



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Parallelizing Machine Learning as a service for the end-user

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Availability:

This version is available at: <https://hdl.handle.net/11585/709029> since: 2019-12-18

Published:

DOI: <http://doi.org/10.1016/j.future.2019.11.042>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Daniela Loreti, Marco Lippi, Paolo Torroni. (2020). Parallelizing Machine Learning as a service for the end-user. *Future Generation Computer Systems*, 105: 275-286

The final published version is available online at:

<https://doi.org/10.1016/j.future.2019.11.042>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Parallelizing Machine Learning as a Service for the End-User

Daniela Loreti^{a,*}, Marco Lippi^b, Paolo Torroni^a

^a*DISI – University of Bologna*

^b*DISMI – University of Modena and Reggio Emilia*

Abstract

As Machine Learning (ML) applications are becoming ever more pervasive, fully-trained systems are made increasingly available to a wide public, allowing end-users to submit queries with their own data, and to efficiently retrieve results. With increasingly sophisticated such services, a new challenge is how to scale up to ever growing user bases. In this paper, we present a distributed architecture that could be exploited to parallelize a typical ML system pipeline. We propose a case study consisting of a text mining service, and discuss how the method can be generalized to many similar applications. We demonstrate the significance of the computational gain boosted by the distributed architecture by way of an extensive experimental evaluation.

Keywords: Machine Learning as a Service, Parallelization, MapReduce

1. Introduction

In the last decade, ML has undoubtedly become one of the hottest topics in computer science. By exploiting big data collections, ML algorithms are now being implemented and deployed on a large scale across countless application domains, including health-care, transportation, speech analysis, computer vision, market analysis, life sciences, and many others [1].

Recently, ML applications have been moving to the cloud, in order to exploit high performance parallel and distributed computing, which has given rise to the concept of Machine Learning as a Service (MLaaS) [2]. This usually refers to the availability of online platforms and frameworks, which

*Corresponding author

has enabled to implement in the cloud all the customary stages of a ML pipeline. Such stages include, for instance, input pre-processing, feature extraction, and in particular the training phase, which is usually the most expensive from a computational and memory consumption viewpoint [3].

In this paper, we turn our attention to a particular aspect of MLaaS, that of deploying and parallelizing systems that have already been trained, and need to be made available to (possibly many) end-users. Although not computationally expensive as the training phase, prediction and forecasting tasks may be nonetheless burdensome in terms of resources, especially when results must be delivered within a close deadline, or when the service has to be made available to a wide public. In this setting, an infrastructure for MLaaS should support large-scale distributed batch processing, as well as run-time stream processing. Besides, scalability and fault tolerance are fundamental requirements. In this regard, since its first formulation in 2004, the MapReduce [4] distributed programming model has gained significant diffusion in the big data research community. This success is mainly due to its simple yet powerful and intrinsically parallelizable paradigm. Several distributed computing engines have been proposed since then, to support the development of distributed programs for batch processing of very large data collections, providing autonomous fault-tolerant mechanisms and run-time infrastructure scaling capabilities. The latest evolution of these frameworks [5, 6, 7] offers stream processing support, alongside with the more traditional batch processing, and leverages different programming paradigms, in addition to MapReduce. All that together greatly simplifies the implementation of efficient distributed applications for the analysis of big data flows.

Our aim with this paper is to show how a MapReduce-inspired programming paradigm can be used to improve the performance and scalability of the ML pipeline, by parallelizing the customary steps of the prediction process and supporting the development of ready-to-use pre-trained services for the end user. These are our main contributions:

- A structural characterization of systems addressing the prediction task of different ML applications, from a MapReduce-inspired parallelization viewpoint.
- The detailed discussion of a concrete application of the outlined structure in the natural language processing domain, specifically in the area of argumentation mining, together with an empirical evaluation of the performance and scalability of the approach.

- A discussion of the issues and challenges that must be addressed when parallelizing the identified general ML pipeline, as well as the features that might induce significant enhancements in performance and desirability of the offered ML service.

The paper is structured as follows. Section 2 discusses related work, introducing the concept of MLaaS and its interpretations in various domains of computer science. Section 3 describes scenarios where the proposed parallel architecture could play a role. Section 4 focuses on the proposal of MLaaS for the end user by presenting a case study in the area of argumentation mining. Section 5 illustrates the parallel architecture. Section 6 presents its empirical evaluation, whereas Section 7 discusses the challenges the argumentation mining case study has helped bringing up. Section 8 concludes.

2. Related Work

MLaaS is a phrase found in various areas of computer science, where it is used to refer to various concepts. A significant body of literature focuses on the description and analysis of platforms that implement a whole ML pipeline. That may include, for instance, the capability to perform data pre-processing and feature selection, to choose the best-performing classifier, to train a model, and to predict the outcomes on query data. For example, Yao et al. [3] compare the performance and complexity of several solutions for building MLaaS applications implementing the entire ML pipeline. Similarly, Chan et al. [8] describe the distributed architecture exploited in ML applications in Uber, with a focus on model training and features selection. A complete architecture for MLaaS is also described by Ribeiro et al. [2] who present a specific analysis on three ML classifiers (multi-layer perceptrons, support vector machines, K -nearest neighbors). In each of these works, however, the parallelization of the final prediction stage is only superficially addressed, or ignored altogether, whereas we believe that a ML tool provided as-a-service to a wide public of end-users cannot disregard the parallelization of this last step (although it is generally less computationally expensive than the previous ones). In particular, we argue that the prediction stages of different ML applications have common characteristics, which make the adoption of a MapReduce-oriented approach [9] particularly suitable for the purpose. The need for a ML tool as a service through the adoption of MapReduce is also envisaged by other scholars. For instance, Chan et al. [10] and Baldominos et al. [11] do so while focusing on the features that such a tool should

expose, rather than on techniques to obtain scalability. The application of MapReduce to the ML pipeline comes with the great advantages brought by distributed computing architectures, which allow the developer to focus on the implementation of its parallel program while disregarding lower-level architectural and infrastructural details, such as the coordination between nodes, the employment of heterogeneous hardware in the same data-center [12, 13] or even the use of multiple cloud data-centers [14, 15, 16, 17].

Another line of research that uses the MLaaS term, focuses instead on the parallelization of training algorithms for ML systems. For example, a multicore implementation for the training of many ML systems, which exploits the MapReduce paradigm, is presented by Chu et al. [18]. Sergeev et al. [19] present an interesting framework to enable faster and easier distributed training in TensorFlow [20]. Tamano et al. [21] illustrate an approach to job scheduling optimization for MapReduce tasks in ML applications. The challenges and opportunities of exploiting MLaaS in the context of the Internet of Things are discussed by Assem et al. [22], again with a focus on the aspects related to training and classifier selection.

Fewer strands of work are instead dedicated to the performance of parallel algorithms for already trained and deployed ML models. Xu et al. [23] propose a software architecture that encompasses model deployment for real-time analytics. Their emphasis is on the processing of big data collections, including RESTful web services for data wrapping and integration, dynamic model training, and up-to-date prediction and forecasting. The resulting architecture relays on Sparks MLlib library to offer a set of already-parallelized, popular ML functions to the end-user. Given the particular focus of their study, Xu et al. do not propose a general methodology for distributing the ML pipeline, and in particular for parallelizing the prediction task of possibly more sophisticated applications and leave it to the developer to address such a problem. Similarly, Harnie et al. [24] use the Apache Spark technology [7] to achieve the desired scalability in cheminformatics applications, which is also a choice we made in our work. However, as their main target is obtaining performance enhancements through the parallelization of the scientific application at hand, they do not provide a general methodology for parallelizing prediction services. Nonetheless, it can be observed how the software implementation proposed by Harnie et al. (composed of two initial mapping tasks and a subsequent aggregation by reduction) can be easily framed in the methodology we propose (see Section 3). The focus of other works such as that by Hanzlik et al. [25] is on model deployment, and in particular on

problems related to model stealing and reverse engineering. In the envisaged scenario, the ML model is made available to the client for offline prediction on sensitive data that cannot leave the clients data center. The authors devote great attention to the performance of the security-enhanced forecasting step, but do not address the parallelization of such a task.

The problem of parallelizing sophisticated artificial intelligence-based reasoning engines has been studied by Loreti et al. [26, 27] in the domain of business process compliance monitoring. However, the focus of such study lies in the input data-set partitioning strategies, whereas the present study aims to identify common patterns in various ML tasks, which could lead back to a MapReduce-inspired approach.

3. Prediction as a Service

The typical pipeline of a ML system to parallelize and distribute has the following characteristics. A collection of data is given as input, either as a set of batches, or as a continuous stream. As Figure 1 illustrates, a first stage of computation is present, where all the instances of the input dataset need independent processing. Such independence enables the distribution of the computation load across many nodes. This set up is common to many application domains. For instance, a great deal of natural language processing tasks consider independent input elements such as single sentences, paragraphs, or documents. Among such tasks we shall mention sentiment

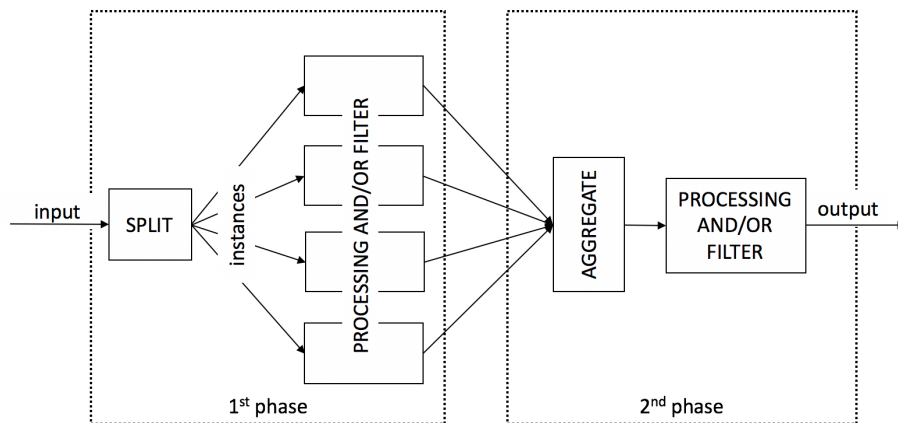
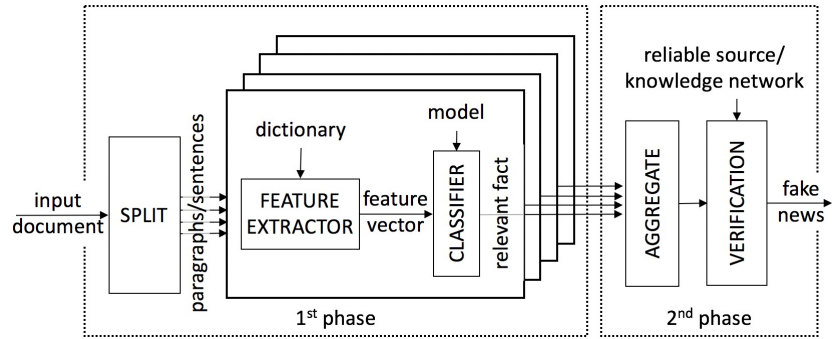


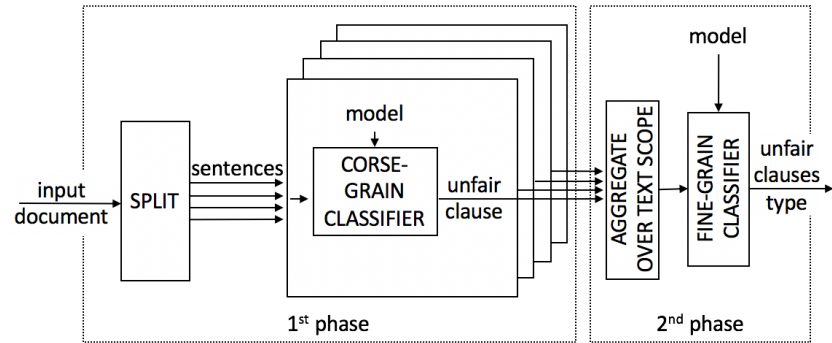
Figure 1: Generic pipeline of steps of a ML application

analysis, whose goal is to assign a *sentiment* to a given segment of text, being it a tweet, a post on a social network, or a comment to a newspaper article. Document categorization aims to *classify* a piece of text into one or more semantic categories. Fake news detection, as well as many other tasks, aim to *detect* sentences or text portions with certain characteristics [28]. The same can be said of many important tasks in computer vision. For example, the goal of image tagging or image classification is to assign a label (or a set of labels) to a given image. In video surveillance [29], as well as in many other video processing tasks, the computation is often carried out at the level of single frames, or small batches of frames. In other domains, such as bioinformatics, cheminformatics, or genomics data analysis, several different predictors can be applied to input data, such as protein or DNA sequences, so as to classify different properties of single sequence elements [30].

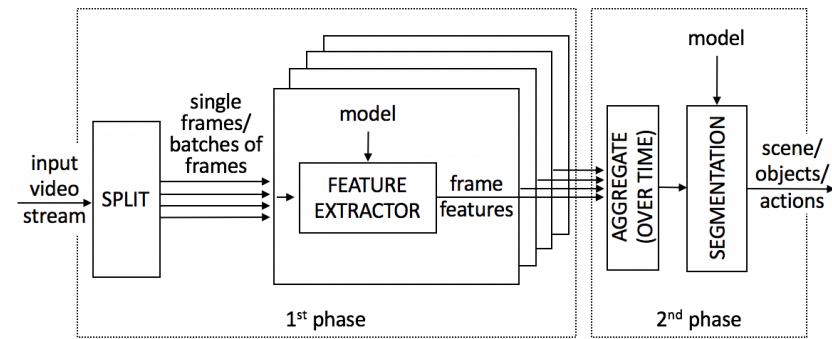
The output of this first stage of the pipeline is often used as input to a second stage. In general, the first stage could be considered a sort of filtering, or detection phase, whereas the second stage works as a sort of aggregation phase, where the instances detected during the first round are matched and compared with one another. That is the case in some natural language processing setups, where the text segments selected during the first phase are processed by a clustering algorithm, such as topic modeling, or by a further categorization stage. For instance, single sentences could be first classified according to coarse-grained categories, and then classification could be further developed into fine-grained categories. That happens with fact-checking systems, where candidate fake news items are first selected during an initial processing phase, and then a fact-checking algorithm is applied to a collection of instances found in the first phase [31], possibly relaying on a trustworthy source. Fig 2a sketches the steps of this pipeline. In some consumer-oriented applications [32], such as the detection of unfair contract clauses [33], a first stage could identify sentences expressing potentially unfair clauses through a coarse-grained classifier (see Figure 2b), whereas a second stage could additionally recognize the category of unfairness among a set of pre-defined cases (e.g., arbitration, limitation of liability, etc.), usually on the basis of the context provided by other relevant sentences within a certain document scope. In bioinformatics, the output of the predictors developed during the first stage are often aggregated into a higher-level set of predictions, that combine the information coming from the different classifiers [24]. In computer vision, video segmentation is typically performed on frame groups, resulting from a first processing stage [34]. Typically, as shown in Fig. 2c, the frames



(a)



(b)



(c)

Figure 2: Examples of pre-trained ML applications: detection of fake news (Fig. 2a), unfair clauses in legal documents (Fig. 2b), and video segmentation. The steps of these pipelines recall the MapReduce programming paradigm.

are initially analyzed to extract low-level features such as brightness, contrast, etc. and later aggregated over a temporal basis to extract higher-level information (e.g., scenes, objects, actions, etc.).

The general pipeline described so far recalls the principles of the MapReduce programming paradigm [9]. MapReduce is a well-known technique for simplifying the parallel execution of software, whereby the input data is partitioned into an arbitrary number of slices, each exclusively processed by a *mapper* task emitting intermediate results in the form of key/value pairs. The pairs are then passed on to other tasks, called *reducers*, which are in charge of merging together the values associated with the same key, and emitting the final result. As Figure 1 illustrates, the first phase of the general ML pipeline could indeed be modeled as a mapping task, where each instance undergoes the same processing to emit a key/value pair. The corresponding reduction task is then carried out in the second phase of the pipeline, where the outputs of the first phase are merged and processed together.

Programs reformulated according to the MapReduce model can be automatically parallelized by means of a big data analytics platform [35], which can turn a collection of computing nodes into a distributed infrastructure able to automatically spread (and balance) the execution tasks across the data center. Most of these platforms supply mechanisms to scale up or down the cluster on demand and perform fault detection and recovery from failures at runtime. All these features seem particularly beneficial to a ML tool provided as-a-service to a wide public.

In an effort to ground the general discussion and illustrate in concrete how MLaaS can be enabled by a MapReduce-oriented approach, the section that follows is devoted to a particular case study. The case study is a natural language processing application in the area of argumentation mining [36], consisting of an argument component identification task followed by an argument structure prediction task, whose goal is to identify links and relations between argument components.

4. A Case Study

Argumentation mining [36] defines the task of extracting arguments from unstructured text by automated analysis methods. Such a task is usually, but not necessarily, referred to a specific genre, such as legal texts, persuasive essays, scientific literature, etc. One notable and increasingly popular application area where argumentation mining plays a key role is that of debating

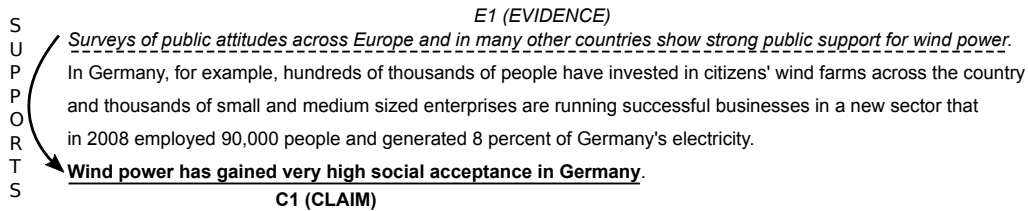


Figure 3: Example of the output of an argumentation mining system, taken from the IBM Wikipedia corpus [40]. Given a text, this system detects evidence (in italics), claims (in bold), and their support links.

technologies [37], where the data to be processed may consist of “static” Wikipedia pages as well as streamed audio signals. Owing also to the variety of addressed genres, the argumentation mining literature offers several alternative definitions of argument [38], with varying degrees of sophistication. For the aims of this work, and without loss of generality, we shall refer to a common and rather basic claim/premise argument model [39], whereby an argument is composed by an assertion, or statement (the *claim*) supported by one or more *premises*, and the inference that connects claim and premises.

Argumentation mining includes several sub-tasks, spanning from claim and evidence detection (i.e., the identification of argument components), to attribution (i.e., detection of the authorship of an argument), to increasingly challenging tasks such as the prediction of the relations between arguments (*argument structure prediction*), the inference of implicit argument components (*enthymemes*), and so forth. Figure 3 offers an example of a possible output produced by an argumentation mining system.

One such system is Mining ARGuments frOm Text (MARGOT) [41]. MARGOT exploits a combination of ML and natural language processing techniques in order to perform argumentation mining on unstructured texts of various genres. In particular, MARGOT exploits a Support Vector Machine (SVM)-based method for context-independent claim detection [42] using tree kernels, and extends its application to evidence detection, and to the identification of their boundaries.

The initial version of MARGOT follows a pipeline of subsequent stages, by initially segmenting the text into sentences, then performing the detection of argumentative sentences, that is sentences that include argument components (claims or evidence), and finally identifying the boundaries of each argument component. The current MARGOT prototype [43], also available as a web

server,¹ adopts the “traditional”, sequential architecture. It analyzes each input sentence individually, thereby processing the document as a whole in a sequential manner, in order for each sentence to undergo all the subsequent steps of the pipeline. It does not address argument structure prediction. It is worthwhile noticing that MARGOT’s training phase is executed offline, once and for all. It results in two trained models, one for claim detection and one for evidence detection, which are then deployed for use of the online server.

In this work, taking MARGOT as a case study, we present a more advanced version of the system, where the claim and evidence detection stage is followed by the application of a further SVM-based step, aimed to detect the existing support links between all possible pairs of claims and evidence found in each document. The offline training of this additional stage resulted in a third model, which was deployed for our experimental analysis.

Following the model introduced in Section 3, we describe MARGOT’s pipeline as composed of two subsequent phases. A first phase, shown in Figure 4, processes each sentence of the input file using the Stanford Parser [44], a largely successful third-party software package. MARGOT uses the Stanford Parser to obtain a first set of features, that is, the trees encoding the grammatical structure of a sentence, known in the literature as *constituency parse trees*. From the same sentence, MARGOT also extracts the bag-of-words feature vector, which represents a binary encoding of its words. The

¹See <http://margot.disi.unibo.it>

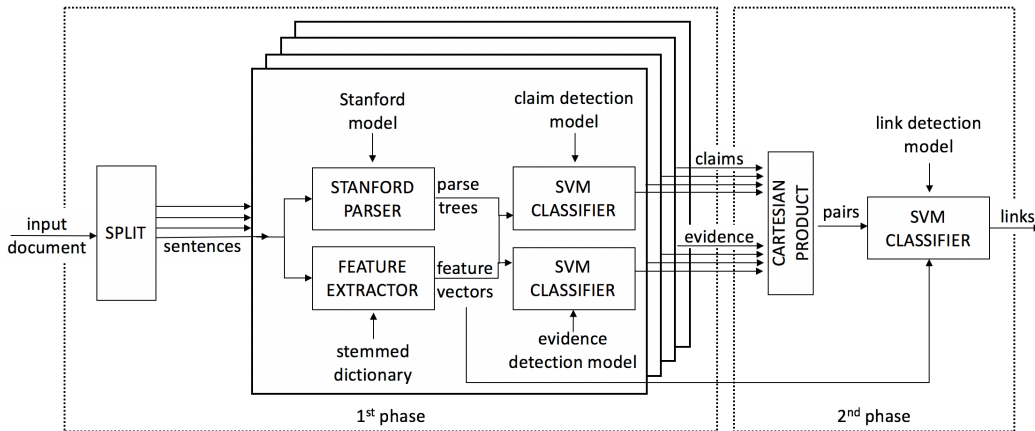


Figure 4: The pipeline of steps implemented in MARGOT

two sets of features are passed on to two different classifiers employing the SubSet Tree Kernel (SSTK). The aim of these classifiers is to identify claims and evidence, respectively. All and only the sentences containing the identified claims/evidence are then sent to the second phase of the pipeline, which considers all the possible *(claim, evidence)* pairs in each file and calls another SVM-based classifier to detect the possible links between the two (link detection).

Given the complexity of the analysis to be carried out, especially when large files are mined, the sequential execution of MARGOT’s pipeline can be highly time- and resource-consuming. The problem is exacerbated when the number of arguments detected in the first phase is large, since this entails to consider, for each input file, the Cartesian product of large sets of the detected premises and claims. Furthermore, as we envisage the future necessity of argumentation mining analysis as a service, it is likely that such a service would be required to consider streams of input text instead of documents already materialized in a certain disk location. The run-time nature of stream processing further amplifies the need for scalable and reliable architectures to support argumentation mining as a service.

Nonetheless, as prescribed by the general model in Figure 1, the parser and the feature extractor of the first phase can process each sentence independently from the others, whereas the Cartesian product in the second phase operates as a pair-wise sentence aggregator. This observation suggests a way to distribute the computational load of MARGOT’s pipeline on a network of computing nodes, leveraging a MapReduce-oriented approach and an engine for large-scale batch and stream processing.

5. A Parallel Architecture

Among the existing variety of MapReduce-oriented engines for large-scale data processing, some offer the possibility to analyze batches of documents already materialized in a certain location [45]; others only deal with the processing of data flows [5, 46, 47]; whereas a restricted number offer the possibility – particularly desirable for the current work – to operate on both batches and streams [6, 7]. For the purposes of this work, and without loss of generality, we will describe how MARGOT can be re-implemented to be automatically executed on a distributed infrastructure using the facilities provided by Apache Spark [7]. Apache Spark has become increasingly popular in the last years also because it allows a developer to write her application

in several different languages, without forcing her to think in terms of only map and reduce operators. Its good performance [48, 49, 50] and resilience to faults [51] has been empirically verified by various studies. We will first consider the case of large batches of input documents, and later refine the algorithm in order to accommodate streams as well.

5.1. MARGOT for batch processing

In case of batch processing, all the documents to be analyzed are already present in a certain (centralized or distributed) disk location, and the data analytics infrastructure is in charge of spreading (and balancing) the computation load on the available nodes.

As Figure 5 illustrates, in the first phase of MARGOT’s pipeline the files to be analyzed are split into sentences. A collection of $(key, value)$ pairs is thus produced, where the key is the file name and the value is the sentence found.

Then, the core operations of the first phase are performed on each sentence independently by applying a map function. This operation extracts the parse tree and the bag-of-words feature vector, and passes them in input to two third-party classifiers that emit, for each sentence, a claim and an

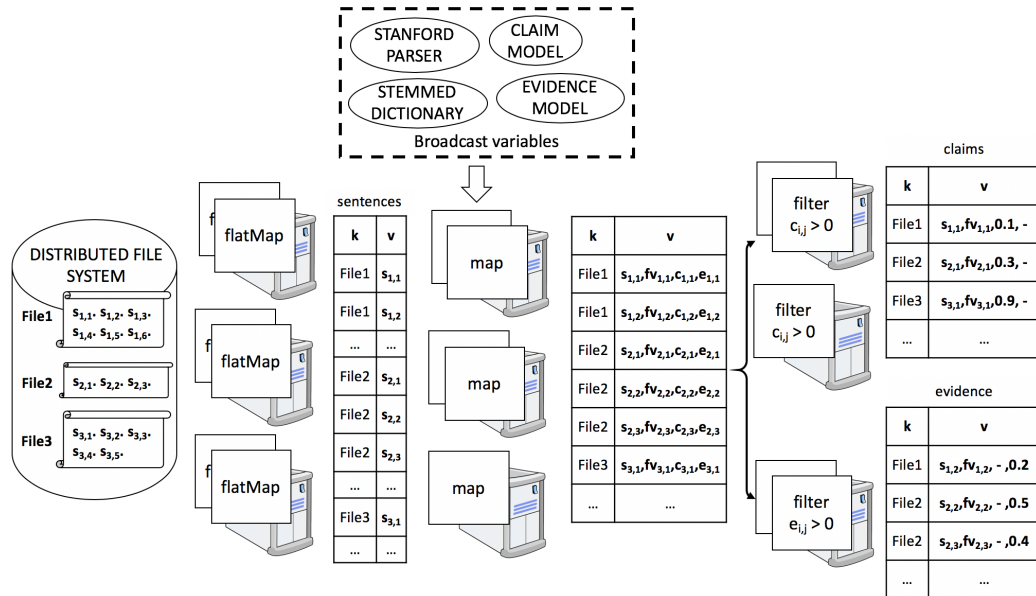


Figure 5: Distributed version of the first phase of MARGOT, working on batch input.

evidence score, respectively. As these functions require sizable models (the parsing model, the stemmed dictionary and the claim/evidence models) to operate on each sentence, we apply the general suggestion of loading such models once at the beginning of the computation. The objects produced are sent to all the computing nodes, by leveraging the concept of *broadcast variable* offered by Apache Spark i.e., an immutable shared variable which is cached on each worker node of the Spark cluster. The output of the map function is again a $(key, value)$ pair with the same key (the file name), but with the values (which hosted just the sentence in the input) now enriched by the phrase feature vector, claim and evidence score.

Finally, two filters are applied to select only claims and premises. Indeed, two different collections of pairs are produced: one containing elements with positive claim score, and another with positive evidence score.

The details of the first phase implementation on a distributed MapReduce-oriented platform are presented in Listing 1 following an Apache Spark-inspired approach with lambda functions. It is worthwhile underlining that, since the SVM-based classifiers are realized by third-party C software [52] that cannot be directly converted into a Spark broadcast variable, we attempt to minimize the initial overhead of loading these external software by resorting to a `mapPartitions` function. Differently from `map` and `mapValues` functions, which are executed for each sentence (as in line 10), `mapPartitions` operates on a collection of sentences, which is a partition of the whole sentences in the input files. The size of this partition is optimized by the underlying infrastructure, based on the number of available cores. As a consequence, the call to external SVM-based classifiers and the loading of the models in lines 16 and 17 is not repeated for each sentence (which would be highly inefficient) but for each group of sentences the infrastructure has partitioned the input in. We shall remark that this solution could be applied to other MLaaS applications too, in order to deal with the frequent challenge of reducing the overhead of third-party software invocation.

Since the result of `mapPartitions` is later accessed by two different filters, cluster-wide caching (line 29) is employed to avoid a duplicated computation of the same `mapPartitions` stage.

The collections computed during the first phase are then merged together in the second phase of MARGOT, as shown in Figure 6. In particular, a distributed `join` operation is performed, to obtain a collection of all the possible $(claim, evidence)$ pairs in each file. The subsequent map function considers each pair individually. Indeed, parallelizing the steps following the

Listing 1: Distributed implementation of MARGOT’s first phase (batch mode).

```

Input: inputDir /*directory containing the input documents*/, stanfordModel /*model for the
Stanford Parser*/, stemmedDict /*dictionary of stems for feature extraction*/, claimModel,
evidenceModel /*SVM classification models*/
Output: claims, evidence /*sentences containing claims/evidences*/
1 stanfordParser = LexicalizedParser.loadModel(stanfordModel)
2 sparkContext.broadcast(stanfordParser)
3 sparkContext.broadcast(stemmedDict)
4 sparkContext.broadcast(claimModel)
5 sparkContext.broadcast(evidenceModel)
6 phrases = sparkContext
7   .wholeTextFiles(inputDir) //emit (fileName,fileContent) pairs
8   .flatMapValues(fileContent =>
9     fileContent.split("[!?!?"])) //split into sentences and emit (fileName,sentence) pairs
10  .mapValues(sentence => {
11    tree = stanfordParser.apply(sentence) // generate parsing tree
12    fv = FeatureExtractor
13      .createBagOfWords(tree,dmStemmed) // generate feature vector
14    (sentence, tree, fv) })
15  .mapPartitions((partition) => {
16    svmC = SvmClassify.load(claimModel)
17    svmE = SvmClassify.load(evidenceModel)
18    outList = List()
19    l = partition.toList
20    l.foreach( x => { //compute claim and evidence score of each sentence in partition
21      filename = x._1
22      sentence = x._2._1
23      tree = x._2._2
24      fv = x._2._3
25      claimScore = svmC.getScore(tree, fv)
26      evidScore = svmE.getScore(tree, fv)
27      outList=(filename,(sentence,fv,claimScore,evidScore))::outList })
28    outList.iterator }) //emit a partition with the elements of the novel list
29  .cache
30 claims = phrases.filter(_._2._3>0) //only sentences with claimScore > 0
31 evidence = phrases.filter(_._2._4>0) //only sentences with evidScore > 0

```

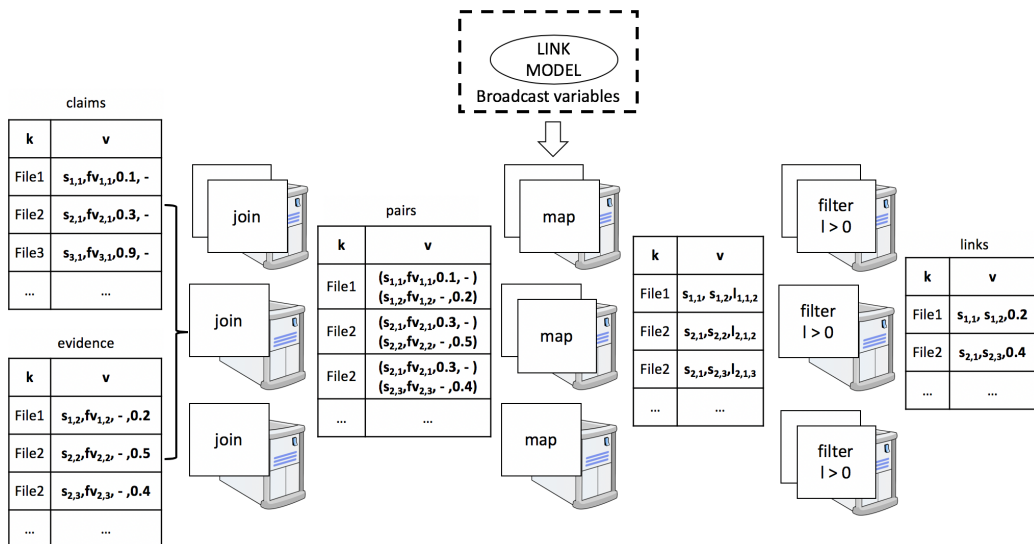



Figure 6: Distributed version of the second phase of MARGOT, working on batch input.

aggregation – whenever possible – should help to speed up the computation. Inside each map function, the link model and the pair of feature vectors in each record are employed by a third-party SVM-based classifier to predict a link score, indicating whether the claim and evidence in each pair are linked. As reported in Listing 2, similarly to claim/evidence detection, a **mapPartitions** function (line 11) is actually performed at this stage, aiming to minimize the initial overhead of loading the external software. Finally, only elements with positive link scores are maintained (line 37 of Listing 2), as these represent the final output of the argumentation mining algorithm.

5.2. MARGOT for stream processing

The distributed algorithm presented in the previous subsections can accommodate streams of input text after only minor modifications. In particular, those operations conducted on each sentence independently, such as map and filter, can be performed on streams and batches alike, whereas the semantics of functions that merge together different collections of key-value pairs, such as the join operation in the second phase of the pipeline, needs further refining for stream processing.

Link detection on a batch of documents implicitly entails a natural definition of “scope” which corresponds with the document at hand. Input files should thus be considered independently, so as to identify the claims and

Listing 2: Distributed implementation of MARGOT’s second phase (batch mode).

```

1 Input: claims, evidence /*sentences containing claims/evidences*/, linkModel /*SVM
   classification model*/
2 Output: links /* (claim,evidence) pairs containing an argumentative link*/
3
5 sparkContext.broadcast(linkModel)
7 pairs = claims.join(evidence)
9 links = pairs
11 .mapPartitions((partition) => {
13     svmL = SvmClassify.load(linkModel)
15     outList = List()
17     l = partition.toList
19     l.foreach( x => { //compute link score of each sentence in partition
21         filename = x._1
23         claim = x._2._1._1
25         evid = x._2._2._1
27         fvC = x._2._1._2
29         fvE = x._2._2._2
31         linkScore = svmC.getScore(fvC, fvE)
33         outList = (filename, (claim, evid, linkScore)) :: outList })
35     outList.iterator })
37 .filter(_. _2._3 > 0) //only pairs with linkScore > 0

```

evidence therein, enabling link detection on a per-file basis, as shown in the previous section. When dealing with a stream of data instead, two different semantics for link detection are possible.

- *Scope file* pairing: if the stream itself contains the indication of the input file transmitted, we might be asked to pair claims with evidence in each file, with a semantics similar to the one adopted for batch processing, thus keeping track of the currently identified claims/evidence in each file along the stream.
- *Scope window* pairing: if there is no explicit or implicit concept of document/file, a natural choice would be to detect the links on a sliding window of sentences in the input stream, that is, to find a connection between claims and evidence separated by at most n other sentences in the stream. This sort of “locality principle” especially holds true in such contexts as argumentation mining, where claims and supporting evidence are typically near to each other in the input text [53].

The implementation of the second phase of MARGOT for stream processing, presented in Listing 3 enables our system to accommodate both

Listing 3: Distributed implementation of MARGOT’s second phase (streaming mode).

```

Input: claims, evidence /*sentences containing claims or evidences*/, linkModel /*SVM
        classification model*/, mode /*file-wise or window-wise*/, winDim /*dimension of the join
        window for window-wise link detection*/
Output: links /* (claim,evidence) pairs containing an argumentative link*/

5  if (mode==file-wise){
7    //join all claims and evidences in each file
9    claimsCollection = claims
11   .updateStateByKey(updateClaimCollection) //emits a stateful collection of all the claims
        encountered since the beginning of the streamed document
13   pairs=claimsCollection.join(evidence)
15 }else{
17   //join on a specific window dimension
19   pairs = claims
21     .window(winDim)
23     .join(evidences.window(winDim))
25 }
27 links = pairs
29 .mapPartitions((partition) => {
31   svmL = SvmClassify.load(linkModel)
33   outList : List[(String, (String, String),Double)] = List()
35   l = partition.toList
37   l.foreach( x => { //compute link score of each sentence in partition
39     filename = x._1
41     claim = x._2._1._1
43     evid = x._2._2._1
45     fvC = x._2._1._2
47     fvE = x._2._2._2
49     linkScore = svmC.getScore(fvC, fvE)
51     outList = (filename, (claim, evid, linkScore)) :: outList })
53   outList.iterator })
55 .filter(_._2._3>0) //only pairs with linkScore > 0

```

semantics. We do not report the details of the first phase, since it is rather similar to that of batch processing.

To perform link detection on a sliding window of the input stream (lines 17 to 23), the flows of claim and evidence collections are sliced using a basic **window** operation before performing the join. In this way, only the claims and evidence inside each window are paired. To perform link detection with scope file, instead, a stateful operation is performed to maintain a growing collection of claims encountered in each streamed file (line 11). This set of past claims is joined with each new evidence detected in the stream (line 13).

6. Empirical Analysis

The objective of the analysis we are going to discuss here is gaining a quantitative understanding of the performance enhancement that can be obtained when a ML task is distributed on a network of computing nodes. We are not interested here in evaluating the accuracy of the ML methods themselves, since that measure would be independent of the architecture being parallel or otherwise, and anyway it has been studied in previous works [41, 42]. Indeed, MARGOT here works as a case study for investigating the scalability of a MLaaS application, in case of both batch and stream processing.

6.1. Simulation setup

We evaluate the performance of the proposed distributed system for argumentation mining on a cluster of 126 physical nodes. One of these machines, configured as a Spark master, coordinates the work of the others 125 slaves. All the computers are equipped with 8 CPUs, 16GB of RAM, and a 400GB hard disk. The nodes are interconnected by a 100Mbit/s bandwidth local network. We evaluate the distributed version of MARGOT on a collection of input files downloaded from the Project Gutenberg web site.² As we need to operate on texts containing a significant number of claims and evidence, common novels are not suitable as input. We therefore restrict our attention to essays in English language. The considered dataset includes 50 files, yielding 466,483 total sentences. The complete source code of the distributed version of MARGOT is available on GitHub [54].

²<https://www.gutenberg.org>

6.2. Evaluation approach

We run separate experiments to evaluate the performance of the system in the two execution modes: with batch input documents and text streams.

To investigate batch processing, we stored the documents downloaded from Project Gutenberg into a Hadoop Distributed File System (HDFS) [55], which automatically slices and distributes the files on the network of computing nodes. When studying the system for stream processing, instead, we must consider an input flow of text with a certain rate, measured in bytes per second. In that case, it is crucial for the system to be able to perform all the steps in the pipeline, while keeping up with the input rate. If the computation is slower than the input flow, not only will the system introduce an increasing delay in the time to emit the output, but also the buffer area employed by Spark to temporarily store the data waiting to be analyzed may eventually become saturated. The Spark Streaming module treats the flow by periodically slicing it into portions called micro-batches, which are later distributed on the network and separately processed on each node. The period of micro-batch slicing is a configurable parameter. As a general recommendation large micro-batch period helps to keep up with high input rates at the price of an increased latency in the results. Because our goal is not to evaluate the performance of the Spark Streaming’s micro-batch processing mechanism, but to evaluate the scalability of the system, we fixed the batch time to 100 seconds for all the tests on streams.

Real-world stream processing services usually experience increases and decreases of the input rate over the day. In order to evaluate the proposed MLaaS application we progressively increase the input rate during each test, so as to identify the maximum input rate that the system can sustain before it starts falling behind.

In both streaming and batch cases, we conduct three scalability tests:

- *Test 1* – experiment with increasing input size (i.e., increasing file dimension for batches, and larger window for streams). The objective is to determine the scalability of the overall MLaaS application.
- *Test 2* – experiment with increasing number of (*key, value*) pairs emitted by the first phase (i.e., the number of emitted claims and evidence that would be later joined and processed in the second phase). The objective is to study the effect of the aggregation step bottleneck on the performance.

- *Test 3* – experiment with increasing number of support vectors in the employed SVM model. The objective is to study the impact of computationally demanding ML tasks.³

We shall remark that *Tests 1* and *2* are independent of the specific ML application considered, whereas *Test 3* has been specifically conceived in the context of MARGOT, because it is strictly related to the kind of operations conducted in its pipeline. However, it is of general interest, since SVM classifiers are widely popular due to their excellent performance in a large variety of tasks.

6.3. Results

Concerning *Test 1*, Figure 7 illustrates the scalability of the distributed system by increasing amounts of input data. In particular, the plot on the left (Figure 7a) shows the time required to process datasets of different size using batch processing with increasing numbers of computing nodes. The size of datasets are reported in Table 1.

As desired, the total execution time greatly benefits from the introduction of additional nodes. The most significant improvements are observed between 1 and 50 nodes. After 50 nodes, the cost of distributing the tasks on the

³Because the features of each sentence must be compared with all the vectors, it is well-known that for an SVM classifier larger models yield longer computation times.




Figure 7: System performance by varying size/rate of the input, and computing nodes. Batch mode (left) and stream mode (right).

Dataset	Sentences	Claims	Evidence	Links
DS1	9,783	1,244	2,739	4,489
DS2	67,917	8,050	16,267	9,827
DS3	233,254	13,597	76,173	29,027
DS4	466,483	27,193	152,345	58,464

Table 1: Details of the datasets employed in batch processing evaluation of Figure 7a

network balances off the benefits yielded by the additional computational resources. Furthermore, when dealing with small input files such as the “DS1” series, a slight performance loss is observed as the number of nodes increases from 25 to 50 and above. That could be the effect of the overhead generated by partitioning and distributing in the network small amounts of data.

The other plot (Figure 7b) illustrates the system’s performance with text streams. The y-axis reports the maximum input rate in bytes per second that the system can tolerate without falling behind. The graph has been plotted by periodically increasing the input rate and checking when the processing time of each micro-batch in the stream started to exceed the configured micro-batch period. Hence, the higher is the curve in the figure, the better is the system performance. We have made several experiments by varying the windows size. For example, with $w = 5,000s$ we indicate the performance of the system when claims and evidence are joined over a window of 5,000 seconds, that is fifty times bigger than the micro-batch period. In this case, when employing 125 nodes, the system cannot keep up with an input frequency higher than 600 bytes/s. If we assume the average sentence length to be 200 characters, this means that the system cannot process more than 3 sentences per second. Although such a performance may seem unimpressive at a first glance, we should consider the sheer number of claim/evidence pairs to be analyzed by the link classifier. In particular, with a 600 bytes/s input rate and a 100s micro-batch period, a 5,000 seconds window contains around 15,000 sentences.

Figure 7b also reports the performance in the “scope-file” series, when no window is employed, but instead the claim/evidence join is executed on a per-file basis. The performance are worse than those obtained with a window of size 100s (up to around 900 sentences per window processed on 125 nodes) and better than 1,000s (around 72,000 sentences per window). As the average

Model	Dimension (MB)	Support Vectors
M1	6.7	7,085
M2	17.3	18,604
M3	29.9	30,363

Table 2: Details of the model employed in *Test 3*

number of sentences in each file is 9,000, the position of the “scope-file” curve between series “w=100s” and “w=1,000s” appears reasonable.

Figure 8 illustrates the results of *Test 2*. In order to obtain different amounts of elements to be analyzed in the second phase, in *Test 2* we artificially varied the filtering thresholds that MARGOT uses to identify claims and evidence at the end of the first phase. The names of the series (5%, 35%, 65% and 90%) report the percentage of input sentences that reach the second phase. As expected, both batch (Figure 8a) and stream (Figure 8b) processing reveal a significant effect of this parameter on the processing time and the maximum input rate. This confirms the bottleneck effect of the aggregation step in the ML pipeline. Nonetheless, the scalability trend is evidently maintained for all the series in the graphs.

Figure 9 illustrates the results of *Test 3*, whose aim was to study the impact of the number of support vectors in the link prediction model on the system’s performance. Table 2 summarizes the details of each model.




Figure 8: Scalability of the system by increasing amount of $(key, value)$ pairs emitted by the first phase. Left: batch processing; right: stream processing.




Figure 9: System performance by varying complexity of the model. Left: batch processing; right: stream processing.

As the SVM classifier checks the features of each $(claim, evidence)$ pair with all the support vectors in the model, one could imagine that larger models cause longer computation times. Instead, rather surprisingly, we observed that the effect of this parameter on the overall batch (Figure 9a) and stream processing (Figure 9b) performance is insignificant and that scalability is not visibly affected.

7. Discussion

Our empirical analysis indicates a promising scalable behaviour in all the considered scenarios. In addition, our study was instrumental in identifying challenges and opportunities that could arise in general from the parallelization of other ML services.

7.1. Challenges of ML pipeline parallelization

Independently of the particular ML application, the implementation of the pipeline initially depicted in Figure 1 on a network of cooperating computing nodes presents three major architectural and technological challenges. First of all, while the first phase enjoys the benefits of a highly parallelizable structure, the aggregation step in the second phase needs careful consideration, to prevent it from turning into a bottleneck. Indeed, combining several records together is an expensive operation in distributed environments because it causes massive data shuffle over the network. This is actually a

known issue for many MapReduce applications. In case of MLaaS, it may be crucial to apply, if possible, a filtering operation aimed to reduce the input space for the aggregation function. For instance, one could anticipate some steps that conceptually may belong further down in the pipeline, in an effort to narrow down the data emitted at the end of the first phase.

A second challenge has to do with the trained model supervised or semi-supervised ML applications commonly use in order to process the input data and provide predictions. When a distributed engine is employed, each computing node must be provided a copy of such a model alongside with the input data. However, depending on the application and the underlying ML technologies, the dimension of the trained model can be significant. Nonetheless, if the learning process is not continuous, meaning that the model evolution is limited to the training phase, such a model is stable during the whole prediction process. Accordingly, the model could be distributed to all computing nodes at the beginning of the computation once and for all (e.g., thanks to mechanisms such as the broadcast variables of Apache Spark), and then employed by each node independently to process/filter its portion of the input.

A final technological challenge has to do with the third-party software ML applications often employ at different stages of the pipeline. The integration of third-party methods into the framework of a distributed data processing engine might not be a straightforward operation. Ad-hoc solutions (such as Spark's *mapPartitions*) might be required, for instance, to limit the number of calls to external processes executing the third-party software.

7.2. Other aspects of MLaaS on large-scale data processing engines

The implementation of MLaaS on large-scale data processing engines not only offers architectural and technological challenges, but it also often requires making some detailed choices whose impact on the performance of the offered service may be crucial. For example, although the aggregation step of the second phase works by collecting together the output of the previous processing, performing the subsequent processing/filtering step on a single node is not mandatory, and indeed it is generally worth avoiding. If a certain degree of parallelism is possible (because for example the instances emitted by the aggregation step can be considered independently for one another, as was for MARGOT's claim/evidence pairs), a good practice would be to split the second phase too into tasks that can be carried on concurrently. Modern big data analytics engines support the implementation of distributed programs not strictly limited to a map and a reduce phase. More complex variants

and combinations of the two are possible. As far as MLaaS is concerned, a further map step following the aggregation could contribute to boosting the performance of the second phase.

Furthermore, when the processing/filtering steps involve a large trained model, the task of loading such model can be time consuming and should be performed carefully. Consider for example a natural language processing system whose first phase classifies the sentences of a text according to the features of a large trained model. In such a system it is certainly not recommended to load the model for each sentence. A much preferable solution would be to load the model once, and then use it to analyze a consistent number of phrases together. Nonetheless, grouping too many sentences together reduces the degree of parallelism of the pipeline step, thus slowing down the computation. Depending on the specific predictive task to be carried out and on the degree of parallelism allowed by the underlying infrastructure, a trade-off must be found between the need to avoid unnecessarily repeating costly operations (such as model loading) and that of splitting the work to speed up the computation.

A similar conundrum regards the features extracted from the input data which might be used in multiple subsequent steps of the pipeline. Feature extraction may be a costly operation, in which case it would be worthwhile performing it only once, and then hand over the features to the following steps until they are no longer needed. This is indeed the solution adopted by MARGOT, where the features extracted in the first phase are passed along the pipeline together with the relevant sentence. However, one must consider that between each pair of steps of the pipeline a data shuffle over the network might occur. Therefore, if feature extraction produces large outputs, the shuffling of such data on a network with limited bandwidth might cause poor performance. In these cases, there is a trade-off to consider, between recomputing large features when needed and shuffling them between nodes.

Finally, a common feature of MLaaS applications is the requirement to accommodate a variety of input modes, in particular document batches and data streams. From a developer's standpoint, a shift of perspective is unavoidable when passing from a ML application working on input files that are already materialized and stored in a specific location (i.e., batch processing) to the analysis of a flow of input data (i.e, stream processing). Nonetheless, some relatively recent MapReduce-oriented platforms [6, 7] allow this shift at the price of slight changes in the application implementation. Since the first phase of our reference ML pipeline operates in theory on each input instance

independently, it could operate on batches and streams alike. The second phase instead is likely to be dependent on the input mode, because it focuses on the aggregation and processing of the previous step’s output. Indeed, for stream processing aggregation entails the need to specify not only the collections of data to be merged, but also the period over which such operation must be performed. For example, aggregation could be based on the last n occurrences in the flow, on a specific time window, or on all the data received so far. Each option would have a different meaning from the others and may produce completely different results. In the case of batch processing, instead, the aggregation task has a less varied range of semantics, because there are no concepts like “time window” or “arrival instant” of an input entry.

8. Conclusions

A growing number of ML applications are being deployed as ready-to-use, already trained services for the end-users. This calls for implementing distributed architectures able to scale up such services to broader user communities, and larger data collections.

In this paper, we presented a distributed architecture inspired by the MapReduce paradigm, which could be used to parallelize the prediction phase of a typical ML pipeline. We conducted experimental results on a real-world text mining application case study. We also discussed how the methodology is general enough to be applied to many other different scenarios. We considered both batch and stream processing, and studied the performance gain that can be achieved by this architecture under many angles.

An interesting open challenge is how to effectively extend this architecture to accommodate ML applications dealing with structured data, such as sequences, trees, and graphs. In that case, since the relations between the input data are as relevant as the data themselves, it is unlikely that the first phase could be implemented through a simple split operation, followed by independent map processes. Instead, we expect that a more elaborated slicing procedure would be needed at the beginning of the pipeline, in order to correctly partition the data across the nodes. A technique used to achieve a desired level of parallelism with input sequences of data in another application domains involved data replication [27]. The idea was to divide sequences into slices and replicating the data on the extremities, before assigning each slice to a different nodes. The identification of patterns for the split step that can be applied and reused in case of more elaborated input

structures would be an interesting subject for future investigation. An even more challenging setting would be the case where the input data has a highly connected structure, hindering any kind of split and re-partition of the work. Then a completely different approach could be explored: instead of slicing the data, one could provide each node with the whole dataset, but only a portion of the trained model. In this way, each machine of the distributed architecture could conduct a lightweight prediction analysis on all the input. Like in the pipeline described in this work, the results of those analyses will have to be conveniently aggregated in the following phase.

Acknowledgements

This work has been partially supported by the H2020 Project AI4EU (g.a. 825619).

References

- [1] T. J. Sejnowski, *The deep learning revolution*, MIT Press, 2018.
- [2] M. Ribeiro, K. Grolinger, M. A. Capretz, Mlaas: Machine learning as a service, in: *Machine Learning and Applications (ICMLA)*, 2015 IEEE 14th International Conference on, IEEE, pp. 896–902.
- [3] Y. Yao, Z. Xiao, B. Wang, B. Viswanath, H. Zheng, B. Y. Zhao, Complexity vs. performance: empirical analysis of machine learning as a service, in: *Proceedings of the 2017 Internet Measurement Conference*, ACM, pp. 384–397.
- [4] J. Dean, S. Ghemawat, MapReduce: Simplified data processing on large clusters, in: E. A. Brewer, P. Chen (Eds.), *6th Symposium on Operating System Design and Implementation (OSDI 2004)*, San Francisco, California, USA, December 6-8, 2004, USENIX Association, 2004, pp. 137–150.
- [5] Apache Storm, <http://storm.apache.org> (accessed October 3, 2019), 2019.
- [6] Apache Flink, <https://flink.apache.org> (accessed October 3, 2019), 2019.

- [7] Apache Spark, <http://spark.apache.org> (accessed October 3, 2019), 2019.
- [8] L. E. Li, E. Chen, J. Hermann, P. Zhang, L. Wang, Scaling machine learning as a service, in: International Conference on Predictive Applications and APIs, pp. 14–29.
- [9] J. Dean, S. Ghemawat, MapReduce: simplified data processing on large clusters, *Commun. ACM* 51 (2008) 107–113.
- [10] S. Chan, T. Stone, K. P. Szeto, K. H. Chan, Predictionio: a distributed machine learning server for practical software development, in: Q. He, A. Iyengar, W. Nejdl, J. Pei, R. Rastogi (Eds.), 22nd ACM International Conference on Information and Knowledge Management, CIKM’13, San Francisco, CA, USA, October 27 - November 1, 2013, ACM, 2013, pp. 2493–2496.
- [11] A. Baldominos, E. Albacete, Y. Sáez, P. Isasi, A scalable machine learning online service for big data real-time analysis, in: 2014 IEEE Symposium on Computational Intelligence in Big Data, CIBD 2014, Orlando, FL, USA, December 9-12, 2014, IEEE, 2014, pp. 112–119.
- [12] M. Zaharia, A. Konwinski, A. D. Joseph, R. H. Katz, I. Stoica, Improving mapreduce performance in heterogeneous environments, in: R. Draves, R. van Renesse (Eds.), 8th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2008, December 8-10, 2008, San Diego, California, USA, Proceedings, USENIX Association, 2008, pp. 29–42.
- [13] D. Cheng, J. Rao, Y. Guo, C. Jiang, X. Zhou, Improving performance of heterogeneous mapreduce clusters with adaptive task tuning, *IEEE Trans. Parallel Distrib. Syst.* 28 (2017) 774–786.
- [14] G. Antoniu, A. Costan, J. Bigot, F. Desprez, G. Fedak, S. Gault, C. Pérez, A. Simonet, B. Tang, C. Blanchet, R. Terreux, L. Bougé, F. Briant, F. Cappello, K. Keahey, B. Nicolae, F. Suter, Scalable data management for map-reduce-based data-intensive applications: a view for cloud and hybrid infrastructures, *IJCC* 2 (2013) 150–170.

- [15] F. J. Clemente-Castelló, B. Nicolae, K. Katrinis, M. M. Rafique, R. Mayo, J. C. Fernández, D. Loreti, Enabling big data analytics in the hybrid cloud using iterative mapreduce, in: [56], pp. 290–299.
- [16] D. Loreti, A. Ciampolini, Mapreduce over the hybrid cloud: A novel infrastructure management policy, in: [56], pp. 174–178.
- [17] D. Loreti, A. Ciampolini, A hybrid cloud infrastructure for big data applications, in: 17th IEEE International Conference on High Performance Computing and Communications, HPCC 2015, 7th IEEE International Symposium on Cyberspace Safety and Security, CSS 2015, and 12th IEEE International Conference on Embedded Software and Systems, ICESS 2015, New York, NY, USA, August 24-26, 2015, IEEE, 2015, pp. 1713–1718.
- [18] C.-T. Chu, S. K. Kim, Y.-A. Lin, Y. Yu, G. Bradski, K. Olukotun, A. Y. Ng, Map-reduce for machine learning on multicore, in: Advances in neural information processing systems, pp. 281–288.
- [19] A. Sergeev, M. D. Balso, Horovod: fast and easy distributed deep learning in tensorflow, CoRR abs/1802.05799 (2018).
- [20] TensorFlow, <https://www.tensorflow.org> (accessed October 3, 2019), 2019.
- [21] H. Tamano, S. Nakadai, T. Araki, Optimizing multiple machine learning jobs on MapReduce, in: Cloud Computing Technology and Science (CloudCom), 2011 IEEE Third International Conference on, IEEE, pp. 59–66.
- [22] H. Assem, L. Xu, T. S. Buda, D. OSullivan, Machine learning as a service for enabling internet of things and people, Personal and Ubiquitous Computing 20 (2016) 899–914.
- [23] D. Xu, D. Wu, X. Xu, L. Zhu, L. Bass, Making real time data analytics available as a service, in: Quality of Software Architectures (QoSA), 2015 11th International ACM SIGSOFT Conference on, IEEE, pp. 73–82.
- [24] D. Harnie, M. Saey, A. E. Vapirev, J. K. Wegner, A. Gedich, M. Steijaert, H. Ceulemans, R. Wuyts, W. De Meuter, Scaling machine learning

- for target prediction in drug discovery using apache spark, *Future Generation Computer Systems* 67 (2017) 409–417.
- [25] L. Hanzlik, Y. Zhang, K. Grosse, A. Salem, M. Augustin, M. Backes, M. Fritz, Mlcapsule: Guarded offline deployment of machine learning as a service, *arXiv preprint arXiv:1808.00590* (2018).
 - [26] D. Loreti, F. Chesani, A. Ciampolini, P. Mello, Distributed compliance monitoring of business processes over MapReduce architectures, in: W. Binder, V. Cortellessa, A. Koziolok, E. Smirni, M. Poess (Eds.), *Companion Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering, ICPE 2017, L’Aquila, Italy, April 22-26, 2017*, ACM, 2017, pp. 79–84.
 - [27] D. Loreti, F. Chesani, A. Ciampolini, P. Mello, A distributed approach to compliance monitoring of business process event streams, *Future Generation Comp. Syst.* 82 (2018) 104–118.
 - [28] C. D. Manning, C. D. Manning, H. Schütze, *Foundations of statistical natural language processing*, MIT press, 1999.
 - [29] C. Steger, M. Ulrich, C. Wiedemann, *Machine vision algorithms and applications*, John Wiley & Sons, 2018.
 - [30] P. Baldi, S. Brunak, F. Bach, *Bioinformatics: the machine learning approach*, MIT press, 2001.
 - [31] G. L. Ciampaglia, P. Shiralkar, L. M. Rocha, J. Bollen, F. Menczer, A. Flammini, Computational fact checking from knowledge networks, *PLoS ONE* (2015) 1–13.
 - [32] M. Lippi, G. Contissa, F. Lagioia, H.-W. Micklitz, P. Pałka, G. Sartor, P. Torroni, Consumer protection requires artificial intelligence, *Nature Machine Intelligence* (2019).
 - [33] M. Lippi, P. Pałka, G. Contissa, F. Lagioia, H.-W. Micklitz, G. Sartor, P. Torroni, Claudette: an automated detector of potentially unfair clauses in online terms of service, 2019.
 - [34] A. M. Tekalp, *Digital video processing*, Prentice Hall Press, 2015.

- [35] D. Singh, C. K. Reddy, A survey on platforms for big data analytics, *Journal of Big Data* 2 (2014) 8.
- [36] M. Lippi, P. Torroni, Argumentation mining: State of the art and emerging trends, *ACM Trans. Internet Techn.* 16 (2016) 10:1–10:25.
- [37] S. Mirkin, G. Moshkovich, M. Orbach, L. Kotlerman, Y. Kantor, T. Lavee, M. Jacovi, Y. Bilu, R. Aharonov, N. Slonim, Listening comprehension over argumentative content, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 719–724.
- [38] A. Peldszus, M. Stede, From argument diagrams to argumentation mining in texts: A survey, *IJCINI* 7 (2013) 1–31.
- [39] D. N. Walton, What is reasoning? what is an argument?, *The Journal of Philosophy* 87 (1990) 399–419.
- [40] R. Rinott, L. Dankin, C. A. Perez, M. M. Khapra, E. Aharoni, N. Slonim, Show me your evidence—an automatic method for context dependent evidence detection, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 440–450.
- [41] M. Lippi, P. Torroni, MARGOT: A web server for argumentation mining, *Expert Syst. Appl.* 65 (2016) 292–303.
- [42] M. Lippi, P. Torroni, Context-independent claim detection for argument mining, in: Q. Yang, M. Wooldridge (Eds.), *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, Buenos Aires, Argentina, July 25-31, 2015, AAAI Press, 2015, pp. 185–191.
- [43] Mining ARGuments frOm Text (MARGOT), <http://margot.disi.unibo.it> (accessed October 3, 2019), 2019.
- [44] C. D. Manning, M. Surdeanu, J. Bauer, J. R. Finkel, S. Bethard, D. McClosky, The stanford corenlp natural language processing toolkit, in: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014*, June 22-27, 2014, Baltimore, MD,

- USA, System Demonstrations, The Association for Computer Linguistics, 2014, pp. 55–60.
- [45] Apache Hadoop, <https://hadoop.apache.org/> (accessed October 3, 2019), 2019.
- [46] Apache Samza, <http://samza.apache.org> (accessed October 3, 2019), 2019.
- [47] Google Cloud Dataflow, <https://cloud.google.com/dataflow/> (accessed October 3, 2019), 2019.
- [48] J. Veiga, R. R. Expósito, X. C. Pardo, G. L. Taboada, J. Touriño, Performance evaluation of big data frameworks for large-scale data analytics, in: J. Joshi, G. Karypis, L. Liu, X. Hu, R. Ak, Y. Xia, W. Xu, A. Sato, S. Rachuri, L. H. Ungar, P. S. Yu, R. Govindaraju, T. Suzumura (Eds.), 2016 IEEE International Conference on Big Data, BigData 2016, Washington DC, USA, December 5-8, 2016, IEEE, 2016, pp. 424–431.
- [49] S. Chintapalli, D. Dagit, B. Evans, R. Farivar, T. Graves, M. Holderbaugh, Z. Liu, K. Nusbaum, K. Patil, B. Peng, P. Poulosky, Benchmarking streaming computation engines: Storm, flink and spark streaming, in: 2016 IEEE International Parallel and Distributed Processing Symposium Workshops, IPDPS Workshops 2016, Chicago, IL, USA, May 23-27, 2016, IEEE Computer Society, 2016, pp. 1789–1792.
- [50] J. Samosir, M. Indrawan-Santiago, P. D. Haghighi, An evaluation of data stream processing systems for data driven applications, in: M. Connolly (Ed.), International Conference on Computational Science 2016, ICCS 2016, 6-8 June 2016, San Diego, California, USA, volume 80 of *Procedia Computer Science*, Elsevier, 2016, pp. 439–449.
- [51] M. A. Lopez, A. G. P. Lobato, O. C. M. B. Duarte, A performance comparison of open-source stream processing platforms, in: 2016 IEEE Global Communications Conference, GLOBECOM 2016, Washington, DC, USA, December 4-8, 2016, IEEE, 2016, pp. 1–6.
- [52] A. Moschitti, Efficient convolution kernels for dependency and constituent syntactic trees, in: J. Fürnkranz, T. Scheffer, M. Spiliopoulou (Eds.), Machine Learning: ECML 2006, 17th European Conference on

Machine Learning, Berlin, Germany, September 18-22, 2006, Proceedings, volume 4212 of *Lecture Notes in Computer Science*, Springer, 2006, pp. 318–329.

- [53] S. Eger, J. Daxenberger, I. Gurevych, Neural end-to-end learning for computational argumentation mining, in: R. Barzilay, M. Kan (Eds.), Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers, Association for Computational Linguistics, 2017, pp. 11–22.
- [54] ParallelMargot, <https://doi.org/10.5281/zenodo.3251016> (accessed October 3, 2019), 2019.
- [55] Hadoop Distributed File System, https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html (accessed October 3, 2019), 2019.
- [56] I. Raicu, O. F. Rana, R. Buyya (Eds.), 8th IEEE/ACM International Conference on Utility and Cloud Computing, UCC 2015, Limassol, Cyprus, December 7-10, 2015, IEEE Computer Society, 2015.