

Large-Scale Learning from Data Streams with Apache SAMOA

Nicolas Kourtellis¹, Gianmarco De Francisci Morales², and Albert Bifet³

¹ Telefonica Research, Spain, nicolas.kourtellis@telefonica.com

² Qatar Computing Research Institute, Qatar, gdfm@acm.org

³ LTCI, Télécom ParisTech, France, albert.bifet@telecom-paristech.fr

Abstract. Apache SAMOA (Scalable Advanced Massive Online Analysis) is an open-source platform for mining big data streams. Big data is defined as datasets whose size is beyond the ability of typical software tools to capture, store, manage, and analyze, due to the time and memory complexity. Apache SAMOA provides a collection of distributed streaming algorithms for the most common data mining and machine learning tasks such as classification, clustering, and regression, as well as programming abstractions to develop new algorithms. It features a pluggable architecture that allows it to run on several distributed stream processing engines such as Apache Flink, Apache Storm, and Apache Samza. Apache SAMOA is written in Java and is available at <https://samoa.incubator.apache.org> under the Apache Software License version 2.0.

1 Introduction

Big data are “data whose characteristics force us to look beyond the traditional methods that are prevalent at the time” [18]. For instance, social media are one of the largest and most dynamic sources of data. These data are not only very large due to their fine grain, but also being produced continuously. Furthermore, such data are nowadays produced by users in different environments and via a multitude of devices. For these reasons, data from social media and ubiquitous environments are perfect examples of the challenges posed by big data.

Currently, there are two main ways to deal with these challenges: streaming algorithms and distributed computing (e.g., MapReduce). Apache SAMOA aims at satisfying the future needs for big data stream mining by combining the two approaches in a single platform under an open source umbrella [9].

Data mining and machine learning are well established techniques among Web and social media companies to draw insights from data coming from ubiquitous and social environments. Online content analysis for detecting aggression [6], stock trade volume prediction [5], online spam detection [7], recommendation [12], and personalization [10] are just a few of the applications made possible by mining the huge quantity of data available nowadays. Just think of Facebook’s relevance algorithm for the news feed for a famous example.

In order to cope with Web-scale datasets, data scientists have resorted to *parallel and distributed computing*. MapReduce [11] is currently the de-facto standard programming paradigm in this area, mostly thanks to the popularity of Hadoop⁴, an open source implementation of MapReduce. Hadoop and its ecosystem (e.g., Mahout⁵) have proven to be an extremely successful platform to support the aforementioned process at web scale.

However, nowadays most data are generated in the form of a stream, especially when dealing with social media. Batch data are just a snapshot of streaming data obtained in an interval (window) of time. Researchers have conceptualized and abstracted this setting in the *streaming model*. In this model, data arrive at high speed, one instance at a time, and algorithms must process it in one pass under very strict constraints of space and time. Extracting knowledge from these massive data streams to perform dynamic network analysis [19] or to create predictive models [1], and using them, e.g., to choose a suitable business strategy, or to improve healthcare services, can generate substantial competitive advantages. Many applications need to process incoming data and react on-the-fly by using comprehensible prediction mechanisms (e.g., card fraud detection) and, thus, streaming algorithms make use of probabilistic data structures to give fast and approximate answers.

On the one hand, MapReduce is not suitable to express streaming algorithms. On the other hand, traditional sequential online algorithms are limited by the memory and bandwidth of a single machine. *Distributed stream processing engines* (DSPEs) are a new emergent family of MapReduce-inspired technologies that address this issue. These engines allow to express parallel computation on streams, and combine the scalability of distributed processing with the efficiency of streaming algorithms. Examples include Storm⁶, Flink⁷, Samza⁸, and Apex⁹.

Alas, currently there is no common solution for mining big data streams, that is, for running data mining and machine learning algorithms on a distributed stream processing engine. The goal of Apache SAMOA is to fill this gap, as exemplified by Figure 1(left).

2 Description

Apache SAMOA (SCALABLE ADVANCED MASSIVE ONLINE ANALYSIS) is a platform for mining big data streams [8]. For a simple analogy, think of Apache SAMOA as Mahout for streaming. As most of the rest of the big data ecosystem, it is written in Java.

Apache SAMOA is both a framework and a library. As a framework, it allows the algorithm developer to abstract from the underlying execution engine, and

⁴ <http://hadoop.apache.org>

⁵ <http://mahout.apache.org>

⁶ <http://storm.apache.org>

⁷ <http://flink.apache.org>

⁸ <http://samza.apache.org>

⁹ <https://apex.apache.org>

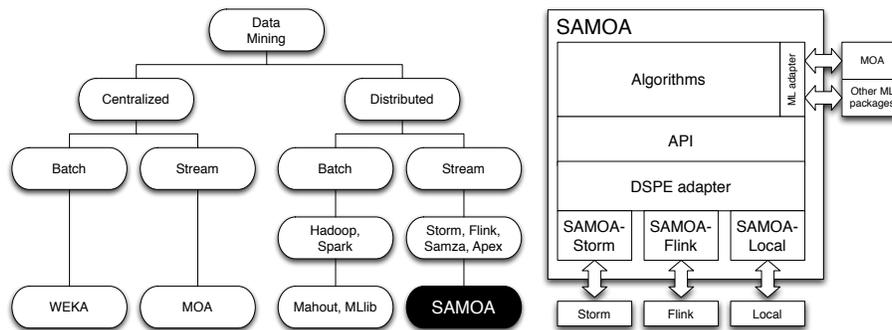


Fig. 1: (Left) Taxonomy of data mining and machine learning tools. (Right) High level architecture of Apache SAMOA.

therefore reuse their code on different engines. It features a pluggable architecture that allows it to run on several distributed stream processing engines such as Storm, Flink, Samza, and Apex. This capability is achieved by designing a minimal API that captures the essence of modern DSPEs. This API also allows to easily write new bindings to port Apache SAMOA to new execution engines. Apache SAMOA takes care of hiding the differences of the underlying DSPEs in terms of API and deployment.

As a library, Apache SAMOA contains implementations of state-of-the-art algorithms for distributed machine learning on streams. For classification, Apache SAMOA provides a Vertical Hoeffding Tree (VHT), a distributed streaming version of a decision tree. For clustering, it includes an algorithm based on CluStream. For regression, HAMR, a distributed implementation of Adaptive Model Rules. The library also includes meta-algorithms such as bagging and boosting [26]. The platform is intended to be useful for both research and real world deployments.

Related Work. We identify two frameworks that belong to the category of distributed streaming machine learning: Jubatus and StormMOA. Jubatus¹⁰ is an example of distributed streaming machine learning framework. It includes a library for streaming machine learning such as regression, classification, recommendation, anomaly detection and graph mining. It introduces the local ML model concept which means there can be multiple models running at the same time and they process different sets of data. Using this technique, Jubatus achieves scalability via horizontal parallelism in partitioning data. We test horizontal parallelism in our experiments, by implementing a horizontally scaled version of the hoeffding tree. Jubatus establishes tight coupling between the machine learning library implementation and the underlying distributed stream processing engine (SPE). The reason is Jubatus builds and implements its own

¹⁰ <http://jubat.us/en>

custom distributed SPE. In addition, Jubatus does not offer any tree learning algorithm, as all of its models need to be linear by construction.

StormMOA¹¹ is a project to combine MOA with Storm to satisfy the need of scalable implementation of streaming ML frameworks. It uses Storm’s Trident abstraction and MOA library to implement OzaBag and OzaBoost[21]. Similarly to Jubatus, StormMOA also establishes tight coupling between MOA (the machine learning library) and Storm (the underlying distributed SPE). This coupling prevents StormMOA’s extension by using other SPEs to execute the machine learning library. StormMOA only allows to run a single model in each Storm bolt (processor). This characteristic restricts the kind of models that can be run in parallel to ensembles.

3 High Level Architecture

We identify three types of Apache SAMOA users:

1. Platform users, who use available ML algorithms without implementing new ones.
2. ML developers, who develop new ML algorithms on top of Apache SAMOA and want to be isolated from changes in the underlying SPEs.
3. Platform developers, who extend Apache SAMOA to integrate more DSPEs into Apache SAMOA.

There are three important design goals of Apache SAMOA:

1. **Flexibility** in term of developing new ML algorithms or reusing existing ML algorithms from other frameworks.
2. **Extensibility** in term of porting Apache SAMOA to new DSPEs.
3. **Scalability** in term of handling ever increasing amount of data.

Figure 1(right) shows the high-level architecture of Apache SAMOA which attempts to fulfill the aforementioned design goals. The *algorithm* layer contains existing distributed streaming algorithms that have been implemented in Apache SAMOA. This layer enables platform users to easily use the existing algorithm on any DSPE of their choice.

The *application programming interface* (API) layer consists of primitives and components that facilitate ML developers when implementing new algorithms. The *ML-adapter* layer allows ML developers to integrate existing algorithms in MOA or other ML frameworks into Apache SAMOA. The API layer and ML-adapter layer in Apache SAMOA fulfill the flexibility goal since they allow ML developers to rapidly develop algorithms.

Next, the *DSPE-adapter* layer supports platform developers in integrating new DSPEs into Apache SAMOA. To perform the integration, platform developers should implement the *samoa-SPE* layer as shown in Figure 1(right). Currently Apache SAMOA is equipped with four adapters: the *samoa-Storm* adapter for

¹¹ <http://github.com/vpa1977/stormmoa>

Storm, the *samoa-Samza* adapter for Samza, the *samoa-Flink* adapter for Flink, and the *samoa-Apex* adapter for Apex. To satisfy the extensibility goal, the DSPE-adapter layer decouples DSPEs and ML algorithms implementations in Apache SAMOA, so that platform developers are able to easily integrate more DSPEs.

The last goal, scalability, implies that Apache SAMOA should be able to scale to cope ever increasing amount of data. To fulfill this goal, Apache SAMOA utilizes modern DSPEs to execute its ML algorithms. The reason for using modern DSPEs such as Storm, Flink, Samza and Apex in Apache SAMOA is that they are designed to provide horizontal scalability to cope with Web-scale streams.

4 System Design

An algorithm in Apache SAMOA is represented by a directed graph of nodes that communicate via messages along streams which connect pairs of nodes. Borrowing the terminology from Storm, this graph is called a *Topology*. Each node in a Topology is a *Processor* that sends messages through a *Stream*. A Processor is a container for the code implementing the algorithm. A Stream can have a single source but several destinations (akin to a pub-sub system). A Topology is built by using a *TopologyBuilder*, which connects the various pieces of user code to the platform code and performs the necessary bookkeeping in the background. The following code snippet builds a topology that joins two data streams in Apache SAMOA:

```
TopologyBuilder builder = new TopologyBuilder();
Processor sourceOne = new SourceProcessor();
builder.addProcessor(sourceOne);
Stream streamOne = builder
    .createStream(sourceOne);

Processor sourceTwo = new SourceProcessor();
builder.addProcessor(sourceTwo);
Stream streamTwo = builder
    .createStream(sourceTwo);

Processor join = new JoinProcessor();
builder.addProcessor(join)
    .connectInputShuffle(streamOne)
    .connectInputKey(streamTwo);
```

A *Task* is an execution entity, similar to a job in Hadoop. A Topology is instantiated inside a Task to be run by Apache SAMOA. An example of a Task is *PrequentialEvaluation*, a classification task where each instance is used for testing first, and then for training.

A message or an event is called *Content Event* in Apache SAMOA. As the name suggests, it is an event which contains content that needs to be processed by the processors. Finally, a *Processing Item* is a hidden physical unit of the

topology and is just a wrapper of Processor. It is used internally, and it is not accessible from the API.

5 Machine Learning Algorithms

In Apache SAMOA there are currently three types of algorithms performing basic machine learning functionalities such as classification via a decision tree (VHT), clustering (CluStream) and regression rules (AMR).

The Vertical Hoeffding Tree (VHT) [20] is a distributed extension of the VFDT [13]. VHT uses vertical parallelism to split the workload across several machines. Vertical parallelism leverages the parallelism across attributes in the same example, rather than across different examples in the stream. In practice, each training example is routed through the tree model to a leaf. There, the example is split into its constituting attributes, and each attribute is sent to a different Processor instance that keeps track of sufficient statistics. This architecture has two main advantages over one based on horizontal parallelism. First, attribute counters are not replicated across several machines, thus, reducing the memory footprint. Second, the computation of the fitness of an attribute for a split decision (via, e.g., entropy or information gain) can be performed in parallel. The drawback is that in order to get good performances, there must be sufficient inherent parallelism in the data. That is, the VHT works best for high-dimensional data.

Apache SAMOA includes a distributed version of CluStream, an algorithm for clustering evolving data streams. CluStream keeps a small summary of the data received so far by computing micro-clusters online. These micro-clusters are further refined to create macro-clusters by a micro-batch process, which is triggered periodically. The period is configured via a command line parameter (e.g., every 10 000 examples).

For regression, Apache SAMOA provides a distributed implementation of Adaptive Model Rules [23]. The algorithm, HAMR, uses a hybrid of vertical and horizontal parallelism to distribute AMRules on a cluster.

Apache SAMOA also includes adaptive implementations of ensemble methods such as bagging and boosting. These methods include state-of-the-art change detectors such as as ADWIN, DDM, EDDM, and Page-Hinckley [15]. These meta-algorithms are most useful in conjunction with external single-machine classifiers, which can be plugged in Apache SAMOA. For instance, open-source connectors for MOA [4] are provided separately by the Apache SAMOA-MOA package.¹²

The following listing shows how to download, build, and run Apache SAMOA.

```
# download and build SAMOA
git clone http://git.apache.org/incubator-samoa.git
cd incubator-samoa
mvn package
```

¹² <https://github.com/samoa-moa/samoa-moa>

```

# download the Forest Cover Type dataset
wget "http://downloads.sourceforge.net/project/moa-datastream
/Datasets/Classification/covtypeNorm.arff.zip"
unzip "covtypeNorm.arff.zip"

# run SAMOA in local mode
bin/samoa local target/SAMOA-Local-0.4.0-SNAPSHOT.jar "
  PrequentialEvaluation -l classifiers.ensemble.Bagging -s
  (ArffFileStream -f covtypeNorm.arff) -f 100000"

```

6 Vertical Hoeffding Tree

We explain the details of the *Vertical Hoeffding Tree* [20], which is a data-parallel, distributed version of the Hoeffding tree. The *Hoeffding tree* [13] (a.k.a. VFDT) is a streaming decision tree learner with statistical guarantees. In particular, by leveraging the Chernoff-Hoeffding bound [16], it guarantees that the learned model is asymptotically close to the model learned by the batch greedy heuristic, under mild assumptions. The learning algorithm is very simple. Each leaf keeps track of the statistics for the portion of the stream it is reached by, and computes the best two attributes according to the splitting criterion. Let ΔG be the difference between the value of the functions that represent the splitting criterion of these two attributes. Let ϵ be a quantity that depends on a user-defined confidence parameter δ , and that decreases with the number of instances processed. When $\Delta G > \epsilon$, then the current best attribute is selected to split the leaf. The Hoeffding bound guarantees that this choice is the correct one with probability at least $1 - \delta$.

In this section, first, we describe the parallelization and the ideas behind our design choice. Then, we present the engineering details and optimizations we employed to obtain the best performance.

6.1 Vertical Parallelism

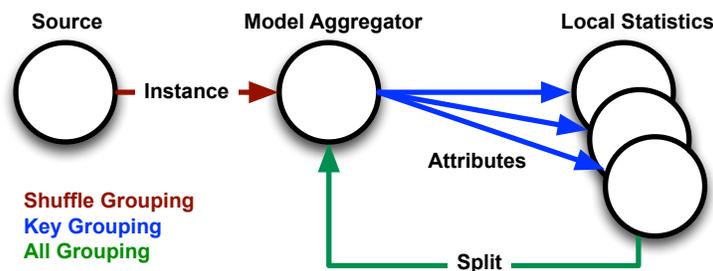


Fig. 2: High level diagram of the VHT topology.

Data parallelism is a way of distributing work across different nodes in a parallel computing environment such as a cluster. In this setting, each node executes the same operation on different parts of the dataset. Contrast this definition with task parallelism (aka pipelined parallelism), where each node executes a different operator, and the whole dataset flows through each node at different stages. When applicable, data parallelism is able to scale to much larger deployments, for two reasons: (i) data has usually much higher intrinsic parallelism that can be leveraged compared to tasks, and (ii) it is easier to balance the load of a data-parallel application compared to a task-parallel one. These attributes have led to the high popularity of the currently available DSPEs. For these reasons, we employ data parallelism in the design of VHT.

In machine learning, it is common to think about data in matrix form. A typical linear classification formulation requires to find a vector x such that $A \cdot x \approx b$, where A is the data matrix and b is a class label vector. The matrix A is $n \times m$ -dimensional, with n being the number of data instances and m being the number of attributes of the dataset.

There are two ways to *slice* the data matrix to obtain data parallelism: by row or column. The former is called *horizontal parallelism*, the latter *vertical parallelism*. With horizontal parallelism, data instances are independent from each other, and can be processed in isolation while considering all available attributes. With vertical parallelism, instead, attributes are considered independently from each other.

The fundamental operation of the algorithm is to accumulate statistics n_{ijk} (i.e., counters) for triplets of attribute i , value j , and class k , for each leaf l of the tree. The counters for each leaf are independent, so let us consider the case for a single leaf. These counters, together with the learned tree structure, constitute the state of the VHT algorithm.

Different kinds of parallelism distribute the counters across computing nodes in different ways. With horizontal parallelism [3], the instances are distributed randomly, thus multiple instances of the same counter can exist on several nodes. On the other hand, when using vertical parallelism, the counters for one attribute are grouped on a single node.

This latter design has several advantages. First, by having a single copy of the counter, the memory requirements for the model are the same as in the sequential version. In contrast, with horizontal parallelism a single attribute may be tracked on every node, thus the memory requirements grow linearly with the parallelism level. Second, by having each attribute tracked independently, the computation of the split criterion can be performed in parallel by several nodes. Conversely, with horizontal partitioning the algorithm needs to (centrally) aggregate the partial counters before being able to compute the splitting criterion.

Of course, the vertically-parallel design has also its drawbacks. In particular, horizontal parallelism achieves a good load balance more easily, even though solutions for these problems have recently been proposed for vertical parallelism as well [24, 25]. In addition, if the instance stream arrives in row-format, it needs to be transformed in column-format, and this transformation generates

additional CPU overhead at the source. Indeed, each attribute that constitutes an instance needs to be sent independently, and needs to carry the class label of its instance. Therefore, both the number of messages and the size of the data transferred increase. Nevertheless, the advantages of vertical parallelism outweigh its disadvantages for several real-world settings.

6.2 Algorithm Structure

Table 1: Definitions and symbols.

Description	Symbol
Source receiving/distributing instances	S
Training instance from S	E
Model Aggregator	MA
Current state of the decision tree in MA	VHT_tree
Local Statistic	LS
Counter for attribute i , value j , class k	n_{ijk}
Instances seen at leaf l	n_l
Information gain or entropy of attribute i in leaf l	$\overline{G}_l(X_i^{local})$

We are now ready to explain the structure of the VHT algorithm. In general, there are two main parts to the Hoeffding tree algorithm: *sorting* the instances through the current model, and accumulating *statistics* of the stream at each leaf node. This separation offers a neat cut point to modularize the algorithm in two separate components. We call the first component *model aggregator*, and the second component *local statistics*. Figure 2 presents a visual depiction of the algorithm, and specifically, of its components and how the data flow among them. Also, Table 1 summarizes a list of components used in the rest of the algorithm description.

The model aggregator holds the current model (the tree) produced so far in a Processor node. Its main duty is to receive the incoming instances and sort them to the correct leaf. If the instance is unlabeled, the model predicts the label at the leaf and sends it downstream (e.g., for evaluation). Otherwise, if the instance is labeled it is also used as training data. The VHT decomposes the instance into its constituent attributes, attaches the class label to each, and sends them independently to the following stage, the *local statistics*. Algorithm 1 shows a pseudocode for the model aggregator.

The local statistics contain the counters n_{ijk} for a set of attribute-value-class triplets. Conceptually, the local statistics can be viewed as a large distributed table, indexed by leaf id (row), and attribute id (column). The value of the cell represents a set of counters, one for each pair of attribute value and class. The local statistics accumulate statistics on the data sent by the model aggregator. Pseudocode for the update function is shown in Algorithm 2.

Algorithm 1 *MA*: VerticalHoeffdingTreeInduction(E, VHT_tree)

Require: E wrapped in **instance** content event

Require: VHT_tree in *MA*

- 1: Use VHT_tree to sort E into a leaf l
 - 2: Send **attribute** content events to *LSs*
 - 3: Increment n_l
 - 4: **if** $n_l \bmod n_{min} = 0$ **and** not all instances seen at l belong to the same class **then**
 - 5: Add l into the list of splitting leaves
 - 6: Send **compute** content event with the id of leaf l to all *LSs*
 - 7: **end if**
-

Algorithm 2 *LS*: UpdateLocalStatistic($attribute, local_statistic$)

Require: $attribute$ is an **attribute** content event

Require: $local_statistic$ is the *LS* in charge of $attribute$, could be implemented as $Table \langle leaf_id, attribute_id \rangle$

- 1: Update $local_statistic$ with data in $attribute$: attribute value, class value and instance weights
-

In Apache SAMOA, we implement vertical parallelism by connecting the model to the statistics via key grouping. We use a composite key made by the leaf id and the attribute id. Horizontal parallelism can similarly be implemented via shuffle grouping on the instances themselves.

Leaf splitting. Periodically, the model aggregator tries to see if the model needs to evolve by splitting a leaf. When a sufficient number of instances n_{min} have been sorted through a leaf, and not all instances that reached l belong to the same class (line 4, Algorithm 1), the aggregator sends a broadcast message to the statistics, asking to compute the split criterion for the given leaf id. The statistics Processor gets the table corresponding to the leaf, and for each attribute compute the splitting criterion in parallel (an information-theoretic function such as information gain or entropy). Each local statistic then sends back to the model the top two attributes according to the chosen criterion, together with their scores ($\overline{G}_l(X_i^{local}), i = a, b$; Algorithm 3).

Subsequently, the model aggregator (Algorithm 4) simply needs to compute the overall top two attributes received so far from the available statistics, apply the Hoeffding bound (line 4), and see whether the leaf needs to be split (line 5). The algorithm also computes the criterion for the scenario where no split takes places (X_\emptyset). Domingos and Hulten [13] refer to this inclusion of a no-split scenario with the term *pre-pruning*. The decision to split or not is taken after a time has elapsed, as explained next.

By using the top two attributes, the model aggregator computes the difference of their splitting criterion values $\Delta\overline{G}_l = \overline{G}_l(X_a) - \overline{G}_l(X_b)$. To determine whether the leaf needs to be split, it compares the difference $\Delta\overline{G}_l$ to the Hoeffding bound $\epsilon = \sqrt{\frac{R^2 \ln(1/\delta)}{2n_l}}$ for the current confidence parameter δ (where R is the range of possible values of the criterion). If the difference is larger than the bound ($\Delta\overline{G}_l >$

Algorithm 3 *LS*: ReceiveComputeMessage(*compute*, *local_statistic*)

Require: *compute* is a **compute** content event

Require: *local_statistic* is the *LS* in charge of *attribute*, could be implemented as

Table \langle *leaf_id*, *attribute_id* \rangle

- 1: Get leaf l ID from **compute** content event
 - 2: For each attribute i that belongs to leaf l in local statistic, compute $\bar{G}_l(X_i)$
 - 3: Find X_a^{local} , i.e., the attribute with highest \bar{G}_l based on the local statistic
 - 4: Find X_b^{local} , i.e., the attribute with second highest \bar{G}_l based on the local statistic
 - 5: Send X_a^{local} and X_b^{local} using **local-result** content event to model-aggregator PI via **computation-result** stream
-

ϵ), then X_a is the best attribute with high confidence $1 - \delta$, and can therefore be used to split the leaf. If the best attribute is the no-split scenario (X_\emptyset), the algorithm does not perform any split. The algorithm also uses a tie-breaking τ mechanism to handle the case where the difference in splitting criterion between X_a and X_b is very small. If the Hoeffding bound becomes smaller than τ ($\Delta\bar{G}_l < \epsilon < \tau$), then the current best attribute is chosen regardless of the values of $\Delta\bar{G}_l$.

Two cases can arise: the leaf needs splitting, or it does not. In the latter case, the algorithm simply continues without taking any action. In the former case instead, the model modifies the tree by splitting the leaf l on the selected attribute, replacing l with an internal node (line 6), and generating a new leaf for each possible value of the branch (these leaves are initialized using the class distribution observed at the best attribute splitting at l (line 8)). Then, it broadcasts a **drop** message containing the former leaf id to the local statistics (line 10). This message is needed to release the resources held by the leaf and make space for the newly created leaves. Subsequently, the tree can resume sorting instances to the new leaves. The local statistics creates a new table for the new leaves lazily, whenever they first receive a previously unseen leaf id. In its simplest version, while the tree adjustment is performed, the algorithm drops the new incoming instances. We show in the next section an optimized version that buffers them to improve accuracy.

Messages. During the VHT execution several types of events are sent and received from the different parts of the algorithm, as summarized in Table 2.

Table 2: Types of content events used during the execution of VHT algorithm.

Name	Parameters	From	To
instance	\langle attr 1, ..., attr m, class C \rangle	S	MA
attribute	\langle attr <i>id</i> , attr value, class C \rangle	MA	LS $id = \langle$ leaf <i>id</i> + attr <i>id</i> \rangle
compute	\langle leaf <i>id</i> \rangle	MA	All LS
local-result	\langle $\bar{G}_l(X_a^{local}), \bar{G}_l(X_b^{local})$ \rangle	LS_{id}	MA
drop	\langle leaf <i>id</i> \rangle	MA	All LS

Algorithm 4 *MA*: Receive(*local_result*, *VHT_tree*)

Require: *local_result* is an **local-result** content event**Require:** *VHT_tree* in *MA*

- 1: Get correct leaf l from the list of splitting leaves
 - 2: Update X_a and X_b in the splitting leaf l with X_a^{local} and X_b^{local} from *local_result*
 - 3: **if** *local_results* from all *LSs* received or time out reached **then**
 - 4: Compute Hoeffding bound $\epsilon = \sqrt{\frac{R^2 \ln(1/\delta)}{2n_l}}$
 - 5: **if** $X_a \neq X_\emptyset$ **and** $(\bar{G}_l(X_a) - \bar{G}_l(X_b) > \epsilon$ **or** $\epsilon < \tau)$ **then**
 - 6: Replace l with a split-node on X_a
 - 7: **for all** branches of the split **do**
 - 8: Add new leaf with derived sufficient statistic from split node
 - 9: **end for**
 - 10: Send **drop** content event with id of leaf l to all *LSs*
 - 11: **end if**
 - 12: **end if**
-

6.3 Evaluation

In our experimental evaluation of the VHT method, we aim to study the following questions:

- Q1:** How does a centralized VHT compare to a centralized hoeffding tree with respect to accuracy and throughput?
- Q2:** How does the vertical parallelism used by VHT compare to horizontal parallelism?
- Q3:** What is the effect of number and density of attributes?
- Q4:** How does discarding or buffering instances affect the performance of VHT?

Experimental setup. In order to study these questions, we experiment with five datasets (two synthetic generators and three real datasets), five different versions of the hoeffding tree algorithm, and up to four levels of computing parallelism. We measure classification accuracy during, and at the end of the execution, and throughput (number of classified instances per second). We execute each experimental configuration ten times and report the average of these measures.

Synthetic datasets. We use synthetic data streams produced by two random generators: one for dense and one for sparse attributes.

- **Dense attributes** are extracted from a random decision tree. We test different number of attributes, and include both categorical and numerical types. The label for each configuration is the number of categorical-numerical used (e.g, 100-100 means the configuration has 100 categorical and 100 numerical attributes). We produce 10 differently seeded streams with 1M instances for each tree, with one of two balanced classes in each instance, and take measurements every 100k instances.

- **Sparse attributes** are extracted from a random tweet generator. We test different dimensionalities for the attribute space: 100, 1k, 10k. These attributes represent the appearance of words from a predefined bag-of-words. On average, the generator produces 15 words per tweet (size of a tweet is Gaussian), and uses a Zipf distribution with skew $z = 1.5$ to select words from the bag. We produce 10 differently seeded streams with 1M tweets in each stream. Each tweet has a binary class chosen uniformly at random, which conditions the Zipf distribution used to generate the words.

Real datasets. We also test VHT on three real data streams to assess its performance on benchmark data.¹³

- (*elec*) Electricity: 45312 instances, 8 numerical attributes and 2 classes.
- (*phy*) Particle Physics: 50000 instances, 78 numerical attributes and 2 classes.
- (*covtype*) CovtypeNorm: 581012 instances, 54 numerical attributes and 7 classes.

Algorithms. We compare the following versions of the hoeffding tree algorithm.

- **moa**: This is the standard Hoeffding tree in MOA.
- **local**: This algorithm executes VHT in a local, sequential execution engine. All split decisions are made in a sequential manner in the same process, with no communication and feedback delays between statistics and model.
- **wok**: This algorithm discards instances that arrive during a split decision. This version is the vanilla VHT.
- **wk(z)**: This algorithm sends instances that arrive during a split decision downstream. It also adds instances to a buffer of size z until full. If the split decision is taken, it replays the instances in the buffer through the new tree model. Otherwise, it discards the buffer, as the instances have already been incorporated in the statistics downstream.
- **sharding**: Splits the incoming stream horizontally among an ensemble of Hoeffding trees. The final prediction is computed by majority voting. This method is an instance of horizontal parallelism applied to Hoeffding trees. It creates an ensemble of hoeffding trees, but each tree is built with a subset of instances split horizontally, while using all available attributes.

Experimental configuration. All experiments are performed on a Linux server with 24 cores (Intel Xeon X5650), clocked at 2.67GHz, L1d cache: 32kB, L1i cache: 32kB, L2 cache: 256kB, L3 cache: 12288kB, and 65GB of main memory. On this server, we run a Storm cluster (v0.9.3) and zookeeper (v3.4.6). We use Apache SAMOA v0.4.0 (development version) and MOA v2016.04 available from the respective project websites.

We use several parallelism levels in the range of $p = 2, \dots, 16$, depending on the experimental configuration. For dense instances, we stop at $p = 8$ due to memory constraints, while for sparse instances we scale up to $p = 16$. We disable

¹³ <http://moa.cms.waikato.ac.nz/datasets/>,
<http://osmot.cs.cornell.edu/kddcup/datasets.html>

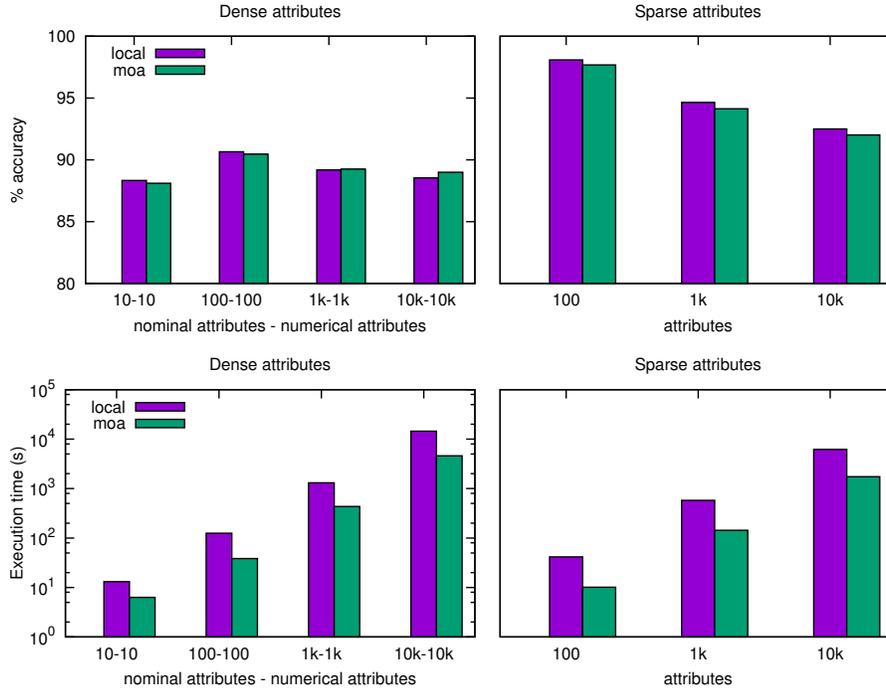


Fig. 3: Accuracy and execution time of VHT executed in local mode on Apache SAMOA compared to MOA, for dense and sparse datasets.

model replication (i.e., use a single model aggregator), as in our setup the model is not the bottleneck.

Accuracy and time of VHT local vs. moa. In this first set of experiments, we test if VHT is performing as well as its counterpart hoeffding tree in MOA. This is mostly a sanity check to confirm that the algorithm used to build the VHT does not affect the performance of the tree when all instances are processed sequentially by the model. To verify this fact, we execute VHT local and MOA with both dense and sparse instances. Figure 3 shows that VHT local achieves the same accuracy as MOA, even besting it at times. However, VHT local always takes longer than MOA to execute. Indeed, the local execution engine of Apache SAMOA is optimized for simplicity rather than speed. Therefore, the additional overhead required to interface VHT to DSPEs is not amortized by scaling the algorithm out. Future optimized versions of VHT and the local execution engine should be able to close this gap.

Accuracy of VHT local vs. distributed. Next, we compare the performance of VHT local with VHT built in a distributed fashion over multiple processors

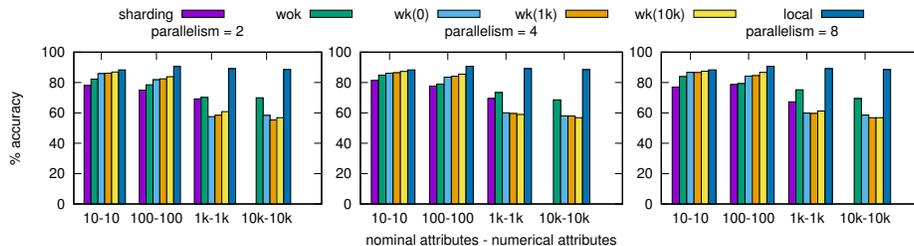


Fig. 4: Accuracy of several versions of VHT (local, **wok**, **wk(z)**) and sharding, for dense datasets.

for scalability. We use up to $p = 8$ parallel statistics, due to memory restrictions, as our setup runs on a single machine. In this set of experiments we compare the different versions of VHT, **wok** and **wk(z)**, to understand what is the impact of keeping instances for training after a model’s split. Accuracy of the model might be affected, compared to the local execution, due to delays in the feedback loop between statistics and model. That is, instances arriving during a split are classified using an older version of the model compared to the sequential execution. As our target is a distributed system where independent processes run without coordination, this delay is a characteristic of the algorithm as much as of the distributed SPE we employ.

We expect that buffering instances and replaying them when a split is decided would improve the accuracy of the model. In fact, this is the case for dense instances with a small number of attributes (i.e., around 200), as shown in Figure 4. However, when the number of available attributes increases significantly, the load imposed on the model seems to outweigh the benefits of keeping the instances for replaying. We conjecture that the increased load in computing the splitting criterion in the statistics further delays the feedback to compute the split. Therefore, a larger number of instances are classified with an older model, thus, negatively affecting the accuracy of the tree. In this case, the additional load imposed by replaying the buffer further delays the split decision. For this reason, the accuracy for VHT **wk(z)** drops by about 30% compared to VHT local. Conversely, the accuracy of VHT **wok** drops more gracefully, and is always within 18% of the local version.

VHT always performs approximatively 10% better than sharding. For dense instances with a large number of attributes (20k), sharding fails to complete due to its memory requirements exceeding the available memory. Indeed, sharding builds a full model for each shard, on a subset of the stream. Therefore, its memory requirements are p times higher than a standard hoeffding tree.

When using sparse instances, the number of attributes per instance is constant, while the dimensionality of the attribute space increases. In this scenario, increasing the number of attributes does not put additional load on the system. Indeed, Figure 5 shows that the accuracy of all versions is quite similar, and

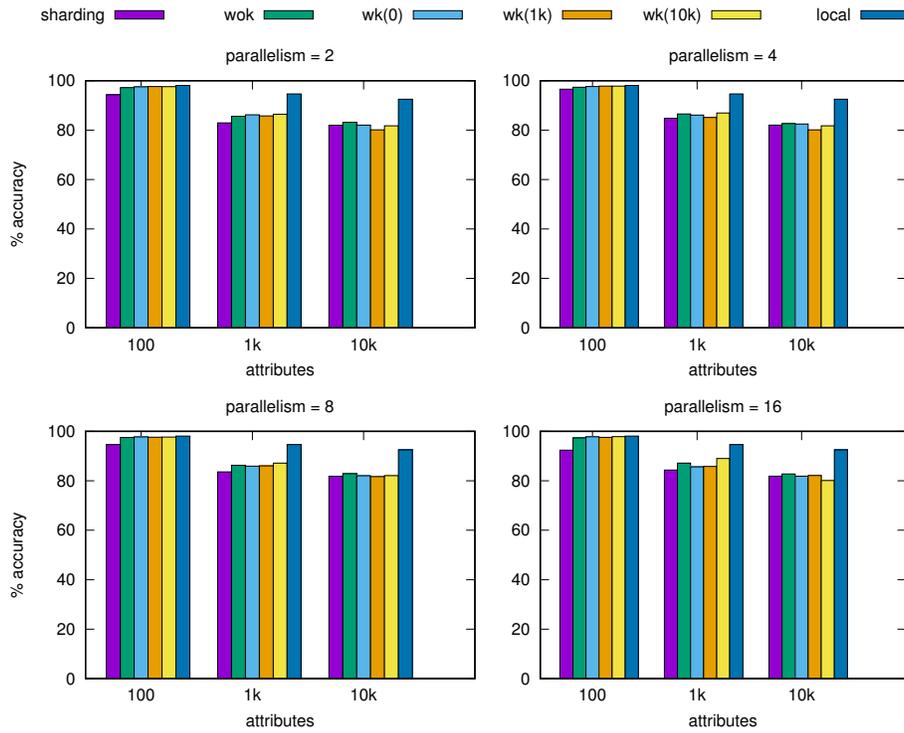


Fig. 5: Accuracy of several versions of VHT (local, **wok**, **wk(z)**) and sharding, for sparse datasets.

close to the local one. This observation is in line with our conjecture that the overload on the system is the cause for the drop in accuracy on dense instances.

We also study how the accuracy evolves over time. In general, the accuracy of all algorithms is rather stable, as shown in Figures 6 and 7. For instances with 10 to 100 attributes, all algorithms perform similarly. For dense instances, the versions of VHT with buffering, $\mathbf{wk}(\mathbf{z})$, outperform \mathbf{wok} , which in turn outperforms sharding. This result confirms that buffering is beneficial for small number of attributes. When the number of attributes increases to a few thousand per instance, the performance of these more elaborate algorithms drops considerably. However, the VHT \mathbf{wok} continues to perform relatively well and better than sharding. This performance, coupled with good speedup over MOA (as shown next) makes it a viable option for streams with a large number of attributes and a large number of instances.

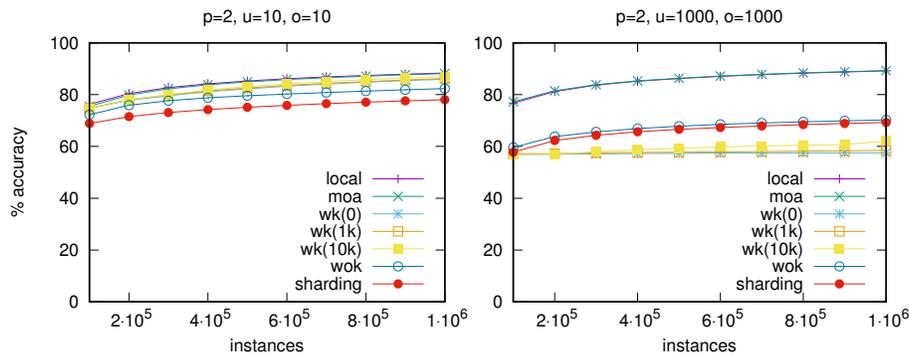


Fig. 6: Evolution of accuracy with respect to instances arriving, for several versions of VHT (local, \mathbf{wok} , $\mathbf{wk}(\mathbf{z})$) and sharding, for dense datasets.

Speedup of VHT distributed vs. moa. Since the accuracy of VHT $\mathbf{wk}(\mathbf{z})$ is not satisfactory for both types of instances, next we focus our investigation on VHT \mathbf{wok} . Figure 8 shows the speedup of VHT for dense instances. VHT \mathbf{wok} is about 2-10 times faster than VHT local and up to 4 times faster than MOA. Clearly, the algorithm achieves a higher speedup when more attributes are present in each instance, as (i) there is more opportunity for parallelization, and (ii) the implicit load shedding caused by discarding instances during splits has a larger effect. Even though sharding performs well in speedup with respect to MOA on small number of attributes, it fails to build a model for large number of attributes due to running out of memory. In addition, even for a small number of attributes, VHT \mathbf{wok} outperforms sharding with a parallelism of 8. Thus, it is clear from the results that the vertical parallelism used by VHT offers better scaling behavior than the horizontal parallelism used by sharding.

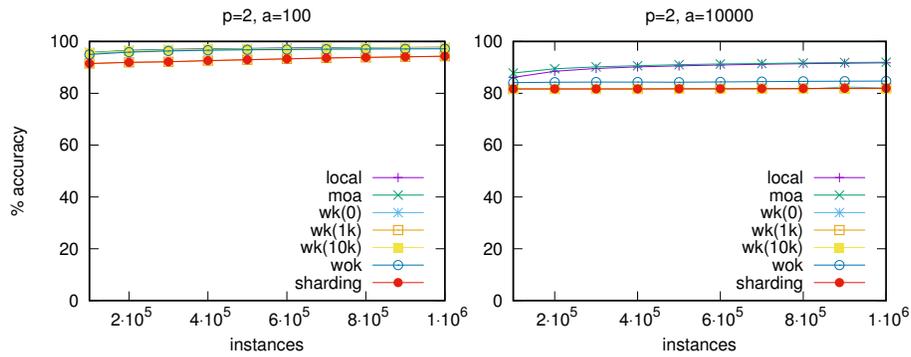


Fig. 7: Evolution of accuracy with respect to instances arriving, for several versions of VHT (local, **wok**, **wk(z)**) and sharding, for sparse datasets.

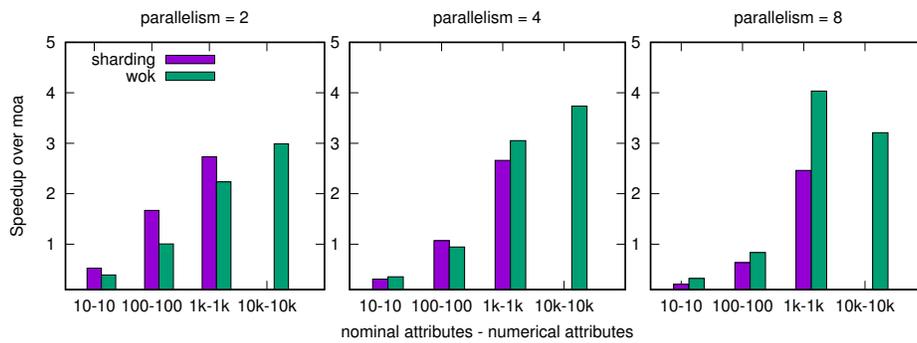


Fig. 8: Speedup of VHT **wok** executed on Apache SAMOA compared to MOA for dense datasets.

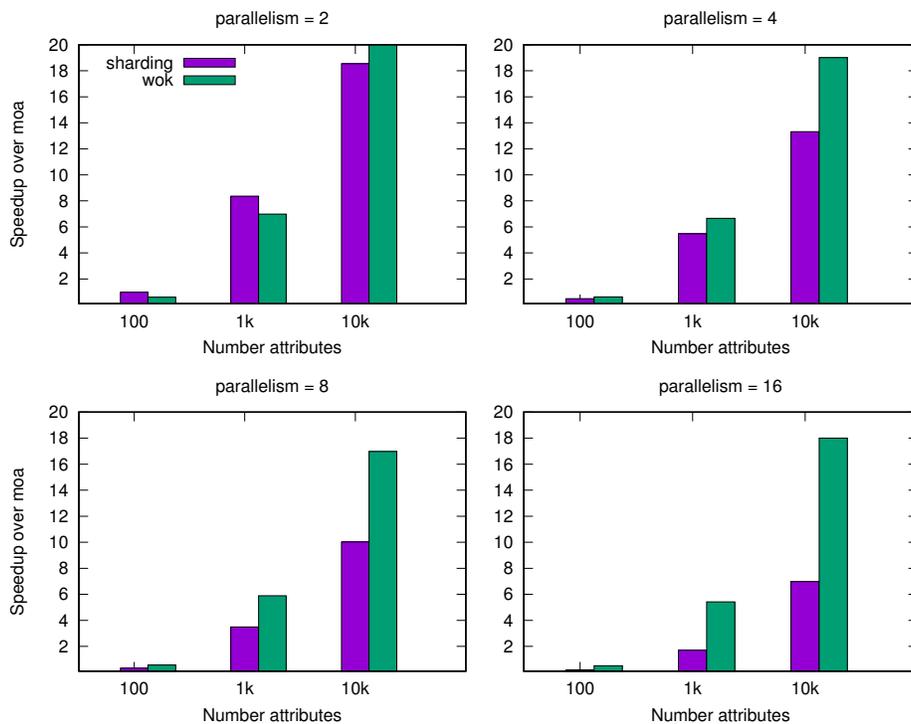


Fig. 9: Speedup of VHT **wok** executed on Apache SAMOA compared to MOA for sparse datasets.

When testing the algorithms on sparse instances, as shown in Figure 9, we notice that VHT **wok** can reach up to 60 times the throughput of VHT local and 20 times the one of MOA (for clarity we only show the results with respect to MOA). Similarly to what observed for dense instances, a higher speedup is observed when a larger number of attributes are present for the model to process. This very large, superlinear speedup ($20\times$ with $p = 2$), is due to the aggressive load shedding implicit in the **wok** version of VHT. The algorithm actually performs consistently less work than the local version and MOA.

However, note that for sparse instances the algorithm processes a constant number of attributes, albeit from an increasingly larger space. Therefore, in this setup, **wok** has a constant overhead for processing each sparse instance, differently from the dense case. VHT **wok** outperforms sharding in most scenarios and especially for larger numbers of attributes and larger parallelism.

Increased parallelism does not impact accuracy of the model (see Figure 4 and Figure 5), but its throughput is improved. Boosting the parallelism from 2 to 4 makes VHT **wok** up to 2 times faster. However, adding more processors does not improve speedup, and in some cases there is a slowdown due to additional communication overhead (for dense instances). Particularly for sparse instances,

parallelism does not impact accuracy which enables handling large sparse data streams while achieving high speedup over MOA.

Performance on real-world datasets. Tables 3 and 4 show the performance of VHT, either running in a local mode or in a distributed fashion over a storm cluster of a few processors. We also test two different versions of VHT: **wok** and **wk(0)**. In the same tables we compare VHT’s performance with MOA and sharding.

The results from these real datasets demonstrate that VHT can perform similarly to MOA with respect to accuracy and at the same time process the instances faster. In fact, for the larger dataset, *covtypeNorm*, VHT **wok** exhibits 1.8 speedup with respect to MOA, even though the number of attributes is not very large (54 numeric attributes). VHT **wok** also performs better than sharding, even though the latter is faster in some cases. However, the speedup offered by sharding decreases when the parallelism level is increased from 2 to 4 shards.

Table 3: Average accuracy (%) for different algorithms, with parallelism level (p), on the real-world datasets.

dataset	MOA		VHT				Sharding	
	local	wok p=2	wok p=4	wk(0) p=2	wk(0) p=4	p=2	p=4	
elec	75.4	75.4	75.0	75.2	75.4	75.6	74.7	74.3
phy	63.3	63.8	62.6	62.7	63.8	63.7	62.4	61.4
covtype	67.9	68.4	68.0	68.8	67.5	68.0	67.9	60.0

Table 4: Average execution time (seconds) for different algorithms, with parallelism level (p), on the real-world datasets.

Dataset	MOA		VHT				Sharding	
	local	wok p=2	wok p=4	wk(0) p=2	wk(0) p=4	p=2	p=4	
elec	1.09	1	2	2	2	2	2.33	
phy	5.41	4	3.25	4	3	3.75	4	
covtype	21.77	16	12	12	13	12	11	

6.4 Summary

In conclusion, our VHT algorithm has the following performance traits. We learned that for a small number of attributes, it helps to buffer incoming instances that can be used in future decisions of split. For larger number of attributes, the load in the model can be high and larger delays can be observed in the integration of the feedback from the local statistics into the model. In this

case, buffered instances may not be used on the most up-to-date model and this can penalize the overall accuracy of the model.

With respect to a centralized sequential tree model (MOA), it processes dense instances with thousands of attributes up to $4\times$ faster with only 10–20% drop in accuracy. It can also process sparse instances with thousands of attributes up to $20\times$ faster with only 5–10% drop in accuracy. Also, its ability to build the tree in a distributed fashion using tens of processors allows it to scale and accommodate thousands of attributes and parse millions of instances. Competing methods cannot handle these data sizes due to increased memory and computational complexity.

7 Distributed AMRules

Decision rule learning is a category of machine learning algorithms whose goal is to extract a set of decision rules from the training data. These rules are later used to predict the unknown *label* values for test data. A rule is a logic expression of the form:

IF *antecedent* **THEN** *consequent*

or, equivalently, $head \leftarrow body$, where *head* and *body* correspond to the *consequent* and *antecedent*, respectively.

The body of a rule is a conjunction of multiple clauses called *features*, each of which is a condition on an attribute of the instances. Such conditions consist of the identity of an attribute, a threshold value and an operator. For instance, the feature “ $x < 5$ ” is a condition on attribute x , with threshold value 5 and operator *less-than* ($<$). An instance is said to be *covered* by a rule if its attribute values satisfy all the features in the rule body. The head of the rule is a function to be applied on the covered instances to determine their label values. This function can be a constant or a function of the attributes of the instances, e.g., $ax + b \leftarrow x < 5$.

AMRules is an algorithm for learning regression rules on streaming data. It incrementally constructs the rule model from the incoming data stream. The rule model consists of a set of *normal* rules (which is empty at the beginning), and a *default* rule. Each normal rule is composed of 3 parts: a *body* which is a list of features, a *head* with information to compute the prediction for those instance covered by the rule, and *statistics* of past instances to decide when and how to add a new feature to its body. The default rule is a rule with an empty *body*.

For each incoming instance, AMRules searches the current rule set for those rules that cover the instance. If an instance is not covered by any rule in the set, it is considered as being covered by the default rule. The heads of the rules are first used to compute the prediction for the instance they cover. Later, their statistics are updated with the attribute values and label value of the instance. There are two possible modes of operation: ordered and unordered. In ordered-rules mode, the rules are checked according to the order of their creation, and only the first rule is used for prediction and then updated. In unordered-rules

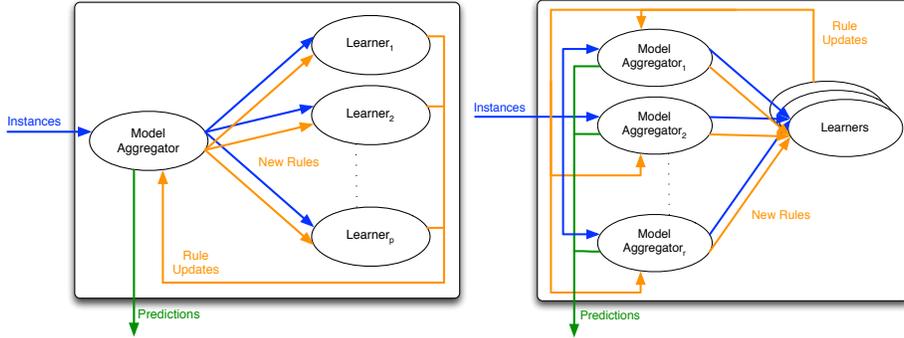


Fig. 10: (Left) Vertical AMRules (VAMR). (Right) AMRules with multiple horizontally parallelized Model Aggregators.

mode, all covering rules are used and updated. In this work, we focus on the former which is more often used albeit more challenging to parallelize.

Each rule tries to expand its body after it receives N_m updates. In order to decide on the feature to expand, the rule incrementally computes a standard deviation reduction (SDR) measure [17] for each potential feature. Then, it computes the *ratio* of the second-largest SDR value over the largest SDR value. This *ratio* is used with a high confidence interval ϵ computed using the Hoeffding bound [16] to decide to expand the rule or not: if $ratio + \epsilon < 1$, the rule is expanded with the feature corresponding to the largest SDR value. Besides, to avoid missing a good feature when there are two (or more) equally good ones, rules are also expanded if the Hoeffding bound ϵ falls below a threshold. If the default rule is expanded, it becomes a *normal* rule and is added to the rule set. A new default rule is initialized to replace the previous one.

Each rule records its prediction error and applies a modified version of the Page-Hinkley test [22] for streaming data to detect changes. If the test indicates that the cumulative error has exceeded a threshold, the rule is evicted from the rule set. The algorithm also employ outlier detection to check if an instance, although being covered by a rule, is an anomaly. If an instance is deemed as an anomaly, it is treated as if the rule does not cover it and is checked against other rules. The following sections describe two possible strategies to parallelize AMRules that are implemented in Apache SAMOA.

7.1 Vertical Parallelism

In AMRules, each rule can evolve independently, as its expansion is based solely on the statistics of instances it covers. Also, searching for the best feature among all possible ones in an attempt to expand a rule is computationally expensive.

Given these observations, we decide to parallelize AMRules by delegating the training process of rules to multiple *learner processors*, each of which handles

only a subset of the rules. Besides the learners, a *model aggregator processor* is also required to filter and redirect the incoming instances to the correct learners. The aggregator manages a set of simplified rules that have only head and body, i.e., do not keep statistics. The bodies are used to identify the rules that cover an instance, while the heads are used to compute the prediction. Each instance is forwarded to the designated learners by using the ID of the covering rule. At the learners, the corresponding rules' statistics are updated with the forwarded instance. This parallelization scheme guarantees that the rules created are the same as in the sequential algorithm. Figure 10(left) depicts the design of this vertically parallelized version of AMRules, or Vertical AMRules (VAMR for brevity).

The model aggregator also manages the statistics of the default rule, and updates it with instances not being covered by any other rule in the set. When the default rule is expanded and adds a new rule to the set, the model aggregator sends a message with the newly added rule to one of the learners, which is responsible for its management. The assignment of a rule to a learner is done based on the rule's ID via key grouping. All subsequent instances that are covered by this rule are forwarded to the same learner.

At the same time, learners update the statistics of each corresponding rule with each processed instance. When enough statistics have been accumulated and a rule is expanded, the new feature is sent to the model aggregator to update the body of the rule. Learners can also detect changes and remove existing rules. In such an event, learners inform the model aggregator with a message containing the removed rule ID.

As each rule is replicated in the model aggregator and in one of the learners, their bodies in model aggregator might not be up-to-date. The delay between rule expansion in the learner and model update in the aggregator depends mainly on the queue length at the model aggregator. The queue length, in turn, is proportional to the volume and speed of the incoming data stream. Therefore, instances that are in the queue before the model update event might be forwarded to a recently expanded rule which no longer covers the instance.

Coverage test is performed again at the learner, thus the instance is dropped if it was incorrectly forwarded. Given this additional test, and given that rule expansion can only increase the selectivity of a rule, when using unordered rules the accuracy of the algorithm is unaltered. However, in ordered-rules mode, these temporary inconsistencies might affect the statistics of other rules because the instance should have been forwarded to a different rule.

7.2 Horizontal Parallelism

A bottleneck in VAMR is the centralized model aggregator. Given that there is no straightforward way to vertically parallelize the execution of the model aggregator while maintaining the order of the rules, we explore an alternative based on horizontal parallelism. Specifically, we introduce multiple replicas of the model aggregator, so that each replica maintains the same copy of the rule set but processes only a portion of the incoming instances.

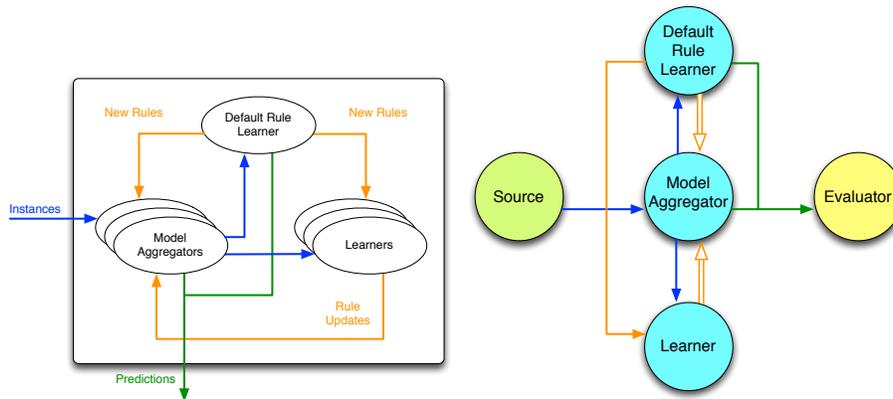


Fig. 11: (Left) Hybrid AMRules (HAMR) with multiple Model Aggregators and separate Default Rule Learner. (Right) Prequential evaluation task for HAMR. Single lines represent single messages (*key grouping, shuffle grouping*) while double lines represent broadcast messages (*all grouping*).

Horizontally parallelized model aggregator. The design of this scheme is illustrated in Figure 10(right). The basic idea is to extend VAMR and accommodate multiple model aggregators into the design. Each model aggregator still has a rule set and a default rule. The behavior of this scheme is similar to VAMR, except that each model aggregator now processes only a portion of the input data, i.e., the amount of instances each of them receives is inversely proportional to the number of model aggregators. This affects the prediction statistics and, most importantly, the training statistics of the default rules.

Since each model aggregator processes only a portion of the input stream, each default rule is trained independently with different portions of the stream. Thus, these default rules evolve independently and potentially create overlapping or conflicting rules. This fact also introduces the need for a scheme to synchronize and order the rules created by different model aggregators. Additionally, at the beginning, the scheme is less reactive compared to VAMR as it requires more instances for the default rules to start expanding. Besides, as the prediction function of each rule is adaptively constructed based on attribute values and label values of past instances, having only a portion of the data stream leads to having less information and potentially lower accuracy. We show how to address these issues next.

Centralized rule creation. In order to address the issues with distributed creation of rules, we move the default rule in model aggregators to a specialized default rule learner processor. With the introduction of this new component, some modifications are required in the model aggregators, but the behavior of the learners is still the same as in VAMR. However, as a result, all the model aggregators are in synch.

As the default rule is now moved to the designated learner, those instances that are not covered by any rules are forwarded from the model aggregators to this learner. This specialized learner updates its statistics with the received instances and, when the default rule expands, it broadcasts the newly created rule to the model aggregators. The new rule is also sent to the assigned learner, as determined by the rule’s ID.

The components of this scheme are shown in Figure 11(left), where this scheme is referred as Hybrid AMRules (HAMR), as it is a combination of vertical and horizontal parallelism strategies.

7.3 Evaluation

We evaluate the performance of the 2 distributed implementations of AMRules, i.e., VAMR and HAMR, in comparison to the centralized implementation in MOA¹⁴ (MAMR).

Evaluation methodology. We plug VAMR and HAMR into a *prequential evaluation task* [14], where each instance is first used to test and then to train the model. This evaluation task includes a *source* processor which provides the input stream and a *evaluator* processor which records the rate and accuracy of prediction results. The final task for HAMR is depicted in Figure 11(right). The parallelism level of the model is controlled by setting the number of learners p and the number of model aggregators r . The task for VAMR is similar but the default rule learner is excluded and model aggregator’s parallelism level is always 1. Each task is repeated for five runs.

Datasets. We perform the same set of experiments with 3 different datasets, i.e., *electricity*, *airlines* and *waveform*.

- *Electricity*: A dataset from the UCI Machine Learning Repository [2] which records the electricity power consumption (in watt-hour) of a household from December 2006 to November 2010. The dataset contains more than 2 millions 12-attribute records.
- *Airlines*: A dataset recording the arrival delay (in seconds) of commercial flights within the USA in year 2008¹⁵. It contains more than 5.8 millions records, each with 10 numerical attributes.
- *Waveform*: A dataset generated using an artificial random generator. To generate an instance, it picks a random waveform among the 3 available ones, and 21 attribute values for this instance are generated according to the chosen waveform. Another 19 noise attributes are generated and included in the new instance, making a total of 40 attributes for each instance. The label value is the index of the waveform (0, 1, or 2). Although this dataset does not fit perfectly to the definition of a regression task, it allows us to test our implementations with a high number of numerical attributes.

¹⁴ <http://moa.cms.waikato.ac.nz>

¹⁵ http://kt.ijs.si/elena_ikonovska/data.html

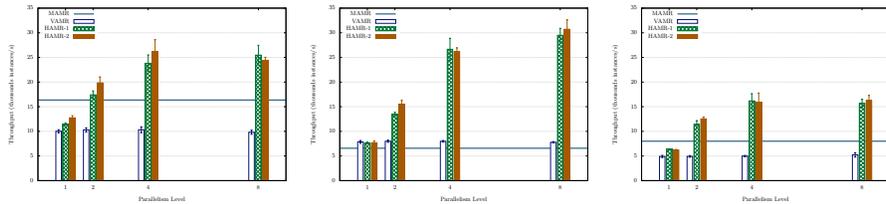


Fig. 12: Throughput of distributed AMRules with *electricity* (left), *airlines* (mid), and *waveform* (right).

Setup. The evaluation is performed on an OpenStack¹⁶ cluster of 9 nodes, each with 2 Virtual CPUs @ 2.3GHz and 4GB of RAM. All the nodes run Red Hat Enterprise Linux 6. The distributed implementations are evaluated on a Samza¹⁷ cluster with Hadoop YARN 2.2¹⁸ and Kafka 0.8.1.¹⁹ The Kafka brokers coordinate with each other via a single-node ZooKeeper 3.4.3²⁰ instance. The replication factor of Kafka’s streams in these experiments is 1. The performance of MAMR is measured on one of the nodes in the cluster.

Throughput. The throughput of several variants of AMRules is shown in Figure 12 for each dataset examined. HAMR-1 and HAMR-2 stand for *HAMR with 1 learner* and *HAMR with 2 learners*. The parallelism levels of VAMR represents the number of its learners, while in the case of HAMR it represents the number of model aggregators.

With *electricity* and *waveform*, the communication overhead in VAMR (to send instances from model aggregator to learners and evaluator) exceeds the throughput gained from delegating the training process to the learners and results in a lower overall throughput compared to MAMR’s. However, with *airlines*, the performance of VAMR is better than MAMR. To verify that the training process for *airlines* is more computationally intensive than the other two datasets and, thus, VAMR is more effective for this dataset, we compare the statistics of rules and predicates creation for the 3 datasets. The number of rules created, rules removed, and features created (when a rule is expanded) by MAMR with the 3 datasets are presented in Table 5. By subtracting the total number of rules removed from the total number of rules created, we can have an estimation of the average number of rules in the model for each dataset. A higher number of rules in the model and a higher number of features created suggest that the model is more complex and it takes more processing power to search for the best new feature when a rule attempts to expand.

Although VAMR can perform better than MAMR for computationally intensive datasets, its throughput does not change with different parallelism level.

¹⁶ <http://www.openstack.org>

¹⁷ <http://samza.incubator.apache.org>

¹⁸ <http://hadoop.apache.org>

¹⁹ <http://kafka.apache.org>

²⁰ <http://zookeeper.apache.org>

This is due to the bottleneck at the model aggregator. The processing of each instance at the model aggregator consists of 3 steps: finding the covering rules, forwarding the instance to the corresponding learners and applying covering rules to predict the label value of the instance. Since the complexity of these 3 steps for an instance is constant, the final throughput is unaffected by the parallelism level.

The throughput of HAMR-1 and HAMR-2 exhibit a better scalability compared to VAMR. Up to parallelism level of 4, the throughput increases almost linearly with the number of model aggregators. However, there is no or little improvement when this number is increased from 4 to 8. As we measure this throughput at a single evaluator, we suspect that the bottleneck is in the maximum rate the evaluator can read from the output streams of the model aggregators and default rule learner. To investigate this issue, we plot the maximum throughput of HAMR against the size of messages from model aggregators and default rule learner to evaluator in Figure 13. The values of throughput of a single-partition Samza stream with messages of size 500B, 1000B and 2000B are used to compute the linear regression line (reference line) in the figure. The message size for different datasets is shown in Table 5.

As reading is expected to be faster than writing in Samza and Kafka, the maximum rate the evaluator in HAMR can read from multiple stream partitions is expected to be higher than the throughput represented by the reference line. This fact is reflected in Figure 13 as the maximum throughput of HAMR for the 3 datasets constantly exceeds the reference line. However, the difference between them is relatively small. This result is a strong indicator that the bottleneck is the maximum reading rate of the evaluator. If there is no need to aggregate the result from different streams, this bottleneck can be eliminated.

Accuracy. We evaluate accuracy of the different implementations of AMRules in terms of *Mean Absolute Error* (MAE) and *Root Mean Square Error* (RMSE). Figure 14, 15 and 16 show the MAE and RMSE for the three datasets, normalized by the range of label values in each dataset.

Most of the figures show that error values of distributed implementations presents very small fluctuations around the corresponding MAMR error line. However, there is a significant increment in the value of RMSE in HAMR with

Table 5: Statistics for features and rules (with MAMR) for different datasets.

	Electricity	Airlines	Waveform
Instances	2 049 280	5 810 462	1 000 000
# Attributes	12	10	40
Result message size (B)	891	764	1446
# Rules created	1203	2501	270
# Rules removed	1103	1040	51
Avg. # Rules	100	1461	219
# Features created	1069	10 606	1245

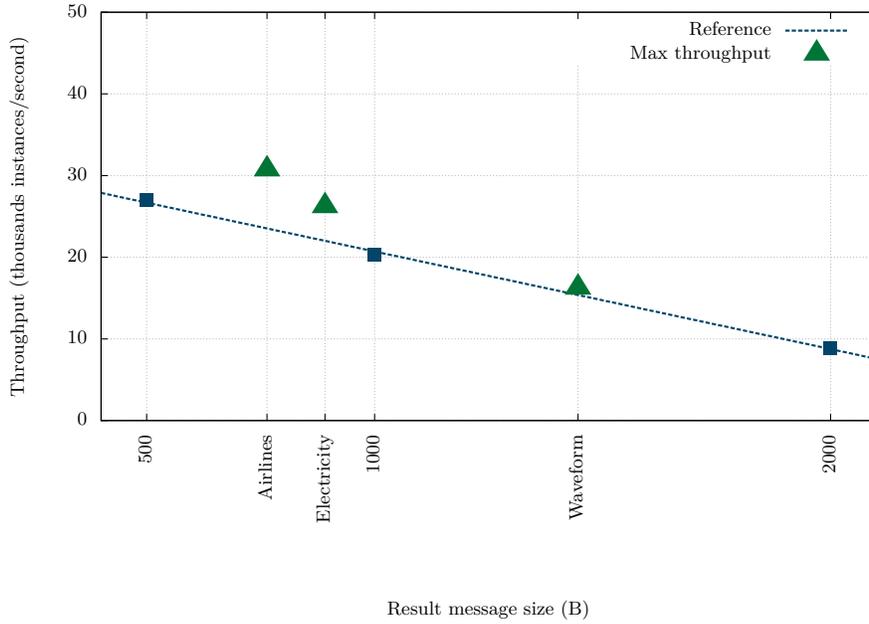


Fig. 13: Maximum throughput of HAMR vs message size.

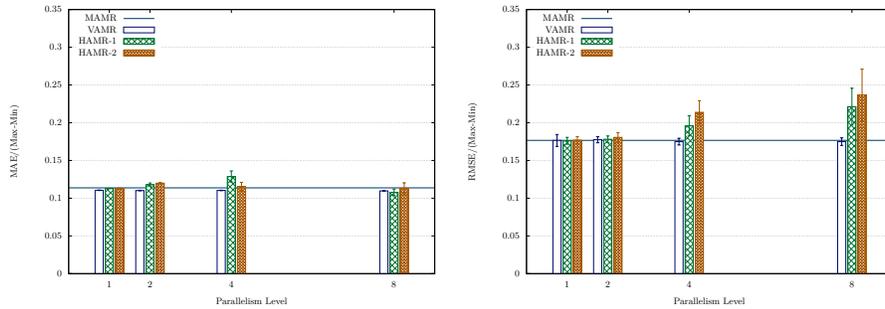


Fig. 14: MAE (left) and RMSE (right) of distributed AMRules with *electricity* dataset.

electricity dataset when the number of model aggregators is increased to 4 or 8. Moreover, larger variances, i.e., *standard error* is greater than 5% of the average MAE or RMSE, are also observed in the case of higher parallelism levels ($p \geq 4$) of HAMR. The probable cause is that when a rule in the model aggregators is out-of-sync with the corresponding one in the learners, i.e., model aggregators are not yet updated with a newly created rule, the number of instances that use an outdated model is multiplied by the throughput.

Table 6: Memory consumption of MAMR for different datasets.

Dataset	Memory consumption (MB)	
	Avg.	Std. Dev.
Electricity	52.4	2.1
Airlines	120.7	51.1
Waveform	223.5	8

Memory. The reference memory consumption of MAMR for different datasets is presented in Table 6. This table shows that *waveform* consumes the largest amount of memory, followed by *airlines* and then *electricity*.

Table 7 reports the memory consumption (*average* and *standard deviation*) of model aggregator and learner processors of VAMR for the 3 datasets with different parallelism levels. First of all, we notice a high memory consumption at the model aggregator. However this amount does not differ much for different datasets. This suggests that, at the model aggregator, there is a constant overhead in memory usage but the memory consumption due to the growing decision model is small.

Second, the memory consumption per learner decreases as more learners are used. As there is definitely some memory overhead for each learner instance, there is no significant reduction of memory usage per learner when the parallelism level goes from 4 to 8. However, the result indicates that we can spread a large decision model over multiple learners and make it possible to learn very large models whose size exceeds the memory capacity of a single machine.

8 Conclusions

We presented the Apache SAMOA platform for mining big data streams. The platform supports the most common machine learning tasks such as classifica-

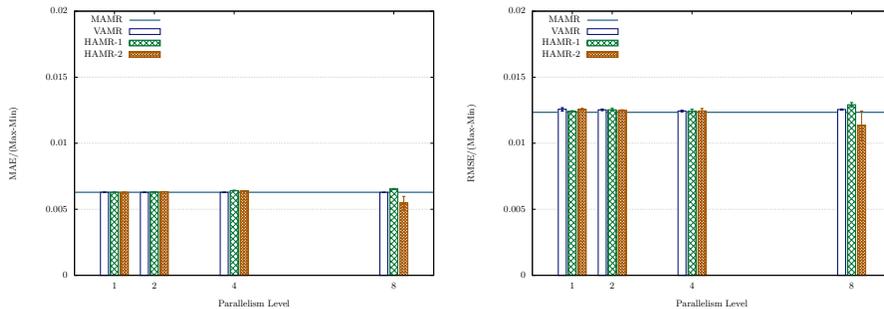


Fig. 15: MAE (left) and RMSE (right) of distributed AMRules with *airlines* dataset.

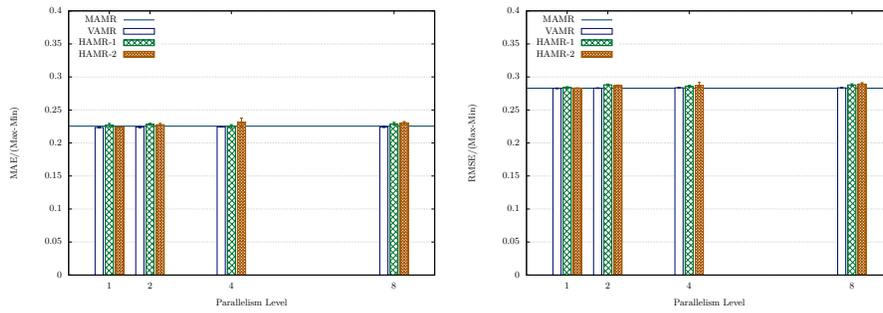


Fig. 16: MAE (left) and RMSE (right) of distributed AMRules with *waveform* random generator.

tion, clustering, and regression. It also provides a simple API for developers that allows to implement distributed streaming algorithms easily.

Apache SAMOA is already available and can be found online at <https://samoa.incubator.apache.org>. The website includes a wiki, an API reference, and a developer's manual. Several examples of how the software can be used are also available. The code is hosted on GitHub. Finally, Apache SAMOA is released as open source software under the Apache Software License v2.0.

Table 7: Memory consumption of VAMR for different datasets and parallelism levels.

Dataset	Parallelism	Memory Consumption (MB)			
		Model Aggregator		Learner	
		Avg.	Std. Dev.	Avg.	Std. Dev.
Electricity	1	266.5	5.6	40.1	4.3
	2	264.9	2.8	23.8	3.9
	4	267.4	6.6	20.1	3.2
	8	273.5	3.9	34.7	2.9
Airlines	1	337.6	2.8	83.6	4.1
	2	338.1	1.0	38.7	1.8
	4	337.3	1.0	38.8	7.1
	8	336.4	0.8	31.7	0.2
Waveform	1	286.3	5.0	171.7	2.5
	2	286.8	4.3	119.5	10.4
	4	289.1	5.9	46.5	12.1
	8	287.3	3.1	33.8	5.7

Bibliography

- [1] C. C. Aggarwal. *Data Streams: Models and Algorithms*. Springer, 2007.
- [2] K. Bache and M. Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- [3] Y. Ben-Haim and E. Tom-Tov. A Streaming Parallel Decision Tree Algorithm. *JMLR*, 11:849–872, Mar. 2010. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1756006.1756034>.
- [4] A. Bifet, G. Holmes, R. Kirkby, and B. Pfahringer. MOA: Massive Online Analysis. *The Journal of Machine Learning Research*, 11:1601–1604, 2010.
- [5] I. Bordino, N. Kourtellis, N. Laptev, and Y. Billawala. Stock trade volume prediction with Yahoo Finance user browsing behavior. In *30th International Conference on Data Engineering (ICDE)*, pages 1168–1173. IEEE, 2014.
- [6] D. Chatzakou, N. Kourtellis, J. Blackburn, E. De Cristofaro, G. Stringhini, and A. Vakali. Mean Birds: Detecting Aggression and Bullying on Twitter. In *9th International Conference on Web Science (WebSci)*. ACM, 2017.
- [7] C. Chen, J. Zhang, X. Chen, Y. Xiang, and W. Zhou. 6 million spam tweets: A large ground truth for timely Twitter spam detection. In *International Conference on Communications (ICC)*. IEEE, 2015.
- [8] G. De Francisci Morales. SAMOA: A Platform for Mining Big Data Streams. In *RAMSS: 2nd International Workshop on Real-Time Analysis and Mining of Social Streams @WWW*, 2013.
- [9] G. De Francisci Morales and A. Bifet. SAMOA: Scalable Advanced Massive Online Analysis. *JMLR: Journal of Machine Learning Research*, 2014.
- [10] G. De Francisci Morales, A. Gionis, and C. Lucchese. From chatter to headlines: harnessing the real-time web for personalized news recommendation. In *5th ACM International Conference on Web Search and Data Mining (WSDM)*, pages 153–162. ACM, 2012.
- [11] J. Dean and S. Ghemawat. MapReduce: Simplified Data processing on Large Clusters. In *6th Symposium on Operating Systems Design and Implementation (OSDI)*, pages 137–150. USENIX Association, 2004.
- [12] R. Devooght, N. Kourtellis, and A. Mantrach. Dynamic matrix factorization with priors on unknown values. In *21st International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pages 189–198. ACM, 2015.
- [13] P. Domingos and G. Hulten. Mining high-speed data streams. In *6th International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 71–80, 2000.
- [14] J. Gama, R. Sebastião, and P. P. Rodrigues. On evaluating stream learning algorithms. *Machine Learning*, 90(3):317–346, 2013.
- [15] J. Gama, I. Zliobaite, A. Bifet, M. Pechenizkiy, and A. Bouchachia. A Survey on Concept Drift Adaptation. *ACM Computing Surveys*, 46(4), 2014.

- [16] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963. URL <http://amstat.tandfonline.com/doi/abs/10.1080/01621459.1963.10500830>.
- [17] E. Ikonomovska, J. Gama, and S. Džeroski. Learning model trees from evolving data streams. *Data Mining and Knowledge Discovery*, 23(1):128–168, 2011.
- [18] A. Jacobs. The Pathologies of Big Data. *Communications of the ACM*, 52(8):36–44, Aug. 2009.
- [19] N. Kourtellis, F. Bonchi, and G. De Francisci Morales. Scalable online betweenness centrality in evolving graphs. *IEEE Transactions on Knowledge and Data Engineering*, pages 2494–2506, 2015.
- [20] N. Kourtellis, G. De Francisci Morales, and A. Bifet. VHT: Vertical Hoeffding Tree. In *4th IEEE International Conference on Big Data (BigData)*, 2016.
- [21] N. C. Oza and S. Russell. Online Bagging and Boosting. In *Artificial Intelligence and Statistics*, pages 105–112. Morgan Kaufmann, 2001.
- [22] E. Page. Continuous inspection schemes. *Biometrika*, pages 100–115, 1954.
- [23] A. Thu Vu, G. De Francisci Morales, J. Gama, and A. Bifet. Distributed Adaptive Model Rules for Mining Big Data Streams. In *2nd IEEE International Conference on Big Data (BigData)*, 2014.
- [24] M. A. Uddin Nasir, G. De Francisci Morales, D. Garcia-Soriano, N. Kourtellis, and M. Serafini. The Power of Both Choices: Practical Load Balancing for Distributed Stream Processing Engines. In *31st International Conference on Data Engineering (ICDE)*, 2015.
- [25] M. A. Uddin Nasir, G. De Francisci Morales, N. Kourtellis, and M. Serafini. When Two Choices Are not Enough: Balancing at Scale in Distributed Stream Processing. In *32nd International Conference on Data Engineering (ICDE)*, 2016.
- [26] T. Vasiloudis, F. Beligianni, and G. De Francisci Morales. BoostVHT: Boosting Distributed Streaming Decision Trees. In *26th ACM International Conference on Information and Knowledge Management (CIKM)*, 2017.