

# Binary classification in unstructured space with hypergraph case-based reasoning

Alexandre Quemy

IBM Krakow Software Lab, Cracow, Poland

Faculty of Computing, Poznań University of Technology, Poznań, Poland



## HIGHLIGHTS

- A new method for binary classification in unstructured space is proposed.
- The method is well calibrated: its support can directly be interpreted as probability.
- Without data assumption or ad-hoc metric, it competes with more specialized methods.
- It ranks second after Neural Network compared to 9 other methods on 7 datasets.
- It gives consistently good results without hyperparameter tuning or feature engineering.

## ARTICLE INFO

### Article history:

Received 11 June 2018

Received in revised form 26 February 2019

Accepted 7 March 2019

Available online 18 March 2019

### Keywords:

Binary classification

Hypergraph

Case-based reasoning

## ABSTRACT

Binary classification is one of the most common problem in machine learning. It consists in predicting whether a given element belongs to a particular class. In this paper, a new algorithm for binary classification is proposed using a hypergraph representation. The method is agnostic to data representation, can work with multiple data sources or in non-metric spaces, and accommodates with missing values. As a result, it drastically reduces the need for data preprocessing or feature engineering. Each element to be classified is partitioned according to its interactions with the training set. For each class, a seminorm over the training set partition is learnt to represent the distribution of evidence supporting this class.

Empirical validation demonstrates its high potential on a wide range of well-known datasets and the results are compared to the state-of-the-art. The time complexity is given and empirically validated. Its robustness with regard to hyperparameter sensitivity is studied and compared to standard classification methods. Finally, the limitation of the model space is discussed, and some potential solutions proposed.

© 2019 Elsevier Ltd. All rights reserved.

## 1. Introduction

In many real-life situations, one tries to take a decision based on previous *similar* situations. Each situation is described by a certain amount of relevant information, either collected by an expert, or automatically e.g. by some sensors. Those situations share similarities on which to make analogies or counter-examples in order to take new decisions. Conversely, in general, if two situations do not share any common characteristic, then they are totally independent, i.e. it is impossible to infer one's outcome from the other one. The purpose of supervised machine learning algorithms is to exploit the available information and interactions between past cases or examples in order to build a model or infer the key rules to take correct decisions.

E-mail address: [aquemy@pl.ibm.com](mailto:aquemy@pl.ibm.com).

<https://doi.org/10.1016/j.is.2019.03.005>

0306-4379/© 2019 Elsevier Ltd. All rights reserved.

Due to the large variety of concrete situations that can be reduced to binary classification, it is one of the most studied problem in machine learning. In this paper, we investigate the problem of binary prediction under a supervised setting.

This paper **contributes** to binary classification with a new algorithm called Hypergraph Case-Based Reasoning (HCBR). The idea is to create a hypergraph where each element of the training set is a hyperedge and vertices are represented by the features describing the elements. The intersections between edges create a partition, unique to a hypergraph. For each case, we model the support as a convex combination of the elements of this partition. Each of those elements is valued according to its importance w.r.t. the set of all the hyperedges it belongs to and their respective labels.

The paper is an extension of [1] with a focus on separating the abstract framework from the specific implementation used for the experiments. On top of the previous comparison to the

best literature results, we extended the empirical validation with a comparison between HCBR and nine alternative methods. We also added a validation on unstructured datasets. A study of learning curves and model spaces, highlighted some properties of HCBR and the limitations of the current model space, thus shaping priority axes for future work. Among other refinements, matricial formulation for the model space and model selection is provided to ease efficient implementation. More evidence concerning model calibration is added.

The plan of the paper is as follows: in Section 2, we present the problem of binary classification and related work. In particular, in Section 2.1, we present the well-understood and unified framework to solve binary classification. Section 2.2 briefly introduces other approaches to classification. Section 2.3 highlights the necessity of working in unstructured spaces and presents related work w.r.t. metric learning and data wrangling. A formulation of binary classification in an abstract space of information is proposed in Section 2.5. The main contribution of this paper is Section 3 which defines HCBR framework. The rest of the paper focuses on empirical validation: Section 4 presents empirical results on seven structured datasets from the UC Irvine Machine Learning Repository (UCI)<sup>1</sup> and the LIBSVM,<sup>2</sup> while Section 5 is dedicated to unstructured datasets for text classification. Section 6 studies important properties, namely the computational time, learning curves and hyperparameters usage. In Section 7, we discuss the current limitations and possible model space extensions. The paper ends in Section 8 with a discussion about the results, future work, and possible improvements.

This paper is accompanied by Additional Material.<sup>3</sup>

## 2. Binary classification and related work

Before introducing the problem of binary classification, we present the notations used throughout this paper. Vectors are denoted in bold and small case (e.g.  $\mathbf{x}$ ) and their components in small case (e.g.  $x_i$ ). A collection of vectors is denoted in bold and large case (e.g.  $\mathbf{X}$ ).  $|\mathbf{x}|$  (resp.  $|\mathbf{X}|$ ) denotes the cardinal of the vector  $\mathbf{x}$  (resp. the collection  $\mathbf{X}$ ) while  $\|\mathbf{x}\|$  is its norm. The scalar product is denoted by  $\langle \cdot, \cdot \rangle$ . For a matrix  $A = (a_{ij})$ ,  $a_i$  is the  $i$ th row vector and  $a_j$  the  $j$ th column vector.

In machine learning, the problem of classification consists in finding a mapping from an input vector space  $\mathcal{X}$  to a discrete decision space  $\mathcal{Y}$  using a set of examples. The binary classification problem is a special case such that  $\mathcal{Y}$  has only two elements. It is often viewed as an approximation problem s.t. we want to find an estimator  $\bar{J}$  of an unknown mapping  $J$  available only through a sample called *training set*. A training set  $(\mathbf{X}, \mathbf{y})$  consists of  $N$  input vectors  $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$  and their associated correct class  $\mathbf{y} = \{y_i = J(\mathbf{x}_i)\}_{i=1}^N$ .

Let  $\mathcal{J}(\mathcal{X}, \mathcal{Y})$  be the class of mappings from  $\mathcal{X}$  to  $\{-1, 1\}$ , or simply  $\mathcal{J}$  if there is no ambiguity. A machine learning algorithm to solve binary classification is an application  $\mathcal{A} : \mathcal{X}^N \times \mathcal{Y}^N \rightarrow \mathcal{J}$  capable of providing a good approximation for any  $J \in \mathcal{J}$  under some assumptions on the *quality* of the training set. In practice, it is not reasonable to search directly in  $\mathcal{J}$  and some assumptions on the “shape” of  $J$  are made s.t.  $\bar{J} = \mathcal{A}(\mathbf{X}, \mathbf{y})$  belongs to a *hypothesis space* or *model space*  $\mathcal{H} \subset \mathcal{J}$ . This restriction implies not only that the exact mapping  $J$  is not always reachable but might also not be approximated correctly by any element of  $\mathcal{H}$ . The choice of the model space is thus crucial as it should be large enough to represent fairly complex functions and small enough to easily find the best available approximation.

In general, a robust classification algorithm must be able to approximate correctly any possible mapping. The problem of finding such algorithm consists in minimizing the generalization error for all possible mappings. Formally, it consists in solving:

$$\min_{\mathcal{A}} \sum_{J \in \mathcal{J}} \int_{\mathcal{X}} \|J(\mathbf{x}) - \bar{J}(\mathbf{x})\| \mu(\mathbf{x}) d\mathbf{x} \quad (1)$$

where  $\mu$  is a probability measure over  $\mathcal{X}$ .

In practice, the generalization error cannot be computed, and the set of possible mappings is unreasonably large. For those reasons, we aim at minimizing the empirical classification error on a reasonably large set of datasets  $\mathcal{D} = \{(\mathbf{X}_1, \mathbf{y}_1), \dots, (\mathbf{X}_K, \mathbf{y}_K)\}$ , i.e.

$$\min_{\mathcal{A}} \sum_{(\mathbf{X}, \mathbf{y}) \in \mathcal{D}} \sum_{(\mathbf{x}, y) \in (\mathbf{X}, \mathbf{y})} \mathbb{I}_{\{y \neq \bar{J}(\mathbf{x})\}}. \quad (2)$$

To highlight the differences between standard classification problem and our approach, we briefly present linear models for classification followed by the challenges of classification in unstructured spaces. For an overview of theoretical results on classification, we refer the reader to [2].

### 2.1. Linear binary classification

The problem of binary classification is commonly studied with  $\mathcal{X} = \mathbb{R}^M$ . Many popular classification approaches such as SVM [3], perceptron [4] or logistic regression [5] define the model space as the set of  $M$ -hyperplanes. A  $M$ -hyperplane is uniquely defined by a vector  $\mathbf{w} \in \mathbb{R}^M$  and a bias  $w_0 \in \mathbb{R}$ , and is formulated by

$$h_{\mathbf{w}}(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle + w_0 = 0 \quad (3)$$

The *homogeneous* notation consists in adding  $w_0$  to  $\mathbf{w}$  by rewriting  $\mathbf{x}$  such that  $\mathbf{x} = (1, x_1, \dots, x_M)$ . The hyperplane equation (3) is then expressed by  $h_{\mathbf{w}}(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle$ . A hyperplane separates  $\mathbb{R}^M$  into two regions, and thus, can be used as a discriminative rule s.t.

$$\bar{J}_{\mathbf{w}}(\mathbf{x}) = \begin{cases} 1 & h_{\mathbf{w}}(\mathbf{x}) > 0 \\ -1 & h_{\mathbf{w}}(\mathbf{x}) \leq 0 \end{cases} \quad (4)$$

Then, given a training set  $(\mathbf{X}, \mathbf{y})$ , the classification problem is equivalent to finding the best hyperplane s.t. it minimizes a certain *loss function* over the training set [6]:

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \sum_{i=1}^N \ell(\mathbf{w}, \mathbf{x}_i) + \lambda R(\mathbf{w}) \quad (5)$$

where  $R(\mathbf{w})$  is *regularization* term to prevent overfitting (usually  $\|\mathbf{w}\|_2^2$  or  $\|\mathbf{w}\|_1$ ) and  $\lambda > 0$  a hyperparameter controlling the effect of regularization. Note that (5) is a parametric problem due to the choice of model space. Several losses functions exists and are generally based on the *margin* of  $\mathbf{x}_i$  which is defined by  $m(\mathbf{w}, \mathbf{x}_i) = J(\mathbf{x}_i)h_{\mathbf{w}}(\mathbf{x}_i)$ . The margin represents the distance of a vector  $\mathbf{x}_i$  to the hyperplane defined by  $h_{\mathbf{w}}$ . It is positive if  $\mathbf{x}_i$  is correctly classified, negative otherwise. For the most known, the 0–1 loss, hinge loss and log loss are defined by

$$\begin{aligned} \ell_{01}(\mathbf{w}, \mathbf{x}) &= \mathbb{I}_{\{m(\mathbf{w}, \mathbf{x}) \leq 0\}} \\ \ell_{\text{hinge}}(\mathbf{w}, \mathbf{x}) &= \max(0, 1 - m(\mathbf{w}, \mathbf{x})) \\ \ell_{\log}(\mathbf{w}, \mathbf{x}) &= \ln(1 + e^{-m(\mathbf{w}, \mathbf{x})}) \end{aligned} \quad (6)$$

The Perceptron algorithm uses the 0–1 loss, while SVM minimizes the hinge loss and the Logistic Regression the log loss.

<sup>1</sup> <http://archive.ics.uci.edu/ml/index.php>.

<sup>2</sup> <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>.

<sup>3</sup> <http://www.cs.put.poznan.pl/events/2018-IS-special-issue.html>.

## 2.2. Other approaches to binary classification

The method developed in this paper is discriminative and fits the mathematical framework introduced above. Therefore, it models  $p(y|\mathbf{x})$ . There exists other framework to handle binary classification. We briefly introduce them as they will be compared to HCBR in Section 4.

Generative techniques such as Naive Bayes or Bayesian Networks [7] estimate the joint probability  $p(\mathbf{x}, y)$ . They make assumptions on the data distribution (e.g. a mixture of gaussians) and estimate the unknown parameters to generate predictions.

Another popular family of classification techniques are based on trees [8,9]. For instance, Random Forest [10,11] is one of the most widely used method due to its good performances [12] and several extensions have been developed [13,14]

Last but not least, Deep Learning is a family of techniques based on stacking layers of neurons whose weights are adjusted using gradient back-propagation [15]. It represents the state-of-the-art in classification in multiple domains [16]: vision, audio and natural language processing to name few.

## 2.3. On the necessity and problems of unstructured spaces

In machine learning, and broadly speaking in mathematics, the term structured refers to objects that are not living in  $\mathbb{R}^n$ . In particular, there is no obvious metric between objects. Graphs are an example of structured object and the right notion of distance between two graphs depends on the considered problem.

In the database and data science community, the usage is different. A dataset is structured if it can be easily mapped into pre-defined fields or represented by a relational database or tabular-like schema. By opposition, any other dataset is unstructured or semi-structured (raw text, BoW, JSON, XML, graph, etc.). A structured dataset does not guarantee a norm on the rows because some columns might not be expressed in a space with an obvious metric (e.g. categorical variables).

A key element to apply classification methods to datasets is a (meaningful) metric, independently of whether or not the data are stored in a structured format. Thus, we qualify of unstructured a space without obvious or informative metric.<sup>4</sup>

A tremendous amount of data is collected but never used. The data lakes turn into *data swamps*, notably because the natural form of data is *messy*: not sanitized, in different formats, without informative metric, etc. Most of the time, solving optimally a concrete problem requires using data from multiple sources, thus multiplying the above-mentioned problems. This highlights the necessity to develop algorithms capable to work in any unstructured space. In this paper, we present a method working for any unstructured space, allowing to relax usual constraints on the input space:

1. **non-metric space**: there is no metric embedded with the data space,
2. **non-homogeneous cardinality**: the number of features per element is not fixed a priori,
3. **representation agnostic**: the concrete *representation* of features does not influence the model.

To handle the problem of **non-metric spaces**, one can use *metric learning* methods. We present those methods in Section 2.3.1 and discuss how they do not handle the two other points and differ from the method presented in this paper.

<sup>4</sup> A discrete set can always be indexed, however, in this case, the usual distance in  $\mathbb{N}$  is uninformative, therefore not defining a structure.

A special yet very common case of **non-homogeneous cardinality** is some missing values. As the data may be missing for different reasons, there exists different ways of handling them that we briefly present in Section 2.3.2. However, none of them is as straightforward as having a classification method that works for non-homogeneous cardinality.

Regarding **representation agnosticity**, the representation is very often imposed by the classification method itself and when the dataset does not respect those requirements, it has to be transformed. Of course, exploiting the specificities of features often leads to better results but also requires manual expertise that blocks a wider adoption by end-users.

### 2.3.1. Metric learning

Choosing an appropriate pairwise metric to measure distance between two points is crucial in the success of classification algorithms [17]. Learning a metric consists in finding a projection  $f$  from an initial space to a Euclidian space s.t. for any elements  $\mathbf{x}$  and  $\mathbf{x}'$ ,  $d(\mathbf{x}, \mathbf{x}') = \|f(\mathbf{x}) - f(\mathbf{x}')\|$ . The metric should reflect the semantic difference between objects. If the elements are non-numerical, a metric can be viewed as a proxy to represent the elements in a vector space s.t. it becomes possible to apply standard classification methods. However, as reported by [18], the literature mostly focused on numerical data, i.e. when the elements already belong to a vector space.

The most common setting for metric learning is to find a parametrization  $M$  of the Mahalanobis pseudometric  $d_M(\mathbf{x}, \mathbf{x}') = \sqrt{(\mathbf{x} - \mathbf{x}')^T M (\mathbf{x} - \mathbf{x}')}$  using the training set  $\mathbf{X}$  and under the constraint that  $M$  is Positive Semi-Definite (PSD). Using  $M$  as the identity recovers the usual Euclidian metric, whereas setting  $M$  as the covariance matrix of  $\mathbf{X}$  gives the original Mahalanobis distance [19]. Most methods aim at finding the parameters that best agree with the following three sets of relations on the examples:

- $\mathcal{S} = \{(\mathbf{x}_i, \mathbf{x}_j) : \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ are similar}\}$ , (must-link)
- $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{x}_j) : \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ are dissimilar}\}$ , (cannot-link)
- $\mathcal{R} = \{(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k) : \mathbf{x}_i \text{ more similar to } \mathbf{x}_j \text{ than to } \mathbf{x}_k\}$ . (relative)

However, supervised methods usually derived those sets (implicitly or explicitly) from the examples and a proper notion of neighborhood.

We now present the approaches with similarities to the one introduced in this paper. For more detailed surveys on metric learning, we refer the reader to [18,20].

LMNN [21,22] is one the most popular Mahalanobis metric learning techniques and many other linear methods are based on it. The constraint sets  $\mathcal{S}$  and  $\mathcal{R}$  are defined by a notion of neighborhood:

- $\mathcal{S} = \{(\mathbf{x}_i, \mathbf{x}_j) : y_i = y_j \text{ and } \mathbf{x}_j \text{ belongs to the } k\text{-nearest neighbors of } \mathbf{x}_i\}$ ,
- $\mathcal{R} = \{(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k) : (\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}, y_i \neq y_k\}$ .

In the original work, the neighborhood is defined using the Euclidian distance, and thus, assumes the elements lives in a vector space. To avoid this, the method presented in this paper defines the neighborhood based on set intersections to derive  $\mathcal{D}$  and  $\mathcal{R}$ . The PSD matrix  $M$  is found by solving the following convex optimization problem:

$$M^* = \arg \min_{M \in \mathbb{S}_+^n} (1 - \mu) \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} d_M^2(\mathbf{x}_i, \mathbf{x}_j) + \mu \sum_{i,j,k} \varepsilon_{i,j,k} \quad (7)$$

$$\text{s.t. } d_M^2(\mathbf{x}_i, \mathbf{x}_k) - d_M^2(\mathbf{x}_k, \mathbf{x}_j) \geq 1 - \varepsilon_{i,j,k}, \forall (\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k) \in \mathcal{R}$$

$$\varepsilon_{i,j,k} > 0$$

with  $\mu \in [0, 1]$  a tradeoff parameter, and  $\varepsilon_{i,j,k}$  some slack variables.

OASIS [23] learns a bilinear similarity metric of the form  $d_M(\mathbf{x}, \mathbf{x}') = \mathbf{x}^T M \mathbf{x}'$  without the PSD constraint on  $M$ . It can define similarity between instances of different dimensions such as Bag-of-Words. Relaxing the PSD constraint allows the author to use an efficient online Passive-Aggressive algorithm [24] to solve for  $M$ :

$$\begin{aligned} M^t &= \arg \min_{M, \varepsilon} \frac{1}{2} \|M - M^{t-1}\|_{\mathcal{F}}^2 + C\varepsilon \\ \text{s.t. } &1 - d_M^2(\mathbf{x}_i, \mathbf{x}_j) + d_M^2(\mathbf{x}_i, \mathbf{x}_k) \leq \varepsilon \\ &\varepsilon > 0 \end{aligned} \quad (8)$$

with  $C$  a regularization factor and  $\|\cdot\|_{\mathcal{F}}$  the Frobenius norm. Similarly, HCBR works with input vectors of different dimensions. The iterative form of the optimization program is close to the training phase introduced in Section 3.4

SLLC [25] focuses on learning a bilinear similarity matrix as a quadratic constraint program defined by

$$\min_{M \in \mathbb{R}^{n \times n}} \frac{1}{n} \sum_{i=1}^n \ell_{\text{hinge}}(1 - y_i \frac{1}{\gamma |S|} \sum_{\mathbf{x}_j \in \mathcal{D}} y_j d_M(\mathbf{x}_i, \mathbf{x}_j)) + \beta \|M\|_{\mathcal{F}}^2 \quad (9)$$

where  $\mathcal{D}$  is a sample from the training set,  $\gamma$  a margin parameter and  $\beta$  a regularization parameter. The objective of SLLS is to find a metric s.t. the elements of one class are in average more similar than the elements of the other class by a margin  $\gamma$ . The idea behind HCBR is similar except that it learns one seminorm per class such that we can impose to calibrate the distance (i.e. the distance is a confidence measure) over the training set rather than having an average distance parameter. Also, the seminorms are parametrized by a vector and not a matrix.

MMDA [26] learns the Euclidian distance between  $W^T \mathbf{x}_i$  and  $W^T \mathbf{x}_j$  and thus is a linear method. To do so, it learns  $k$  projection hyperplanes  $\{\mathbf{w}_r\}_{r=1}^k$  by solving the following optimization program:

$$\begin{aligned} \min_{W, \mathbf{b}, \varepsilon_r} &\frac{1}{2} \sum_{r=1}^k \|\mathbf{w}_r\|_2^2 + \frac{C}{n} \sum_{r=1}^k \sum_{i=1}^n \varepsilon_{ri} \\ \text{s.t. } &\ell_{\text{hinge}}(\mathbf{w}_r^T \mathbf{x}_i + b_i) \leq 1 - \varepsilon_{ri} \\ &\varepsilon_{ri} > 0 \\ &W^T W = Id \end{aligned} \quad (10)$$

where  $\varepsilon_r$  are slack variables to penalize margins and  $C$  a regularization parameter. HCBR adopts the opposite direction: it gives explicitly a projection per  $\mathbf{x}_i$  and tries to find a vector  $\mu$  such that the sum of hinge loss  $\ell_{\text{hinge}}(\mathbf{w}_i^T \mu)$  is minimized over the training set.

LSMD [27] learns a non-linear projection  $G_W$  parametrized by  $W$  s.t.  $d(\mathbf{x}, \mathbf{x}') = \|G_W(\mathbf{x}) - G_W(\mathbf{x}')\|_1$  is small when the  $\mathbf{x}$  and  $\mathbf{x}'$  are in the same class and large otherwise. The parameter  $W$  represents the weights of a neural network used to learn the projections. HCBR also learns a non-linear projection but such that the sign indicates the class and the distance to 0 the confidence in this class. Therefore, a small distance between two elements does not indicate the same class but the same level of confidence.

The method proposed in this paper is not a pairwise metric learning method. However, many similarities exist as it needs a proxy to express non-numerical data in a numerical decision space s.t. a simple decision rule well performs. What is learnt is a seminorm per class, representing the distribution of discriminative information to support this class w.r.t. the interactions between the training set examples. The pairwise metric can be

artificially defined using the pseudo-norms but is never directly used in the current state of the method.

According to the usual taxonomy of metric learning methods [18,20], HCBR can be viewed as a fully-supervised, non-linear and global<sup>5</sup> metric learner algorithm dedicated to classification on non-vector space. The dimensionality reduction is implicitly defined by the hypergraph representation. The literature work on non-numerical data mostly focused on distance between strings [28] and trees [29,30], with a problem of combinatorial explosion. The scalability study in Section 3.6 shows this is not the case with HCBR.

### 2.3.2. Data wrangling and missing data

Data wrangling consists in selecting, transforming and curating data that are unstructured or not suitable for the selected algorithm. It is commonly accepted that up to 80% of data scientists time is spent on data wrangling [31]. On top of the time consumption, the whole preprocessing step has a huge impact on the final model quality. For instance, in [32], the authors showed that the accuracy obtained by Neural Network, SVM and Decision Trees are significantly impacted by data scaling, sampling and encoding. For a more comprehensive view on data processing impact, we refer the reader to [33].

Data wrangling offers several challenges notably dealing with missing data and outliers or combining data from multiple sources in an automated way [34,35]. In this section, we present the work related to those challenges. However, we believe that another way to tackle data wrangling challenges is to design models that are less sensitive to data preprocessing, missing data or data representation. Not only it decreases the time allocation needed to create efficient data pipelines, but it also decreases the CPU time by shrinking the data pipeline itself. For instance, the method proposed in this paper does not need to remove outliers because they will be marginalized by the model selection itself. As HCBR works with any unstructured space, it becomes easy to combine data from multiple sources. Despite intelligent data transformation might increase the overall performances, we show in Section 4.3.2 that it is not a requirement to perform well.

Modeling datasets in a latent space using meta-features allows to predict the impact of preprocessing operators on the model accuracy [36]. To merge multiple data sources semi-automated tools exists such a Clio [37] or Schema Mapper [38]. For knowledge discovery, SeeDB [39] automatically generates useful visualization of data relations that can help to perform feature selection. Recently, [40] proposed a fully automated way of building data pipelines using standard hyperparameter tuning techniques. The approach still involves a large CPU time overhead.

The most common technique to deal with missing value is replacing the value (imputation) using the mode, average or a value obtained by a model, over the available data. The imputation might be done using a  $k$ -nearest neighbors algorithm [41], local regression [42] or Random Forest [43]. In general, machine learning approaches tend to outperform traditional statistical methods [44]. To handle values under missing at random (MAR) settings, one can use multiple imputations [45,46]. However, this technique requires to generate several variations of the training set and build several models that are then averaged s.t. once again, it requires a large computational effort.

Even when values are missing completely at random (MCAR), there are scalability issues with practical machine learning. Assume some values are MCAR (e.g. due to a defective sensor

<sup>5</sup> In this paper, we use the notion of locality in a different meaning than in the metric learning community. The locality of our model expresses the fact the model cannot be used outside of a certain neighborhood, but the metric has been defined using the entire training set.



for some time), or that new pieces of information are available after an initial model was built (e.g. adding new sensors). In the case of linear classification, e.g. in  $\mathbb{R}^3$ , if an element has only 2 components it does not describe a point but a plane and as a result, there is no way to separate it with a (hyper)plane.

Conversely, consider a model built in  $\mathbb{R}^2$  with some additional information available afterward, s.t. the classification instance now lives in  $\mathbb{R}^3$ . Mathematically, there is no guarantee that the model in two dimensions is even close to the model built from scratch in three dimensions. For instance, if the points are not linearly separable in two dimensions but are in three dimensions, there is no chance for the projection of the plane into the subspace of two dimensions to be the model found while working only in this subspace. Despite multiple algorithms have an *online* counterpart, as far as we know, there is no significant work on having this *horizontal* scalability (as opposed to vertical scalability). One step toward horizontal scalability is to work with unstructured spaces as proposed in this paper.

#### 2.4. Classification and hypergraph

Hypergraphs generalize graphs and can be used to represent higher order relations while graphs are limited to binary relations. A hypergraph is defined by a set of vertices and a collection of hyperedges where each hyperedge is a subset of this set of vertices. Therefore, a graph is a special case of hypergraph for which each hyperedge contains only two vertices. We will formally introduce hypergraphs in Section 3.1. Recently hypergraphs have been used as data representation, and some classification algorithms on hypergraph have been proposed. A vast majority of approaches models the objects to classify as the set of vertices and constructs the hyperedges as the representation of a metric. This conventional approach is known as *neighborhood-based* hypergraph. The metric relies on some assumptions on the data or a specific representation (e.g. Zernike moment and histogram of oriented gradient to measure the distance between images in [47]) and for each vertex, a hyperedge is created to represent its  $k$ -nearest neighbors [48].

The problem of classification on hypergraph consists in labeling some unlabeled vertices given a training set such that all vertices in a same hyperedge have the same label. As all the vertices are known a priori, the problem is part of transductive learning. To learn the labels, the standard approach is to minimize a cost function based on a hypergraph equivalent of a graph Laplacian [47,49] with a structural risk:

$$C(\mathbf{x}) = \mathbf{x}^t \Delta \mathbf{x} + \mu \|\mathbf{x} - \mathbf{y}\|^2 \quad (11)$$

where  $\Delta$  is the hypergraph Laplacian,  $\mu > 0$  a regularization factor and  $\|\cdot\|$  a norm. The vector  $\mathbf{y}$  represents the initial labels for all vertices with  $y_i = 0$  for unlabeled vertices, a negative (resp. positive) value for label -1 or 1.

On the contrary, the **method proposed in this paper** models the elements to classify as the hyperedges and the vertices as the different components of those elements. As far as we know, there is no previous work that uses this modeling choice. In addition, it does not require knowing all the elements before building the model: our approach is inductive. More important, as most previous work consists in building metrics based on the feature representation, it obviously conflicts with our goal of agnosticity described in the previous section.

#### 2.5. Binary classification in unstructured spaces

With all the considerations of Section 2.3 in mind, we reformulate the problem of binary classification for unstructured space. The only *ad-hoc* operation we allow is checking if a particular

feature belongs to the element to classify, and by extension we allow elementary set operations. We consider a countable space  $\mathbb{F}$  and define  $\mathcal{X} = 2^{\mathbb{F}}$ . By abuse of notation, we call  $\mathbf{x} \in \mathcal{X}$  an “vector” despite  $\mathcal{X}$  is not a vector space. We also refer to  $\mathbf{x}$  as a *case*. In practice, it is very likely that only a subset of  $2^{\mathbb{F}}$  may appear (for instance if two features encode two contradictory propositions or if every case has the same number of features). The real class for any input vector  $\mathbf{x}$  of  $2^{\mathbb{F}}$  is given by the mapping:

$$\begin{aligned} J: 2^{\mathbb{F}} &\rightarrow \{-1, 1\} \\ \mathbf{x} &\mapsto J(\mathbf{x}) \end{aligned} \quad (12)$$

Assuming the unknown mapping takes value from the powerset of  $\mathbb{F}$  allows us not to have to know  $\mathbb{F}$  at all. Concretely, the elements of  $\mathbb{F}$  can take different forms depending on the data: a pair “variable = value”, a word, a proposition that is true or false, etc. The only minimal requirement is that there exists a way to index the data elements. In the experimental validation performed in Section 4, the datasets are structured (with some missing values) and thus, the features are all represented by a pair “variable = value”. In Section 5, the features are represented by “word = occurrences” if a word appears at least once. However, as  $\mathbb{F}$  can be seen as a subset of  $\mathbb{N}$ , the method presented below would work on datasets mixing representations (descriptive features, Bag-of-Words, etc.).

Notice that in this paper we do not consider *uncertainty*: if two situations are described the same in  $\mathbb{F}$ , then they have the same label.

### 3. Hypergraph case-based reasoning

In this section, we introduce our main contribution with a **new framework for binary classification in unstructured spaces** called Hypergraph Case-Based Reasoning (HCBR). The presentation is broken down into five steps. First, we present in Section 3.1 how to represent our training set as a hypergraph and how to project any element of  $2^{\mathbb{F}}$  onto it. Those elements allow us to formally define the model space in Section 3.2. Section 3.3 is dedicated to parameter estimations while Section 3.4 focuses on the training. Finally, Section 3.5 is dedicated to the decision rule refinement and hyperparameters. Section 3.6 provides the time complexity of the main phases of the algorithm.

#### 3.1. Representation and projection

Before defining the projection operator used by HCBR to make predictions, we recall the definition of a hypergraph. For additional results on hypergraphs, we refer the reader to [50].

**Definition 3.1 (Hypergraph).** A hypergraph is defined by  $H = (V, \mathbf{X})$  with  $V$  a set of vertices,  $\mathbf{X}$  the hyperedges s.t.  $\forall \mathbf{x} \in \mathbf{X}, \mathbf{x} \subseteq V$ .

A hypergraph can be viewed as a collection of subsets  $\mathbf{X}$  of a given set of vertices  $V$ . It is sometimes convenient to define a hypergraph solely by a collection of sets. In this case, the set of vertices, denoted  $V_{\mathbf{X}}$ , is implicitly defined as the union of edges.

A training set  $\mathbf{X}$  can be seen as a hypergraph  $H = (\mathbb{F}, \mathbf{X})$ , i.e. such that each example is a hyperedge (Fig. 1). In practice, we do not need to know  $\mathbb{F}$  as we can always use the implicit hypergraph  $H = (\mathbb{F}_{\mathbf{X}}, \mathbf{X})$  where  $\mathbb{F}_{\mathbf{X}}$  is the restriction of  $\mathbb{F}$  to the features present in the sample  $\mathbf{X}$ . For structured datasets, the elements of  $\mathbb{F}_{\mathbf{X}}$  are all the distinct pairs “variable=value” in  $\mathbf{X}$ . For any hypergraph  $H = (\mathbb{F}_{\mathbf{X}}, \mathbf{X})$ , there exists a unique partition  $\mathcal{E}_H = \{\mathbf{e}_i\}_{i=1}^M$ ,  $\forall 1 \leq i \leq M$ ,  $\mathbf{e}_i \subseteq \mathbb{F}_{\mathbf{X}}$  defined by the intersections of its edges as illustrated by Fig. 1.

The projection of a case over a hypergraph returns the elements of  $\mathcal{E}_H$  it intersects with.




Fig. 1. The family  $\mathcal{E} = \{\mathbf{e}_i\}_i$  forms the partition obtained by the union of the projection of cases and represents how the three cases share information.

**Definition 3.2 (Projection Over a Hypergraph).** The projection operator  $\pi_H$  over a hypergraph  $H = (V, \mathbf{X})$  for any  $A \subseteq V$  is defined by  $\pi_H(A) = \{\mathbf{e} \in \mathcal{E}_H \mid \mathbf{e} \cap A \neq \emptyset\}$ .

**Example 1.** Each element of  $\pi_H(\mathbf{x})$  is a (sub)set of features. For instance, in Fig. 1,  $\pi_H(\mathbf{x}_1) = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_6\}$  and in Fig. 2, the projection of  $\mathbf{x}$  (in yellow) on  $H$  is the whole partition  $\mathcal{E}_H$  as  $\mathbf{x}$  intersects with every  $\mathbf{e} \in \mathcal{E}_H$ .

We call *discretionary features* of  $\mathbf{x}$  (w.r.t.  $H$ ) the (possibly empty) set of features that are not in  $V_{\mathbf{x}}$ , denoted  $D_{\mathbf{x}}$ . It can be interpreted as the features of  $\mathbf{x}$  that do not belong to any hyperedge. When adding hyperedges, the discretionary features of  $\mathbf{x}$  are new pieces of information that were never encountered before.

**Example.** Considering the hypergraph composed of  $\mathbf{x}_1$  and  $\mathbf{x}_2$  as illustrated by Fig. 1, the set of discretionary features of  $\mathbf{x}_3$  is  $\mathbf{e}_4$ . In Fig. 2, the yellow case  $\mathbf{x}$  has no discretionary feature: all its features are present at least in one example.

For any set of features  $\mathbf{x} \subseteq \mathbb{F}$ , we define  $d_H(\mathbf{x}) = \{\mathbf{x}' \in \mathbf{X} \mid \mathbf{x} \cap \mathbf{x}' \neq \emptyset\}$  the set of edges sharing some features with  $\mathbf{x}$ . In particular, the set  $d_H$  can be split into  $d_H^{(1)}$  and  $d_H^{(-1)}$  depending on the label of the case, and defined by  $d_H^{(l)}(\mathbf{x}) = \{\mathbf{x}' \in d_H(\mathbf{x}) \mid J(\mathbf{x}') = l\}$ . Note that if  $\mathbf{x} \notin \mathbf{X}$  and  $|d_H(\mathbf{x})| = 0$ , then  $\mathbf{x} = D_{\mathbf{x}}$ , i.e. the case  $\mathbf{x}$  is in relation with no case in  $\mathbf{X}$ . In the hypergraph literature,  $|d(\mathbf{x})|$  is called the *degree* of a hyperedge and its domain of definition is restricted to  $\mathbf{X}$  while here, it is extended to  $2^{\mathbb{F}}$ . Finally, to use matricial notations, we define the vectors  $\mathbf{d}^l = (|d_H^{(l)}(\mathbf{x}_i)|)_{i=1}^N$ .

From now, we consider only the specific hypergraph generated by the set of examples  $\mathbf{X}$ . For the sake of readability, we remove the subscript  $H$ .

### 3.2. Model space

As discussed in Section 2.3, we relaxed some implicit constraints on the input vector space. As a counterpart, HCBBR relies on the assumption that if two input vectors  $\mathbf{x}$  and  $\mathbf{x}'$  do not share any feature, they are *independent* i.e.  $\mathbf{x}$  cannot help to understand the correct class of  $\mathbf{x}'$  and vice versa. This limitation comes from the fact there is no metric on  $\mathbb{F}$  to determine a distance between two elements s.t. we can rely only on intersections. As a result, HCBBR produces only *local* models because if a new input vector is independent of all examples, it is impossible to generate a prediction. On the contrary, a hyperplane model is *global* in a sense that it can produce a prediction for any element of  $\mathbb{R}^M$ . We discuss and propose a solution to this issue in Section 7.3.

An example of concrete situation for which such assumption is natural is a justice trial. Cases are composed of some elements,

and the correct label is the result of a reasoning that can possibly use analogies or counter-examples with a set of past cases on top of the legal texts. However, if a judge or lawyer wants to use  $\mathbf{x}$  to justify the outcome of  $\mathbf{x}'$ ,  $\mathbf{x}'$  must have similarities with  $\mathbf{x}$ .

First, we give the **intuition** behind our model space. For each element  $\mathbf{e}_i$  in the projection, we value (1) its importance in the case  $\mathbf{x}$  and (2) its support toward a specific class given the whole hypergraph. Intuitively, (1) consists in answering the question: *how important is  $\mathbf{e}_i$  w.r.t.  $\mathbf{x}$ ? or what is the potential for analogies or counter-examples between  $\mathbf{x}$  and the other examples also containing  $\mathbf{e}_i$ ?* The purpose of (2) is to value how important is  $\mathbf{e}_i$  with regards to the other elements of  $\mathcal{E}$ . For instance, to explain the case  $\mathbf{x}_1$  in Fig. 1, we will use only the elements of the projection  $\pi_H(\mathbf{x}_1) = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_6\}$ . Without prior information, we will use the size of  $\mathbf{e}_i$  in  $\mathbf{x}$  to measure their importance  $\mathbf{x}$ . The way of measuring the support with regard to the whole hypergraph as well as an intuitive interpretation will be given in Section 3.3.

Let us now **formally** define the model space. Given the hypergraph  $H = (\mathbb{F}, \mathbf{X})$  defined by a training set  $\mathbf{X}$  and its associated partition  $\mathcal{E} = \{\mathbf{e}_i\}_{i=1}^M$ , the relation between an example  $\mathbf{x}$  and its class is modeled by

$$\begin{cases} s_{w,\mu}(\mathbf{x}_j) &= \sum_{i=1}^M w(\mathbf{e}_i, \mathbf{x}_j) \mu(\mathbf{e}_i, \mathbf{X}, \mathbf{y}) \\ \sum_{i=1}^M w(\mathbf{e}_i, \mathbf{x}_j) &= 1 \quad \forall 1 \leq j \leq N \\ \sum_{i=1}^M \mu(\mathbf{e}_i, \mathbf{X}, \mathbf{y}) &= 1 \end{cases} \quad (13)$$

where  $w(\mathbf{e}_i, \mathbf{x}_j) \geq 0$  models the importance of  $\mathbf{e}_i$  in  $\mathbf{x}_j$  and  $\mu(\mathbf{e}_i, \mathbf{X}, \mathbf{y})$  the support of  $\mathbf{e}_i$  for class 1 w.r.t. whole training set. For this reason, we call  $\mu$  the *intrinsic strength* of  $\mathbf{e}_i$ . The above-mentioned assumption implies that if  $\mathbf{e}_i \cap \mathbf{x} = \emptyset$  then  $w(\mathbf{e}_i, \mathbf{x}) = 0$ . For readability, we write  $s$  in place of  $s_{w,\mu}$ .

The classification rule consists in selecting the class with the total highest support:

$$\bar{J}_{w,\mu}(\mathbf{x}) = \begin{cases} 1 & s_{w,\mu}(\mathbf{x}) > 0 \\ -1 & s_{w,\mu}(\mathbf{x}) \leq 0 \end{cases} \quad (R1)$$

The classification problem consists then in finding the couple  $(w, \mu)$  that minimizes the classification error:


$$(w^*, \mu^*) = \arg \min_{w, \mu} \sum_{(\mathbf{x}, \mathbf{y}) \in (\mathbf{X}, \mathbf{Y})} \mathbb{I}_{\{\bar{J}_{w,\mu}(\mathbf{x}) \neq \mathbf{y}\}} \quad (14)$$

The problem described by (14) is a functional problem, in general much harder than parametric ones such as (5), and without restrictions on the search space it may not be reasonable to look for a solution. For this reason, and for the rest of this paper, we will fix  $w$  a priori. As we do not formulate any particular assumption on the feature space and allow only basic set operations, a natural yet trivial way to model the importance of  $\mathbf{e}_i$  in  $\mathbf{x}_j$  is to set  $w(\mathbf{e}_i, \mathbf{x}_j) = \frac{|\mathbf{x}_j \cap \mathbf{e}_i|}{|\mathbf{x}_j|}$ . However, if one has some information about the impact of some features in some particular cases,  $w$  might be redefined to integrate this prior knowledge.

Once the method to calculate  $w$  and  $\mu$  is specified, the model can be expressed matricially by

$$\begin{cases} \mathbf{s} &= W\mu \quad W \in \mathcal{M}(\mathbb{R})_{N \times M}, \mu \in \mathbb{R}^M \\ \|\mathbf{w}_j\|_1 &= 1 \quad \forall 1 \leq j \leq N \\ \|\mu\|_1 &= 1 \end{cases} \quad (15)$$

By setting  $d_{\mu}(\mathbf{x}, \mathbf{x}') = |s(\mathbf{x}) - s(\mathbf{x}')|$ , we can see the problem as learning a metric parametrized by  $\mu$  s.t. the training set elements are properly classified. Rather than learning directly  $\mu$ , we split it into  $\mu^{(1)}$  and  $\mu^{(-1)}$ , representing the distribution of the support toward class  $-1$  and  $1$  over  $\mathcal{E}$ .



**Fig. 2.** The projection of  $\mathbf{x}$  (in yellow) on  $H$  is the set of  $\mathbf{e}_i$  intersecting with it. On the right, the graph represents the projection elements  $\{\mathbf{e}_i\}_i$  and their respective connections to the cases  $\{\mathbf{x}_i\}_i$ , in particular,  $D_{\mathbf{x}} = \emptyset$ . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

To fit a model capable of generalization, it is not enough to simply select  $\mu$  s.t. it minimizes the classification error. Assuming  $W$  is fixed and  $\mathbf{s}$  set to some arbitrary values s.t. each example is correctly classified,  $\mu$  can be obtained by using the Moore–Penrose pseudo-inverse of  $W$ . The Moore–Penrose pseudo-inverse generalizes the inverse for non-square matrix. It always exists and is unique. Depending on  $\mathbf{s}$  and  $W$ , it might be impossible to classify correctly all the elements using  $\mu = W^+\mathbf{s}$ . However, by construction of the Moore–Penrose pseudo-inverse,  $W^+\mathbf{s}$  is a least-square minimizer of  $\mathbf{s} = W\mu$ . If this system has many solutions,  $W^+\mathbf{s}$  is the solution with the lowest norm  $\|\cdot\|_2$ .

By artificially setting  $\mathbf{s}$  to obtain  $\mu$  with the Moore–Penrose inverse, the constructed model holds no discriminative information and there is no chance it can provide good results on new cases. For this reason,  $\mu$  has to be designed such that  $\mathbf{s}$  is *calibrated*, i.e. the higher  $\mathbf{s}$  is, the more confident the model is about the prediction. Ideally, among all the elements with e.g. a (normalized) strength close to 0.8, 80% of them should be correctly classified.

To summarize, after defining a general model space to solve the binary classification problem in unstructured space, we made additional assumptions to reduce the problem of model selection to find a vector  $\mu$  s.t. the classification rule (R1) provides a good solution to the original problem (1). For this purpose, we will proceed in three steps:

- **Step 1:** Define  $\mu$  to capture as much information as possible from  $\mathcal{E}$  for any  $\mathbf{X}$  (Section 3.3).
- **Step 2:** Train the model to adjust  $\mu$  on a specific  $\mathbf{X}$  using the classification rule (R1) (Section 3.4).
- **Step 3:** Refine the classification rule (R1) to take into account the local nature of the model (Section 3.5). The refined classification rule (R2) is presented in Section 3.5.

A summary and high level view of HCBR is given by Algorithm 1.

---

**Algorithm 1** HCBR (High level view)

---

- 1: Build  $H$  and  $\mathcal{E}$  from  $\mathbf{X}$ .
  - 2: Calculate  $w$  and  $\mu$  on  $\mathcal{E}$ .
  - 3: Adjust  $\mu$  with training algorithm 2 on  $\mathbf{X}$  using rule (R1)
  - 4: **for** each  $\mathbf{x}$  in the test set **do**
  - 5:   Calculate the projection  $\pi(\mathbf{x})$ .
  - 6:   Calculate the support  $s(\mathbf{x})$  using the projection.
  - 7:   Predict using the refined rule (R2).
  - 8: **end for**
- 

While SVM aims at separating the data using a single hyperplane in the original input vector space, HCBR tries to explain the decision for each case by a convex combination expressed in a lower dimensional space  $\mathcal{E}$ .

### 3.3. Step 1 – Model selection

In this section, we define how to concretely select  $\mu$  for a given hypergraph. Ultimately,  $\mu$  is a measure over  $\mathbb{F}_{\mathbf{X}}$  directly, and cannot be interpreted as a probability:  $\mathbb{F}_{\mathbf{X}} \in 2^{\mathbb{F}_{\mathbf{X}}}$  but by definition of Eq. (13),  $\mu(\mathbb{F}_{\mathbf{X}}) = 1$ . However, there is absolutely no reason to think that  $J(\mathbb{F}_{\mathbf{X}}) = 1$ . Our definition of  $\mu$  tries to answer the following question: *knowing a certain amount of information materialized by the intersection family  $\mathcal{E}$ , where are located the features holding more discriminative information than the others?* Once again, with further information about the features, a prior might be introduced to select  $\mu$ .

For  $\mathbf{x}$  and  $\mathbf{x}'$  in  $2^{\mathbb{F}}$ , a basic quantitative measure on the importance of  $\mathbf{x}'$  w.r.t.  $\mathbf{x}$  can be expressed by  $\frac{|\mathbf{x} \cap \mathbf{x}'|}{|\mathbf{x}|}$ , i.e. how much  $\mathbf{x}'$  overlaps with  $\mathbf{x}$ . This measure is a sort of *potential* for an analogy with  $\mathbf{x}$ . Potential because it does not account for the individual importance of the features and simply holds the idea that the larger is a subset of features in a case, the higher is the chance it holds important features to decide the outcome.

Let us consider  $\mathcal{E}$  and an example  $\mathbf{x} \in \mathbf{X}$ .

**Definition 3.3** (*Intrinsic Strength w.r.t.  $\mathbf{x}$* ). The intrinsic strength of  $\mathbf{e} \in \mathcal{E}$  w.r.t.  $\mathbf{x} \in \mathbf{X}$  is defined by

$$\forall l \in \{-1, 1\}, \bar{S}^{(l)}(\mathbf{e}, \mathbf{x}) = \frac{|d^{(l)}(\mathbf{e})| \frac{|\mathbf{x} \cap \mathbf{e}|}{|\mathbf{x}|}}{\sum_{\mathbf{e}_j \in \mathcal{E}} |d^{(l)}(\mathbf{e}_j)| \frac{|\mathbf{x} \cap \mathbf{e}_j|}{|\mathbf{x}|}} \quad (16)$$

$$= \frac{|d^{(l)}(\mathbf{e})| |\mathbf{x} \cap \mathbf{e}|}{\sum_{\mathbf{e}_j \in \mathcal{E}} |d^{(l)}(\mathbf{e}_j)| |\mathbf{x} \cap \mathbf{e}_j|}$$

Matricially, for any  $\mathbf{e}_i$  and  $\mathbf{x}_i$ ,

$$\bar{S}_{i,j}^{(l)} = \frac{d_i^{(l)} w_{ji}}{\langle \mathbf{d}^{(l)}, \mathbf{w}_{\cdot i} \rangle} \quad (17)$$

In particular, for a given  $\mathbf{x} \in \mathbf{X}$ ,  $\bar{S}^{(l)}(\mathbf{e}_i, \mathbf{x}) = 0$  if  $\mathbf{e}_i$  is not part of the projection of  $\mathbf{x}$  on  $H$ .

The more  $\mathbf{e}_i$  belongs to many cases with the same class  $l$  and the higher  $\bar{S}^{(l)}(\mathbf{e}_i, \mathbf{x})$  is. Conversely, for a fixed number of cases, the more  $\mathbf{e}_i$  describes  $\mathbf{x}$ , the higher  $\bar{S}^{(l)}(\mathbf{e}_i, \mathbf{x})$  is. As  $\forall \mathbf{e}_i \in \mathcal{E}$ ,  $|d(\mathbf{e}_i)| > 0$ , either we have  $\bar{S}^{(1)}(\mathbf{e}_i, \mathbf{x}) \neq 0$  or  $\bar{S}^{(-1)}(\mathbf{e}_i, \mathbf{x}) \neq 0$ . We have  $\bar{S}^{(l)}(\mathbf{e}_i, \mathbf{x}) = 0$  only when all the cases in which  $\mathbf{e}_i$  results of are labeled with the opposite class  $\bar{l}$ . For  $\bar{S}^{(l)}(\mathbf{e}_i, \mathbf{x}) = 1$ , it needs both the unanimity of labels for the cases in which  $\mathbf{e}_i$  belongs to and that  $\mathbf{e}_i = \mathbf{x}$ . The relation  $\mathbf{e}_i = \mathbf{x}$  implies that  $\mathbf{x}$  does not share any feature with any other example or that  $\mathbf{x}$  is included into another example.

**Definition 3.4** (Intrinsic Strength w.r.t. a Hypergraph  $H$ ). The intrinsic strength of  $\mathbf{e} \in \mathcal{E}$  w.r.t.  $H = (\mathbb{F}_X, \mathbf{X})$  is defined by

$$\begin{aligned} \forall l \in \{-1, 1\}, S^{(l)}(\mathbf{e}) &= \frac{|\mathbf{e}|}{\sum_{\mathbf{e}' \in \mathcal{E}} |\mathbf{e}'|} \sum_{\mathbf{x} \in d^{(l)}(\mathbf{e})} \bar{S}^{(l)}(\mathbf{e}, \mathbf{x}) \\ &= \frac{|\mathbf{e}|}{|\mathbb{F}_X|} \sum_{\mathbf{x} \in d^{(l)}(\mathbf{e})} \bar{S}^{(l)}(\mathbf{e}, \mathbf{x}) \end{aligned} \quad (18)$$

Matricially, for any  $\mathbf{e}_i$ ,

$$S_i^{(l)} = \frac{|\mathbf{e}_i|}{|\mathbb{F}_X|} \|S_i^{(l)}\|_1 \quad (19)$$

The more  $\mathbf{e}$  belongs to several cases, the more information it represents to support a class or another. As  $\mathcal{E}$  represents the sets of features that appear all the time together, we favor the larger  $\mathbf{e} \in \mathcal{E}$  as they hold more information to explain a decision. The normalized version is defined by:

$$\forall l \in \{-1, 1\}, \mu^{(l)}(\mathbf{e}) = \frac{S^{(l)}(\mathbf{e})}{\sum_{\mathbf{e}' \in \mathcal{E}} S^{(l)}(\mathbf{e}')} \quad (20)$$

Finally, the measure  $\mu$  is simply defined by the difference between the strength of both classes:

$$\mu(\mathbf{e}) = \mu^{(1)}(\mathbf{e}) - \mu^{(-1)}(\mathbf{e}) \quad (21)$$

which can be expressed matricially for a given  $\mathbf{e}_i$  by

$$\mu_i = \frac{|\mathbf{e}_i|}{|\mathbb{F}_X|} \left[ \frac{S_i^{(1)}}{\|S^{(1)}\|_1} - \frac{S_i^{(-1)}}{\|S^{(-1)}\|_1} \right] \quad (22)$$

**Example (Numerical Example).** Consider the hypergraph in Fig. 1 made of  $\mathbf{x}_1$ ,  $\mathbf{x}_2$  and  $\mathbf{x}_3$  arbitrarily labeled with resp. 1,  $-1$  and 1. Arbitrarily, we assume the cardinal of the elements of  $\mathcal{E}$  to be  $\#\mathbf{e} = (2, 1, 2, 3, 1, 2, 3)$  s.t. the cardinal of cases are  $\#\mathbf{x} = (7, 8, 7)$  and  $|\mathbb{F}_X| = 14$ . The values of  $|d^{(l)}(\mathbf{e})|$  can be summarized by the vectors  $\mathbf{d}^{(-1)} = (0, 0, 1, 0, 1, 1, 1)$  and  $\mathbf{d}^{(1)} = (1, 2, 2, 1, 1, 1, 0)$ . Let us calculate  $S^{(-1)}(\mathbf{e}_3)$ :

$$\begin{aligned} \bar{S}^{(-1)}(\mathbf{e}_3, \mathbf{x}_1) &= \frac{2 \times 2}{2 \times 1 + 1 \times 2 + 2 \times 2 + 1 \times 2} = \frac{4}{10} \\ \bar{S}^{(-1)}(\mathbf{e}_3, \mathbf{x}_2) &= \frac{2 \times 2}{2 \times 2 + 1 \times 1 + 1 \times 2 + 0 \times 3} = \frac{4}{7} \\ \bar{S}^{(-1)}(\mathbf{e}_3, \mathbf{x}_3) &= \frac{2 \times 2}{2 \times 1 + 2 \times 2 + 3 \times 1 + 1 \times 1} = \frac{4}{10} \end{aligned}$$

which we interpret as  $\mathbf{e}_3$  being responsible for  $\frac{4}{10}$  of the support toward class 1 in  $\mathbf{x}_1$  and  $\mathbf{x}_3$ , while  $\frac{4}{7}$  for  $\mathbf{x}_2$ . This leads to

$$S^{(-1)}(\mathbf{e}_3) = \frac{2}{14} \left[ \frac{4}{10} + \frac{4}{7} + \frac{4}{10} \right] \simeq 0.1959$$

Similarly, we compute the support for each  $\mathbf{e}$  and both labels. We summarize this into the following vectors:

$$\begin{aligned} \mathbf{S}^{(1)} &\simeq (0.0286, 0.0286, 0.1959, 0.0643, 0.0173, 0.0694, 0.0000) \\ \mathbf{S}^{(-1)} &\simeq (0.0000, 0.0000, 0.2024, 0.0000, 0.0327, 0.1071, 0.0000) \end{aligned}$$

After normalization, we obtain the intrinsic strength:

$$\mu \simeq (0.0707, 0.0707, 0.0060, 0.1591, -0.0345, -0.0818, -0.1901)^T$$

Let us evaluate the model on the three examples:

$$\begin{aligned} s(\mathbf{x}_1) &= \frac{2}{7}\mu(\mathbf{e}_1) + \frac{1}{7}\mu(\mathbf{e}_2) + \frac{2}{7}\mu(\mathbf{e}_3) + \frac{2}{7}\mu(\mathbf{e}_6) \simeq 0.0086 \\ s(\mathbf{x}_2) &= \frac{2}{8}\mu(\mathbf{e}_3) + \frac{1}{8}\mu(\mathbf{e}_5) + \frac{2}{8}\mu(\mathbf{e}_6) + \frac{3}{8}\mu(\mathbf{e}_7) \simeq -0.0946 \\ s(\mathbf{x}_3) &= \frac{1}{7}\mu(\mathbf{e}_2) + \frac{2}{7}\mu(\mathbf{e}_3) + \frac{3}{7}\mu(\mathbf{e}_4) + \frac{1}{7}\mu(\mathbf{e}_5) \simeq 0.0751 \end{aligned}$$

As a result,  $\mathbf{x}_1$  and  $\mathbf{x}_3$  are labeled 1 and  $\mathbf{x}_2$  is labeled  $-1$ . All three cases are correctly labeled. The highest support is given for case  $\mathbf{x}_2$  and  $\mathbf{x}_3$  while the support for  $\mathbf{x}_1$  is one order of magnitude lower than for  $\mathbf{x}_3$ . This is because the discretionary features of  $\mathbf{x}_3$  are larger while the intersection with  $\mathbf{x}_2$  is lower than for  $\mathbf{x}_1$  ( $\frac{3}{8}$  of  $\mathbf{x}_3$  against  $\frac{4}{7}$  of  $\mathbf{x}_1$ ).

Consider a new case  $\mathbf{x}$  as described in Fig. 2. Its support is given by  $s(\mathbf{x}) = \sum_{\mathbf{e} \in \pi(\mathbf{x})} w(\mathbf{e}, \mathbf{x})\mu(\mathbf{e})$  with  $\pi(\mathbf{x}) = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_5, \mathbf{e}_6, \mathbf{e}_7\}$ .

It is impossible for  $\mathbf{x}$  to be classified as 1 because the highest support would be for a maximal intersection with  $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$  and minimal for  $\mathbf{e}_5, \mathbf{e}_6$  and  $\mathbf{e}_7$  s.t.  $s(\mathbf{x}) = \frac{1}{8}(2\mu(\mathbf{e}_1) + \mu(\mathbf{e}_2) + 2\mu(\mathbf{e}_3) + \mu(\mathbf{e}_5) + \mu(\mathbf{e}_6) + \mu(\mathbf{e}_7)) \simeq -0.0103 < 0$ . It can be explained by the fact that the support for 1 is provided by a larger set of features (11 features versus (8)). On top of that, the intersections between positive cases ( $\mathbf{e}_2$  and  $\mathbf{e}_3$ ) are too small (1 for  $\mathbf{e}_2$  compared to e.g. 3 for  $\mathbf{e}_7$ ) or include also negative cases ( $\mathbf{e}_3$ ).

**Example (Car Accident).** We are interested in knowing if a driver is responsible for an accident involving a pedestrian. From the past cases and the considered case, we collected seven facts: (1) the driver was not speeding, (2) the driver was speeding, (3) the pedestrian was outside the crosswalk, (4) the driver was drunk, (5) the accident happened at night, (6) the driver was a young driver, (7) the accident happened on the highway, (8) the crosswalk light was red.

To simplify, we assume two past cases:  $\mathbf{x}_1 = \{2, 4, 5, 7\}$ ,  $\mathbf{x}_2 = \{1, 3, 5, 6\}$  with  $y_1 = 1, y_2 = -1$ . We are interested in  $y_3$  knowing  $\mathbf{x}_3 = \{2, 5, 6, 8\}$ . The partition is given by  $\mathcal{E} = \{\mathbf{e}_1 = \{2, 4, 7\}, \mathbf{e}_2 = \{5, 6\}, \mathbf{e}_3 = \{1, 3\}\}$ . In particular,  $\pi(\mathbf{x}_3) = \{\mathbf{e}_1, \mathbf{e}_2\}$ .

We assume  $\mu^{(1)} = (0.75, 0.25, 0)$  and  $\mu^{(-1)} = (0, 0.3, 0.7)$  s.t. after normalization  $\mu = (0.5, -0.03, -0.46)$ .

$$\begin{aligned} s(\mathbf{x}_3) &= w(\mathbf{e}_1, \mathbf{x}_3)\mu_1 + w(\mathbf{e}_2, \mathbf{x}_3)\mu_2 \\ &= \frac{1}{4} \cdot 0.5 - \frac{1}{2} \cdot 0.03 = 0.11 > 0 \end{aligned}$$

The model concludes that the driver is guilty. The fact the crosswalk light was red is not taken into account. The decision can be explained mostly by the features of  $\mathbf{e}_1$ . The elements of  $\mathbf{e}_2$  are not discriminative enough to reverse the judgment. If the driver's lawyer would highlight the fact his client is not responsible because of the outcome of  $\mathbf{x}_2$  and the fact it shares many similarities, the defense could argue that the main reason  $y_2 = -1$  holds in  $\mathbf{e}_3$ . By building a chain of analogies and counter-examples, each decision can be explained w.r.t. past cases.

### 3.4. Step 2 – Training

At this stage, it is already possible to generate predictions for new cases. However, the intrinsic strength vector calculated on the hypergraph might not be perfectly accurate on the training set because of the lack of information contained in the training set (or some limitations on the model space itself that we will discuss in Section 6). In this section, we give an algorithm to adjust the intrinsic strength in order to correct the initial estimation.

Once the model is built, it can be evaluated on the training set. Analogously to SVM, we define the margin as the distance to the correct class, i.e.  $m(w, \mu, \mathbf{x}) = J(\mathbf{x})s_{w, \mu}(\mathbf{x})$ . To improve the pertinence of the strength of the elements of  $\mathcal{E}$ , we use the iterative algorithm described by Algorithm 2 to minimize the total margin over the training set  $\mathbf{X}$ .

The order in which points are considered is fixed in the current implementation. When a decision is incorrect for  $\mathbf{x}$ , the algorithm modifies each element of the projection by lowering its strength



**Algorithm 2** Model training**Input:**

- $\mathbf{X}$ : training set
- $\mathbf{y}$ : correct labels for  $\mathbf{X}$
- $k$ : number of training iterations
- $\mu^{(1)}, \mu^{(-1)}$ : weights calculated with (3.4)

**Output:**

- Modified vectors  $\mu^{(1)}, \mu^{(-1)}$

```

1: for  $k$  iterations do
2:   for  $\mathbf{x}_i \in \mathbf{X}$  do
3:      $\bar{y}_i \leftarrow \bar{J}(\mathbf{x}_i)$ 
4:     if  $\bar{y}_i \neq y_i$  then
5:       for  $\mathbf{e} \in \pi(\mathbf{x}_i)$  do
6:          $\mu^{(y_i)}(\mathbf{e}) \leftarrow \mu^{(y_i)}(\mathbf{e}) + w(\mathbf{e}, \mathbf{x}_i)|\mu(\mathbf{e})|$ 
7:          $\mu^{(\bar{y}_i)}(\mathbf{e}) \leftarrow \mu^{(\bar{y}_i)}(\mathbf{e}) - w(\mathbf{e}, \mathbf{x}_i)|\mu(\mathbf{e})|$ 
8:       end for
9:     end if
10:  end for
11: end for

```

for the wrong class and increasing it for the proper class. The margin is split between the element of the projection w.r.t. their respective weight in  $\mathbf{x}$  i.e.  $w(\mathbf{e}, \mathbf{x})$ . If a case  $\mathbf{x}$  is wrongly classified, it is due to the cases intersecting with it. Indeed, if  $\mathbf{x}$  was not intersecting with any other example, its projection would be itself, and its support toward the wrong class would be 0 and positive for the real class. In other words,  $\mathbf{x}$  would be correctly classified. Thus, the idea is not to directly bring the support of  $\mathbf{x}$  to the correct class but to gradually adjust the weights s.t. the neighbors are modified enough for  $\mathbf{x}$  to be correctly classified. In particular, it is sensitive to the order in which the cases are considered: a modification in the strength of any  $\mathbf{e} \in \mathcal{E}$  impacts all cases in which it appears and potentially changes the predicted class for those cases.

The update rule being a contracting mapping because  $|\mu(\mathbf{e})| < 1$  at the initial step and  $w(\mathbf{e}, \mathbf{x}_i) < 1$  unconditionally, Algorithm 2 is guaranteed to converge. A more formal justification can be found in the Additional Material, Section 1. However, too many iterations may lead to overfitting  $\mu$  to the training set. Empirical experiments suggest that the result after one to five training iterations is (near-)optimal (see Section 6).

### 3.5. Step 3 – Decision rule refinement

This step is not mandatory in a sense that the model can be built and already generates predictions. The hyperparameters introduced in this section can help either to increase prediction accuracy (see Section 4.3.1) or control the risk associated to a prediction (see Section 6.3).


The measure  $\mu$  is defined as the difference of support for both classes. Thus, by linearity we can rewrite

$$\begin{aligned} s(\mathbf{x}) &= \sum_{i=1}^M w(\mathbf{e}_i, \mathbf{x})\mu^{(1)}(\mathbf{e}_i) - \sum_{i=1}^M w(\mathbf{e}_i, \mathbf{x})\mu^{(-1)}(\mathbf{e}_i) \\ &= s^{(1)}(\mathbf{x}) - s^{(-1)}(\mathbf{x}) \end{aligned} \quad (23)$$

This form is convenient because we can control how much evidence we need to support a specific class using the following constraints and a family  $\eta$  of four hyperparameters:

$$s^{(-1)}(\mathbf{x}) > \max\left(\frac{\bar{\eta}_{-1}}{1 - \bar{\eta}_{-1}}s^{(-1)}(\mathbf{x}), \eta_{-1}\right) \geq 0 \quad (C_0)$$

$$s^{(1)}(\mathbf{x}) > \max\left(\frac{\bar{\eta}_1}{1 - \bar{\eta}_1}s^{(1)}(\mathbf{x}), \eta_1\right) \geq 0 \quad (C_1)$$



**Fig. 3.** Representation of the updated decision rule (R2) in the extended decision space.

with  $\eta_{-1}, \eta_1 \in \mathbb{R}^+$  and  $\bar{\eta}_{-1}, \bar{\eta}_1 \in [0, 1]$ . The constraints on  $\eta_{-1}$  and  $\eta_1$  define a minimal amount of support respectively toward class  $-1$  and  $1$  while  $\bar{\eta}_{-1}$  and  $\bar{\eta}_1$  requires the support toward a class to be significantly higher than the support for the other class. As  $\mu$  is normalized over  $\mathcal{E}$ , the value of  $\eta_{-1}$  and  $\eta_1$  must be set w.r.t. the hypergraph. On the contrary,  $\bar{\eta}_{-1}$  and  $\bar{\eta}_0$  can be set independently of the hypergraph.

Those constraints may be used to design a decision rule for new cases depending on the application or the dataset. The most generic decision rule is as follows:

$$\tilde{J}(\mathbf{x}) = \begin{cases} 1 & s(\mathbf{x}) > 0 \quad \text{and } C_1 \\ -1 & s(\mathbf{x}) \leq 0 \quad \text{and } C_0 \\ l_1 & s(\mathbf{x}) > 0 \quad \text{and not } C_1 \\ l_{-1} & s(\mathbf{x}) \leq 0 \quad \text{and not } C_0 \end{cases} \quad (R2)$$

where  $l_{-1}, l_0$  are two labels. A representation is given by Fig. 3. Those hyperparameters are intended to model the “burden of proof”. For instance, in a trial, one is assumed innocent until proven guilty which implies the support for the class “guilty” must be *beyond a reasonable doubt* (where the term reasonable is defined by the jurisprudence of the applicable country). In case  $\eta_{-1} = \eta_1 = \bar{\eta}_{-1} = \bar{\eta}_1$  (and  $l_{-1} = 0$  and  $l_1 = 1$ ), then the decision rule is equivalent to the original one defined by (R1).

### 3.6. Time complexity

**Model building.** Given  $\mathbf{X} \in (2^{\mathcal{F}})^N$ , constructing  $\mathcal{E}_H$  can be done in  $\mathcal{O}(\sum_{\mathbf{x} \in \mathbf{X}} |\mathbf{x}|)$  by using a Partition Refinement data structure [51].

Given  $\mathbf{x} \in \mathbf{X}$ , calculating the family  $\{\bar{S}(\mathbf{e}, \mathbf{x})\}_{\mathbf{e} \in \mathcal{E}_H}$  can be done in  $\mathcal{O}(|\mathbf{x}|)$  by asking for each feature of  $\mathbf{x}$  the  $\mathbf{e}$  it belongs to and maintaining the size of each  $\mathbf{e}$  during the construction of  $\mathcal{E}_H$ . Thus, calculating  $\{\bar{S}(\mathbf{e}, \mathbf{x})\}_{\mathbf{e} \in \mathcal{E}_H}$  for all  $\mathbf{x} \in \mathbf{X}$  can be done in  $\mathcal{O}(\sum_{\mathbf{x} \in \mathbf{X}} |\mathbf{x}|)$ .

On  $m$ -uniform hypergraphs (when all cases are described with  $m$  features), it becomes  $\mathcal{O}(mN)$ .

Calculating  $\{S(\mathbf{e})\}_{\mathbf{e} \in \mathcal{E}_H}$  and  $\mu$  can be done in  $\mathcal{O}(|\mathcal{E}_H|)$  because it requires to iterate over  $\mathcal{E}_H$ . An obvious upper bound on  $|\mathcal{E}_H|$  is  $|\mathbb{F}_{\mathbf{X}}|$  i.e. the number of vertices in the hypergraph. The worst-case cardinal of  $\mathcal{E}_H$  is when each  $\mathbf{x} \in \mathbf{X}$  intersects with all the others and none of them is strictly a subset of any other. Thus,  $|\mathcal{E}_H| \leq \min(2^N - 1, |\mathbb{F}_{\mathbf{X}}|)$ .

**Learning phase.** For each wrongly classified  $\mathbf{x} \in \mathbf{X}$ , a training iteration requires at most  $\mathcal{O}(|\mathbf{x}|)$  steps (maximal cardinal for  $\pi(\mathbf{x})$ ). The worst-case scenario is when the model wrongly classifies every  $\mathbf{x} \in \mathbf{X}$ . Thus, the learning phase worst-case complexity is  $\mathcal{O}(k \sum_{\mathbf{x} \in \mathbf{X}} |\mathbf{x}|)$  and on  $m$ -uniform hypergraphs it becomes  $\mathcal{O}(kmN)$ .

**Model query.** For a case  $\mathbf{x} \in 2^{\mathcal{F}}$ , the projection can be done in  $\mathcal{O}(|\mathbf{x}|)$ . Calculating the classification rule also requires at most  $\mathcal{O}(|\mathbf{x}|)$  (maximal cardinal for  $\pi(\mathbf{x})$ ).

## 4. Experiments on structured datasets

In this section, we validate HCBR on well-known structured datasets. Two series of experiments are performed. The first one compares HCBR to the best results from the literature and includes hyperparameter tuning. The second one focuses on the robustness, i.e. the capacity to deliver good performances on a wide range of datasets without spending time on feature engineering, data preprocessing or hyperparameter tuning.

Some specific elements such as the learning curves or computing times are discussed in Section 6 dedicated to intrinsic performances and properties.

### 4.1. Data and method

We used seven structured datasets for binary classification. All of them are available either from the UCI Machine Learning Repository<sup>6</sup> or provided with LIBSVM<sup>7</sup>: `adult`, `breasts`, `heart`, `mushrooms`, `phishing`, `skin` and `splice`. For each dataset, the original features (`name=value`) are converted into a unique identifier and the union of all such identifiers constitutes the information set  $\mathbb{F}$ .

The datasets are described by Table 1. The minimal, maximal and average size give information about the case sizes (notice `adult`, `heart` and `mushrooms` datasets are missing values). The column `unique` reports the cardinal of  $\mathbb{F}$ . Two datasets have at least one real-valued attribute as indicated by the column “Real”. Three datasets (`adult`, `breasts` and `skin`) are highly imbalanced.

For both experiments, we saved the confusion matrix obtained over all runs and after each prediction. From this confusion matrix, we calculated standard performance indicators: accuracy, recall, specificity, precision, negative prediction value,  $F_1$ -score and Matthews correlation coefficient. Denoting by TP the number of true positives, TN the true negatives, FP the false positives and FN the false negative, the accuracy, the  $F_1$ -score and Matthews correlation coefficient (MCC) are defined by:

$$\text{ACC} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}$$

$$F_1 = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}}$$

$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}}$$

The accuracy,  $F_1$ -score and MCC respectively belongs to  $[0, 1]$ ,  $[0, 1]$  and  $[-1, 1]$ . The closer to 1, the better it is.  $F_1$ -score and MCC take into account false positive and false negatives. Furthermore, MCC has been shown to be more informative than other metrics derived from the confusion matrix [52], in particular with imbalanced datasets.

#### 4.1.1. Experiment 1 – Literature comparison

The objective is to compare HCBR performance to the best results from the literature. As most studies do not report  $F_1$ -score nor MCC, we will base our comparison on the accuracy.

We used a 10-fold cross-validation with stratified sample to preserve the original dataset prevalence. The number of training steps  $k$  was adjusted with a manual trial-and-error approach.

To measure the impact of hyperparameters on HCBR, the runs have been performed twice (with the same seed). The first

campaign was done with the family  $\eta$  sets to  $\eta_{-1} = \eta_1 = \bar{\eta}_{-1} = \bar{\eta}_1 = 0$ . The second one was done with a nested 10-fold cross-validation in order to tune hyperparameters. The training set was itself divided into 10 folds. The first nine folds were used to build and train the model, the last one used as validation set to find the optimal values for the family  $\eta$ . On the test set, we used for  $\eta$  the average values found over the 10 validation sets.

It is difficult to know in advance the range of values for the components of  $\mu$  such that setting a grid is not convenient. Instead, we define a hyperparameter space that depends on the strength expressed in the decision defined by Fig. 3. Consider the upper half-plane of this space, i.e. the points with positive support. The hyperparameters can define a vertical split, horizontal split or another (open) half-plane with a certain margin from the axes. We performed an exhaustive search for the family  $\eta$ . We describe here the procedure for cases with a positive support i.e. to find  $\eta_1$  and  $\bar{\eta}_1$ . The procedure is symmetric for negative support.

1. For each point  $(x, y)$  from the validation set, we considered three points:  $p_1 = (x + \varepsilon, y + \bar{\varepsilon})$ ,  $p_2 = (x + \varepsilon, 0)$  and  $p_3 = (0, y + \bar{\varepsilon})$  where  $\varepsilon$  and  $\bar{\varepsilon}$  are positive or negative depending on whether  $(x, y)$  is correctly classified or not and set to half the distance to the immediate neighbor of  $(x, y)$ .
2. For each point of those three points, we recalculated the accuracy when  $(\eta_1, \bar{\eta}_1)$  is sets to  $p_i$
3. We choose the values that return the highest accuracy.

#### 4.1.2. Experiment 2 – Robustness comparison

The state-of-the-art results obtained per dataset required a lot of efforts in terms of feature engineering, as well as model selection or hyperparameter tuning. As discussed in Section 2.3, it represents nowadays most data scientists work and CPU time. If well-calibrated models with high performances on their domain of application is undeniably useful for practical usage, being able to solve a problem on a large variety of instances without effort is also of a vital importance, especially in the industry where the end-user might not be specialized in data science and machine learning. One main-challenge of large scale machine learning systems is thus to achieve the compromise between ready-to-deploy models and good performance metrics.

There exist tools such as auto-sklearn [53] to automatically tune hyperparameters and select the best algorithm in a given portfolio. However, this approach still requires an extensive CPU time which might be prohibitive in case end-users need to build classification models over numerous datasets. For instance, one may think to cloud offers such as IBM Cloud<sup>8</sup> or Amazon AWS<sup>9</sup> that create for their customers thousands of models from scratch every day. Having models that give good results with limited CPU time dedicated to feature engineering and hyperparameter is thus an important competitive advantage. On top of that, Automated Machine Learning is far from being widely adopted in the industry where hyperparameter tuning might not even be done for several reasons, as noted by [12] and confirmed by our experience.

Therefore, we would like to quantify how good HCBR can perform on different datasets without feature engineering or hyperparameter tuning. To measure this robustness, we performed a 10-fold cross-validation for each of the seven selected datasets using nine standard classification methods: AdaBoost, k-Nearest

<sup>6</sup> <https://archive.ics.uci.edu/ml/index.php>.

<sup>7</sup> <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>.

<sup>8</sup> <https://www.ibm.com/cloud/>.

<sup>9</sup> <https://aws.amazon.com/>.

**Table 1**  
Datasets description.

	Cases	Total features	Unique	Min. size	Max. size	Avg. size	Real	Prev.
adult	32561	418913	118	10	13	12.87	No	0.7586
breasts	699	5512	80	8	8	8	No	0.3338
heart	270	3165	344	12	13	12.99	Yes	0.5107
mushrooms	8124	162374	106	20	20	20	No	0.4804
phishing	11055	319787	808	29	29	29	No	0.5562
skin	245057	734403	768	3	3	3	Yes	0.2075
splice	3175	190263	237	60	60	60	No	0.5164

Neighbors, Linear SVM, Radius-Based Function (RBF) SVM, Decision Tree, Random Forest, Neural Network and Quadratic Discriminant Analysis (QDA). The implementation is provided by Scikit-Learn [54]. No hyperparameter tuning was performed and the default values of parameters were used.

#### 4.2. Previous work on the datasets

To compare the results of the proposed method, we explored for each dataset the best results from the literature. The results are summarized in Table 2. In [55], 5 rule-based classification techniques dedicated to medical databases are compared and achieve at best 95.85% and 82.96% accuracy resp. on breast, and heart datasets. Comparing bayesian approaches, [56] demonstrated 97.35% (breast) and 83.00% (heart) accuracy. A 5 layers neural network with fuzzy inference rules achieved 87.78% on heart [57] while a k-NN algorithm reached 99.96% on mushrooms [58]. The best alternative among 6 rules-based classification methods achieved 95.84% on breast and 100.00% on mushroom [59]. Using 80% of phishing as training set, an adaptive neural network achieved an average accuracy of 93.76% (among 6 alternatives) with the best run at 94.90% [60]. Still on phishing, 7881507 proposes to combine several classifiers and reaches 97.75% accuracy for the best hybrid model (and demonstrates 97.58% for Random Forest classifier). On adult, the comparison of several classifiers (naive bayes, decision tree, ...) demonstrated at most 86.25% accuracy [61] while a Support Vector Machine approach reached 85.35% [62]. On splice, a method using Fuzzy Decision Trees [63] reaches 94.10% accuracy and a neural network combined to boosting [64] 97.54%. On breast, Support Vector Machine approaches reached resp. 96.87%, 98.53%, 99.51% accuracy [65–67], 99.26% and 97.36% for neural network based techniques [68,69], 98.1% for a bayesian network method [70], or 94.74% using Decision Trees [71]. On skin, catak2017 reports 98.94% accuracy against 99.68% for Decision Tree based method [72]. The best result, as far as we know, is 99.92%, obtained by a Generalized Linear Model [73].

#### 4.3. Results

The integrality of the data used for the experiments, as well as the scripts to transform them and analyze the results are available within the HCBR Github repository<sup>10</sup> s.t. the whole experimental campaign starting from the raw data can be easily reproduced.

##### 4.3.1. Literature comparison

The average confusion matrix obtained for each dataset is showed in the Additional Material, Section 2. The performance

**Table 2**

Previous literature results measured as the highest accuracy obtained by the authors.

Dataset	Ref.	Type	Accuracy
adult	[61]	Many classifiers	86.25%
	[62]	SVM	85.35%
		<b>HCBR (tuned)</b>	<b>82.90%</b>
		<b>HCBR</b>	<b>82.06%</b>
breast	[67]	SVM	99.51%
	[68]	Neural Net	99.26%
	[66]	SVM	98.53%
	[70]	Bayes	98.1%
		<b>HCBR (tuned)</b>	<b>97.83%</b>
	[69]	Neural Net	97.36%
	[56]	Bayes	97.35%
		<b>HCBR</b>	<b>96.96%</b>
	[65]	SVM	96.87%
	[55]	Rule-based	95.85%
[59]	Rule-based	95.84%	
[71]	Decision Tree	94.74%	
heart		<b>HCBR (tuned)</b>	<b>90.77%</b>
	[57]	Neural Network + Rule-based	87.78%
		<b>HCBR</b>	<b>85.77%</b>
	[56]	Bayes	83.00%
[55]	Rule-based	82.96%	
mushrooms	[59]	Rule-Based	100.00%
		<b>HCBR</b>	<b>100.00%</b>
	[58]	k-NN	99.96%
phishing	[74]	Ensemble	97.75%
	[74]	Random-Forest	97.58%
		<b>HCBR (tuned)</b>	<b>96.82%</b>
		<b>HCBR</b>	<b>96.05%</b>
	[60]	Neural Net	94.90%
skin	[73]	Generalized Linear Model	99.92%
	[72]	Decision Tree	99.68%
	[64]	Neural Network + Boosting	98.94%
		<b>HCBR (tuned)</b>	<b>98.68%</b>
	<b>HCBR</b>	<b>98.65%</b>	
splice	[64]	Neural Network + Boosting	97.54%
		<b>HCBR (tuned)</b>	<b>95.09%</b>
	[63]	<b>HCBR</b> (fuzzy) Decision Tree	<b>94.10%</b>

indicators are reported in Table 3. The proposed algorithm performs very well on a wide range of datasets as reported by the Additional Material, Section 2 and Table 3.

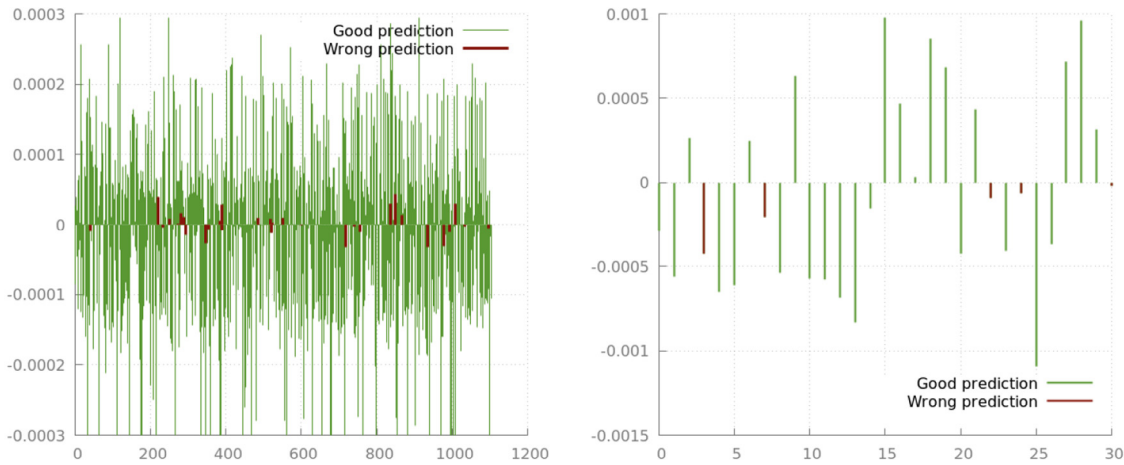
*Without hyperparameter tuning.* The accuracy is contained in a range from 0.8206 (adult) to 1 (mushrooms) while the  $F_1$ -score is bounded by 0.8653 (heart) and 1 (mushrooms). On adult, the accuracy is only 6% higher than the prevalence, i.e. a baseline model consisting in returning 1 for any point would be only 6% worse. This relatively poor performance in learning the underlying decision mapping is better reflected by the Matthews correlation coefficient of 0.51.

The false positives and false negatives are equilibrated for each dataset, despite a huge variation in the prevalence (between 20% and 64%, cf. Table 1) which is a desirable property as it is known to be a problem for many machine learning algorithms [75].


<sup>10</sup> <https://github.com/aquemy/HCBR>.

**Table 3**  
Average performances obtained with a 10-fold cross-validation.

	Accuracy (std dev.)	Recall	Specificity	Precision	Neg. Pred. Value	F <sub>1</sub> score	MCC
adult	0.8206 (0.0094)	0.8832	0.6233	0.8808	0.6290	0.8820	0.5081
	0.8290 (0.0063)	0.9008	0.6029	0.8773	0.6029	0.8889	0.5194
breasts	0.9696 (0.0345)	0.9691	0.9676	0.9479	0.9844	0.9575	0.9344
	0.9783 (0.0204)	0.9553	0.9910	0.9833	0.9910	0.9691	0.9526
heart	0.8577 (0.0943)	0.8695	0.8437	0.8699	0.8531	0.8653	0.7178
	0.9077 (0.0659)	0.9310	0.8783	0.9060	0.8783	0.9184	0.8126
mushrooms	1.0000 (0.0000)	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
phishing	0.9605 (0.0081)	0.9680	0.9514	0.9615	0.9590	0.9647	0.9199
	0.9682 (0.0067)	0.9689	0.9672	0.9741	0.9672	0.9715	0.9355
skin	0.9865 (0.0069)	0.9608	0.9932	0.9736	0.9898	0.9672	0.9587
	0.9868 (0.0062)	0.9740	0.9900	0.9618	0.9900	0.9679	0.9596
splice	0.9443 (0.0124)	0.9478	0.9398	0.9450	0.9441	0.9463	0.8884
	0.9509 (0.0108)	0.9478	0.9544	0.9577	0.9544	0.9527	0.9018



**Fig. 4.** Difference between the weight assigned to both classes for each decision on phishing and splice (average). Similar results are observed for all datasets.




**Fig. 5.** Histogram of decisions depending on the strength for phishing and splice. In blue the correctly classified elements, in red the wrongly classified ones. The false positives and false negatives are concentrated around 0. Similar results are observed for all datasets. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The support is a metric of confidence for the prediction as illustrated in Fig. 4. In general, the wrongly classified cases have a smaller difference between the evidence for each class. This can also be observed in Fig. 5.

HCBR performs the best on mushrooms and heart datasets. For the rest, the accuracy is slightly lower than the best results from the literature (adult 82.90% against 86.25%, breast 97.83% against 99.51%, phishing 96.82% against 97.75%, splice 95.09% against 97.54%, skin 98.68% against 99.92%). We explain this by at least two factors. First, the best methods on a given dataset

are often dedicated to this dataset with *ad-hoc* or engineered parts which is not the case of HCBR. Secondly, the learning curves study in Section 6.2.1 reveals that the model space is not complex enough. HCBR performed better than Bayes classifier in two thirds of cases. Bayes classifier performs better on breast by approximately 1% which represents less than one case wrongly classified. Similar results are observed with Decision Trees. However, the 1% difference on skin represents an average of 7 cases misclassified in comparison in favor of Bayes. It performs better than Rule-based approaches (or gives similar results on mushrooms with





**Fig. 6.** Illustration of the hyperparameter optimization on `splice` (one specific run, zoomed in the origin). The dashed lines are defined by averaging the best parameters obtained on the test set (10 folds). For negative support, it turned 4 FN into FP and 1 TN into FP compared to the version without hyperparameter tuning. For positive support, it turned one FP into TN. The hyperparameter values are not optimal as a better result could have been achieved.

an accuracy of (1) in the four considered references on three different datasets. Except for `heart` and `phishing`, Neural Network returns better results (0.46 more cases with correct classification in average for `breast`, 71 for `skin` and almost 10 for `splice`). Last, SVM gives better results in all three cases, but appear only as best results in two datasets.

*With hyperparameter tuning.* An illustration of parameters tuning on a real instance is depicted by Fig. 6. With hyperparameter tuning, the accuracy lower bound is obtained by the same dataset (0.8290 with `adult`). The lower bound on  $F_1$ -score is now obtained by `adult` with 0.8889. On `heart`, the accuracy gains 5 percent point (pp) which represents a 35% error reduction. This allowed HCBR to rank first. On `breast`, the gain represents 0.87pp which is about 29% error improvement and a higher accuracy than [56] and [69]. For the other datasets, the accuracy variation does not result in a better rank and the error improvement lies between 2% for `skin` and 19% for `phishing`.

We explain the limited effect of hyperparameters by the fact that for small support values, there can be several cases with the exact same support but different outcomes. Two different cases have the same support if the elements of their projections are weighted the same despite being composed of different features. In order to be more discriminative, one can increase the model space complexity s.t. the support is obtained by more variables which decreases the probability of collisions. See Section 6.2 for a study of the model space limitations and possible remedies.

*Discussion.* Even if HCBR does not rank first in most cases, we see at least three reasons to use it in practice. First, it provides consistently good results on all datasets without a need to tune any hyperparameter, without correcting the imbalanced datasets or feature transformation that would require expert knowledge. In other words, it requires absolutely no domain knowledge or data science expertise to deploy and use. Second, HCBR works directly in unstructured spaces allowing combining data from multiple sources without tedious transformations. This has a considerable advantage in practice where the information per case is rarely structured by default. Last but not least, HCBR provides local explanation: for each case and each group of features  $e_i$  in this case, HCBR provides not only the support for  $e_i$  but also all the cases that participated in increasing the support. In the medical domain, the practitioner may check, for instance, the top

**Table 4**


Average MCC and rank obtained with several methods (Scikit-Learn implementation). The column  $\Delta_{HCBR}$  represents the relative difference in MCC w.r.t. HCBR.

Method	MCC	Rank	$\Delta_{HCBR}$
Neural Network	0.8914	1	5.68%
<b>HCBR</b>	0.8435	2	–
RBF SVM	0.8267	3	2.00%
Decision Tree	0.8066	4	4.37%
AdaBoost	0.8063	5	4.41%
k-NN	0.7859	6	6.82%
Linear SVM	0.7858	7	6.84%
QDA	0.7358	8	12.76%
Random Forest	0.7237	9	14.20%
Naive Bayes	0.6953	10	17.56%

**Table 5**

Average accuracy and rank obtained with several methods (Scikit-Learn implementation). The column  $\Delta_{HCBR}$  represents the relative difference in accuracy w.r.t. HCBR.

Method	Accuracy	Rank	$\Delta_{HCBR}$
<b>HCBR</b>	0.9360	1	–
Neural Network	0.9354	2	0.06%
RBF SVM	0.9247	3	1.16%
AdaBoost	0.9207	4	1.60%
Linear SVM	0.9115	5	2.62%
Decision Tree	0.9056	6	3.25%
k-NN	0.9011	7	3.81%
Random Forest	0.8903	8	4.88%
QDA	0.8802	9	6.27%
Naive Bayes	0.8235	10	13.59%



**Fig. 7.** Matthew Correlation Coefficient depending on the training set size for `skin`. See the Additional Material, Section 5 for other datasets.

two groups of features, and for each, the five most influencing past cases to find a reasonable justification to the prediction.

#### 4.3.2. Robustness comparison

The average MCC achieved per method over the datasets is reported in Table 4. The details per dataset can be found in the Additional Material, Section 3. The counterpart for the accuracy is provided by Table 5 and the Additional Material, Section 4. Additionally, we displayed the evolution of MCC with the training set size in Fig. 7.

The highest MCC value is obtained by Neural Network with 0.8914 followed by HCBR with 0.8435. Neural Network improves HCBR result by 5.68% while HCBR improves the result of all other methods from 2% to 17.56%. The lowest MCC score is obtained by

Naive Bayes with 0.6953, mostly because it performs poorly on some datasets (0.2493 on *adult*, 0.5292 on *phishing* or 0.7600 on *skin*). In general, Neural Network and HCBR are the only two methods whose ranks remain consistent across all datasets, close to the first and second rank. On the contrary, methods like *k*-Nearest Neighbors or Decision Tree performed very well on one dataset (*skin*, resp. *splice*). Surprisingly, Random Forest not only perform poorer than Decision Tree but also performs worse than most methods. Naive Bayes and QDA perform very poorly in general which is less surprising knowing the assumptions behind those methods that are likely to be unrealistic on real datasets.

Regarding accuracy, HCBR is the best performing algorithm followed by Neural Network. In general, the ranking is consistent with the one obtained for MCC.

In other words, HCBR and Neural Network have shown to be more versatile than the other approaches on this selected set of classification instances and without parameter tuning. We have no doubt that the results obtained with other methods can be highly improved by a proper hyperparameter selection, in particular for Random Forest or even *k*-Nearest Neighbors. However, this time-consuming operation can be avoided with HCBR that provides the best compromise between ready-to-use and good performances.

## 5. Experiments on unstructured datasets for text classification

In this section we study the performance of HCBR on unstructured datasets and show it performs better in most cases compared to the reference study that uses SVM.

### 5.1. Data and method

We used the European Court of Human Rights dataset provided by the authors of [76]. This dataset for binary classification is broken down into three articles, namely Article 3, Article 6 and Article 8. For each article, the authors collected some judgments (250, 80 and 254 respectively for Article 3, 6 and 8) s.t. the prevalence is 50%. For each document, they isolated different sections (Procedure, Circumstances, Relevant Law, Law) and excluded the verdict section. The authors do not provide the raw text but the Bag-of-Words representation obtained by keeping the 2000 most common *N*-grams for  $N \in \{1, 2, 3, 4\}$ . The authors used a linear SVM after applying a TF-IDF schema on the dataset to predict whether an article has been violated.

On top of those Bag-of-Words representations, the authors provide the representation of a judgment in a topic space. This topic space is obtained by calculating the matrix of cosine distance between the documents and keeping the 30 top-components after applying a spectral clustering. Finally, the authors include two additional parts defined as the mean vector of other sections **after** applying the TF-IDF transformation. Facts is the mean vector of Circumstances and Relevant Law, and Full as mean of all the other parts. Thus, they are continuous contrarily to the other sections. In this work, we stacked the representations to obtain Facts and Full as HCBR does not handle properly continuous variables.

It has to be noted that for some parts, there are empty cases. For the Law part with Article 3, there are 162 empty cases (64.8% of the total casebase) with a prevalence of 17%, for the Article 6, 52 cases are empty with 17% prevalence and for the Article 8, 146 cases are empty for a 21% prevalence. This is not compatible with the hypothesis s.t. if two cases are described by the same features, they must have the same outcome used both by HCBR and SVM.

Therefore, the results are biased for both methods and could be improved by a better preprocessing step.<sup>11</sup>

For each of the seven sections (Full, Procedure, Circumstances, Relevant law, Facts, Law, Topics), we performed a 10-fold cross-validation and reported the accuracy.

### 5.2. Results

Table 6 summarizes the results. The accuracy is improved in 14 cases out of 21, deteriorated in 5 and remained unchanged in 2 cases. The improvement ranges from 1 to 20pp while the deterioration ranges from 1pp to 8pp.

All sections have seen their average accuracy unchanged or increased except for Topics. This is not surprising knowing Topics is made out of 30 continuous variables which are not correctly handled by HCBR. Surprisingly, the conclusions we can draw from this experiment are quite opposed to those of the reference study [76]. The full text was outperformed by using only the section Circumstances while here, taking the full text returns a much higher accuracy. The section with the best predictive power was Circumstances. It is now Relevant Law. The section Law was the worst predictor, it now one of the best, notably thanks to a gain of 20pp on Article 3.

In the original study, the best performances are obtained on Sections Topics and Topics and Circumstances with an accuracy of 0.78, 0.84 and 0.78 respectively for Article 3, 6 and 8. Thus, HCBR performed 2% lower on Article 3 and 6 and 1% better on Article 8. Knowing continuous variables are not properly handled yet, it is encouraging. Also, it turns out the best section for HCBR is the full text, which implies that in practice there is no need for feature engineering or time-consuming operations such as splitting the text into subsections. We are confident that HCBR could perform better with a larger number of tokens in the bag-of-word representation and let this for future work.

## 6. Intrinsic performances and properties

In this section, we validate the time complexity (Section 6.1) and discuss the model space limitations of HCBR (Section 6.2). We show how the hyperparameters can be used to control the confusion matrix in Section 6.3.

### 6.1. Computation time


We generated a casebase of *N* cases of size *m* s.t. case *i* is described by  $\{i, \dots, i+m\}$  i.e., each case is partitioned into *m* elements (one discretionary feature). This is the worst-case scenario in terms of the size of  $\mathcal{E}$  if  $m < N$  because the family grows exponentially in function of *m* or *N*. We split the computation time into constructing the hypergraph (and determining the intersection family) and calculating the strength of the partition. The results are illustrated in Fig. 8. By increasing *N* with a fixed *m*, the partition grows exponentially and thus, it is expected to have an exponential curve for the strength computation. On the contrary, building the hypergraph can be done in linear time when *m* fixed. When *N* is fixed and *m* increases, constructing the hypergraph is still doable in linear time as expected. Interestingly, calculating the strength has two phases: if  $m \leq N$ , increasing *m* exponentially increases the time (because  $\mathcal{E}$  exponentially increases) but if  $m > N$ , increasing *m* cannot result in an exponential growth in the computation time (because  $\mathcal{E}$  grows linearly).

<sup>11</sup> Ironically, this supports our claim that shortening the data preprocessing steps is necessary in order to avoid such hard-to-notice problems.

**Table 6**

Accuracy obtained by HCBR on the European Court of Human Rights dataset, depending on the article and the judgment section. The columns  $\Delta\text{acc}$  represent the difference with the reference study [76]. The color indicates if the result has been improved (green), deteriorated (red) or unchanged (white).

	Article 3	$\Delta\text{acc}$	Article 6	$\Delta\text{acc}$	Article 8	$\Delta\text{acc}$	Average	$\Delta\text{acc}$
Full	.76	+.06	.83	+.01	.77	+.05	.79	+.04
Procedure	.67	-	.81	-	.72	+.01	.73	-
Circumstances	.67	-.01	.81	-.01	.72	-.05	.73	-
Relevant law	.71	+.02	.86	+.08	.76	+.04	.78	+.05
Facts	.73	+.03	.76	-.04	.72	+.04	.74	+.01
Law	.76	+.20	.80	+.12	.73	+.11	.76	+.14
Topics	.70	-.08	.83	+.02	.74	-.02	.76	-.02



**Fig. 8.** On the left, computation time to build the model (hypergraph construction + strength calculation) depending on  $N$  ( $m = 10$ ), and on the right, depending on  $m$  ( $N = 100$ ). The case  $i$  is described by  $\{i, \dots, i + m\}$  s.t. each case is partitioned into  $m$  elements (one discretionary feature).

## 6.2. Model space limitations

We now study the learning curves to show the model space limitations and propose some extensions left for future work. We also discuss the tradeoff between model locality and generalization.

### 6.2.1. Learning curves

The learning curves are useful to study the limit of the model space on different datasets. It consists in plotting the accuracy in function of the training set size for both the training and the test sets. For the training set, it is expected to observe an accuracy starting close to 1 and decreasing fast to reach a plateau. A low stationary accuracy value indicates a high bias in the model or/and an irreducible error contained in the dataset, such as noisy or uninformative features. Oppositely, the accuracy on the test sets starts close to 0 as the training set is small and should increase until a plateau which is very often expected to be lower than the accuracy of the training set. If the training set curve converges toward a much higher value than the test set curve, then the model has a large variance. In other words, to achieve a good bias–variance tradeoff, the accuracy of both curves should converge toward more or less the same value, expected as high as possible.

The learning curves are calculated as the average over 10 runs with random splits and are shown in Figs. 9 and 10. The first remark concerns the variance. On all datasets, the variance is very low and the accuracy on the test set extremely close to accuracy of the training set. Despite a relatively low accuracy, HCBR seems to reach a good bias–variance compromise on heart suggesting a more complex model space might help. In general, the remark applies on the four datasets displayed in Fig. 9 as the observed error rate results from bias and not only from irreducible error. Indeed, the literature comparison provided by Table 2 proves that the accuracy could be improved.

For phishing and skin, we arrive at the same conclusion. However, for those datasets, a heuristic is used during the prediction phase. We discuss its implication on the learning curves in the Additional Material, Section 6. Finally, the learning curve of adult in Fig. 10 shows once again a small variance but high bias.

The analysis of the learning curves indicates that the main limitation of HCBR lies in a large bias. This may come either from the model space complexity or the model selection method that does not properly fit the parameters. The number of parameters in the model is equal to the cardinality of  $\mu$ , itself proportional to the number of atoms in the training set. However, a closer look at how a decision is taken reveals that each case has its own *small model* as a convex combination of its features. In other words, the real number of parameters to model the decision for a given case never exceed its number of features. This might not be enough to represent fairly complex functions. To discard the second hypothesis, we conducted additional investigations on the model selection in the next section.

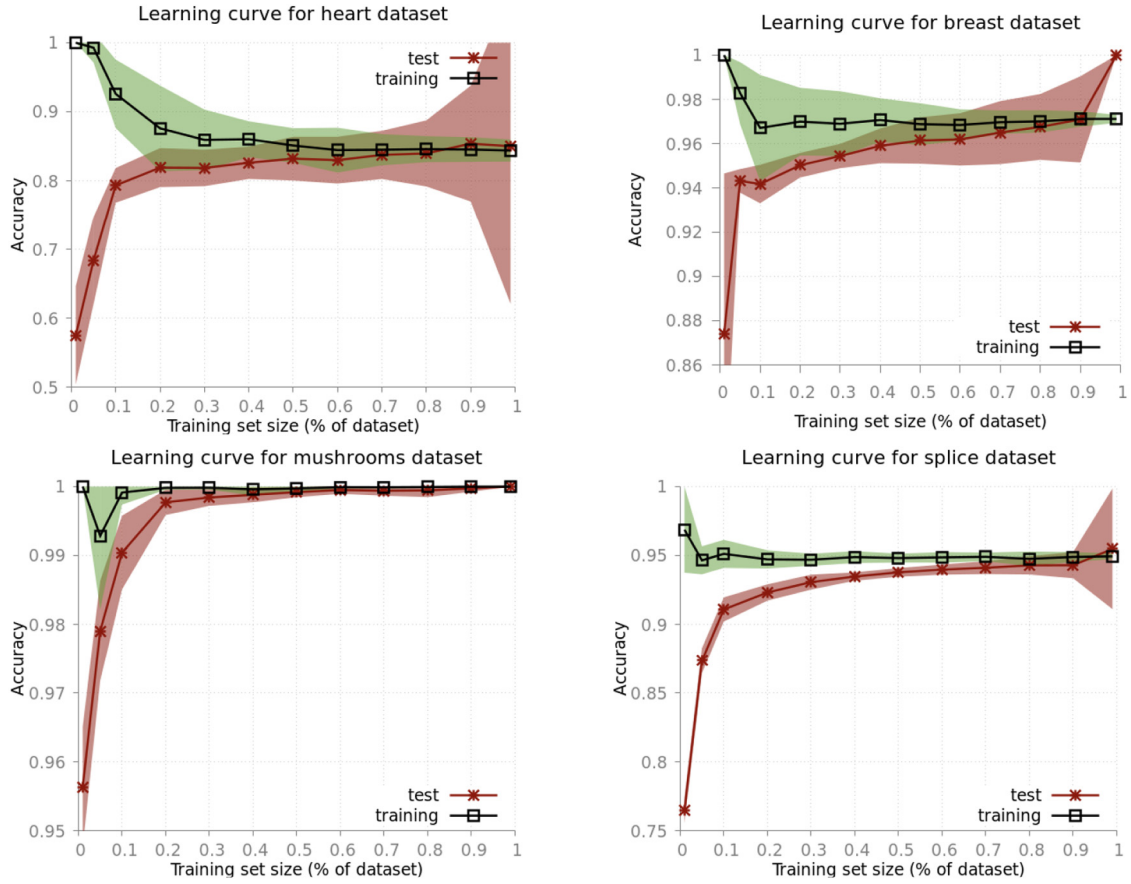
### 6.2.2. Assessing model space limitations

We are interested in understanding whether it is possible for HCBR to overfit the training set or at least improve significantly the accuracy or MCC on the training set. We consider the matricial formulation of the support,


$$\mathbf{s} = W\mu \quad (24)$$

and we are interested in knowing if there exists a  $\mu$  s.t.  $\mathbf{s}$  leads to classify correctly all the examples. As we already verified that the initial  $\mu$  provides better results than baseline models, we would like to perturbate as little as possible  $\mu$ . For this, we calculate  $\mathbf{k}$  s.t.  $\mathbf{s} + \mathbf{k}$  would correctly classify all the examples. We are looking for  $\delta$  s.t.,

$$\mathbf{s} + \mathbf{k} = W(\mu + \delta) \quad (25)$$



**Fig. 9.** Learning curves for heart, breast, mushrooms and splice datasets. Despite a bias and/or irreducible error observed on heart, breast and splice. The model variance appears very low on all datasets. Conversely, the bias and/or irreducible error ranges from very low on mushrooms, relatively low on breast and splice, to high for heart.



**Fig. 10.** Learning curve for adult. On the right, the training set size is restricted from 0 to 30% of the dataset.

$$\Leftrightarrow \mathbf{s} + \mathbf{k} = W\mu + W\delta \tag{26}$$

$$\Leftrightarrow \mathbf{k} = W\delta \tag{27}$$

$$\Rightarrow \delta = W^+\mathbf{k} + [I - W^+W]\mathbf{w}, \forall \mathbf{w} \in \mathbb{R}^N \tag{28}$$

with  $W^+$  the Moore–Penrose pseudo-inverse of  $W$ . When  $W$  is not of full rank, the solutions of the undetermined system are given for any vector  $\mathbf{w} \in \mathbb{R}^N$ , however, it can be shown that  $\delta = W^+\mathbf{k}$  is a least-square minimizer, i.e.  $\forall \mathbf{x} \in \mathbb{R}^M, \|\mathbf{W}\mathbf{x} - \mathbf{k}\|_2 \geq \|W\delta - \mathbf{k}\|_2$ . In particular, if  $\|W\delta - \mathbf{k}\|_2 = 0$ , then all the elements would be correctly classified.

Of course, it might not be possible to obtain  $\delta$  s.t.  $\|W\delta - \mathbf{k}\|_2 = 0$ , but it does imply that there exist no couple  $(\mathbf{k}, \delta)$  s.t. the accuracy is 1. Solving directly for all such couples seem to be a difficult problem without additional assumptions and thus, we adopted a slightly different approach: instead of fixing  $\mathbf{k}$  a priori, we formulated the problem as optimizing the Matthew Correlation Coefficient:

$$\delta^* = \max_{\delta \in \mathbb{R}^M} \text{MCC}(\delta, \mu, W) - c\|\delta\|_2^2 \tag{29}$$

where MCC is the Matthew Correlation Coefficient associated to  $\mathbf{s} = W\delta$  and  $c$  a regularization factor. The regularization factor



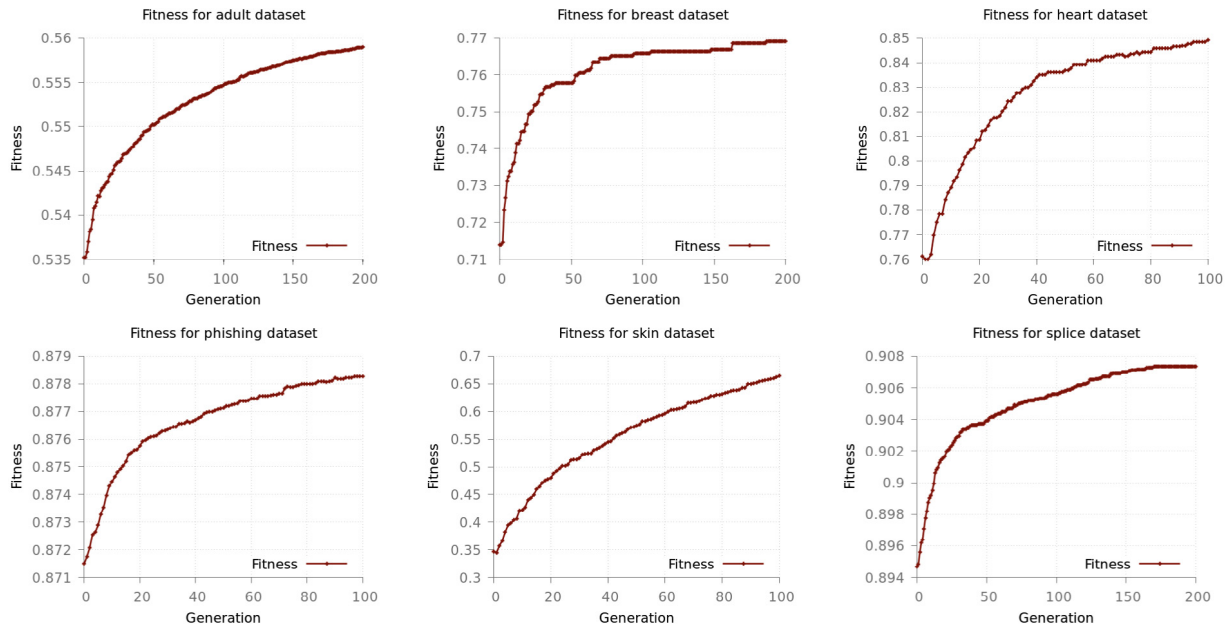


Fig. 11. Maximal fitness value in the population in function of generations for adult, breast, heart, phishing, skin and splice. The fitness is lower than the MCC by definition of (29).

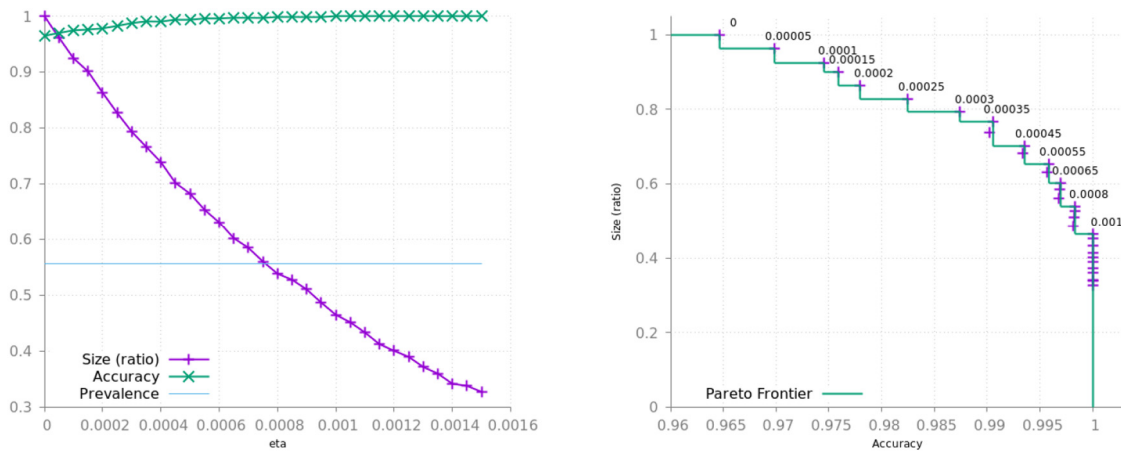


Fig. 12. Influence of  $\eta$  on phishing dataset.

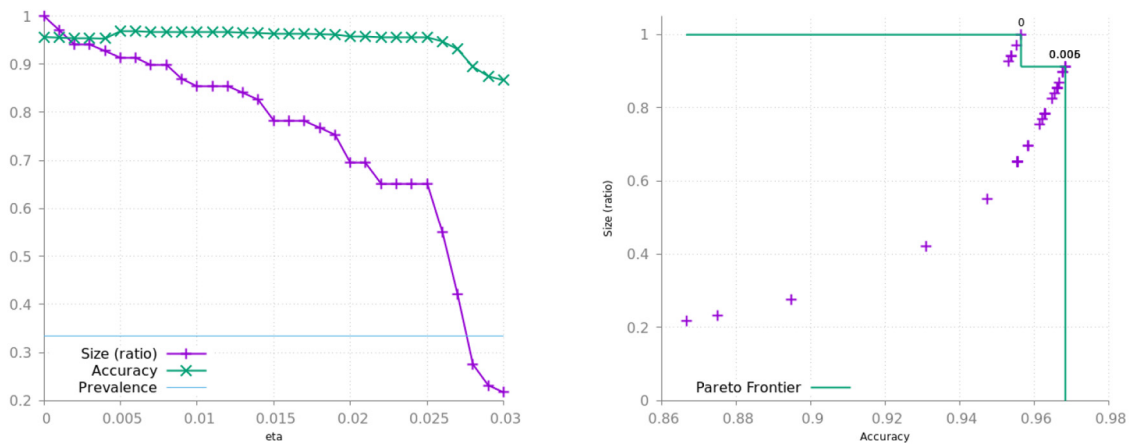


Fig. 13. Influence of  $\eta$  on breast dataset.

translates the idea that a smaller perturbation to obtain the same MCC is better than a larger one. We arbitrarily set  $c$  to 0.1. Despite

further work could be needed to determine a more tailored value, the conclusions drawn from this section would remain valid.


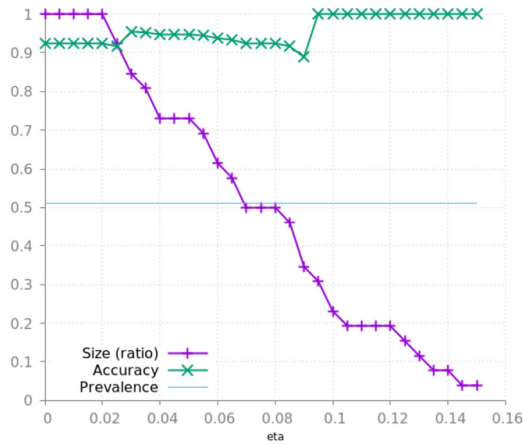


Fig. 14. Influence of  $\eta$  on heart dataset.

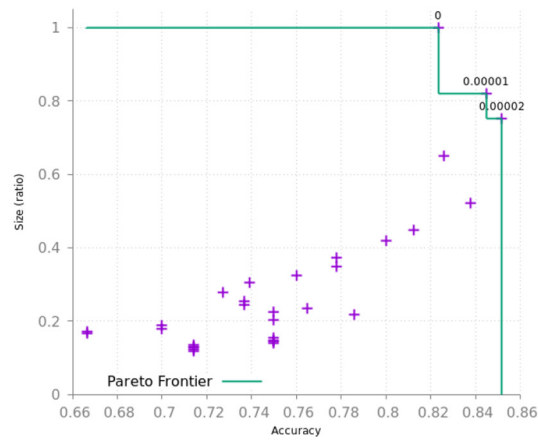
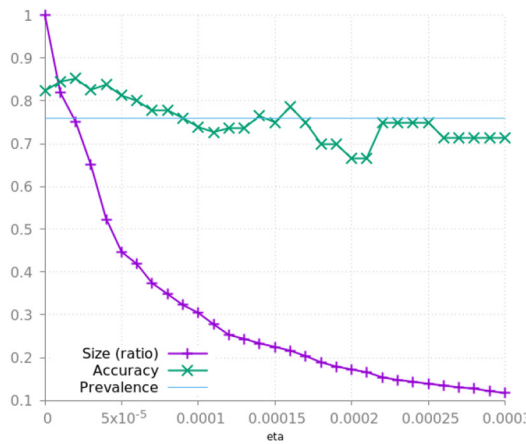


Fig. 15. Influence of  $\eta$  on adult dataset.

To solve (29), we used a  $(\mu + \lambda)$  genetic algorithm with an evolution strategy. Each individual is made of two vectors:  $\delta \in \mathbb{R}^M$  and a  $\nu \in \mathbb{R}^M$  representing the mutation probability of each corresponding component of  $\delta$ . The mutation operator is a centered gaussian perturbation and the crossover a 2-points crossover. The implementation is provided by DEAP [77].

We performed 10 runs of a 10-fold cross-validation with random split and compared to the result obtained without the optimization process. The dataset `mushrooms` has been discarded as HCBR reached an accuracy of 1. We set the population to 100, and for each dataset, we adjusted the number of generations and the standard deviation of the gaussian mutation manually (see the Additional Material, Section 7). To set the standard deviation of the mutation, we used  $\sigma = \frac{\mu^-}{\alpha}$  where  $\alpha$  is a factor determined empirically, and  $\mu^-$  defined by  $\min_{i,j} \mu_i - \mu_j$  i.e. the minimal difference between two components of  $\mu$ . The rationale behind is that, once again, we would like to slightly perturbate  $\mu$  and it is reasonable to think that a perturbation should be small enough not to directly switch the estimation of two elements of  $\mu$ . To conclude, we used a Wilcoxon signed-rank test at 5% and 1% on the test MCC obtained with and without the optimization process. There are three possible scenarios: (1) if the MCC can be improved on the training set and on the test set, it means the problem comes from the model selection method (estimation of  $\mu$  and training) and results might be improved without changing the model space, (2) if the MCC can be improved on the training set but remains the same or deteriorate on the test set, the

model space can represent the training set but overfits, and thus, should be extended (3) if the MCC cannot be improved even on the training set, then the model space is definitely not capable of representing properly the underlying decision mapping and should be extended.

A summary of the results are provided in Table 7 and the evolution of the fitness by Fig. 11. A look at Fig. 11 confirms that the genetic algorithm converged or was closed to converged and was able to optimize the cost function. In all cases except `skin`, the optimization procedure succeeded to find a vector  $\delta$  that yields a better MCC on the test sets. However, the improvement is quantitatively different from a dataset to another. For instance, on `heart` the absolute difference in MCC is 16% (relative difference: 23.64%) while on `phishing` the difference is barely 1%. In general, the higher is the initial MCC and the less the improvement is visible.

The variations of MCC on the test set are mitigated. The procedure returns significant changes in only two cases (`adult`) and (`skin`) for which one is improved (`adult`) and one deteriorated (`skin`). It indicates that the model selection and training phase described in Sections 3.3 and 3.4 return one of the best MCC achievable within the model space. More precisely, the vector  $\mu$  represents one of the best support approximation in its neighborhood. Notice that for `skin`, the optimization process deteriorated the MCC on the training set. A smaller mutation factor might help, however we believe this would barely change the results and thus entail the conclusion on the model space limitation.

**Table 7**

Results obtained on solving (29). The columns *initial MCC*,  $\Delta$  *MCC training* and  $\Delta$  *MCC test* represent respectively the initial MCC on the training set, the difference of MCC obtain on the training with and without the genetic algorithm, the difference of MCC obtain on the test sets with and without the genetic algorithm. *WSR  $r\%$*  indicates the result of the Wilcoxon signed-rank test at  $r\%$  risk (yes for significant difference in the sample median, no otherwise).

Dataset	Initial MCC	$\Delta$ MCC training	$\Delta$ MCC test	WSR 5%	WSR 1%
adult	0.5190	0.0400	0.0084	Yes	Yes
breasts	0.9360	0.0312	-0.0025	No	No
heart	0.6912	0.1634	-0.0023	No	No
phishing	0.8690	0.0093	0.0002	No	No
skin	0.8432	-0.0242	-0.0118	Yes	Yes
splice	0.8866	0.0208	0.0006	No	No

### 6.3. Hyperparameters $\eta$ to control the accuracy

We showed in Section 4.3.1 that the hyperparameters can increase the overall performances while the impact is limited by the model space. In this experiment, we show how  $\eta$  can be used to control the accuracy by specifying a threshold on the risk associated to a prediction.

We used a 90–10 split and set  $\eta_{-1} = \eta_1$  to ease the visualization. Instead of using the decision function defined by (R2), we did not produce a prediction if the constraints  $C_1$  or  $C_0$  were not respected. It can be viewed as creating a third class *unknown* for which we consider HCBR cannot produce a decision. We measured the accuracy and the test set size ratio for which a prediction has been produced for different values of  $\eta := \eta_{-1} = \eta_1$ . If the model correctly approximates the underlying mapping function  $J$ , increasing  $\eta$  should increase the accuracy while the test set ratio should remain high. Additionally, we plot the test set ratio in function of the accuracy and calculate the Pareto frontier<sup>12</sup> which represents the best compromises accuracy/ratio. The closer the points are to (1, 1) the better it is. A Pareto frontier consisting of (1, 1) represents the perfect model (e.g. reached on mushroom dataset). Figs. 12–15 provides the result for the best and worst two datasets. Fig. 16 shows all of the four Pareto frontiers.

As expected, the results are better on phishing and breast. On phishing, breast and heart, the accuracy globally increases with  $\eta$  while on heart the accuracy slightly decreases indicating poor influence of the hyperparameters and model.

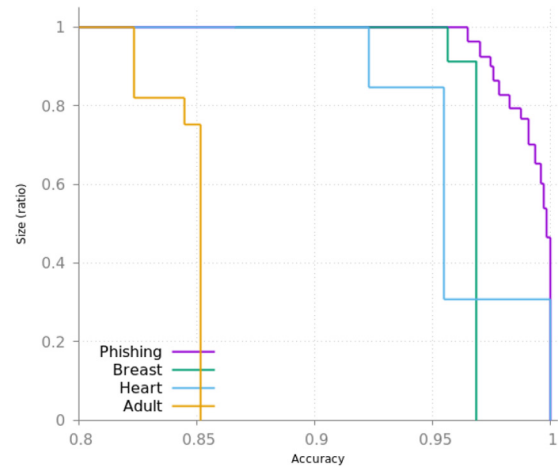
Notice that for certain values of  $\eta$  it is possible to reach 100% accuracy with heart (sacrificing over 70% of the dataset) while it is not with breast. Also, for high values of  $\eta$ , we observe a fall in accuracy for breast. We suspect those two phenomena to appear because we used the same value for  $\eta_0$  and  $\eta_1$ .

## 7. Discussion

### 7.1. On the hypergraph representation

The reader may have observed that the model space has a purely set-theoretic interpretation, and no hypergraph-specific property is used so far. Another possibility would be to use a bipartite graph with the first class being the cases of  $\mathbf{X}$  and the second the elements of  $\mathcal{E}$ . We justify viewing the method from a hypergraph perspective by three axes currently being investigated:

<sup>12</sup> Points s.t. improving one component would deteriorate the other one.



**Fig. 16.** Pareto Frontiers comparison.

- **Model space extension.** The current model space is not complex enough (see Section 6.2). An extension using hyperpaths is proposed in Section 7.2. Sets or graphs are not suitable to manipulate such space.
- **Justification.** Most justification techniques provide a justification about the model itself (e.g. a decision tree<sup>13</sup>) or provide hints about each decisions under the form of weights for each variable [78]. In particular, [79] unifies the literature and formulates the justification as learning a simple model. With HCBR we are exploring the possibility to generate explanations tailored for an element depending on its neighbors, in a case-based reasoning fashion. For this, we need a notion of neighborhood given by the hyperpaths. This is not possible with a set approach, feasible with graphs but less natural than with hypergraphs.
- **Performances.** Hypergraphs have computational advantages over graphs. Indeed, hypergraphs can be represented as graphs using clique expansion technique. However, increasing hyperedges cardinality leads to a larger increase in the graph counterpart [80].

### 7.2. On the model space extension

It is now clear that the model space is too limited. The number of parameters to describe a case is lower than  $m$ , the cardinal of the partition  $\mathcal{E}$ , and at most  $|\mathcal{X}|$ . The number of parameters to describe the whole datasets is  $m$ .

To increase the model space complexity, there are two paths to explore in future work:

1. **Increasing the number of parameters per case.** Currently, the support of  $\mathbf{x}$  is modeled by a linear combination over the partition of its features s.t. the interactions between the elements of the partition are not taken into account. A reasonable way to go is to use some combinations of those elements. For instance, if  $\mathbf{x}$  is partitioned into  $e_1, e_2$  instead of modeling the support as  $s(\mathbf{x}) = w(e_1, \mathbf{x})\mu(e_1) + w(e_2, \mathbf{x})\mu(e_2)$ , we would have  $s(\mathbf{x}) = w(e_1, \mathbf{x})\mu(e_1) + w(e_2, \mathbf{x})\mu(e_2) + w(e_{1 \cup 2}, \mathbf{x})\mu(e_{1 \cup 2})$ . This would allow to be more discriminative and solve the problem of cases having the same support.

<sup>13</sup> <https://github.com/andosa/treinterpreter>.

**2. Increasing the number of parameters of the whole model.** A hyperpath of length  $k$  is a sequence of hyperedges  $(\mathbf{x}_1, \dots, \mathbf{x}_k)$  such that  $\forall i \in \{1, \dots, k-1\}, \mathbf{x}_i \cap \mathbf{x}_{i+1} \neq \emptyset$ . Instead of modeling the support as a combination of the elements  $\mathbf{e}_i$  belonging to a specific case  $\mathbf{x}_i$ , it could be extended to include the elements  $\mathbf{e} \in \mathcal{E}$  belonging to the neighbor cases where a neighbor is a case that can be reached by a hyperpath of length  $k$ . Thus, the model space proposed in this paper would be the particular case with  $k = 0$ .

Naturally, it raises many questions, notably how to extend  $w$  and  $\mu$ . The second point seems to be the most interesting because it really requires the hypergraph representation and cannot be interpreted using set theory.

### 7.3. On the model locality

The generalization capacity of HCBR depends on the number of intersections between the examples. Conversely, if a new case does not intersect with the examples, it is impossible to generate a prediction. Therefore, the locality property depends on how much the examples cover the feature space.

Consider that it is possible to choose  $n$  examples of  $k$  features in a space with  $K$  features and  $K \gg nk$ . There is a tradeoff between covering as much space as possible and having enough intersections to construct a meaningful model. The first extreme configuration consists in maximizing the space cover: there is no intersection between cases, the accuracy during the training phase is 1 but the generalization capacity is null. The probability that a new case will intersect with some examples is maximized. The second extreme configuration maximizes the intersections between the examples: all cases intersect with each other, therefore minimizing the space cover and thus the probability that a new case will intersect with the training set. The accuracy on the training set might not be 1 but the generalization capacity is high compared to the previous case.

For configurations close to the first extreme, it is possible to generate more intersections by using clustering or discretization in the original feature space. For instance, “variable\_1 = v\_1” and “variable\_1 = v\_2” could be encoded the same if both values belong to the same cluster. This requires slightly more feature engineering that previously stated. For the second case, the problem lies in the training set itself. It is not specific to HCBR, but to the fact the training set is not representative of the underlying distribution that generates the cases to classify. In both cases, acquiring more data can help.

In the Additional Material, Section 8, we discuss a way of redefining the decision function to take into account this locality property.

## 8. Conclusion

This paper presented HCBR, a method for binary classification in unstructured space. The method can be seen as learning a metric that optimizes the classification score on the training set. Contrarily to most classification methods or metric learning methods, HCBR does not require to work in vector space and is agnostic to data representation. Therefore, it allows combining information from multiple sources by simply *stacking* the information. It does not require transforming the data to obtain satisfactory results.

The general framework introduced in Section 3.2 is instantiated in Sections 3.3 and 3.4 where the support is determined using all the interactions between the hyperedges. Beyond this generic implementation, one can imagine different model selection methods using some assumptions or prior on the data.

HCBR has been tested on seven well-known structured datasets and demonstrated similar accuracy when compared to the best results from the literature, with and without hyperparameter tuning. We showed that the model is properly calibrated. Additionally, we performed a comparison with nine alternative methods to find out HCBR, along with Neural Network, outperforms in average with a constant good result. Those experiments showed that HCBR can easily be used and deployed in practice, as it lowers the requirement for feature engineering, data pre-processing and hyperparameter tuning, i.e. the most consuming operations in practical machine learning nowadays.

In Section 5, we tested HCBR on unstructured datasets and showed it improves the accuracy in most cases compared to reference study.

We empirically validated the worst-case complexity. Finally, we studied the properties and limitations of the model space. We showed that the model selection procedure provides one of the best possible performance within the model space. Hence, further work will focus on extending the model space as proposed in Section 7.2.

This proof of concept raises many questions and offer many improvement axes. First, it seems relatively easy to extend the method to several classes, with a linear increase of the computation time. As calculating the class support represents most of the computational effort, working on an approximation of the main measure should be investigated. The solution may come from exploring the feature selection capacity of HCBR. It may be possible to remove from the partition some elements that are not discriminative enough, reducing the computation time.

Additionally, we plan to investigate explanation generation about each prediction, using the link between cases in a similar way a lawyer may use past cases to make analogies or counter-examples. We also work on an online version of HCBR where the hypergraph is constructed case after case, including forgetting some old cases (which would allow handling non-stationary environment). It seems possible not only to add new examples dynamically, but also some vertices (i.e. adding some pertinent information to some cases) without generating the whole model from scratch.

Last but not least, we would like to answer some questions: can we provide some quality bounds depending on the initial hypergraph configuration w.r.t. the number of intersections and space cover? How to handle continuous values without discretization?

## Acknowledgment

The author warmly thanks Dr. Jean-François Puget, IBM Analytics, for his useful suggestions and advice in order to improve this paper.

## References

- [1] A. Quemy, Binary classification with hypergraph case-based reasoning, in: Proc. Int. Workshop Des. Optim. Lang. Anal. Process. Big Data, 2018.
- [2] S. Boucheron, O. Bousquet, G. Lugosi, Theory of classification: A survey of some recent advances, *ESAIM Probab. Stat.* 9 (2005) 323–375.
- [3] V. Vapnik, *The nature of Statistical Learning Theory*, Springer science & business media, 2013.
- [4] F. Rosenblatt, The perceptron: a probabilistic model for information storage and organization in the brain, *Psychol. Rev.* 65 (6) (1958) 386.
- [5] D.R. Cox, The regression analysis of binary sequences, *J. R. Stat. Soc.* (1958) 215–242.
- [6] Y. Lin, A note on margin-based loss functions in classification, *Statist. Probab. Lett.* 68 (1) (2004) 73–82.
- [7] N. Friedman, D. Geiger, M. Goldszmidt, Bayesian network classifiers, *Mach. Learn.* 29 (2–3) (1997) 131–163.
- [8] W.N. Venables, B.D. Ripley, Tree-based methods, in: *Modern Applied Statistics with S*, Springer, 2002, pp. 251–269.



- [9] L. Breiman, *Classification and Regression Trees*, Routledge, 2017.
- [10] T.K. Ho, Random decision forests, in: Proc. Int. Conf. Doc. Anal. Recognit, vol. 1, IEEE, 1995, pp. 278–282.
- [11] L. Breiman, Random forests, *Mach. Learn.* 45 (1) (2001) 5–32.
- [12] R. Couronné, P. Probst, A.-L. Boulesteix, Random forest versus logistic regression: a large-scale benchmark experiment, *BMC Bioinformatics* 19 (1) (2018) 270.
- [13] P. Geurts, D. Ernst, L. Wehenkel, Extremely randomized trees, *Mach. Learn.* 63 (1) (2006) 3–42.
- [14] J.H. Friedman, Greedy function approximation: a gradient boosting machine, *Ann. Statist.* (2001) 1189–1232.
- [15] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436.
- [16] J. Schmidhuber, Deep learning in neural networks: An overview, *Neural Netw.* 61 (2015) 85–117.
- [17] J.V. Davis, B. Kulis, P. Jain, S. Sra, I.S. Dhillon, Information-theoretic metric learning, in: Proc. Int. Conf. Mach. Learn, ACM, 2007, pp. 209–216.
- [18] A. Bellet, A. Habrard, M. Sebban, A survey on metric learning for feature vectors and structured data, CoRR abs/1306.6709 (2013).
- [19] P.C. Mahalanobis, On the generalised distance in statistics, *Proc. Nat. Inst. Sci. India* 2 (1936) 49–55.
- [20] F. Wang, J. Sun, Survey on distance metric learning and dimensionality reduction in data mining, *Data Min. Knowl. Discov.* 29 (2) (2015) 534–564.
- [21] K.Q. Weinberger, J. Blitzer, L.K. Saul, Distance metric learning for large margin nearest neighbor classification, *Adv. Neural Inf. Process. Syst.* (2006) 1473–1480.
- [22] K.Q. Weinberger, L.K. Saul, Distance metric learning for large margin nearest neighbor classification, *J. Mach. Learn. Res.* 10 (Feb) (2009) 207–244.
- [23] G. Chechik, V. Sharma, U. Shalit, S. Bengio, Large scale online learning of image similarity through ranking, *J. Mach. Learn. Res.* 11 (2010) 1109–1135.
- [24] K. Crammer, O. Dekel, J. Keshet, S. Shalev-Shwartz, Y. Singer, Online passive-aggressive algorithms, *J. Mach. Learn. Res.* 7 (Mar) (2006) 551–585.
- [25] A. Bellet, A. Habrard, M. Sebban, Similarity learning for provably accurate sparse linear classification, in: Proc. Int. Conf. Mach. Learn, United Kingdom, 2012.
- [26] A. Kocsor, K. Kovács, C. Szepesvári, Margin maximizing discriminant analysis, in: Eur. Conf. Mach. Learn, Springer, 2004, pp. 227–238.
- [27] S. Chopra, R. Hadsell, Y. LeCun, Learning a similarity metric discriminatively, with application to face verification, in: *Comput. Vision Pattern Recognit*, vol. 1, IEEE, 2005, pp. 539–546.
- [28] J. Oncina, M. Sebban, Learning stochastic edit distance: application in handwritten character recognition, *Pattern Recognit.* 39 (2006) 1575–1587.
- [29] N. Dalvi, P. Bohannon, F. Sha, Robust web extraction: An approach based on a probabilistic tree-edit model, in: Proc. Int. Conf. Manage. Data, in: SIGMOD, ACM, New York, NY, USA, 2009, pp. 335–348.
- [30] M. Bernard, A. Habrard, M. Sebban, Learning stochastic tree edit distance, in: *Eur. Conf. Mach. Learn*, Springer, 2006, pp. 42–53.
- [31] M. Gessell, F. Scheepers, N. Nguyen, R. van Kessel, R. van der Starre, Governing and managing big data for analytics and decision makers.
- [32] S.F. Crone, S. Lessmann, R. Stahlbock, The impact of preprocessing on data mining: An evaluation of classifier sensitivity in direct marketing, *European J. Oper. Res.* 173 (3) (2006) 781–800.
- [33] T. Dasu, T. Johnson, *Exploratory Data Mining and Data Cleaning*, vol. 479, John Wiley & Sons, 2003.
- [34] T. Furche, G. Gottlob, L. Libkin, G. Orsi, N. Paton, Data Wrangling for Big Data: Challenges and Opportunities, OpenProceedings.org, 2016.
- [35] S. Kandel, J. Heer, C. Plaisant, J. Kennedy, F. van Ham, N.H. Riche, C. Weaver, B. Lee, D. Brodbeck, P. Buono, Research directions in data wrangling: Visualizations and transformations for usable and credible data, *Inf. Vis.* 10 (4) (2011) 271–288.
- [36] B. Bilalli, A. Abelló, T. Aluja-Banet, On the predictive power of meta-features in openml, *Int. J. Appl. Math. Comput. Sci.* 27 (4) (2017) 697–712.
- [37] L.M. Haas, M.A. Hernández, H. Ho, L. Popa, M. Roth, Clío grows up: from research prototype to industrial tool, in: Proc. Int. Conf. Manage. Data, ACM, 2005, pp. 805–810.
- [38] G.G. Robertson, M.P. Czerwinski, J.E. Churchill, Visualization of mappings between schemas, in: Proc. SIGCHI Conf Hum. Factors Comput. Syst., ACM, 2005, pp. 431–439.
- [39] M. Vartak, S. Rahman, S. Madden, A. Parameswaran, N. Polyzotis, See db: efficient data-driven visualization recommendations to support visual analytics, *Proc. VLDB Endow.* 8 (13) (2015) 2182–2193.
- [40] A. Quemy, Data pipeline selection and optimization, in: Proc. Int. Workshop Des. Optim. Lang. Anal. Process. Big Data, 2019.
- [41] G.E. Batista, M.C. Monard, et al., A study of k-nearest neighbour as an imputation method, *HIS* 87 (251–260) (2002) 48.
- [42] C. Preda, G. Saporta, M.H.B.H. Mbarek, The NIPALS algorithm for missing functional data, *Rev. Roumaine Math. Pures Appl.* 55 (4) (2010) 315–326.
- [43] D.J. Stekhoven, P. Bühlmann, Missforest—non-parametric missing value imputation for mixed-type data, *Bioinf.* 28 (1) (2011) 112–118.
- [44] M. Kuhn, K. Johnson, *Applied Predictive Modeling*, Vol. 26, Springer, 2013.
- [45] P. Royston, et al., Multiple imputation of missing values, *Stata J.* 4 (3) (2004) 227–241.
- [46] S. Van Buuren, *Flexible Imputation of Missing Data*, Chapman and Hall/CRC, 2018.
- [47] S. Lifan, G. Yue, Z. Xibin, W. Hai, G. Ming, S. Jiaguang, Vertex-weighted hypergraph learning for multi-view object classification, in: Proc. Int. Joint Conf. Artif. Intell., 2017, pp. 2779–2785.
- [48] Y. Huang, Q. Liu, S. Zhang, D.N. Metaxas, Image retrieval via probabilistic hypergraph ranking, in: Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit, 2010, pp. 3376–3383.
- [49] D. Zhou, J. Huang, B. Schölkopf, Learning with hypergraphs: Clustering, classification, and embedding, in: B. Schölkopf, J.C. Platt, T. Hoffman (Eds.), *Adv. Neural Inf. Process. Syst.* MIT Press, 2007, pp. 1601–1608.
- [50] C. Berge, *Hypergraphs: Combinatorics of Finite Sets*, Vol. 45, Elsevier, 1984.
- [51] R. Paige, R.E. Tarjan, Three partition refinement algorithms, *SIAM J. Comput.* 16 (6) (1987) 973–989.
- [52] D. Chicco, Ten quick tips for machine learning in computational biology, *BioData Min.* 10 (1) (2017).
- [53] M. Feurer, A. Klein, K. Eggensperger, J. Springenberg, M. Blum, F. Hutter, Efficient and robust automated machine learning, in: C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, R. Garnett (Eds.), *Adv. Neural Inf. Process. Syst.* Curran Associates, Inc., 2015, pp. 2962–2970.
- [54] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [55] R.P. Datta, S. Saha, Applying rule-based classification techniques to medical databases: an empirical study, *Int. J. Bus. Intel. Syst. Eng.* 1 (1) (2016) 32–48.
- [56] L. Jiang, Learning instance weighted naive bayes from labeled and unlabeled data, *J. Intell. Inf. Syst.* 38 (1) (2012) 257–268.
- [57] A.M. Sagir, S. Sathasivam, A hybridised intelligent technique for the diagnosis of medical diseases, *Pertanika J. Sci. Tech.* 25 (2) (2017).
- [58] S. Das, Filters, wrappers and a boosting-based hybrid for feature selection, in: Proc. Int. Conf. Mach. Learn, 2001, pp. 74–81.
- [59] W. Hadi, G. Issa, A. Ishtaiwi, ACRPRISM: Associative classification based on PRISM algorithm, *Inform. Sci.* 417 (Supplement C) (2017) 287–300.
- [60] F. Thabtah, R.M. Mohammad, L. McCluskey, A dynamic self-structuring neural network model to combat phishing, in: *Int. Jt. Conf. Neur. Net*, 2016, pp. 4221–4226.
- [61] G. Kou, Y. Lu, Y. Peng, Y. Shi, Evaluation of classification algorithms using MCDM and rank correlation, *Int. J. Inf. Tech. Decis. Making* 11 (01) (2012) 197–225.
- [62] Y.-J. Lee, O. Mangasarian, SSVN: A smooth support vector machine for classification, *Comput. Optim. Appl.* 20 (1) (2001) 5–22.
- [63] R.B. Bhatt, G. Sharma, A. Dhall, S. Chaudhury, Efficient skin region segmentation using low complexity fuzzy decision tree model, in: *IEEE India Conf*, 2009, pp. 1–4.
- [64] F.Ö. Çatak, Classification with boosting of extreme learning machine over arbitrarily partitioned data, *Soft Comput.* 21 (9) (2017) 2269–2281.
- [65] H.-L. Chen, B. Yang, J. Liu, D.-Y. Liu, A support vector machine classifier with rough set-based feature selection for breast cancer diagnosis, *Expert Syst. Appl.* 38 (7) (2011) 9014–9022.
- [66] K. Polat, S. Güneş, Breast cancer diagnosis using least square support vector machine, *Digit. Signal Process.* 17 (4) (2007) 694–701.
- [67] M.F. Akay, Support vector machines combined with feature selection for breast cancer diagnosis, *Expert Syst. Appl.* 36 (2) (2009) 3240–3247.
- [68] A. Marciano-Cedeño, J. Quintanilla-Domínguez, D. Andina, WBCD breast cancer database classification applying artificial metaplasticity neural network, *Expert Syst. Appl.* 38 (8) (2011) 9573–9579.
- [69] E.D. Übeyli, Implementing automated diagnostic systems for breast cancer detection, *Expert Syst. Appl.* 33 (4) (2007) 1054–1062.
- [70] A. Fallahi, S. Jafari, An expert system for detection of breast cancer using data preprocessing and Bayesian network, *Int. J. Adv. Sci. Tech.* 34 (2011) 65–70.
- [71] J.R. Quinlan, Improved use of continuous attributes in C4.5, *J. Artificial Intelligence Res.* 4 (1996) 77–90.
- [72] M.T. Cazzolato, M.X. Ribeiro, A statistical decision tree algorithm for medical data stream mining, in: Proc. IEEE Int. Symp. Comput.-Based Med. Syst., 2013, pp. 389–392.
- [73] S. Basterrech, A. Mesa, N.-T. Dinh, Generalized linear models applied for skin identification in image processing, in: *Intell. Data Anal. Appl*, Springer, 2015, pp. 97–107.
- [74] M.A.U.H. Tahir, S. Asghar, A. Zafar, S. Gillani, A hybrid model to detect phishing-sites using supervised learning algorithms, in: *Int. Conf. Comput. Sci. Comp. Intell*, 2016, pp. 1126–1133.
- [75] H. He, E.A. Garcia, Learning from imbalanced data, *IEEE Trans. Knowl. Data Eng.* 21 (9) (2009) 1263–1284.

- [76] N. Aletras, D. Tsarapatsanis, D. Preotiuc-Pietro, V. Lampos, Predicting judicial decisions of the european court of human rights: a natural language processing perspective, *PeerJ Comput. Sci.* 2 (2016) e93.
- [77] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau, C. Gagné, DEAP: Evolutionary algorithms made easy, *J. Mach. Learn. Res.* 13 (2012) 2171–2175.
- [78] M.T. Ribeiro, S. Singh, C. Guestrin, Why should I trust you?: Explaining the predictions of any classifier, in: *Proc. Int. Conf. Knowl. Discov. Data Min.*, ACM, 2016, pp. 1135–1144.
- [79] S.M. Lundberg, S.I. Lee, A unified approach to interpreting model predictions, in: *Adv. Neural Inf. Process. Syst.*, 2017, pp. 4765–4774.
- [80] L. Pu, B. Faltings, Hypergraph learning with hyperedge expansion, in: *Jt. Eur. Conf. Mach. Learn. Knowl. Discov. Datab.*, Springer, 2012, pp. 410–425.