

An Information Theory-Based Feature Selection Framework for Big Data Under Apache Spark

Sergio Ramírez-Gallego, Héctor Mouriño-Talín, David Martínez-Rego, Verónica Bolón-Canedo, José Manuel Benítez, *Member, IEEE*, Amparo Alonso-Betanzos, *Senior Member, IEEE*, and Francisco Herrera, *Senior Member, IEEE*

Abstract—With the advent of extremely high dimensional datasets, dimensionality reduction techniques are becoming mandatory. Of the many techniques available, feature selection (FS) is of growing interest for its ability to identify both relevant features and frequently repeated instances in huge datasets. We aim to demonstrate that standard FS methods can be parallelized in big data platforms like Apache Spark so as to boost both performance and accuracy. We propose a distributed implementation of a generic FS framework that includes a broad group of well-known information theory-based methods. Experimental results for a broad set of real-world datasets show that our distributed framework is capable of rapidly dealing with ultrahigh-dimensional datasets as well as those with a huge number of samples, outperforming the sequential version in all the cases studied.

Index Terms—Apache spark, big data, feature selection (FS), filtering methods, high-dimensional.

I. INTRODUCTION

IN THE last few decades, the dimensionality of datasets used for machine learning (ML) or data mining tasks has increased significantly, representing an unprecedented challenge for researchers, given that existing algorithms do not

always respond in a suitably timely way to extremely high dimensionality. In fact, analyzing the datasets posted in the popular libSVM Database [1], it can be observed that maximum data dimensionality was about 62 000 in the 1990s, but increased to some 16 million in the first decade of this century and has further increased to some 29 million in the current decade. In this new scenario, existing learning methods need to be adapted so as to be able to deal with millions of features.

With the advent of very high-dimensional datasets, identifying relevant features has become paramount. Dimensionality reduction techniques can be applied to reduce the dimensionality of the original data and even to improve learning performance [2]–[4]. These dimensionality reduction techniques usually come in two flavors: 1) feature selection (FS) and 2) feature extraction. Each approach has its own merits. Feature extraction techniques combine original features to yield a new set of features, whereas FS techniques remove irrelevant and redundant features. Since FS maintains the original features, it is especially useful for applications where the original features are important for model interpretation and knowledge extraction [5], [6]. For this reason, this model was the focus of this paper.

Existing FS methods are not expected to scale well when dealing with big data due to the fact that efficiency may significantly deteriorate or that the FS approach may even become inapplicable [7]. Scalable distributed programming protocols and frameworks have emerged in the last decade to manage the problem of big data. The first programming model developed was MapReduce [8] along with its open-source implementation Apache Hadoop [9], [10]. More recently, Apache Spark [11], [12] has been presented as a new distributed framework with a fast, general large-scale data processing engine that is popular with ML researchers due to its suitability for iterative procedures.

Likewise, several libraries for approaching ML tasks in big data environments have emerged in recent years. The first such library was Mahout [13], followed up by MLlib [14], built on the Spark system [12]. Thanks to Spark's capacity for in-memory computation that speeds up iterative processes, algorithms developed for this kind of platform have become pervasive in the industry. Although several gold-standard algorithms for ML tasks have been redesigned to incorporate a distributed implementation for big data technologies, this is not yet the case for FS algorithms.

Manuscript received January 7, 2016; accepted February 1, 2017. This work was supported in part by the Spanish National Research Projects under Grant TIN2014-57251-P, Grant TIN-2015-65069-C2-1-R, and Grant TIN2016-81113-R, in part by the BBVA Foundation Project (BigDaPTOOLS) under Grant 75/2016, in part the Andalusian Research Plan under Grant P11-TIC-7765 and Grant P12-TIC-2958, in part by the Xunta de Galicia through Research Project under Grant GRC2014/035, and in part by the FEDER Funds of the European Union. The work of S. Ramírez-Gallego was supported by the FPU Scholarship from the Spanish Ministry of Education and Science under Grant FPU13/00047. The work of D. Martínez-Rego and V. Bolón-Canedo was supported by the Xunta de Galicia under Grant POS-A/2013/196 and Grant ED481B 2014/164-0. This paper was recommended by Associate Editor K. Wolf. (*Corresponding author: Sergio Ramírez-Gallego.*)

S. Ramírez-Gallego and J. M. Benítez are with the Department of Computer Science and Artificial Intelligence, CITIC-UGR, University of Granada, 18071 Granada, Spain (e-mail: sramirez@decsai.ugr.es; j.m.benitez@decsai.ugr.es).

H. Mouriño-Talín, D. Martínez-Rego, V. Bolón-Canedo, and A. Alonso-Betanzos are with the Department of Computer Science, University of Corunna, 15071 A Coruña, Spain (e-mail: h.mtalín@udc.es; dmartinez@udc.es; veronica.bolon@udc.es; ciamparo@udc.es).

F. Herrera is with the Department of Computer Science and Artificial Intelligence, CITIC-UGR, University of Granada, 18071 Granada, Spain, and also with the Faculty of Computing and Information Technology, King Abdulaziz University, Jeddah 21589, Saudi Arabia (e-mail: herrera@decsai.ugr.es).

Digital Object Identifier 10.1109/TSMC.2017.2670926

To date only a simple chi-squared approach¹ and a Random Forest improvement [15] have been proposed to deal with this problem.

This paper aims to fill the existing gap by demonstrating that standard FS methods can be designed for big data platforms that can be usefully applied to big datasets while boosting both performance and accuracy. We propose a new distributed design for an FS generic framework based on information theory [16] that is implemented using the Apache Spark paradigm. To make this adaptation feasible, a wide variety of techniques from the distributed environment have been used, including information caching, data partitioning and replication of relevant variables. Note that adapting this framework to Spark implies a major challenge as it requires deep restructuring of classic algorithms.

To test the effectiveness of our framework, we applied it to a complete set of real-world datasets [up to $O(10^7)$ features and instances]. The results point to competitive performance (in terms of generalization and efficiency) when dealing with datasets that are huge in terms of both number of features and instances. As an illustrative example, we were able to select 100 features in a dataset with 29×10^6 features and 19×10^6 instances in under 60 min (using a 432-core cluster).

The remainder of this paper is organized as follows: Section II provides some background information on FS, big data, the MapReduce programming model and other frameworks. Section III describes the distributed framework we propose for FS in big data. Section IV presents and discusses the experiments carried out. Finally, Section V concludes this paper.

II. BACKGROUND

Below we briefly introduce FS, discuss the advent of big data and its implications and outline the particularities of MapReduce and the models derived from this framework.

A. Feature Selection

FS is a dimensionality reduction technique that tries to remove irrelevant and redundant features from original data. Its goal is to obtain a subset of features that properly describes a given problem with minimum performance degradation in order to obtain simpler and more accurate schemes [2].

We can formally define FS as follows: let e_i be an instance $e_i = (e_{i1}, \dots, e_{in}, e_{iy})$, where e_{ir} corresponds to the r th feature value of the i th sample and e_{iy} corresponds to the value of the output class Y . Let us assume a training set D with m examples, whose instances e_i are formed by a set X of n characteristics or features, and a test set D_t . $S_\theta \subseteq X$ is this defined as a subset of selected features yielded by an FS algorithm.

FS methods can be broadly categorized as follows [17].

- 1) *Wrapper Methods*: It uses an evaluation function dependent on a learning algorithm [18]. They are aimed at optimizing a predictor as part of the learning process.
- 2) *Filtering Methods*: It uses other selection techniques as separability measures or statistical dependences.

They only consider the general characteristics of the dataset, as they are independent of any predictor [19].

- 3) *Embedded Methods*: It uses a search procedure that is implicit in the classifier/regressor [20].

Filtering methods usually result in better generalization due to learning independence. Nevertheless, since they usually select larger feature subsets, they sometimes require the application of a threshold. Regarding complexity, filters are normally less costly than wrappers. When the number of features is large (especially the case of big data), filtering methods need to be used as they are much faster than either of the other approaches.

B. Big Data: Two-Sided Coin

The Internet continues generating quintillions of bytes of data; in 2012, for example, 2.5 exabytes of data were created daily. A solution for the problem of handling large collections of data is therefore becoming increasingly urgent [21]. Exceptional technologies are now required to efficiently collect, maintain, transmit and process large datasets within tolerable time intervals. Extracting relevant information from these collections of data is now one of the most important and complex challenges facing data analytics research, especially since many knowledge extraction algorithms have been rendered obsolete in the face of such vast amounts of data.

Big data, a term coined to describe the exponential growth and availability of data nowadays, has become a problem for classical data analytics. Gartner [22] referred to big data in terms of volume, velocity, and variety, that is, the 3Vs, to which a further 2Vs were subsequently added, namely, veracity and value. Information thus defined requires a radically new approach to large-scale processing. An under-explored but no less important topic is big dimensionality in big data [23]. This phenomenon, also known as the ‘‘curse of big dimensionality,’’ reflects the explosion of features and the combinatorial impact of new large incoming datasets with thousands or even millions of features.

Data scientists have generally focused on only one aspect of big data, namely, the huge number of instances, while paying less attention to the features aspect. Big dimensionality, however, calls for new FS strategies and methods to deal with the feature explosion problem, reflected in many of the best known dataset repositories in computational intelligence (like UCI or libSVM [1], [24]), where the dimensionality of most of the newly added datasets is huge [23] (e.g., almost 30 million for the KDD2010 dataset).

An additional problem is the myriad of feature types and their combinations that are becoming standard in many real-world scenarios. For instance, multimedia content on the Internet, which represents about 60% of total traffic [25], is transmitted in thousands of different formats (audio, video, images, etc.). Another example is natural language processing, where multiple feature types such as words, n -gram templates, etc., are simultaneously used to produce comprehensible and reliable models [26].

Not all features in a problem, however, contribute equally to prediction models and results. FS is thus required, now more

¹<http://spark.apache.org/docs/latest/mllib-feature-extraction.html#chisqselector>

than ever, to ensure fast and cost-effective learning and prediction. Isolating high-value features from a raw set of potentially irrelevant, redundant and noisy features, while complying with measurement and storage requirements, is a key task in big data research. Although some solutions exist that apply FS techniques to large-scale data [27], most of these solutions are applied locally. For high-dimensional data, however, new distributed solutions are necessary.

C. Distributed Programming Models and Frameworks: Apache Hadoop and Apache Spark

The MapReduce framework [8] emerged in 2003 as a revolutionary tool for handling big data. It was designed by Google specifically to generate and process large-scale datasets by automatically processing data in a highly distributed way using large clusters of computers. The framework is in charge of partitioning and managing data, recovering failures, scheduling jobs and communicating, leaving the programmers with a transparent and scalable tool to easily execute tasks on distributed systems.²

MapReduce is based on processing in two phases: 1) mapping (map) and 2) reducing (reduce). The map function transforms key-value pairs into sets of intermediate pairs. The reduce function merges these intermediate pairs with a matching key, with the master node splitting the data into chunks and distributing them across nodes for independent processing (in a divide-and-conquer fashion). Each node next executes the Map function on a given input subset and notifies completion to the master node. The master node then distributes the matching pairs across the nodes according to a key partitioning scheme that combines pairs using the reduce function to form the final output.

Apache Hadoop [9], [10] is an open-source implementation of MapReduce for reliable, scalable, and distributed computing. Despite being a popular open-source implementation of MapReduce, however, Hadoop is not suitable for certain applications, including online or iterative computing, high interprocess communication paradigms, or in-memory computing [29].

In recent years, Apache Spark has been included in the Hadoop ecosystem [11], [12] as a powerful framework that performs faster distributed computing on big data. It does this by using in-memory primitives that allow it to perform 100 times faster than Hadoop for certain applications. The fact that this platform allows user programs to load data into memory and to make repeated queries means that it is particularly well suited for online and iterative processing (especially for ML algorithms). Spark is also versatile in allowing several distributed programming models like Pregel and MapReduce to be implemented.

Spark is based on a distributed data structure called resilient distributed datasets (RDDs). RDDs are an immutable, partitioned set of records that can be generated by either stored data or other RDDs. There are two types of operations for RDDs: 1) *transformation*, which creates a new RDD from

another RDD and 2) *action*, which returns a record/value to the main program after a set of operations is applied. Some examples of transformations are *map* (transform all records) and *filter* (select records), whereas *count* (compute the number of records) and *save* (write output) are examples of actions. The main difference between the two operations are when and how they are applied to data. Transformations, lazily applied, are annotated until an action triggers execution of pending operations in a log called *lineage*, which removes the need for checkpoints while allowing lost partitions to be easily recomputed from the log.

Users can also control other distributed features like persistence and data partitioning. For instance, users can *cache* a dataset to memory for reuse in further iterations. RDDs can also be persisted on disk if we consider that re-execution may be more costly than spilling to disk. Finally, placement of key-based data can be optimized by choosing between different Spark schemes (*range* or *hash*) or even using our own partitioner.

Spark primitives extend MapReduce concepts in such a way as to offer much more complex operations that ease code parallelization. For a full description of Spark operations, see [12]. Here, we outline the most relevant operations for our algorithm.

- 1) *mapPartitions*: Like Map, this operation independently runs a function on each partition, for each of which an iterator of tuples is fetched and another of the same type is generated.
- 2) *groupByKey*: This operation (using a shuffle operation) groups tuples with the same key in a single vector of values.
- 3) *sortByKey*: This operation is a distributed version of merge-sort.
- 4) *broadcast*: This operation, normally used for large permanent variables (such as big hash tables), keeps a read-only copy of a given variable on each node rather than shipping a copy to each task.

Note that *groupByKey* and *sortByKey* imply a complete redistribution of data across the network. This operation, known as shuffling, is complex and costly and so should be avoided if at all possible.

Finally, created as a Spark subproject was a scalable ML library called MLlib [14], [30], composed of common learning algorithms and statistical utilities. Its main functionalities include classification, regression, clustering, collaborative filtering, optimization, and dimensionality reduction (mostly feature extraction).

III. FEATURE SELECTION FILTERING FOR BIG DATA

An information theory-based framework that includes many common FS filtering algorithms has been proposed by Brown *et al.* [16] that proves that algorithms like minimum redundancy-maximum relevance (mRMR) and other algorithms are special cases of conditional mutual information (CMI) when certain specific independence assumptions are made about both the class and the features (further details below). Here, we demonstrate that these criteria are

²For a exhaustive review of MapReduce and similar programming frameworks (see [28]).

not only a sound theoretical formulation, but also fit well with modern big data platforms and allow us to distribute several FS methods and their complexity across a cluster of machines.

We now describe how we redesigned this framework for a distributed paradigm. Our framework contains a generic implementation of several information theory-based FS methods—including mRMR, conditional MI maximization, and joint MI (JMI)—that, furthermore, have been designed to integrate in the Spark MLlib library. The framework can also be extended with other criteria provided by the user as long as it complies with the guidelines proposed by Brown *et al.* [16].

Below we first briefly present this framework and explain its adaptation to the big data environment. We then describe in detail how the selection process and the underlying information theory operations were implemented in a distributed manner using Spark primitives.

A. Information Theory-Based Filter Methods

Information measures tell us how much information has been acquired by the receiver when sent a message [31]. In predictive learning, we associate the message with the output feature in classification.

A commonly used uncertainty function is MI [32], which measures the amount of information one random variable contains about another, in other words, it expresses the reduction in the uncertainty of one random variable due to knowledge of the other variable

$$\begin{aligned} I(A; B) &= H(A) - H(A|B) \\ &= \sum_{a \in A} \sum_{b \in B} p(a, b) \log \frac{p(a, b)}{p(a)p(b)} \end{aligned} \quad (1)$$

where A and B are two random variables with marginal probability mass functions $p(a)$ and $p(b)$, respectively, $p(a, b)$ is the joint mass function and H is the entropy.

MI can likewise be conditioned to a third random variable. Thus, CMI is denoted as

$$\begin{aligned} I(A; B|C) &= H(A|C) - H(A|B, C) \\ &= \sum_{c \in C} p(c) \sum_{a \in A} \sum_{b \in B} p(a, b, c) \log \frac{p(a, b, c)}{p(a, c)p(b, c)} \end{aligned} \quad (2)$$

where C is a third random variable with marginal probability mass function $p(c)$ and where $p(a, c)$, $p(b, c)$ and $p(a, b, c)$ are the joint mass functions.

Filtering methods are based on a quantitative criterion or index, also known as the relevance index or scoring. This index measures the usefulness of each feature for a specific classification problem. Through the relevance of a feature for the class (self-interaction), we can rank features and select the most relevant ones. However, features can also be ranked using a more complex criterion such as whether it is more redundant than another feature (multi-interaction). For instance, redundant features (variables that carry similar information) can be discarded using the MI criterion [33]

$$J_{mifs}(X_i) = I(X_i; Y) - \beta \sum_{X_j \in S} I(X_i; X_j)$$

where $S \subseteq S_\theta$ is the current set of selected features and β is a weight factor. Considered is the MI between each candidate $X_i \notin S$ and the class. Also introduced is a penalty proportional to the redundancy, calculated as the MI between the current set of selected features and each candidate feature.

A wide range of methods have been described in the literature that are built on these information theory measures. To homogenize use of all the criteria, Brown *et al.* [16] proposed a generic expression that allows multiple information theory criteria to be ensembled in a single FS framework, based on a greedy optimization process that assesses features using a simple scoring criterion. Through certain independence assumptions, many criteria can be transformed as linear combinations of the Shannon entropy terms MI and CMI [32]. In some cases, more complex criteria are expressed as non-linear combinations of these terms (e.g., max or min). For a detailed description of the transformation processes (see [16]). The generic formula proposed by Brown *et al.* [16] is as follows:

$$J = I(X_i; Y) - \beta \sum_{X_j \in S} I(X_j; X_i) + \gamma \sum_{X_j \in S} I(X_j; X_i|Y) \quad (3)$$

where γ represents a weight factor for the conditional redundancy component.

The formula can be divided into three components, representing the relevance of a feature X_i , the redundancy between two features X_i and X_j and the conditional redundancy between two features X_i, X_j and the class Y . Through the aforementioned assumptions, many criteria were rewritten by Brown *et al.* [16] to fit the generic formulation in such a way that all the methods could be implemented using a slight variation on this formula. For a comprehensive list of adapted FS methods (see [16]).

B. Feature Selection Filtering Framework for Big Data

We now describe the proposed FS filtering framework for big data using distributed operations, outlining the main changes made to adapt the classical approach to the new big data environment. We also analyze the implications arising from the distributed implementation of (3) and the complexity arising from parallelization of the core operations of this expression, namely, MI and CMI.

Beyond implementation in Spark, we redesigned Brown's framework [16] by making improvements to the classical approach while maintaining certain features.

- 1) *Columnar Transformation*: The access pattern presented by most FS methods is feature-wise, in contrast to many other ML algorithms, which are instance-wise (they operate on rows). Although this may be considered a minor issue, it can significantly degrade performance since the natural way to compute relevance and redundancy in FS methods is normally via columns. This issue is especially important for distributed frameworks like Spark, where the data partitioning scheme has a significant impact on performance.
- 2) *Broadcasting*: Once all features values have been grouped and partitioned into different partitions, data

Algorithm 1 Main FS Algorithm

Input: D Dataset, an RDD of samples.
Input: ns Number of features to select.
Input: $npart$ Number of partitions to set.
Input: $cindex$ Index of the output feature.
Output: S_θ Index list of selected features
 $D_c \leftarrow \text{columnarTransformation}(D, ns, npart)$
 $ni \leftarrow D.nrows; nf \leftarrow D.ncols$
 $REL \leftarrow \text{computeRelevances}(D_c, cindex, ni)$
 $CRIT \leftarrow \text{initCriteria}(REL)$
 $p_{best} \leftarrow CRIT.max$
 $sfeat \leftarrow \text{Set}(p_{best})$
while $|S| < |S_\theta|$ **do**
 $RED \leftarrow \text{computeRedundancies}(D_c, p_{best}.index)$
 $CRIT \leftarrow \text{updateCriteria}(CRIT, RED)$
 $p_{best} \leftarrow CRIT.max$
 $sfeat \leftarrow \text{addTo}(p_{best}, sfeat)$
end while
 $\text{return}(sfeat)$

shift has to be minimal to avoid superfluous network and CPU usage. If the MI process is performed locally in each partition, the overall algorithm will run efficiently (almost linearly). Data shift is minimized by replicating the output feature and the last selected feature in each iteration.

- 3) *Precomputed Data Caching*: The first term in the generic criterion of (3) is relevance, which basically implies calculating MI for all the input features and the output (relevance). This operation is performed once at the start of our algorithm, then cached to be reused in subsequent evaluations of (3). Likewise, subsequent marginal and joint proportions derived from these operations are also kept so as to omit some computations. This also helps isolate redundancy computation by feature as it replicates permanent information in all nodes.
- 4) *Greedy Approach*: Brown *et al.* [16] proposed a greedy search process in which only one feature is selected in each iteration. This approach transforms the quadratic complexity of typical FS algorithms into a more manageable complexity determined by the number of features to select.

1) *Main FS Algorithm*: Algorithm 1 is the main FS algorithm, in charge of deciding which feature to select in a sequential manner. Roughly speaking, it calculates the initial relevance for all the features, and then iterates to select the best features according to (3) and the underlying MI and CMI values.

The first step consists of transforming the data into columnar format. Once the data matrix is transformed, the algorithm obtains the relevance for each feature in X , initializing the criterion value [partial result according to (3)] and creating an initial ranking of the features. Relevance values are saved as part of the previous expression and are reused in subsequent steps to update the criteria. The most relevant feature, p_{best} , is then selected and added to the set $sfeat$, initially empty. The iterative phase begins by calculating MI and CMI for p_{best} , each candidate X_i and Y . The resulting values update the accumulated redundancies (simple and conditional) of the criteria.

Algorithm 2 Function That Transforms Row-Wise Data Into a Column-Wise Format (*columnarTransformation*)

Input: D Dataset, an RDD of samples.
Input: nf Number of features.
Input: $npart$ Number of partitions to set.
Output: Column-wise data (RDD of feature vectors).
1: $D_c \leftarrow$
2: **map partitions** $part \in D$
3: $matrix \leftarrow \text{new Matrix}(nf)(part.length)$
4: **for** $j = 0$ **until** $part.length$ **do**
5: **for** $i = 0$ **until** nf **do**
6: $matrix(i)(j) \leftarrow part(j)(i)$
7: **end for**
8: **end for**
9: **for** $k = 0$ **until** nf **do**
10: $EMIT < k, (part.index, matrix(k)) >$
11: **end for**
12: **end map**
13: $\text{return}(D_c.sortByKey(npart))$

In each iteration, the most relevant candidate features are selected as the new p_{best} and are added to $sfeat$. The loop ends once ns features (where $ns = |S_\theta|$) have been selected or when no more features remain to be selected.

2) *Distributed Operations (Columnar Transformation and MI Computation)*: MI and CMI are undoubtedly the costliest operations to estimate in information theory-based FS. With huge datasets, these operations become impossible to calculate sequentially as the number of combinations grows. Below we describe how these calculations are parallelized as a set of distributed operations (explained in Section II-C). For all the algorithms described below, RDD variables are written in uppercase to distinguish them from ordinary variables.

a) *Columnar transformation*: As mentioned previously, the column-wise format is much more manageable for FS filtering methods than the row-wise format. Algorithm 2 explains this transformation, carried out as the first step in our algorithm. The idea behind this transformation is to transpose the local data matrix provided by each partition. This operation maintains the partitioning scheme without incurring in a high shuffling overhead. Additionally, once data are transformed, they can be cached and reused in the subsequent loop. The result of this operation is a new matrix, with one row per feature, that generates a tuple, where k represents the feature index, $part.index$ is the index of the partition (henceforth the block index) and $matrix(k)$ is the local matrix for this feature block.

In order to benefit from data locality, the algorithm allocates all instances of the same feature to a specific set of partitions (if possible, to just one partition). To do that, the new instances are sorted by key, with the number of partitions limited to $npart$. In subsequent phases, the partitions are mapped with the aim of generating a number of histograms per feature that count the number of occurrences by combination.

Choosing the right number of partitions is important for the next steps. If $npart$ is equal to or less than the number of features, then the number of total histograms per feature will be two at most; otherwise, the total number of histograms can be high, since the same feature can be distributed across

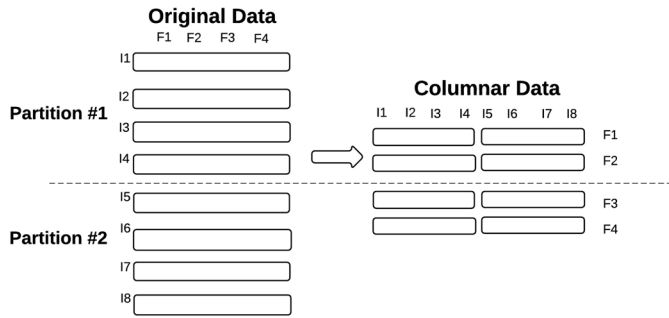


Fig. 1. Columnar transformation scheme. F and I indicate features and instances, respectively. Each rectangle on the left represents a single register in the original dataset and each rectangle on the right represents a transposed feature block in the new columnar format.

Algorithm 3 Compute Mutual Information Between the Set of Features X and Y (*computeRelevances*)

Input: D_C RDD of tuples (index, (block, vector)).

Input: $yind$ Index of Y .

Input: ni Number of instances.

Output: MI values for all input features.

- 1: $ycol \leftarrow D_C.lookup(yind)$
 - 2: $bycol \leftarrow broadcast(ycol)$
 - 3: $counter \leftarrow broadcast(getMaxByFeature(D_C))$
 - 4: $H \leftarrow getHistograms(D_C, yind, bycol, null, null)$
 - 5: $joint \leftarrow getProportions(H, ni)$
 - 6: $marginal \leftarrow getProportions(aggregateByRow(joint), ni)$
 - 7: $return(computeMutualInfo(H, yind, null))$
-

many partitions (more than two). We recommend setting this parameter to, at most, $2 \times$ the number of features.

Fig. 1 depicts this process using a small example with eight instances and four features. This figure shows how the algorithm generates a block for each feature in each partition and how all blocks are sorted by feature in order to place them in the same partitions.

b) Computing relevance: Algorithm 3 describes how to compute relevance (MI) for all the input features and Y [as expressed in (1)]. The design as an initialization method means that all variables in this function can be used in the subsequent algorithms. For example, the number of different values for each feature is first computed and saved as *counter* (to limit the size of histograms).

The main idea behind relevance and redundancy functions is that calculations for each feature are performed independently. This is done by distributing only the single variables (p_{best} and Y) across the cluster and leveraging for the data locality property. The first step consists of collecting all blocks of Y from the data and putting them in a single vector for broadcasting (*bycol*). Histograms for all the candidate features with respect to Y are then calculated in *getHistograms* (explained below). This function is common to both the relevance and redundancy phases. 3-D histograms are computed between all the non-selected features and two secondary variables for redundancy, and between all the nonselected features and one variable for relevance.³ Joint and marginal proportions are generated from

³For relevance, the *null* value is used to represent the lack of the second variable.

Algorithm 4 Compute CMI and MI Between p_{best} , the Set of Candidate Features, and Y (*computeRedundancies*)

Input: D_C RDD of tuples (index, (block, vector)).

Input: $jind$ Index of p_{best} .

Output: CMI values for all input features.

- 1: $jcol \leftarrow D_C.lookup(jind)$
 - 2: $bjcol \leftarrow broadcast(jcol)$
 - 3: $H \leftarrow getHistograms(D_C, jind, bjcol, yind, bycol)$
 - 4: $return(computeMutualInfo(H, jind, yind))$
-

Algorithm 5 Function That Computes 3-D Histograms Between p_{best} , the Set of Candidate Features, and Y for CMI; or Between the Set of Candidate Features, and Y for MI (*getHistograms*)

Input: D_C RDD of tuples (index, (block, vector)).

Input: $jind$ Index of Y or p_{best} .

Input: $yind$ Index of feature Y (can be empty).

Input: $jcol$ Values for Y or p_{best} , a broadcast matrix.

Input: $ycol$ Values for Y , a broadcast matrix (can be empty).

Output: Column-wise dataset (RDD of feature vectors).

- 1: $jsize \leftarrow counter(jind)$
 - 2: $ysize \leftarrow counter(yind)$
 - 3: $H \leftarrow$
 - 4: **map partitions** $part \in partitions$
 - 5: **for** $(k, (block, v)) \leftarrow part$ **do**
 - 6: $isize \leftarrow counter(k)$
 - 7: $m \leftarrow newMatrix(ysize)(isize)(jsize)$
 - 8: **for** $e = 0$ **until** $v.size$ **do**
 - 9: $j \leftarrow jcol(block)(e)$; $y \leftarrow ycol(block)(e)$; $i \leftarrow v(e)$
 - 10: $m(y)(i)(j) += 1$
 - 11: **end for**
 - 12: $EMIT < k, m >$
 - 13: **end for**
 - 14: **end map**
 - 15: $return(H.reduceByKey(sum))$
-

the resulting histograms using matrix operations, that is, by aggregating proportions by row (*marginal*) and by computing joint for joint proportions (*joint*). Using this information we can now obtain the MI value for each candidate feature.

c) Computing redundancy: Simple and conditional redundancy are computed between p_{best} , each candidate feature X_i and Y . Conditional redundancy introduces a third conditional variable (Y), following the formula $I(X_j; X_i|Y)$ (2).

The operation is repeated until we obtain the number of features specified as a parameter. Algorithm 4, an extension of relevance computation (Algorithm 3), depicts this process. Blocks for p_{best} are obtained from the RDD and broadcast to all the nodes. The *getHistograms* function is called up with two variables in order to obtain the histograms for all the candidate features with respect to p_{best} and Y . Note that the vector for Y is already available from the redundancy computation phase. Finally, both types of redundancy are computed using the function that computes MI and CMI (*computeMutualInfo*).

d) Histograms creation: Algorithm 5 computes 3-D histograms for the set of candidate features with respect to p_{best} and Y and subsequently computes MI and CMI. When no conditional variable is provided, this yields histograms whose third dimension is equal to one.

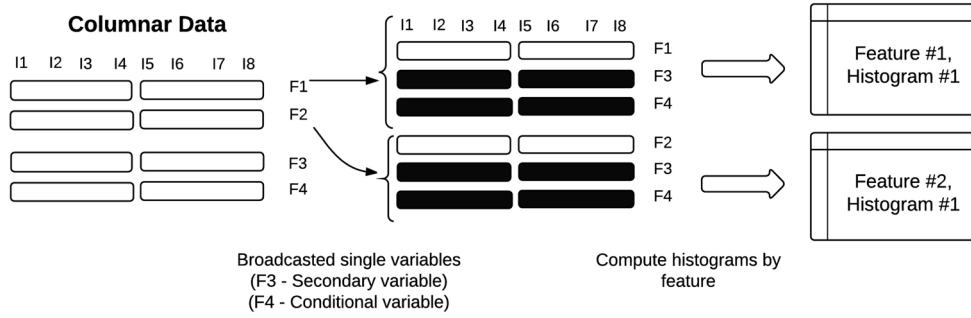


Fig. 2. Histogram creation scheme. F and I indicate features and instances, respectively. The white rectangles each represent a single feature block in the columnar format and the black rectangles represent marginal and joint proportions for single variables, broadcast across the cluster.

The first two lines return the dimensions for the p_{best} and Y variables using *counter* (Algorithm 1). A mapping operation on each partition is launched on the dataset that iterates over the blocks derived from the columnar transformation. Each tuple is formed of a key (the index of a candidate feature) and a value with an index block and the corresponding feature array $[(k, (block, v))]$. For each instance, a matrix is initialized to zero and then incremented by one, according to the value for each combination of p_{best} , X_i , and Y . Single features (p_{best} and Y) are broadcast in the form of matrices, whose first and second axes indicate, respectively, the block index and the partial index of the value in this block. This updating operation is repeated to the end of the feature vector, after which a new tuple is emitted with the feature index as the key and the resulting matrix as the value ($\langle k, m \rangle$). The mapping operation continues with ensuing blocks until the partition is finished. The final histograms are then aggregated by summation. This aggregation process will remain simple as long as the number of histograms per partition is small—which is normally the case, as a single partition usually contains all the blocks for the same feature. This is why it is important, as explained above, to reduce the number of histograms by adjusting the number of partitions.

Fig. 2 depicts the histogram creation process using a simple example, showing how the algorithm generates one histogram for each partition and feature. In this case, the number of partitions corresponds to the number of features, so only one histogram is generated per feature.

e) MI and CMI computations: Algorithm 6 details the process that unifies computation of MI and CMI. The algorithm takes as input the indices of single variables (Y or p_{best} for MI, and both for CMI) and all previously computed histograms. Before launch, the algorithm broadcasts the marginal and joint matrices that correspond to these variables and sends this information to the nodes, as this information is already computed and cannot be computed independently from the previous histograms.

A mapping phase is then launched for each histogram tuple, consisting of a given feature index as the key and a 3-D matrix as the value. The algorithm generates the MI and CMI values for all combinations of histograms and single variables [see (1) and (2)]. First, however, it is necessary to compute the marginal proportions for the set of candidate features as well as the joint proportions between each X_i and Y and

Algorithm 6 Calculate MI and CMI for the Set of Histograms With Respect to Y or p_{best} for MI, and Both Variables for CMI (*computeMutualInfo*)

Input: H Histograms, an RDD of tuples (index, matrix).

Input: *bind* Index of feature p_{best} .

Input: *cind* Index of feature Y (can be empty).

Output: MI and CMI values.

```

1:  $bprob \leftarrow broadcast(marginal(bind))$ 
2:  $cprob \leftarrow broadcast(marginal(cind))$ 
3:  $bcprob \leftarrow broadcast(joint(bind))$ 
4:  $MINFO \leftarrow$ 
5: map  $(k, m) \in H$ 
6:    $aprob \leftarrow computeMarginal(m)$ 
7:    $abprob \leftarrow computeJoint(m, bind)$ 
8:    $acprob \leftarrow computeJoint(m, cind)$ 
9:   for  $c = 0$  until  $getSize(m)$  do
10:    for  $b = 0$  until  $getnRows(m(c))$  do
11:     for  $a = 0$  until  $getnCols(m(c))$  do
12:       $pc \leftarrow cprob(c)$ 
13:       $pabc \leftarrow (m(c)(a)(b)/ninstances)/pc$ 
14:       $pac \leftarrow acprob(c)(a); pbc \leftarrow bcprob(c)(b)$ 
15:       $cmi += conditionalMutualInfo(pabc, pac, pbc, pc)$ 
16:      if  $c == 0$  then
17:         $pa \leftarrow xprob(a)$ 
18:         $pab \leftarrow abprob(a)(b); pb \leftarrow yprob(b)$ 
19:         $mi += mutualInfo(pa, pab, pb)$ 
20:      end if
21:    end for
22:  end for
23: end for
24:  $EMIT < k, (mi, cmi) >$ 
25: end map
26: return  $(MINFO)$ 

```

each X_i and p_{best} (using the matrix operations described in Algorithm 3). Once all joint and marginal proportions are calculated, a loop starts over all combinations to compute the proportion $pabc$ (which comes directly from the histogram) and the final result for each combination. These results are then aggregated to obtain the overall MI and CMI values for each feature.

Fig. 3 depicts the MI process that is launched once all histograms have been computed. Each partition generates a histogram for each feature it contains. Histograms for the same feature are aggregated to obtain a single final histogram. Only marginal and joint proportions that cannot be computed independently from each histogram are broadcast.

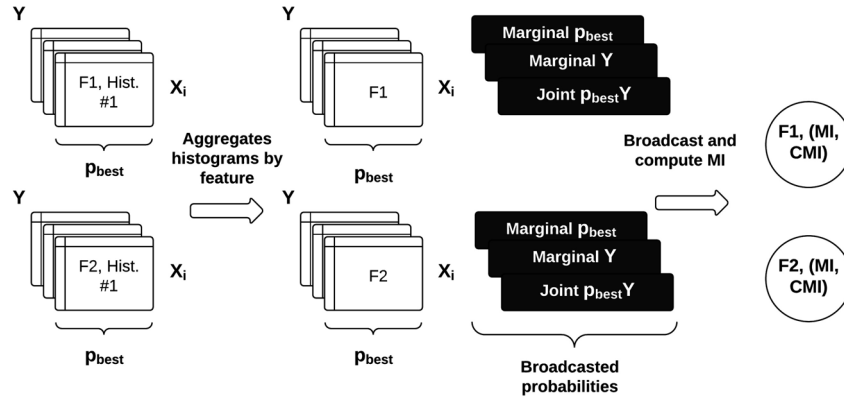


Fig. 3. MI computations. F indicates features and the matrices depict the histograms for each feature. The black rectangles represent broadcast joint and marginal probabilities.

Algorithm 7 Function That Transforms Row-Wise Data Into Column-Wise Data (Sparse Version) (*sparseColumnar*)

Input: D Dataset, an RDD of sparse samples.

Input: $npart$ Number of partitions to set.

Output: Column-wise dataset (feature vectors).

```

1:  $D_c \leftarrow$ 
2: map  $reg \in D$ 
3:    $ireg \leftarrow reg.index$ 
4:   for  $i = 0$  until  $reg.length$  do
5:      $EMIT < reg(i).key, (ireg, reg(i).value) >$ 
6:   end for
7: end map
8:  $D_c \leftarrow D_c.groupByKey(npart).mapValues(vectorize)$ 

```

Using the histograms and broadcast variables, MI and CMI are independently computed per feature.

3) *High-Dimensional and Sparse Versions*: The above algorithms work well with tall-and-skinny data, formed of a small number of features and a large number of instances. However, complexity grows horizontally for sparse datasets characterized by a large number of features and a variable number of instances with an undefined number of nonzero indexed elements. The algorithms affected by this issue are columnar transformation (Algorithm 2) and histogram creation (Algorithm 5). The remaining code remains unchanged except that the structure of data is reduced to a single vector [in the form $(index, vector)$]. Since only one histogram is generated per feature, the block index is excluded from this structure.

Algorithm 7 describes a new model of processing that transposes data directly, generating single vectors for each feature. Sparsity is maintained in sparse vectors while the index of the original instance is used as the key. The algorithm generates a tuple for each value, in such a way that the key is formed of the feature index and the value is formed of the instance index and the value itself. Tuples are all grouped by key to create a single sparse vector (formed of sorted key-value tuples).

As for histogram creation, the partition mapping operation is now replaced by a mapping operation applied to each previously generated feature vector. Algorithm 8, which describes this process, is similar to Algorithm 5, with the caveat that only

Algorithm 8 Function That Computes 3-D Histograms for the Set of Features X_i With Respect to Features p_{best} and Y (Sparse Version) (*sparseHistograms*)

Input: D_c Dataset, an RDD of tuples (Int, (Int, Vector)).

Input: $jind$ Index of Y or p_{best} .

Input: $yind$ Index of feature Y (can be empty).

Input: $jcol$ Values for Y or p_{best} , a broadcast matrix.

Input: $ycol$ Values for Y , a broadcast matrix (can be empty).

Output: Column-wise dataset (RDD of feature vectors).

```

1:  $jsize \leftarrow counter(jind)$ ;  $ysize \leftarrow counter(yind)$ 
2:  $jyhist \leftarrow frequencyMap(jind, yind)$ 
3:  $zhist \leftarrow frequencyMap(yind)$ 
4:  $H \leftarrow$ 
5: map  $(k, v) \in D_c$ 
6:    $isize \leftarrow counter(k)$ 
7:    $m \leftarrow newMatrix(ysize)(isize)(jsize)$ 
8:   for  $e = 0$  until  $v.size$  do
9:      $j \leftarrow jcol(e)$ ;  $y \leftarrow ycol(e)$ 
10:     $i \leftarrow v(e)$ 
11:    if  $j <> 0$  then
12:       $jyhist(j)(y) = jyhist(j)(y) - 1$ 
13:    end if
14:     $m(y)(i)(j) += 1$ 
15:  end for
16:  for  $((j, y), q) \leftarrow jyhist$  do
17:     $m(y)(0)(j) += q$ 
18:  end for
19:  for  $(y, q) \leftarrow yhist$  do
20:     $m(y)(0)(0) += yhist(y) - sum(mat(y))$ 
21:  end for
22:   $EMIT < k, m >$ 
23: end map
24:  $return(H)$ 

```

one histogram is yielded for each feature; hence, no reducing operation is necessary.

Unlike what happens in the dense version, to avoid visiting all possible sparse feature combinations, the matrix generation process is adapted, using, as much as possible, accumulators to compute combinations formed by zeros. As accumulators, the algorithm calculates the class histogram for the conditional variable and the joint class histogram for the parametric variables $jind$ and $yind$. A loop is first started for those combinations in which the first variable (i) is not equal to zero (the procedure is the same as in the dense version). If the second

variable (j) is equal to zero, then the frequency counter (in the joint histogram) for the combination is reduced by one. If j is not equal to zero, then the algorithm completes the matrix with the frequencies in the joint histogram. Finally, when both variables (j and y) are equal to zero, the matrix is updated with the remaining occurrences for all classes.⁴ The outcome is the feature index and the aforementioned matrix.

4) *Algorithm Complexity*: As mentioned earlier, the FS algorithm performs a greedy search which stops once the condition defined as input is reached. Beyond this point, the sequential algorithm is influenced by the set of distributed algorithms/operations as presented above. The distributed primitives used in these algorithms need to be analyzed to check the complexity of the full proposal. Note that the first operation (columnar transformation) is quite time-consuming as it makes great use of the network and memory when shuffling the data (wide dependency). However, once data have a known partition, they can be reused in subsequent phases (leveraging the data locality property). This transformation is performed once at the start or can even be omitted if data are already in columnar format. The list of distributed operations in this algorithm are described as follows.

- 1) *Algorithm 2*: This algorithm starts with a *mapPartitions* operation that transposes the local matrix contained in the partition and emits a tuple for each feature (in a linear distributed order). The total number of tuples is n multiplied by the number of original partitions. These tuples are then shuffled across the cluster and each subset is locally sorted (in a log-linear distributed order).
- 2) *Algorithms 3 and 4*: The first operation for both algorithms is to retrieve a single column (feature) using the *lookup* primitive (in a linear distributed order). Since the data are already partitioned, the operation efficiently only looks at the right partition. This variable is then broadcast to all the nodes, sending a single feature (m values) across the network. The following operations (histograms and MI computations) are described below.
- 3) *Algorithm 5*: This algorithm is a simple *mapReduce* operation whereby the previously generated tuples are transformed to local histograms and then reduced to the final histograms by feature. This mapping operation consists of a linear function ($O(m)$) that fetches the data contained in each local matrix.
- 4) *Algorithm 6*: This operation starts by broadcasting three single values (proportion values). For each feature, three linear operations are launched to compute extra probabilities. The MI values are then computed by fetching the entire 3-D-histogram (in a cubic linear order). Note that the complexity of all these operations is bounded by the cardinality of the features included in the histogram.

IV. EXPERIMENTAL FRAMEWORK AND ANALYSIS

This section describes experiments that evaluate the proposed FS framework applied to a set of real-world problems that are huge in terms of both features and instances.

⁴The class vector is always dense.

TABLE I

SUMMARY DESCRIPTION OF FIVE DATASETS, INDICATING THE NUMBER OF EXAMPLES FOR TRAINING AND TEST SETS (#TRAIN EX., #TEST EX.), THE TOTAL NUMBER OF ATTRIBUTES (#ATTS.), THE NUMBER CLASSES (#CL), AND THE SPARSITY CONDITION (SPARSE)

Data Set	#Train Ex.	#Test Ex.	#Atts.	#Cl.	Sparse
epsilon	400 000	100 000	2000	2	No
dna	79 739 293	10 000 000	200	2	No
ECBDL14	65 003 913	2 897 917	630	2	No
url	1 916 904	479 226	3 231 961	2	Yes
kddb	19 264 097	748 401	29 890 095	2	Yes

A. Datasets and Methods

We used five classification datasets to measure the quality and usefulness of our FS framework implementation for Spark. We classified these datasets into two groups: 1) dense (large number of samples) and 2) sparse (high-dimensional datasets). For the sparse datasets, the high-dimension version described in Section III-B3 was used.

The first dataset ECBDL14, used as a reference dataset at the international GECCO-2014 conference, consists of 631 features (including both numerical and categorical attributes) and 32 million instances. In this binary classification problem the class distribution is imbalanced, with 98% of negative instances. To this imbalanced problem, we applied the MapReduce version of the Random OverSampling (ROS) algorithm [34] (henceforth we will use ECBDL14 to refer to the ROS version). Another dataset used was dna, consisting of 50 000 000 instances and 201 discrete features and created in 2008 for the pascal large scale learning challenge.⁵ Only the training set was used for our experiments; since the test set does not contain the class labels, the training set was used to generate both subsets (using an 80/20 hold-out data split). The ROS technique was also applied to this dataset (henceforth dna) since this problem is also imbalanced between classes. The remaining datasets (epsilon, url, and kddb) come from the LibSVM dataset repository [1]. These datasets and their descriptions can be found in the project's website.⁶ Table I provides summary details of these datasets.

As an FS benchmark method, we used the mRMR algorithm [35] since it is one of the most cited selectors in the literature. Note that the FS algorithm chosen to test performance did not affect the time results yielded by the framework since all criteria were computed in the same way.

For the comparison study, we used support vector machines (SVM) [36], and Naive Bayes [37] as the classifiers. For our experiments, we used the distributed versions of these algorithms implemented in the MLlib library [14]. The parameters of the classifiers, as recommended in the authors' specifications [14], are shown in Table II. For all executions, the datasets were cached in memory as SVM and our method used iterative processes. The level of parallelism (number of

⁵<http://largescale.ml.tu-berlin.de/summary/>

⁶<http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

TABLE II
PARAMETERS OF THE USED CLASSIFIERS

Method	Parameters
Naive Bayes	lambda = 1.0
SVM	stepSize = 1.0, batchFraction = 1.0, regularization = 1.0, iterations = 100
mRMR	level of parallelism = 864

partitions) was set to 864, twice the total number of cores available in the cluster.⁷

A Spark package associated with this paper can be found in the third-party Spark Repository: <http://spark-packages.org/package/sramirez/spark-infotheoretic-feature-selection>. This software was designed for integration in the MLlib Library and has an associated JIRA issue to discuss its integration in this library: <https://issues.apache.org/jira/browse/SPARK-6531>.

We used two common evaluation metrics to assess the quality of the resulting FS schemes: area under the receiver operating characteristic curve (AUROC, henceforth AUC) to evaluate classifier accuracy, and training modeling time to evaluate FS performance.

B. Cluster Configuration

For all the experiments we used a cluster composed of eighteen computing nodes and one master node. The computing nodes had the following characteristics: two processors x Intel Xeon CPU E5-2620, 6 cores per processor, 2.00 GHz, 15 MB cache, QDR InfiniBand Network (40 Gb/s), 2 TB HDD, 64 GB RAM. We used the following configuration for the software: Hadoop 2.5.0-cdh5.3.1 from Cloudera's open-source Apache Hadoop distribution,⁸ HDFS replication factor 2, HDFS default block size 128 MB, Apache Spark and MLlib 1.2.0, 432 cores (24 cores/node), 864 RAM GB (48 GB/node).

Both the HDFS and Spark master processes (the HDFS NameNode and the Spark Master) were hosted in the main node. The NameNode controlled the HDFS and coordinated the slave machines by means of their respective DataNode daemons. The Spark Master controlled all the executors in each worker node. Spark used the HDFS file system to load and save data in the same way as the Hadoop framework.

C. Results Analysis

For the evaluation of the time employed by our implementation to rank the most relevant features, Table III presents the results obtained by our algorithm using different ranking thresholds (number of selected features).

As can be observed, our algorithm yielded competitive results in all cases, irrespective of the number of iterations (represented by the threshold value). Regarding the datasets with the highest data volumes, namely, kddb (ultra-high-dimensionality) and ECBDL14 (with a huge

⁷The Spark creators recommend using 2-4 partitions per core: <http://spark.apache.org/docs/latest/programming-guide.html#parallelized-collections>.

⁸<http://www.cloudera.com/content/cloudera/en/documentation/cdh5/v5-0-0/CDH5-homepage.html>

TABLE III
SELECTION TIME BY DATASET AND THRESHOLD (IN SECONDS)

# Features	kddb	url	dna	ECBDL14	epsilon
10	283.61	94.06	97.83	332.90	111.42
25	774.43	186.22	148.78	596.31	173.39
50	1365.82	333.70	411.84	1084.58	292.07
100	2789.55	660.48	828.35	2420.94	542.05

TABLE IV
SELECTION TIMES (SEQUENTIAL VERSUS DISTRIBUTED APPROACHES, IN SECONDS)

# Examples	Sequential	Distributed	Speedup
1 000 000	450.00	425.56	1.06
5 000 000	2 839.72	508.41	5.59
40 000 000	23 749.77	828.35	28.67
80 000 000	47 646.97	1 597.26	29.83

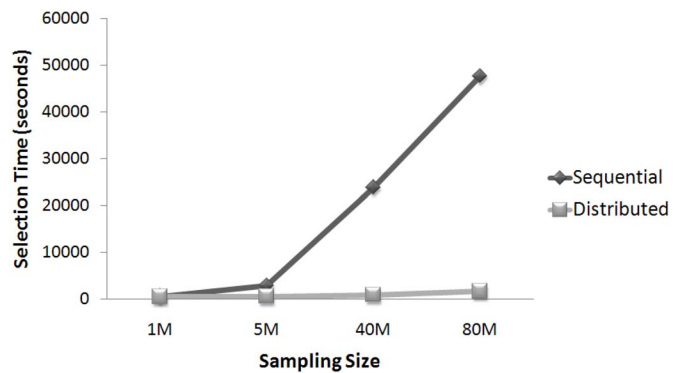


Fig. 4. Selection times (sequential versus distributed approaches).

number of samples), our method was able to rank 100 features in under 60 min.

We conducted a comparison study between our distributed version and the sequential version developed by Brown's laboratory.⁹ Samples from dna were generated with different ratios of instances in order to study the scalability of our approach in comparison with the sequential version.¹⁰ The level of parallelism was set to 200 in the distributed executions as so to facilitate comparison between the distributed version with one core per feature and the sequential version with just a single core.

Table IV and Fig. 4 show the time results for our distributed versions compared to the sequential version. In the latter, the last two values (highlighted in italics) were estimated using linear interpolation since they could not be computed due to memory problems. Table IV indicates that our distributed version outperformed the sequential approach in all cases. What was especially remarkable was that the largest dataset achieved the maximum speedup rate (29.83).

⁹FEAST toolbox (python version): <http://www.cs.man.ac.uk/~gbrown/software/>.

¹⁰The sequential version was executed in one node of our cluster with the aforementioned characteristics.




Fig. 5. Selection time by number of cores in the distributed approach (in s).




Fig. 7. AUC results for SVM using different thresholds.




Fig. 6. AUC results for NaiveBayes using different thresholds.




Fig. 8. NaiveBayes classification times for the training phase using different thresholds (in s).

Finally, as an additional study of scalability we varied the number of cores, using ECBDL14 (the largest dense dataset) as a reference and using the same parameters as in the previous study. Fig. 5, which depicts how our method performed depending on the number of cores (10 to 100), reveals logarithmic behavior as the number of cores increased. Note that for just ten cores not all memory was available.

D. Classification Results Analysis

We also analyzed the usefulness of our FS solution when applied to large-scale classification. Figs. 6 and 7 show the accuracy results for SVM and Naive Bayes using different FS schemes. All the datasets described in Table I were used in this paper except kddb because those classifiers were not designed to work with such great dimensionality.

Fig. 7 points to an important improvement when FS was used for url and epsilon. In contrast, the fact that the impact for dna and ECBDL14 was negligible can be explained by their high imbalance ratio or small number of features. Fig. 6 presents similar results, except that the improvement for the url dataset was much smaller.

Beyond AUC, the time employed to create a classification model is quite important in many large-scale problems. Figs. 8 and 9 show classification times for the training phase for different datasets and thresholds. The results demonstrate that the simplicity and performance of the generated models was improved after applying FS, most especially for SVM, which requires more time for modeling than Naive Bayes.

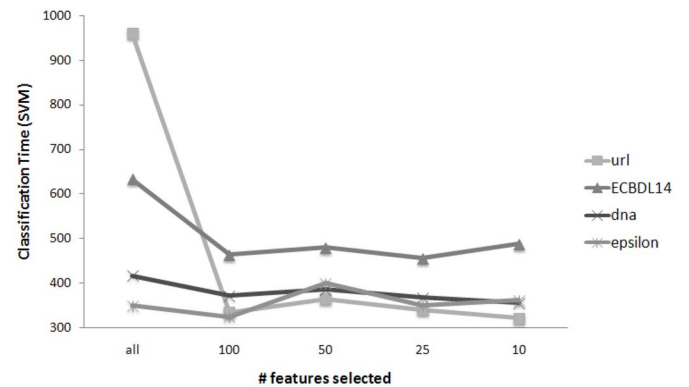


Fig. 9. SVM classification times for the training phase using different thresholds (in s).

The performance results demonstrate that our solution is capable of selecting features in a competitive time interval when applied to datasets that are huge- in both number of instances and features. The results also demonstrate the superiority of our distributed approach to FS over the sequential approach.

Furthermore, using our selection schemes, the classifiers yield better results in most cases, and at least similar results in the other cases. Note that, in all the studied cases, our model was much simpler and faster despite using a small percentage of the original set of features.

V. CONCLUSION

In discussing the problem of processing big data, especially from the perspective of dimensionality, we have highlighted the impact of correctly identifying relevant features in datasets and the corresponding difficulties caused by the combinatorial effects of incoming data growing in terms of both instances and features. Despite the growing interest in dimensionality reduction for big data, few FS methods have been developed to date that can adequately deal with high-dimensionality problems.

Adopting an information theory approach, we adapted a generic FS framework for big data from a proposal by Brown *et al.* [16]. The framework contains implementations of many state-of-the-art FS algorithms, including mRMR and JMI. The adaptation entailed a radical redesign of Brown *et al.*'s framework so as to adapt it to a distributed paradigm. This paper has also resulted in contributing an FS module to the emerging Spark and MLLib platforms, which, to date, included no complex FS algorithm.

Our experimental results demonstrate the usefulness of our FS solution applied to a broad set of large real-world problems. Our solution performed well with 2-D of big data (samples and features) and yielded competitive performance results in dealing with both ultrahigh-dimensionality datasets and datasets with a huge number of samples. Our distributed approach consistently outperformed the sequential version, enabling the resolution of problems that could not be usefully resolved using the classical approach.

Future research will focus on the following.

- 1) Designing new information theory-based approaches for high-speed data streams and also extending these approaches to the handling of drifts in concepts.
- 2) Analyzing the impact of approximative selection on high-dimensional data via faster solutions that do not incur a high penalty on accuracy.
- 3) Designing a new fully automatic FS system that selects the most relevant subset of features from a full set of features, thereby eliminating the need to define the number of features to select in each execution.

REFERENCES

- [1] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, pp. 1–27, 2011. [Online]. Available: <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>
- [2] I. Guyon, S. Gunn, M. Nikravesh, and L. A. Zadeh, *Feature Extraction: Foundations and Applications* (Studies in Fuzziness and Soft Computing). Secaucus, NJ, USA: Springer-Verlag, 2006.
- [3] Z. A. Zhao and H. Liu, *Spectral Feature Selection for Data Mining*. Boca Raton, FL, USA: CRC Press, 2011.
- [4] V. Bolón-Canedo, N. Sánchez-Marño, and A. Alonso-Betanzos, *Feature Selection for High-Dimensional Data* (Artificial Intelligence: Foundations, Theory, and Algorithms). New York, NY, USA: Springer, 2015.
- [5] Q. Wu, Z. Wang, F. Deng, Z. Chi, and D. D. Feng, "Realistic human action recognition with multimodal feature selection and fusion," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 43, no. 4, pp. 875–885, Jul. 2013.
- [6] V. Bolón-Canedo, N. Sánchez-Marño, A. Alonso-Betanzos, J. M. Benítez, and F. Herrera, "A review of microarray datasets and applied feature selection methods," *Inf. Sci.*, vol. 282, pp. 111–135, Oct. 2014.
- [7] V. Bolón-Canedo, N. Sánchez-Marño, and A. Alonso-Betanzos, "Recent advances and emerging challenges of feature selection in the context of big data," *Knowl. Based Syst.*, vol. 86, pp. 33–45, Sep. 2015.
- [8] J. Dean and S. Ghemawat, "Mapreduce: Simplified data processing on large clusters," in *Proc. OSDI*, San Francisco, CA, USA, 2004, pp. 137–150.
- [9] T. White, *Hadoop, The Definitive Guide*. Sebastopol, CA, USA: O'Reilly Media, 2012.
- [10] Apache Hadoop Project. (2017). *Apache Hadoop*. Accessed on Feb. 2017. [Online]. Available: <http://hadoop.apache.org/>
- [11] H. Karau, A. Konwinski, P. Wendell, and M. Zaharia, *Learning Spark: Lightning-Fast Big Data Analytics*. Beijing, China: O'Reilly Media, 2015.
- [12] A. Spark. (2017). *Apache Spark: Lightning-Fast Cluster Computing*. Accessed on Feb. 2017. [Online]. Available: <https://spark.apache.org/>
- [13] Apache Mahout Project. (2017). *Apache Mahout*. Accessed on Feb. 2017. [Online]. Available: <http://mahout.apache.org/>
- [14] A. Spark. (2017). *Machine Learning Library (MLlib) for Spark*. [Online]. Accessed on Feb. 2017. Available: <http://spark.apache.org/docs/latest/mllib-guide.html>
- [15] K. Sun, W. Miao, X. Zhang, and R. Rao, "An improvement to feature selection of random forests on spark," in *Proc. IEEE 17th Int. Conf. Comput. Sci. Eng. (CSE)*, Chengdu, China, 2014, pp. 774–779.
- [16] G. Brown, A. Pocock, M.-J. Zhao, and M. Luján, "Conditional likelihood maximisation: A unifying framework for information theoretic feature selection," *J. Mach. Learn. Res.*, vol. 13, pp. 27–66, Jan. 2012.
- [17] A. L. Blum and P. Langley, "Selection of relevant features and examples in machine learning," *Artif. Intell.*, vol. 97, nos. 1–2, pp. 245–271, 1997.
- [18] R. Kohavi and G. H. John, "Wrappers for feature subset selection," *Artif. Intell.*, vol. 97, nos. 1–2, pp. 273–324, 1997.
- [19] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, Mar. 2003.
- [20] Y. Saeys, I. N. Inza, and P. Larrañaga, "A review of feature selection techniques in bioinformatics," *Bioinformatics*, vol. 23, no. 19, pp. 2507–2517, 2007.
- [21] V. Mayer-Schönberger and K. Cukier, *Big Data: A Revolution That Will Transform How We Live, Work and Think*. London, U.K.: John Murray, 2013.
- [22] D. Laney. (2001). *3D Data Management: Controlling Data Volume, Velocity and Variety*. Accessed on Feb. 2017. [Online]. Available: <http://blogs.gartner.com/doug-laney/files/2012/01/ad949-3D-Data-Management-Controlling-Data-Volume-Velocity-and-Variety.pdf>
- [23] Y. Zhai, Y.-S. Ong, and I. W. Tsang, "The emerging 'big dimensionality,'" *IEEE Comput. Intell. Mag.*, vol. 9, no. 3, pp. 14–26, Aug. 2014.
- [24] M. Lichman. (2013). *UCI Machine Learning Repository*. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [25] W. Roush, *TR10: Peering Into Video's Future*, MIT Technol. Rev., Cambridge, MA, USA, Mar. 2013. [Online]. Available: <http://www.technologyreview.com/Infotech/18284>
- [26] Q. Mao and I. W.-H. Tsang, "Efficient multitemplate learning for structured prediction," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 24, no. 2, pp. 248–261, Feb. 2013.
- [27] T. Wang *et al.*, "FD4C: Automatic fault diagnosis framework for Web applications in cloud computing," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 46, no. 1, pp. 61–75, Jan. 2016.
- [28] A. Fernández *et al.*, "Big data with cloud computing: An insight on the computing environment, mapreduce, and programming frameworks," *Wiley Interdiscipl. Rev. Data Min. Knowl. Disc.*, vol. 4, no. 5, pp. 380–409, 2014.
- [29] J. J. Lin, "MapReduce is good enough? If all you have is a hammer, throw away everything that's not a nail!" *Big Data*, vol. 1, no. 1, pp. 28–37, 2013.
- [30] X. Meng *et al.*, "Mllib: Machine learning in apache spark," *J. Mach. Learn. Res.*, vol. 17, no. 34, pp. 1–7, 2016.
- [31] M. Michael and W.-C. Lin, "Experimental study of information measure and inter-intra class distance ratios on feature selection and orderings," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-3, no. 2, pp. 172–181, Mar. 1973.
- [32] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2006.
- [33] R. Battiti, "Using mutual information for selecting features in supervised neural net learning," *IEEE Trans. Neural Netw.*, vol. 5, no. 4, pp. 537–550, Jul. 1994.
- [34] S. del Río, V. López, J. M. Benítez, and F. Herrera, "On the use of mapreduce for imbalanced big data using random forest," *Inf. Sci.*, vol. 285, pp. 112–137, Nov. 2014.

- [35] H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 8, pp. 1226–1238, Aug. 2005.
- [36] M. A. Hearst, S. T. Dumais, E. Osman, J. Platt, and B. Scholkopf, "Support vector machines," *IEEE Intell. Syst. Appl.*, vol. 13, no. 4, pp. 18–28, Jul./Aug. 1998.
- [37] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*, vol. 3. New York, NY, USA: Wiley, 1973,



Sergio Ramírez-Gallego received the M.Sc. degree in computer science from the University of Jaén, Jaén, Spain, in 2012. He is currently pursuing the Ph.D. degree with the Department of Computer Science and Artificial Intelligence, University of Granada, Spain.

His current research interests include data mining, data preprocessing, big data, and cloud computing.



Héctor Mouriño-Talín received the B.S. degree in computer science from the University of Corunna, A Coruña, Spain, in 2014, and the M.Sc. degree in machine learning from University College London, London, U.K., in 2015.

He currently is a Data Scientist optimizing ad targeting in real time. His current research interests include artificial intelligence and machine learning and, more specifically, reinforcement learning and deep learning and their applications in fields such as bioinformatics or real-time decision making.



David Martínez-Rego received the M.S. degree in computer science from the University of Corunna, A Coruña, Spain, in 2008.

He has been a Research Staff Member, with an FPU grant from the Ministerio de Ciencia e Innovación, with the Department of Computer Science, University of Corunna, since 2009. His current research interests include theoretical works on neural networks, predictive maintenance, and learning optimization.



Verónica Bolón-Canedo received the B.S., M.S., and Ph.D. degrees in computer science from the University of Corunna, A Coruña, Spain, in 2009, 2010, and 2014, respectively.

After a Post-Doctoral Fellowship with the University of Manchester, Manchester, U.K., in 2015, she is currently an Assistant Professor with the Department of Computer Science, University of Corunna. She has extensively published in the area of machine learning and feature selection. She has co-authored one book, one book chapter, and over 50

research papers in international conferences and journals in the above areas.



José Manuel Benítez (M'98) received the M.S. and Ph.D. degrees in computer science from the University of Granada, Granada, Spain.

He is currently an Associate Professor with the Department of Computer Science and Artificial Intelligence, University of Granada. He has published in journals such as the IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS, *Fuzzy Sets and Systems*, the *Journal of Statistical Software*, *Information Sciences*, *Mathware and Soft Computing*, *Neural Networks*, *Applied Intelligence*, the *Journal of Intelligent Information Systems*, the *Artificial Intelligence in Medicine*, *Expert Systems With Applications*, *Computer Methods and Programs in Biomedicine*, *Evolutionary Intelligence*, and *Econometric Reviews*. His current research interests include time series analysis and modeling, distributed/parallel computational intelligence, data mining, and statistical learning theory.



Amparo Alonso-Betanzos (M'87–SM'06) received the Ph.D. degree from the University of Santiago de Compostela, Santiago, Spain, in 1988.

She moved to the Medical College of Georgia, Augusta, GA, USA, for a Post-Doctoral position until 1990. She is currently a Professor with the Department of Computer Science, University of Corunna, A Coruña, Spain. She has published over 11 book chapters, and over 200 articles in journals and international conferences. Her current research interests include decision-making, data science, and feature selection in large-scale datasets.

Dr. Alonso-Betanzos was a recipient of the L'Oréal-UNESCO Award for "Women in Science" in Spain in 1998. She has coordinated and/or participated in over 50 competitive projects at the European, national and regional levels, and on transference projects with different companies. She is an ACM Member and the President of the Spanish Association for Artificial Intelligence.



Francisco Herrera (SM'15) received the M.Sc. and Ph.D. degrees from the University of Granada, Granada, Spain, in 1988 in 1991, respectively, both in mathematics.

He is currently a Professor with the Department of Computer Science and Artificial Intelligence, University of Granada. He has been the supervisor of 40 Ph.D. students. He has published over 300 journal papers that have received over 50 000 Scholar Google citations and an H-index of 112. He has co-authored the books entitled *Genetic Fuzzy Systems* (World Scientific, 2001), *Data Preprocessing in Data Mining* (Springer, 2015), *The 2-Tuple Linguistic Model. Computing With Words in Decision Making* (Springer, 2015), *Multilabel Classification. Problem Analysis, Metrics and Techniques* (Springer, 2016), and *Multiple Instance Learning. Foundations and Algorithms* (Springer, 2016). His current research interests include soft computing (including fuzzy modeling and evolutionary algorithms), information fusion, decision making, biometric, data preprocessing, data science, and big data.

Dr. Herrera was a recipient of the following honors and awards such as, ECCAI Fellow 2009, IFSA Fellow 2013, the 2010 Spanish National Award on Computer Science ARITMEL to the "Spanish Engineer on Computer Science," International Cajastur "Mamdani" Prize for Soft Computing (Fourth Edition, 2010), the IEEE TRANSACTIONS ON FUZZY SYSTEM Outstanding 2008 and 2012 Paper Award (bestowed in 2011 and 2015, respectively), the 2011 Lotfi A. Zadeh Prize Best paper Award of the International Fuzzy Systems Association, the 2013 AEPIA Award to a scientific career in Artificial Intelligence, the 2014 XV Andalucía Research Prize Maimónides (by the regional government of Andalucía), and the Highly Cited Researcher Award from Thomson Reuters in 2014, in the fields of Computer Science and Engineering, respectively. He currently acts as the Editor-in-Chief of the international journals *Information Fusion* (Elsevier) and *Progress in Artificial Intelligence* (Springer). He acts as an Editorial Member of a dozen of journals.