest/joblib.pdf>. (2018).
- [6] Christopher Olston et al. 2017. TensorFlow-Serving: Flexible, High-Performance ML Serving. In *Workshop on ML Systems at NIPS*.
- [7] Daniel Crankshaw et al. 2017. Clipper: A Low-Latency Online Prediction Serving System. In *NSDI 2017*.
- [8] Martin Abadi et al. 2016. TensorFlow: A system for large-scale machine learning. In *OSDI 16*. 265–283.
- [9] Tianqi Chen et al. 2015. MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems. *CoRR* abs/1512.01274 (2015).
- [10] Xiangrui Meng et al. 2016. MMLib: Machine Learning in Apache Spark. *JMLR* 17, 34 (2016), 1–7.
- [11] Jerome H. Friedman. 2000. Greedy Function Approximation: A Gradient Boosting Machine. *Annals of Statistics* 29 (2000), 1189–1232.
- [12] G. Graefe. 1994. Volcano: An Extensible and Parallel Query Evaluation System. *TKDE* 6, 1 (Feb. 1994), 120–135.
- [13] H2O. 2019. <https://github.com/h2oai/h2o-3>. (2019).
- [14] Ruining He and Julian McAuley. 2016. Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering. In *WWW 2016*. 507–517.
- [15] Jupyter. 2018. <http://jupyter.org/>. (2018).
- [16] Guolin Ke et al. 2017. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *NIPS 2017*. 3146–3154.
- [17] Yunseong Lee, Alberto Scolari, Byung-Gon Chun, and et al. 2018. From the Edge to the Cloud: Model Serving in ML.NET. *IEEE Data Eng. Bull.* 41, 4 (2018), 46–53.
- [18] Yunseong Lee, Alberto Scolari, Byung-Gon Chun, and et al. 2018. PRETZEL: Opening the Black Box of Machine Learning Prediction Serving Systems. In *OSDI 2018*. 611–626.
- [19] Matplotlib. 2018. <https://matplotlib.org/>. (2018).
- [20] Wes McKinney. 2011. pandas: a Foundational Python Library for Data Analysis and Statistics. (01 2011).
- [21] Michelangelo. 2018. <http://eng.uber.com/michelangelo/>. (2018).
- [22] Tomas Mikolov et al. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *NIPS 2013*. USA, 3111–3119.
- [23] ML.NET. 2019. <https://github.com/dotnet/machinelearning>. (2019).
- [24] NimbusML. 2019. <https://github.com/Microsoft/NimbusML>. (2019).
- [25] The State of Data Science and Machine Learning. 2017. (2017). <https://www.kaggle.com/surveys/2017/>
- [26] Bureau of Transportation Statistics. 2018. Flight Delay Dataset. (2018). https://www.transtats.bts.gov/Fields.asp?Table_ID=236
- [27] Fabian Pedregosa et al. 2011. Scikit-learn: Machine Learning in Python. *JMLR* 12 (Nov. 2011), 2825–2830.
- [28] PyTorch. 2018. <https://pytorch.org/>. (2018).
- [29] Frank Seide and Amit Agarwal. 2016. CNTK: Microsoft's Open-Source Deep-Learning Toolkit. In *KDD 2016*. 2135–2135.
- [30] Shai Shalev-Shwartz et al. 2011. Pegasos: primal estimated sub-gradient solver for SVM. *Mathematical Programming* 127, 1 (01 Mar 2011), 3–30.
- [31] Mike Stonebraker et al. 2005. C-store: A Column-oriented DBMS. In *Vldb 2005*. 553–564.
- [32] Kenneth Tran et al. 2015. Scaling Up Stochastic Dual Coordinate Ascent. In *KDD 2015*. 1185–1194.
- [33] TransmogriAI. 2018. <https://transmogri.ai/>. (2018).
- [34] S. van der Walt, S. C. Colbert, and G. Varoquaux. 2011. The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science Engineering* 13, 2 (March 2011), 22–30.
- [35] Gyeong-In Yu et al. 2018. Making Classical Machine Learning Pipelines Differentiable: A Neural Translation Approach. *SysML Workshop at NIPS* (2018).
- [36] Matei et al. Zaharia. 2012. Resilient distributed datasets: a fault-tolerant abstraction for in-memory cluster computing. In *NSDI 2012*.
- [37] Zeppelin. 2018. <https://zeppelin.apache.org/>. (2018).
- [38] Martin Zinkevich. 2019. Rules of Machine Learning: Best Practices for ML Engineering. (2019).

REPRODUCIBILITY

In this section we report additional information and insights regarding the experimental evaluation of Section 5. For reproducibility we added the Python scripts (with the employed parameters) we used for ML.NET (i.e., NimbusML with Data Frames), Scikit-learn and H2O.

Criteo. For this scenario we build a pipeline which (1) fills in the missing values in the numerical columns of the dataset; (2) encodes a categorical columns into a numeric matrix using a hash function; and (3) applies a gradient boosting classifier (LightGBM for ML.NET/ NimbusML). The training scripts used for NimbusML, Scikit-learn, and H2O are shown in Figure 10, Figure 11, and Figure 12, respectively.

```

1 import pandas as pd
2 from nimbusml.preprocessing.missing_values import Handler
3 from nimbusml.feature_extraction.categorical import OneHotHashVectorizer
4 from nimbusml import Pipeline
5 from nimbusml.ensemble import LightGbmBinaryClassifier
6
7 def LoadData(fileName):
8     data = pd.read_csv(fileName, header = None, sep='\t',
9         error_bad_lines=False)
10    colnames = [ 'V'+str(x) for x in range(0, data.shape[1])]
11    data.columns = colnames
12    data.iloc[:,14:] = data.iloc[:,14:].astype(str)
13    return data.iloc[:,1:],data.iloc[:,0]
14
15 # Load data
16 dataTrain, labelTrain = LoadData(train_file)
17
18 # Create pipeline
19 nimbus_ppl = Pipeline([
20     Handler(columns = [ 'V'+str(x) for x in range(1, 14)]),
21     OneHotHashVectorizer(columns = [ 'V'+str(x) for x in range(14, 40)]),
22     LightGbmBinaryClassifier(feature = [ 'V'+str(x) for x in range(1, 40)])])
23
24 # Train
25 nimbus_ppl.fit(dataTrain,labelTrain,parallel = n_thread)

```

Figure 10: Criteo in NimbusML with DataFrames.

```

1 import pandas as pd
2 from sklearn.preprocessing import Imputer, FunctionTransformer
3 from sklearn.feature_extraction import FeatureHasher as FeatureHasher
4 from sklearn.ensemble import GradientBoostingClassifier as
5     GradientBoostingClassifier
6 from sklearn.pipeline import Pipeline, make_pipeline, FeatureUnion
7
8 # Helper function for sklearn pipeline to select columns
9 def CategoricalColumn(X):
10     return X[[ 'V'+str(x) for x in range(14,
11         40)]]].astype(str).to_dict(orient = 'records')
12
13 def NumericColumn(X):
14     return X[[ 'V'+str(x) for x in range(1, 14)]]
15
16 def LoadData(fileName):
17     data = pd.read_csv(fileName, header = None, sep='\t',
18         error_bad_lines=False)
19     colnames = [ 'V'+str(x) for x in range(0, data.shape[1])]
20     data.columns = colnames
21     data.iloc[:,14:] = data.iloc[:,14:].astype(str)
22     return data.iloc[:,1:],data.iloc[:,0]
23
24 # Load data
25 dataTrain, labelTrain = LoadData(train_file)
26
27 # Create pipeline
28 imp = make_pipeline(FunctionTransformer(NumericColumn, validate=False),
29     Imputer())
30 hasher = make_pipeline(FunctionTransformer(CategoricalColumn,
31     validate=False), FeatureHasher())
32 sk_ppl = Pipeline([
33     ('union', FeatureUnion(
34         transformer_list=[('hasher',hasher), ('imp',imp)]),
35     ('tree', GradientBoostingClassifier())])
36
37 # Train
38 sk_ppl.fit(dataTrain,labelTrain)

```

Figure 11: Criteo training pipeline in Scikit-learn.

```

1 import h2o
2 from h2o.estimators.gbm import H2OGradientBoostingEstimator
3
4 h2o.init(nthreads=n_thread, min_mem_size='20g')
5
6 # Load data
7 dataTrain = h2o.import_file(train_file,header = 0)
8
9 # Create pipeline
10 response = "C1"
11 features = ["C"+ str(x) for x in range(2,41)]
12 dataTrain[response] = dataTrain[response].asfactor()
13 gbm = H2OGradientBoostingEstimator()
14
15 # Train
16 gbm.train(x=features, y=response, training_frame=dataTrain)

```

Figure 12: Criteo training pipeline implemented in H2O.

Amazon. In this pipeline first we featurize the text column of the dataset and then applies a linear classifier. For the featurization part, we used the TextFeaturizer transform in ML.NET and the TfidfVectorizer in Scikit-learn. The Text Featurizer/Tfidf Vectorizer extracts numeric features from the input corps by producing a matrix of token ngrams counts. As parameters, in both case we used word ngrams of size 1 and 2, and char ngrams of size 1, 2 and 3. H2O implements the Skip-Gram word2vec model [22] as the only text featurizer. The training scripts are shown in Figure 13, Figure 14, and Figure 15.

```

1 import pandas as pd
2 from nimbusml.linear_model import AveragedPerceptronBinaryClassifier
3 from nimbusml.feature_extraction.text import NgramFeaturizer
4 from nimbusml.feature_extraction.text.extractor import Ngram
5 from nimbusml import Pipeline
6
7 def LoadData(fileName):
8     data = pd.read_csv(fileName, sep='\t', error_bad_lines=False)
9     return data.iloc[:,1],data.iloc[:,0]
10
11 # Load data
12 dataTrain, labelTrain = LoadData(train_file)
13
14 # Create pipeline
15 nimbus_ppl = Pipeline([
16     NgramFeaturizer(word_feature_extractor = Ngram(ngram_length = 2),
17         char_feature_extractor = Ngram(ngram_length = 3)),
18     AveragedPerceptronBinaryClassifier()])
19
20 # Train
21 nimbus_ppl.fit(dataTrain,labelTrain,parallel = n_thread)

```

Figure 13: Amazon pipeline implemented in NimbusML with Data Frames.

```

1 import pandas as pd
2 from sklearn.feature_extraction.text import TfidfVectorizer
3 from sklearn.pipeline import Pipeline, make_pipeline, FeatureUnion
4 from sklearn.linear_model import Perceptron
5
6 def LoadData(fileName):
7     data = pd.read_csv(fileName, sep='\t', error_bad_lines=False)
8     return data.iloc[:,1],data.iloc[:,0]
9
10 # Load data
11 dataTrain, labelTrain = LoadData(train_file)
12
13 # Create pipeline
14 wg = make_pipeline(TfidfVectorizer(analyzer='word', ngram_range = (1,2)))
15 cg = make_pipeline(TfidfVectorizer(analyzer='char', ngram_range = (1,3)))
16
17 sk_ppl = Pipeline([
18     ('union', FeatureUnion(
19         transformer_list=[('wordg',wg), ('charg',cg)]),
20     ('ap', Perceptron())])
21
22 # Train
23 sk_ppl.fit(dataTrain,labelTrain)

```

Figure 14: Amazon pipeline implemented in Scikit-learn.

```

1 import h2o
2 from h2o.estimators.word2vec import Word2vecEstimator
3 from h2o.estimators.glm import GeneralizedLinearEstimator
4
5 h2o.init(nthreads=n_thread, min_mem_size='20g')
6
7 # Load data
8 dataTrain = h2o.import_file(train_file, header = 1, sep = "\t")
9
10 def tokenize(sentences):
11     tokenized = sentences.tokenize("\\W+")
12     return tokenized
13
14 # Create featurization pipeline
15 wordsTrain = tokenize(dataTrain["text"])
16 w2v_model = Word2vecEstimator(sent_sample_rate = 0.0, epochs = 5)
17
18 # Train featurizer
19 w2v_model.train(training_frame=wordsTrain)
20
21 # Create model pipeline
22 wordsTrain_vecs = w2v_model.transform(wordsTrain, aggregate_method =
23     "AVERAGE")
24 dataTrain = dataTrain.cbind(wordsTrain_vecs)
25 glm_model = GeneralizedLinearEstimator(family= "binomial", lambda_ = 0,
26     compute_p_values = True)
27
28 # Train
29 glm_model.train(wordsTrain_vecs.names, 'Label', dataTrain)

```

Figure 15: Amazon pipeline implemented in H2O.

Flight Delay. Here we pre-process all the feature columns into numeric values and we compare the performance of a single operator in ML.NET/NimbusML-DV/NimbusML-DF versus Sklearn/H2O without a pipeline (LightGBM vs GradientBoosting). The NimbusML, Scikit-learn and H2O pipelines are reported in Figure 16, Figure 17, and Figure 18, respectively.

Scale-up Experiments For this set of experiments we used the exact same scripts introduced above, except that we properly change the `n_thread` parameter to set the number of used cores.

```

1 import pandas as pd
2 from nimbusml.ensemble import LightGbmRegressor
3
4 # Load data
5 df_train = pd.read_csv(train_file, header = None)
6 df_train.columns = df_train.columns.astype(str)
7
8 # Create pipeline
9 nimbus_m = LightGbmRegressor()
10
11 # Train
12 nimbus_m.fit(df_train.iloc[:,1:], df_train.iloc[:,0], parallel = n_thread)

```

Figure 16: Flight Delay pipeline implemented in NimbusML using the Data Frame API.

```

1 import pandas as pd
2 from sklearn.ensemble import GradientBoostingRegressor
3 from sklearn.metrics import mean_squared_error
4
5 # Load data
6 df_train = pd.read_csv(train_file, header = None)
7 df_train.columns = df_train.columns.astype(str)
8
9 # Create pipeline
10 sklearn_m = GradientBoostingRegressor()
11
12 # Train
13 sklearn_m.fit(df_train.iloc[:,1:], df_train.iloc[:,0])

```

Figure 17: Flight Delay pipeline in Scikit-learn.

```

1 import h2o
2 from h2o.estimators.gbm import H2OGradientBoostingEstimator
3
4 h2o.init(nthreads=n_thread, min_mem_size='20g')
5
6 # Load data
7 dataTrain = h2o.import_file(train_file, header = 0)
8
9 # Create pipeline
10 response = "C1"
11 features = ["C" + str(x) for x in range(2,41)]
12 gbm = H2OGradientBoostingEstimator()
13
14 # Train
15 gbm.train(x=features, y=response, training_frame=dataTrain)

```

Figure 18: Flight Delay pipeline implemented in H2O.