



Article Feature Selection based on the Local Lift Dependence Scale

Diego Marcondes *,† ^(D), Adilson Simonis [†] and Junior Barrera ^{†,‡}

Institute of Mathematics and Statistics, University of São Paulo, São Paulo 05508-090, Brazil; asimonis@ime.usp.br (A.S.); jb@ime.usp.br (J.B.)

- * Correspondence: dmarcondes@ime.usp.br; Tel.: +55-11-97585-6644
- + These authors contributed equally to this work.
- ‡ Junior Barrera was supported by grants 2013/07467-1 and 2015/01587-0, São Paulo Research Foundation (FAPESP).

Received: 11 November 2017; Accepted: 25 January 2018; Published: 30 January 2018

Abstract: This paper uses a classical approach to feature selection: minimization of a cost function applied on estimated joint distributions. However, in this new formulation, the optimization search space is extended. The original search space is the Boolean lattice of features sets (BLFS), while the extended one is a collection of Boolean lattices of ordered pairs (CBLOP), that is (features, associated value), indexed by the elements of the BLFS. In this approach, we may not only select the features that are most related to a variable *Y*, but also select the values of the features that most influence the variable or that are most prone to have a specific value of *Y*. A local formulation of Shannon's mutual information, which generalizes Shannon's original definition, is applied on a CBLOP to generate a multiple resolution scale for characterizing variable dependence, the *Local Lift Dependence Scale* (LLDS). The main contribution of this paper is to define and apply the LLDS to analyse local properties of joint distributions that are neglected by the classical Shannon's global measure in order to select features. This approach is applied to select features based on the dependence between: i—the performance of students on university entrance exams and on courses of their first semester in the university; ii—the congress representative party and his vote on different matters; iii—the cover type of terrains and several terrain properties.

Keywords: feature selection; local lift dependence scale; mutual information; variable dependence; variable selection

1. Introduction

The problem of feature selection is equivalent to the problem of nonparametric estimation of a discrete joint distribution $\mathbb{P}(X, Y)$ from a sample of *n* pairs $\{(x_1, y_1), \ldots, (x_n, y_n)\}$, in which *X* is an m-dimensional real or integer vector of features, i.e., measures of phenomenon characteristics, and *Y* is a natural number, which represents the class of *X*. The problem is to find a subspace of characteristics χ with dimension k < m that permits to properly estimate $\mathbb{P}(\chi, Y)$, emphasizing one of its important characteristics: for example, that χ is a good predictor of *Y*, i.e., for all *x*, the conditional distribution $\mathbb{P}(Y | \chi = x)$ has its mass concentrated around some value $Y = y_x$.

Formally, let *X* be an m-dimensional feature vector and *Y* a single variable. Let χ be a feature vector, whose features are also in *X*, and denote $\mathcal{P}(X)$ as the set of all feature vectors whose features are also in *X*. In this scenario, we define the classical approach to feature selection, in which the search space is the Boolean lattice of features sets (BLFS), as follows.

Definition 1. Given a variable Y, a feature vector X and a cost function $C_Y : \mathcal{P}(X) \to \mathbb{R}^+$ calculated from the estimated joint distribution of $\chi \in \mathcal{P}(X)$ and Y, the classical approach to feature selection consists in finding a subset $\chi \in \mathcal{P}(X)$ of features such that $C_Y(\chi)$ is minimum.

In light of Definition 1, we note that some families of feature selection algorithms may be considered as classical approaches. In fact, according to the taxonomy of feature selection, as presented in [1] for example, feature selection algorithms may be divided into three families, *filters, wrappers* and *embedded* methods, being the last two classical approaches to feature selection. Indeed, in the *wrappers* methods, the feature selection algorithm exists as a wrapper around a learning machine (or induction algorithm), so that a subset of features is chosen by evaluating its performance on the machine [2]. Furthermore, in the *embedded* methods, a subset of features is also chosen based on its performance on a learning machine, although the feature selection and the learning machine cannot be separated [3]. Therefore, both *wrappers* and *embedded* methods satisfy Definition 1, as the performance on the learning machine may be established by a cost function, so that these methods are special cases of the classical approach to feature selection. For more details about these methods see [1–7].

Under this approach for feature selection, a classical choice for the cost function is the estimated mean conditional entropy [8] that measures the mean mass concentration of the conditional distribution. Great mass concentration indicates that the chosen features define equivalence classes with almost homogeneous classifications, hence it is a good choice to represent the complete set of features. For a joint distribution estimated by a sample of size *n*, the curve formed by this cost function applied in a chain of the BLFS has an U shape and is called U-curve. For a small set of features, in the left side of the U-curve, the cost is high for the small amount of features generates large equivalence classes that mix labels, which leads to high entropy. For a large set of features, in the right side of the U-curve, the cost is also high, for severe conditional distribution estimation error leads to high entropy. Therefore, the ideal set of features in this chain is in the U-curve minimum, that is achieved by a set that contains the maximum number of features, whose corresponding distribution estimation is not seriously affected by estimation error. Choosing the best set of features consists of comparing the minimum of all lattice chains. There are some NP-hard algorithms that find the absolute minimum [9,10]. There are also some heuristics that give approximate solutions such as Sequential Forward Selection (SFS) [11], which adds features progressively until it finds a local minimum, and Sequential Forward Floating Selection (SFFS) [11], which, at first, adds features, but after takes some of them out and adds others, trying to improve the first local minimum found.

The main goal of the classical approach is to select the features that are most related to *Y* according to a metric defined by a cost function. Although useful in many scenarios, this approach may not be suitable in some applications in which it is of interest to select not only the features that are most related to *Y*, but also the features' values that most influence *Y*, or that are most *prone* to have a specific value *y* of *Y*. Therefore, it would be relevant to extend the search space of the classical approach to an extended space that also contemplates the range of the features, so that we may select features and subsets of their range. This extended space is a collection of Boolean lattices of ordered pairs (features, associated values) (CBLOP) indexed by the elements of the BLFS. In other words, for each $\chi \in \mathcal{P}(X)$ we have the Boolean lattice that represents the powerset of its range R_{χ} , that is denoted by $\mathcal{P}(R_{\chi})$, and the CBLOP is the collection of these Boolean lattices, i.e., $\{\mathcal{P}(R_{\chi}) : \chi \in \mathcal{P}(X)\}$. If $X = (X_1, X_2)$ are Boolean features, then its CBLOP is as the one in Figure 1. Note that the circle nodes and solid lines form a BLFS, that around each circle node there is an associated Boolean lattice that represents the powerset of R_{χ} , for a $\chi \in \mathcal{P}(X)$, and that the whole tree is a CBLOP (Note that if the features *X* are not Boolean, we also have that $\mathcal{P}(\chi)$ is a Boolean lattice for all $\chi \in \mathcal{P}(X)$, so that the search space of the algorithm is always a CBLOP in this framework, regardless of the features' range).

A downside of this extension is that the sample size needed to perform feature selection at the extended space is greater than the one needed at the associated BLFS, what demands more refined optimal and sub-optimal algorithms in order to select features and subsets of their range. Nevertheless,

the extended space brings advances to the state-of-art in feature selection, as it expands the method to a new variety of applications. As an example of such applications, we may cite market segmentation. Suppose it is of interest to segment a market according to the products that each market share is most prone to buy. Denote *Y* as a discrete variable that represents the products sold by a company, i.e., $\mathbb{P}(Y = y)$ is the probability of an individual of the market buying the product $y \in \{1, ..., p\}$ sold by the company, and *X* as the socio-economic and demographic characteristics of the people that compose the market. In this framework, it is not enough to select the characteristics (features) that are most related to *Y*: we need to select, for each product (value of *Y*), the characteristics $\chi \in \mathcal{P}(X)$ and their values $W \in \mathcal{P}(R_{\chi})$ (the *profile* of the people) that are prone to buy a given product, so that feature selection must be performed on a CBLOP instead of a BLFS.

We call the approach to feature selection in which the search space is a CBLOP *multi-resolution*, for we may choose the features based on a global cost function calculated for each $\chi \in \mathcal{P}(X)$ (low resolution); or choose the features and a subset of their range based on a local cost function calculated for each $\chi \in \mathcal{P}(X)$ and $W \in \mathcal{P}(R_{\chi})$ (medium resolution); or choose the features and a point of their range based on a local cost function calculated for each $\chi \in \mathcal{P}(X)$ and $W \in \mathcal{P}(R_{\chi})$ (medium resolution); or choose the features and a point of their range based on a local cost function calculated for each $\chi \in \mathcal{P}(X)$ and $x \in R_{\chi}$ (high resolution). Formally, the *multi-resolution* approach to feature selection may be defined as follows.

Definition 2. Given a variable Y, a feature vector X and cost functions $C_Y^k : \mathcal{P}(X) \times \mathbb{R}^k \to \mathbb{R}^+, k \in \{1, ..., m\}$, calculated from the estimated joint distribution of $\chi \in \mathcal{P}(X)$ and Y, the multi-resolution approach to feature selection consists in finding a subset $\chi \in \mathcal{P}(X)$ of $k \in \{1, ..., m\}$ features and a $W \in \mathcal{P}(R_{\chi})$ such that $C_Y^k(\chi, W)$ is minimum.

The cost functions $C_Y^k(\chi, W)$ considered in this paper measure the local dependence between $\chi \in \mathcal{P}(X)$ of length $k \in \{1, ..., m\}$ and Y restricted to the subset $W \in \mathcal{P}(R_\chi)$, i.e., for $\chi \in W$. More specifically, our cost functions are based on the *Local Lift Dependence Scale*, a scale for measuring variable dependence in multiple resolutions. In this scale we may measure variable dependence globally and locally. On the one hand, global dependence is measured by a coefficient, that summarizes it. On the other hand, local dependence is measured for each subset $W \in \mathcal{P}(\mathcal{R}_\chi)$, again by a coefficient. Therefore, if the cardinality of \mathcal{R}_χ is N, we have $2^N - 1$ dependence coefficients: one global and $2^N - 2$ local, each one measuring the influence of χ in Y restricted to a subset of \mathcal{R}_χ . Furthermore, the *Local Lift Dependence Scale* also provides a *propensity* measure for each point of the joint range of χ and Y. Note that the dependence is indeed measured in multiple resolutions: globally, for each subset of \mathcal{R}_χ and pointwise.

Thus, in this paper, we extend the classical approach to feature selection in order to select not only the features, but also their values that are most related to *Y* in some sense. In order to do so, we extend the search space of the feature selection algorithm from the BLFS to the CBLOP and use cost functions based on the *Local Lift Dependence Scale*, which is an extension of Shannon's mutual information. The feature selection algorithms proposed in this paper are applied to a dataset consisting of student performances on a university's entrance exam and on undergraduate courses in order to select exam's subjects, and the performances on them, that are most related to undergraduate courses, considering student performance on both. The method is also applied to two datasets publicly available at the UCI Machine Learning Repository [12], namely, the Congressional Voting Records and Covertype datasets. We first present the main concepts related to the *Local Lift Dependence Scale*. Then, we propose feature selection algorithms based on the *Local Lift Dependence Scale* and apply them to solve real problems.

2. Local Lift Dependence Scale

The *Local Lift Dependence Scale* (LLDS) is a scale for measuring the dependence between a random variable Y and a random vector X (also called feature vector) in multiple resolutions. Although consisting of well known mathematical objects, there does not seem to exist any literature that thoroughly defines and study the properties of the LLDS, even though it is highly used in

marketing [13] and data mining [14] (Chapter 10), for example. Therefore, we present an unprecedented characterization of the LLDS, despite the fact that much of it is known in the theory.

The LLDS analyses the raw dependence between the variables, as it does not make any assumption about its kind, nor restrict itself to the study of a specific kind of dependence, e.g., linear dependence. Among LLDS dependence coefficients, there are three measures of dependence, one global and two local, but with different resolutions, that assess variable dependence on multiple levels. The global measure and one of the local are based on well known dependence measures, namely, the Mutual Information and the Kullback-Leibler Divergence. In the following paragraphs we present the main concepts of the LLDS and discuss how they can be applied to the classical and multi-resolution approaches to feature selection. The main concepts are presented for discrete random variables $X = \{X_1, \ldots, X_m\}$ and Y defined on $(\Omega, \mathbb{F}, \mathbb{P})$, with range $R_{X,Y} = R_X \times R_Y$, although, with simple adaptations, i.e., by interchanging probability functions with probability density functions, the continuous case follows from it.

The Mutual Information (MI), proposed by [15], is a classical dependence quantifier that measures the mass concentration of a joint probability function. As more concentrated the joint probability probability function is, the more dependent the random variables are and greater is their MI. In fact, the MI is a numerical index defined as

$$I(X,Y) \coloneqq \sum_{(\mathbf{x},y) \in R_{\mathbf{X},Y}} f(\mathbf{x},y) \log \left(\frac{f(\mathbf{x},y)}{g(\mathbf{x})h(y)}\right)$$

in which $f(x, y) := \mathbb{P}(X = x, Y = y), g(x) := \mathbb{P}(X = x)$ and $h(y) := \mathbb{P}(Y = y)$ for all $(x, y) \in R_{X,Y}$. An useful property of the MI is that it may be expressed as

$$I(\boldsymbol{X},\boldsymbol{Y}) = -\sum_{\boldsymbol{y}\in R_{\boldsymbol{Y}}} h(\boldsymbol{y})\log h(\boldsymbol{y}) + \sum_{(\boldsymbol{x},\boldsymbol{y})\in R_{\boldsymbol{X},\boldsymbol{Y}}} f(\boldsymbol{x},\boldsymbol{y})\log\left(\frac{f(\boldsymbol{x},\boldsymbol{y})}{g(\boldsymbol{x})}\right) \coloneqq H(\boldsymbol{Y}) - H(\boldsymbol{Y}|\boldsymbol{X})$$
(1)

in which H(Y) is the Entropy of Y and H(Y|X) is the Conditional Entropy (CE) of Y given X. The form of the MI in (1) is useful because, if we fix Y, and consider features (X_1, \ldots, X_m) , we may determine which one of them is the most dependent with Y by observing only the CE of Y given each one, as the feature that maximizes the MI is the one that minimizes the CE. In this paper, we consider the normalized MI that is given by

$$\eta_{\mathbf{x}}(Y|R_{\mathbf{X}}) = \frac{\sum\limits_{(\mathbf{x},y)\in R_{\mathbf{X},Y}} f(\mathbf{x},y) \log\left(\frac{f(\mathbf{x},y)}{g(\mathbf{x})h(y)}\right)}{-\sum_{y\in R_{Y}} h(y) \log h(y)} \coloneqq \frac{I(\mathbf{X},Y)}{H(Y)}$$
(2)

,

when $H(Y) \neq 0$. We have that $0 \leq \eta_x(Y|R_X) \leq \frac{\min\{H(Y),H(X)\}}{H(Y)} \leq 1$, that $\eta_x(Y|R_X) = 0$ if, and only if, **X** and Y are independent and that $\eta_x(Y|R_X) = 1$ if, and only if, there exists a function $Q: R_X \to R_Y$ such that $\mathbb{P}(Y = Q(X)) = 1$, i.e., Y is a function of X. The η , MI and CE are equivalent global and general measures of dependence, that summarize to an index a variety of dependence kinds that are expressed by mass concentration.

On the other hand, we may define a LLDS local and general measure of dependence that expands the global dependence measured by the MI into local indexes, enabling a local interpretation of variable dependence. As the MI is an index that measures variable dependence by measuring the mass concentration incurred in one variable by the observation of another, it may only give evidences about the existence of a dependence, but cannot assert what kind of dependence is being observed. Therefore, it is relevant to break down the MI by region, so that it can be interpreted in an useful manner and the

kind of dependence outlined by it may be identified. The *Lift Function* (LF) is responsible for this break down, as it may be expressed as

$$L_{(\mathbf{X},Y)}(\mathbf{x},y) \coloneqq \frac{f(\mathbf{x},y)}{g(\mathbf{x})h(y)} = \frac{f(y|\mathbf{x})}{h(y)}, \quad \forall (\mathbf{x},y) \in R_{\mathbf{X},Y}$$

in which $f(y|x) := \mathbb{P}(Y = y|X = x)$. When there is no doubt about which variables the LF refers to, it is denoted simply by L(x, y). Note that the LF is the exponential of the pontual mutual information (see [16,17] for example for more details).

The MI is the expectation on (X, Y) of the LF logarithm, so that the LF presents locally the mass concentration measured by the MI. As the LF may be written as the ratio between the conditional probability of Y given X and the marginal probability of Y, the main interest in its behaviour is in determining for which points $(x, y) \in R_{X,Y} L(x, y) > 1$ and for which L(x, y) < 1. If L(x, y) > 1then the fact of X being equal to x increases the probability of Y being equal to y, as the conditional probability is greater than the marginal one. Therefore, we say that event $\{X = x\}$ lifts event $\{Y = y\}$ or that instances with profile x are prone to be of class y. In the same way, if L(x, y) < 1, we say that event $\{X = x\}$ inhibits event $\{Y = y\}$, for f(y|x) < h(y). If $L(x, y) = 1, \forall (x, y) \in R_{X,Y}$, then the random variables are independent. Note that the LF is symmetric: $\{X = x\}$ lifts $\{Y = y\}$ if, and only if, $\{Y = y\}$ lifts $\{X = x\}$. Therefore, the LF may be interpreted as X lifting Y or Y lifting X. From now on, we interpret it as X lifting Y, even though it could be the other way around.

An important property of the LF is that it cannot be greater than one nor lesser than one for all points $(x, y) \in R_{X,Y}$. Indeed, if L(x, y) > 1, $\forall (x, y) \in R_{X,Y}$, then $f(y \mid x) > h(y)$, $\forall (x, y) \in R_{X,Y}$, what implies the absurd $1 = \sum_{y \in R_Y} f(y \mid x) > \sum_{y \in R_Y} h(y) = 1$ for $x \in R_X$. With an analogous argument we see that L(x, y) cannot be lesser than one for all $(x, y) \in R_{X,Y}$. Therefore, if there are LF values greater than one, then there must be values lesser than one, what makes it clear that the values of the LF are dependent and that the *lift* is a pointwise characteristic of the joint probability function and not a global property of it. Thus, the study of the LF behaviour gives the full view of the dependence between the variables, without restricting it to a specific kind nor making assumptions about it.

Although the LF presents a wide picture of variable dependence, it may present it in a too high resolution, making it complex to interpret. Therefore, instead of measuring dependence for each point in the range $R_{X,Y}$, we may measure it for a *window* $W \in \mathcal{P}(R_X)$. The dependence between X and Y in the window W, i.e., for $X \in W$, may be measured by the η coefficient defined as

$$\eta_{x}(Y|W) \coloneqq \frac{\sum\limits_{\boldsymbol{x}\in W} g(\boldsymbol{x}) \sum\limits_{\boldsymbol{y}\in R_{Y}} f(\boldsymbol{y}|\boldsymbol{x}) \log \frac{f(\boldsymbol{y}|\boldsymbol{x})}{h(\boldsymbol{y})}}{-\sum\limits_{\boldsymbol{x}\in W} g(\boldsymbol{x}) \sum\limits_{\boldsymbol{y}\in R_{Y}} f(\boldsymbol{y}|\boldsymbol{x}) \log h(\boldsymbol{y})} = \frac{E\left\{D_{KL}(f(\cdot|\boldsymbol{X})||h(\cdot)) \middle| \boldsymbol{X}\in W\right\}}{E\left\{H([Y|\boldsymbol{X}], Y) \middle| \boldsymbol{X}\in W\right\}}$$
(3)

when $E\{H([Y|X], Y) | X \in W\} \neq 0$ and $\mathbb{P}(X \in W) > 0$, in which $D_{KL}(\cdot || \cdot)$ is the Kullback-Leibler divergence [18], $H(\cdot, \cdot)$ is the cross-entropy [19] and H([Y|X], Y) means the cross-entropy between the conditional distribution of Y given X and the marginal distribution of Y. The η coefficient (3) compares the conditional probability of Y given $x, \forall x \in W$, with the marginal probability of Y, so that as greater the coefficient, as *distant* the conditional probability is from the marginal one and, therefore, greater is the influence of the event $\{X \in W\}$ in Y. Note that, analogously to the MI, we may write

$$E\left\{D_{KL}(f(\cdot|\mathbf{X})||h(\cdot))\Big|\mathbf{X}\in W\right\} = \frac{-\sum_{\mathbf{x}\in W}g(\mathbf{x})\sum_{y\in R_{Y}}f(y|\mathbf{x})\log h(y) + \sum_{\mathbf{x}\in W}g(\mathbf{x})\sum_{y\in R_{Y}}f(y|\mathbf{x})\log f(y|\mathbf{x})}{\mathbb{P}(\mathbf{X}\in W)}$$
$$= E\left\{H([Y|\mathbf{X}],Y)\Big|\mathbf{X}\in W\right\} - E\left\{H([Y|\mathbf{X}])\Big|\mathbf{X}\in W\right\}$$

in which H([Y|X]) means the Entropy of the conditional distribution of Y given X, and we have that $0 \le \eta_x(Y|W) \le 1$, that $\eta_x(Y|W) = 0$ if, and only if, $h(y) \equiv f(y|x), \forall x \in W$, and that $\eta_x(Y|W) = 1$

if, and only if, there exists a function $Q : W \to R_Y$ such that $\mathbb{P}(Y = Q(X)|X \in W) = 1$. Observe that the η coefficient of a window is also a local dependence quantifier, although its resolution is lower than that of the LF if the cardinality of W is greater than one. Also note that the η coefficient (3) is a generalization of (2) to all subsets (windows) of R_X , as $W = R_X$ is a window and that the numerator of $\eta_X(Y|W)$ equals $E\{\log L(X,Y) \mid (X,Y) \in W \times \mathcal{R}_Y\}$. It is important to outline that we may *complete* the LLDS with coefficients that are given by $E\{\log L(X,Y) \mid (X,Y) \in W_X \times W_Y\}$ normalized, with $W_X \in \mathcal{P}(\mathcal{R}_X)$ and $W_Y \in \mathcal{P}(\mathcal{R}_Y)$, that measure the dependence between X and Yfor $(X,Y) \in W_X \times W_Y$. However, we do not use coefficients of this type in this paper for in our case the variable Y and its range are always fixed. Note that if the cardinality of $W_X \times W_Y$ is one then $E\{\log L(X,Y) \mid (X,Y) \in W_X \times W_Y\} = \log L(W_X, W_Y)$ so that the scale is indeed *complete*.

The three dependence coefficients presented, when analysed collectively, measure variable dependence in all kinds of resolutions: since the low resolution of the MI, through the middle resolutions of the windows W, until the high resolution of the LF. Indeed, the η coefficients and the LF define a dependence scale in R_X , that we call LLDS, that gives a dependence measure for each subset $W \in \mathcal{P}(R_X)$. This scale may be useful for various purposes and we outline some of them in the following paragraphs.

Potential applications of the Local Lift Dependence Scale

The LLDS, more specifically the LF, is relevant in frameworks in which we want to choose a set of elements, e.g, people, in order to apply some kind of *treatment* to them, obtaining some kind of *response* Y, and are interested in maximizing the number of elements with a given response $y \in R_Y$. In this scenario, given the features X, the LF provides the set of elements that must be chosen, that is the set whose elements have profile $x \in R_X$ such that L(x, y) is greatest. Formally, we must choose elements whose profile is

$$\mathbf{x}_{opt}(y) = \arg \max_{\mathbf{x} \in R_{\mathbf{X}}} L(\mathbf{x}, y).$$

Indeed, if we choose *n* elements randomly from our population, we expect that $n \times \mathbb{P}(Y = y)$ of them will have the desired response. However, if we choose *n* elements from the population of all elements with profile $x_{opt}(y) \in R_X$, then we expect that $n \times \mathbb{P}(Y = y | X = x_{opt}(y))$ of them will have the desired response, what is $[L(x_{opt}(y), y) - 1] \times n$ more elements when comparing with the whole population sampling framework. Observe that this framework is the exact opposite of the classification problem. In the classification problem, we want to classify an instance given its profile $x \in R_X$ into a class $y \in R_Y$, that may be, for example, the class y such that f(y|x) is maximum. On the other hand, in this framework, we are interested in, given a $y \in R_Y$, finding the profile $x_{opt}(y) \in R_X$ such that $f(y|x_{opt}(y))$ is maximum. In the applications section we further discuss the differences between this framework and the classification problem, and how the LLDS may be applied to both.

Furthermore, the η coefficient is relevant in scenarios in which we want to understand the influence of X in Y by region, i.e., for each subset of R_X . As an example of such framework, consider an image in the grayscale, in which $X = (X_1, X_2)$ represents the pixels of the image and Y is the random variable whose distribution is the distribution of the colors in the picture, i.e., $\mathbb{P}(Y = y) = \frac{n_y}{n}$ in which n_y is the number of pixels whose color is $y \in \{1, ..., 255\}$ and n is the total number of pixels in the distribution of $Y|X = (x_1, x_2)$ properly for all $(x_1, x_2) \in R_X$, we may calculate $\eta_x(Y|W), W \in \mathcal{P}(R_X)$, in order to determine the regions that are a representation of the whole picture, i.e., whose color distribution is the same of the whole image, and the regions W whose color distribution differs from that of the whole image. The η coefficient may be useful for identifying textures and recognizing patterns in images.

Lastly, the LLDS may be used for feature selection, when we are not only interested in selecting the features $\chi \in \mathcal{P}(X)$ that are most related to *Y*, but also want to determine the features $\chi \in \mathcal{P}(X)$ whose *levels* $W \in \mathcal{P}(R_{\chi})$ most influence *Y*. In the same manner, we may want to select the features χ whose level $\mathbf{x}_{opt}(y) \in R_{\chi}$ maximizes $L_{(\chi,Y)}(\mathbf{x}_{opt}(y), y)$, for a given $y \in R_Y$, so that we may sample from the population of elements with profile $x_{opt}(y) \in R_{\chi}$ in order to maximize the number of elements of class *y*. Feature selection based on the LLDS is a special case of the classical and multi-resolution approaches to feature selection as presented next.

3. Feature Selection Algorithms based on the Local Lift Dependence Scale

In this section we present the characteristics of feature selection algorithms based on the LLDS. We first outline the special case of the classical approach to feature selection that is based on the LLDS, and then propose *multi-resolution* feature selection algorithms that are also based on the LLDS.

3.1. Classical Feature Selection Algorithm

Let *Y* and $X = (X_1, ..., X_m)$ be random variables. We call the random variables in *X* features and note that $\mathcal{P}(X)$, the set of all feature vectors whose features are also in *X*, may be seen as a BLFS, in which each vector represents a subset of features. In this scheme, feature selection is given by the minimization, in the BLFS, of a cost function applied on the estimated joint probability of a feature vector and *Y*. In fact, the subset of features selected by this approach is given by

$$\boldsymbol{\chi} = \operatorname*{arg\,min}_{\boldsymbol{\chi}^* \in \mathcal{P}(\boldsymbol{X})} C_{Y}(\boldsymbol{\chi}^*)$$

in which $C_Y : \mathcal{P}(\mathbf{X}) \to \mathbb{R}^+$ is a cost function. The estimated error of a predictor Ψ as presented in [20] (Chapter 2), for example, is a classical cost function. Another classical cost function is the CE as defined in (1). A pseudo-code for such algorithm is presented in Algorithm 1. Algorithm 1 is *naive*, performs an exhaustive search on the BLFS and is known to be NP-hard [21]. However, some other algorithms may be applied to find a sub-optimal solution to this problem, as sequential selection algorithms and floating search methods [22–27]. Also, the search space may be restricted to a subspace of $\mathcal{P}(\mathbf{X})$. Nevertheless, there are algorithms, as the *branch-and-bound* [28] and the *u-curve* [9,10,29], that does not perform an exhaustive search, but ensure that the selected subset of features is optimal.

Algorithm 1 Select $\chi \in \mathcal{P}(X)$ that minimizes $C_Y(\chi)$.
Ensure: $c = \infty$ Ensure: $\chi = \emptyset$ 1: for $\chi^* \in \mathcal{P}(X)$ do
2: if $C_Y(\chi^*) < c$ then 3: $c \leftarrow C_Y(\chi^*)$ 4: $\chi \leftarrow \chi^*$ 5: end if 6: end for 7: return χ

As an example of the classical approach to feature selection, suppose that $X = (X_1, X_2)$, in which X_1 and X_2 are Boolean features. Then, the search space $\mathcal{P}(X)$ may be represented by a tree, i.e., a BLFS, as the one displayed in Figure 1, considering only the circle nodes and solid lines. Algorithm 1 may be performed by walking through this tree seeking the minimum of C_Y .



Figure 1. Example of multi-resolution tree for feature selection. The circle nodes and solid lines form a BLFS. The rectangular nodes and dashed lines around each circle node form a Boolean lattice. The whole tree is a CBLOP.

3.2. Multi-resolution Feature Selection based on the Local Lift Dependence Scale

Feature selection based on the LLDS may be performed in three distinct resolutions. As a low resolution approach, we may select the features χ that are most globally related to Y, that are given by

$$\chi = \underset{\chi^* \in \mathcal{P}(X)}{\arg \max} \eta_{\chi^*}(Y | R_{\chi^*}).$$
(4)

Note that, in this resolution, the feature selection approach is the classical one, with $1/\eta$ in (2) as the cost function, i.e., Algorithm 1 may be applied to determine (4) taking C_Y as $1/\eta$. The use of the MI as a cost function in classical feature selection algorithms is quite common in the literature (see [30–32] for example) and is not original of this paper. The search space of (4) may be restricted, sub-optimal algorithms may be applied or the discretization of the continuous features may be performed jointly, so that the subset selected in (4) is not always the subset of all features. In the applications section we show how the continuous features may be discretized jointly.

Increasing the resolution, we may be interested in finding not the features most related to Y, but the features *levels* that most influence Y. In this approach the selected features and their levels are

$$(\boldsymbol{\chi}, W) = \underset{\substack{\boldsymbol{\chi}^* \in \mathcal{P}(\mathbf{X})\\ W^* \in \mathcal{P}(R_{\boldsymbol{\chi}^*})}}{\arg \max \eta_{\boldsymbol{\chi}^*}(\boldsymbol{Y}|W^*)}.$$
(5)

A pseudo-code for this feature selection approach is presented in Algorithm 2. Note that the space in which the exhaustive search is conducted in Algorithm 2, i.e., the CBLOP, is even greater than the one in Algorithm 1. However, optimal algorithms that do not exhaustively search the space, and sub-optimal algorithms, may also be applied in this scenario, saving some computational time. Note that this approach is not suitable for the case in which the features in *X* are continuous, as R_{χ} , $\chi \in \mathcal{P}(X)$, is uncountable, although the continuous features may be discretized. Furthermore, as is further discussed in the applications section, this algorithm is subjected to *overfitting* if the sample size is not relatively great (as is the majority of statistical models and feature selection algorithms), so that it may be needed to restrict its search space to a subset of the CBLOP.

```
Algorithm 2 Select \chi \in \mathcal{P}(X) and W \in \mathcal{P}(R_{\chi}) that maximizes \eta_{\chi}(Y|W).
Ensure: c = 0
Ensure: \chi = \emptyset
Ensure: W = \emptyset
 1: for \chi^* \in \mathcal{P}(X) do
         for W^* \in \mathcal{P}(R_{\chi^*}) do
  2:
             if \eta_{x^*}(Y|W^*) > c then
  3:
                c \gets \eta_{\mathbf{x}^*_*}(Y|W^*)
  4:
  5:
  6:
 7:
             end if
        end for
 8:
 9: end for
10: return (\boldsymbol{\chi}, W)
```

Finally, as a higher resolution approach, we may fix an $y \in R_Y$ and then look for the features levels that maximize the LF at the point *y*. Formally, the selected features and levels are given by

$$(\boldsymbol{\chi}, \boldsymbol{x}_{opt}(\boldsymbol{y})) = \underset{\substack{\boldsymbol{\chi}^* \in \mathcal{P}(\mathbf{X})\\ \boldsymbol{x}^* \in R_{\boldsymbol{\chi}^*}}}{\arg \max L_{(\boldsymbol{\chi}*, \boldsymbol{Y})}(\boldsymbol{x}^*, \boldsymbol{y})}.$$
(6)

continuous features; and it is subjected to overfitting.

Algorithm 3 Select $\chi \in \mathcal{P}(X)$ and $x \in R_{\chi}$ that maximizes $L_{(\chi, Y)}(x, y)$ for some fixed $y \in R_{Y}$.
Ensure: $c = 0$
Ensure: $\chi = \emptyset$
Ensure: $x = \emptyset$
Ensure: $y = y$
1: for $\chi^* \in \mathcal{P}(X)$ do
2: for $x^* \in R_{\chi^*}$ do
3: if $L_{(x^*,Y)}(x^*,y) > c$ then
4: $c \leftarrow L_{(\mathbf{x}^{*}, Y)}(\mathbf{x}^{*}, y)$
5: $\chi \leftarrow \chi^*$
6: $x \leftarrow x^*$
7: end if
8: end for
9: end for
10: return (χ, x)

As an example of a *multi-resolution* approach to feature selection based on the LLDS, suppose again that $X = (X_1, X_2)$ are Boolean features. Then, for all the proposed resolutions, the selection of the features and their levels, i.e., Algorithms 1, 2 and 3, may be performed by walking through the tree (CBLOP) in Figure 1. Indeed, we may calculate the global η at the circle nodes, the η on all windows W at the rectangular nodes and the LF at the leaves, where we may determine its maximum for a fixed value $y \in R_Y$. Therefore, we call a tree as the one in Figure 1 a *multi-resolution* tree for feature selection, where we may apply feature selection algorithms for all the resolutions of the LLDS, i.e., Algorithms 1, 2 and 3.

4. Applications

The *multi-resolution* approach proposed in the previous sections is now applied to three different datasets. First, we apply it to the performances dataset, that consists of student performances on entrance exams and undergraduate courses. Then, we apply the algorithms to two UCI Machine Learning Repository datasets: the Congressional Voting Records and Covertype datasets [12].

4.1. Performances dataset

A recurrent issue in universities all over the world is the framework of their recruitment process, i.e., the manner of selecting their undergraduate students. In Brazilian universities, for example, the recruitment of undergraduate students is solely based on their performance on exams that cover high school subjects, called *vestibulares*, so that knowing which subjects are most related to the performance on undergraduate courses is a matter of great importance to universities admission offices, as it is important to optimize the recruitment process in order to select the students that are most likely to succeed. Therefore, in this scenario, the algorithm presented in the previous sections may be an useful tool in determining the entrance exam subjects, and the performances on them, that are most related to the performance on undergraduate courses, so that students may be selected based on their performance on these subjects.

The recruitment of students to the University of São Paulo is based on an entrance exam that consists of an essay and questions of eight subjects: Mathematics, Physics, Chemistry, Biology, History, Geography, English and Portuguese. The selection of students is entirely based on this exam,

although the *weights* of the subjects differ from one course to another. In the exact sciences courses, as Mathematics, Statistics, Physics, Computer Science and Engineering, for example, the subjects with greater weights are Portuguese, Mathematics and Physics, as those are the subjects that are qualitatively most related to what is taught at these courses. Although weights are given to each subject in a systematic manner, it is not known what subjects are *indeed* most related to the performance on undergraduate courses. Therefore, it would be of interest to measure the relation between the performance on exam subjects and undergraduate courses and, in order to do so, we apply the algorithms proposed on the previous sections.

The dataset to be considered consists of *8,353* students that enrolled in 28 courses of the University of São Paulo between 2011 and 2016. The courses are those of its Institute of Mathematics and Statistics, Institute of Physics and Polytechnic School, and are in general Mathematics, Computer Science, Statistics, Physics and Engineering courses. The variable of interest (*Y*) is the weighted mean grade of the students on the courses they attended in their first year at the university (the weights being the courses credits), and is a number between zero and ten. The features, that are denoted $X = (X_1, X_2, X_3, X_4, X_5, X_6, X_7, X_8, X_9)$, are the performances on each one of the eight entrance exam subjects, that are numbers between zero and one, and the performance on the essay, that is a number between zero and one hundred.

In order to apply the proposed algorithm to the data at hand, it is necessary to conveniently discretize the variables and, to do so, we take into account an important characteristic of the data: the scale of the performances. The scale of the performances, both on the entrance exam and the undergraduate courses, depends on the course and the year. Indeed, the performance on the entrance exam of students of competitive courses is better, as only the students with high performance are able to enrol in these courses. In the same way, the performances differ from one year to another, as the entrance exam is not the same every year and the teachers of the first year courses also change from one year to another, what causes the scale of the grades to change. Therefore, we discretize all variables by tertiles inside each year and course, i.e., we take the tertiles considering only the students of a given course and year. Furthermore, we do not discretize each variable by itself, but rather discretize the variables jointly, by a method based on distance tertiles, as follows.

Suppose that at a step of the algorithm we want to measure the relation between Υ and the features $\chi \in \mathcal{P}(X)$. In order to do so, we discretize Υ by its tertiles inside each course and year, e.g., a student is in the third tertile if he is on the top one third students of his class according to the weighted mean grade, and discretize the performance on χ *jointly*, i.e., by discretizing the distance between the performance of each student on these subjects and zero by its tertiles. Indeed, students whose performance is close to zero have low joint performance on the subjects χ . Therefore, we take the distance between each performance and zero, and then discretize it inside each course and year, e.g., a student is at the first tertile if he is on the bottom students of his class according to his joint performance on the subjects χ . The Mahalanobis distance [33] is used, as it takes into account the variance and covariance of the performance on the subjects χ .

As an example, suppose that we want to measure the relation between the performances on Mathematics and Physics and the weighted mean grade of students that enrolled in the Statistics undergraduate course in 2011 and 2012. In order to do so, we discretize the weighted mean grade by year and the performance on Mathematics and Physics by the Mahalanobis distance between it and zero, also by year, as is displayed in Figure 2. Observe that each year has its own ellipses that partition the performance on Mathematics and Physics in three and the tertile of a student depends on the ellipses of his year. The process used in Figure 2 is extended to the case in which there are more than two subjects and one course. When there is only one subject, the performance is discretized in the usual manner inside each course and year. The LF between the weighted mean grade and the joint performance on Mathematics and Physics is presented in Table 1. From this table we may

search for the maximum lift or calculate the η coefficient for its windows. In this example, we have $\eta_{(M,P)}(Y|R_{(M,P)}) = 0.0387$, in which (M,P) = (Mathematics, Physics).



Figure 2. Discretization of the joint performance on Mathematics and Physics of Statistics students that enrolled at the University of São Paulo in 2011 and 2012 by the tertiles of the Mahalanobis distance inside each year.

Table 1. The Lift Function between the weighted mean grade, discretized by year, and the joint performance on Mathematics and Physics, discretized by the Mahalanobis distance inside each year, of Statistics students that enrolled at the University of São Paulo in 2011 and 2012. The numbers in parentheses represent the quantity of students in each category.

Mathematics and Dhysics	Weigł	nted Mean	Rolativo Eroguonav	
Mathematics and Physics	Tertile 1	Tertile 2	Tertile 3	Kelative Frequency
Tertile 1	0.975 (9)	1.46 (13)	0.563 (5)	0.34
Tertile 2	1.01 (9)	0.935 (8)	1.05 (9)	0.33
Tertile 3	1.01 (9)	0.584 (5)	1.4 (12)	0.33
Relative Frequency	0.342	0.329	0.329	1

The proposed algorithm is applied to the discretized variables using three cost functions. First, we use the η coefficient on the window that represents the whole range of the features in order to determine what are the subjects (features) that are most related to the weighted mean grade, i.e., the features (4). Then, we apply the algorithm using as cost function the η coefficient for all windows in order to determine the subjects performances (features and window) that are most related to the weighted mean grade, i.e., the subjects and performances (5). Finally, we determine what are the subjects and their performance that most lift the weighted mean grade third tertile, i.e., the subjects and performances (6) with y = Tertile 3.

The subjects that are most related to the weighted mean grade, according to the proposed discretization process and the η coefficient (2), are $\chi = (M,P,C,B,Po)$ and $\eta_{\chi}(Y|R_{\chi}) = 0.0354$, in which (M,P,C,B,Po) = (Mathematics,Physiscs,Chemistry,Biology,Portuguese). The LF between the weighted mean grade and χ is presented in Table 2. The features χ are the ones that are in

general most related to the weighted mean grade, i.e., are the output of the classical feature selection algorithm that employs the inverse of the global η coefficient as cost function (Algorithm 1). Therefore, the recruitment of students could be optimized by taking into account only the subjects χ .

Table 2. The Lift Function between the weighted mean grade and the joint performance on $\chi = (Mathematics, Physiscs, Chemistry, Biology, Portuguese). The numbers in parentheses represent the quantity of students in each category.$

Dorformon co in a	Weig	ghted Mean G	Polotivo Ero avon av		
remomence in χ	Tertile 1	Tertile 2	Tertile 3	Relative Frequency	
Tertile 1	1.33 (1277)	1.1 (1018)	0.566 (533)	0.34	
Tertile 2	0.992 (921)	1.06 (951)	0.954 (871)	0.33	
Tertile 3	0.669 (630)	0.848 (775)	1.49 (1377)	0.33	
Relative Frequency	0.339	0.329	0.333	1	

Applying Algorithms 2 and 3 we obtain the same result, that the performance, i.e., window, that is most related to the weighted mean grade and that most *lifts* the third tertile of the weighted mean grade is the third tertile in Mathematics, for which $\eta_M(Y|\{\text{Tertile 3}\}) = 0.0575$ and $L_{(M,Y)}(\text{Tertile 3,Tertile 3}) = 1.51$, in which M = Mathematics. The LF between the weighted mean grade and the performance on Mathematics is presented in Table 3.

Table 3. The Lift Function between the weighted mean grade and the performance on Mathematics. The numbers in parentheses represent the quantity of students in each category.

Portormanco in Mathematico	Weig	ghted Mean G		
renormance in Mathematics	Tertile 1	Tertile 2	Tertile 3	- Kelative Frequency
Tertile 1	1.3 (1398)	1.06 (1111)	0.631 (667)	0.38
Tertile 2	0.935 (843)	1.11 (972)	0.956 (847)	0.32
Tertile 3	0.689 (587)	0.8 (661)	1.51 (1267)	0.30
Relative Frequency	0.339	0.329	0.333	1

The output of the algorithms provides relevant information to the admission office of the University. Indeed, it is now known that the subjects that are most related to the performance on the undergraduate courses are Mathematics, Physics, Chemistry, Biology and Portuguese. Furthermore, in order to optimize the number of students that will succeed in the undergraduate courses, the office must select those that have high performance on Mathematics, as it *lifts* by more than 50% the probability of the student having also a high performance on the undergraduate course, i.e., students with high performance on Mathematics are prone to have high performance on the undergraduate course course. Although the subjects that are most related to the performance on the courses are obtained from the classical feature selection algorithm, only the LLDS outlines what is the performance on the entrance on Mathematics. Therefore, feature selection algorithms based on the LLDS provide more information than the classical feature selection algorithm, as they have a greater resolution and take into account the local relation between the variables.

4.2. Congressional Voting Records dataset

The Congressional Voting Records dataset consists of 435 instances of 16 Boolean features and a Boolean variable that indicates the party of the instance (democrat or republican). The features indicate how the instance voted (yes or no) in the year of 1984 about each one of 16 matters, that are

displayed in Table 4. Algorithm 3 is applied to this dataset in order to determine what are the voting profiles that are most prone to be that of a republican and that of a democrat.

ID	Matter (Feature)
HI	Handicapped infants
WP	Water project cost sharing
AB	Adoption of the budget resolution
PF	Physician fee freeze
SA	El Salvador aid
RG	Religious groups in schools
ST	Anti satellite test ban
AN	Aid to Nicaraguan contras
MM	MX missile
IM	Immigration
SC	Synfuels corporation cutback
ES	Education spending
SR	Superfund right to sue
CR	Crime
DF	Duty Free exports
EA	Export administration act South Africa

Table 4. Features of the Congressional Voting Records dataset.

As the number of instances is relatively small, we perform Algorithm 3 under a restriction that avoids *overfitting*. Indeed, if we apply the algorithm without the restriction, then the chosen profiles are those in which all the instances are of the same party. If there is only a couple of instances with some profile, and all of them are of the same party, then this profile is chosen as a *prone* one for the party. However, we do not know if the profile is really prone, i.e., everybody with it is in fact of the same party, or if the fact of everybody with this profile being of the same party is just a sample deviation. In other words, without the restriction, the estimation error of the LF is too great as some profiles have low frequency in the sample and the feature selection algorithm *overfits*.

Therefore, we restrict the search space to the profiles with a relative frequency in the sample of at least 0.15. In other words, we select the profiles

$$(\boldsymbol{\chi}, \boldsymbol{x}_{opt}(\boldsymbol{y})) = \underset{\substack{\boldsymbol{\chi}^* \in \mathcal{P}(\boldsymbol{X})\\ \boldsymbol{x}^* \in \boldsymbol{R}_{\boldsymbol{\chi}^*}}}{\arg \max} L_{(\boldsymbol{\chi}^*, \boldsymbol{Y})}(\boldsymbol{x}^*, \boldsymbol{y})$$

for $y \in \{democrat, republican\}$, in which $\mathbb{P}(\chi^* = x^*), \chi^* \in \mathcal{P}(X), x^* \in R_{\chi^*}$, is estimated by the relative frequency of the profile. The selected profiles, their LF value and the sample size considered are presented in Table 5. At each iteration of the algorithm, only the instances that have no missing data in the features being considered are taken into account when calculating the LF, so that the sample size of each iteration is not the same.

The profiles with maximum LF lifts by 94% the probability of democrat and by around 165% the probability of republican. This difference in the *lift* is due to the fact that there are more democrats than republicans, so that the probability of democrat is greater and, therefore, cannot be lifted as much as the probability of republican can. The profiles in Table 5 present a wide view of the voting profile of democrats and republicans, what allows an understanding of what differentiates a democrat from a republican regarding their vote.

Table 5. Selected profiles obtained applying Algorithm 3 to the Congressional Voting Records dataset
with the restriction that only the profiles with relative frequency greater than 0.15 are considered.
The instances with missing data were excluded at each iteration of the algorithm, i.e., $L_{(x^*,Y)}(x^*, y)$ is
calculated using only the instances that have all the observations on the features χ^* .

Party	Features (χ)	LF	Profile (<i>x</i>)	Sample Size
	(AB,PF,SA,RG,MM,ES,SR,EA)	1.94	(y,n,n,n,y,n,n,y)	277
	(HI,AB,PF,SA,RG,MM,ES,SR,EA)	1.94	(y,y,n,n,n,y,n,n,y)	275
	(AB,PF,RG,ST,MM,ES,SR,EA)	1.94	(y,n,n,y,y,n,n,y)	279
	(HI,AB,PF,RG,ST,MM,ES,SR,EA)	1.94	(y,y,n,n,y,y,n,n,y)	277
	(AB,PF,SA,RG,ST,MM,ES,SR,EA)	1.94	(y,n,n,n,y,y,n,n,y)	276
democrat	(HI,AB,PF,SA,RG,ST,MM,ES,SR,EA)	1.94	(y,y,n,n,n,y,y,n,n,y)	274
	(AB,PF,SA,RG,MM,ES,SR,CR,EA)	1.94	(y,n,n,n,y,n,n,n,y)	275
	(AB,PF,RG,ST,MM,ES,SR,CR,EA)	1.94	(y,n,n,y,y,n,n,n,y)	276
	(AB,PF,SA,RG,ST,MM,ES,SR,CR,EA)	1.94	(y,n,n,n,y,y,n,n,n,y)	274
	(AB,PF,RG,ST,MM,ES,SR,DF,EA)	1.94	(y,n,n,y,y,n,n,y,y)	269
	(AB,PF,SA,RG,ST,MM,ES,SR,DF,EA)	1.94	(y,n,n,n,y,y,n,n,y,y)	266
	(WP,PF,SC,ES,CR)	2.65	(n,y,n,y,y)	342
	(AB,PF,AN,SC,CR,DF)	2.64	(n,y,n,n,y,n)	369
	(PF,AN,IM,ES,CR,DF)	2.64	(y,n,y,y,y,n)	361
	(PF,AN,SC,CR,DF)	2.64	(y,n,n,y,n)	373
	(AB,PF,AN,SC,ES)	2.63	(n,y,n,n,y)	376
	(HI,AB,PF,AN,SC,ES)	2.63	(n,n,y,n,n,y)	373
	(AB,PF,AN,SC,ES,CR)	2.63	(n,y,n,n,y,y)	368
ropublican	(HI,AB,PF,AN,SC,ES,CR)	2.63	(n,n,y,n,n,y,y)	365
Tepublican	(PF,AN,SC,DF)	2.63	(y,n,n,n)	380
	(AB,PF,AN,SC,DF)	2.63	(n,y,n,n,n)	376
	(PF,AN,IM,ES,DF)	2.63	(y,n,y,y,n)	368
	(AB,PF,AN,SC,ES,DF)	2.63	(n,y,n,n,y,n)	360
	(HI,AB,PF,AN,SC,CR,DF)	2.63	(n,n,y,n,n,y,n)	365
	(PF,AN,SC,ES,CR,DF)	2.63	(y,n,n,y,y,n)	356
	(AB,PF,AN,SC,ES,CR,DF)	2.63	(n,y,n,n,y,y,n)	353
	(HI,AB,PF,AN,SC,ES,CR,DF)	2.63	(n,n,y,n,n,y,y,n)	350

y = yes; n = no.

This application to the Congressional Voting Records dataset shed light on two interesting properties of the LLDS approach to feature selection in its higher resolution. First, this approach is indeed *local*, as we are not interested in selecting the features that best classify the representatives accordingly to their party, but rather the voting *profiles* that are most prone to be that of a democrat or republican. Secondly, the problem treated here is the opposite of the classification problem. Indeed, in the classification problem, we are interested in classifying a representative according to his party, given his voting profile. On the other hand, the problem treated here is the exact opposite: given a party, we want to know what are the profiles of the representatives that are most prone to be of that party. In other words, in the classification problem we want to determine the party given the voting profile, while on the LLDS problem we want to determine the voting profile given the party.

4.3. Covertype dataset

The Covertype dataset consists of *581,012* instances (terrains) of *54* features (10 continuous and 44 discrete) and a variable that indicates the cover type of the terrain (7 types). We apply Algorithms 1, 2 and 3 to select features among the continuous ones that are displayed in Table 6. The features are discretized in the same way they were in the performances dataset: by taking sample quantiles of the Mahalanobis distance between the features and zero. However, we now consider the quantiles 0.2, 0.4, 0.6 and 0.8 as cutting points, i.e., *quintiles*, instead of tertiles.

ID	Feature
Е	Elevation
А	Aspect
S	Slope
HH	Horizontal distance to hydrology
HR	Horizontal distance to roadways
HF	Horizontal distance to fire points
H9	Hillshade 9 a.m.
HN	Hillshade Noon
H3	Hillshade 3 p.m.
VH	Vertical distance to hydrology

Table 6. Features of the Covertype dataset that are considered in this application.

Applying Algorithm 1 we select the features $\chi = (E,HH,HF)$, with a coefficient $\eta_x(Y | R_\chi) = 0.307$ and the LF in Table 7. We see that being in the first quintile of the selected features lifts classes 3, 4, 5 and 6; being in the second quintile lifts classes 2 and 5; being in the third quintile lifts class 2; being in the fourth quintile lifts class 1; and being in the fifth quintile lifts classes 1 and 7. From Table 7 we may interpret the relation between the selected features and the cover type. For example, we see that terrains with cover types 3, 4, 5 and 6 tend to have low joint values in the selected features, while terrains with cover 7 tend to have great joint values in them. This example shows how the proposed approach allows not only to select the features, but also understand why these features were selected, i.e., what is the relation between them and the cover type, by analysing the local dependence between the variables.

Table 7. The Lift Function between the cover type of the terrain and the features Elevation, Horizontal distance to hydrology and Horizontal distance to fire points discretized by the sample quintiles of the Mahalanobis distance to zero. The numbers in parentheses represent the sample size of each category.

(E HH HE)	Cover Type						Rolativo Fraguency	
(E,III,IIF)	1	2	3	4	5	6	7	Relative frequency
Quintile 1	0.0766 (3244)	0.961 (54,473)	4.94 (35,344)	5 (2747)	1.78 (3385)	4.9 (17,010)	0 (0)	0.20
Quintile 2	0.444 (18,816)	1.6 (90,872)	0.0573 (410)	0 (0)	2.98 (5663)	0.103 (357)	0.0205 (84)	0.20
Quintile 3	0.949 (40,195)	1.33 (75,562)	0 (0)	0 (0)	0.234 (445)	0 (0)	0 (0)	0.20
Quintile 4	1.66 (70,427)	0.8 (45,314)	0 (0)	0 (0)	0 (0)	0 (0)	0.112 (461)	0.20
Quintile 5	1.87 (79,158)	0.301 (17,080)	0 (0)	0 (0)	0 (0)	0 (0)	4.87 (19,965)	0.20
Relative Frequency	0.365	0.488	0.0615	0.00473	0.0163	0.0299	0.0353	1

Applying Algorithm 2 to this dataset we obtain the windows displayed in Table 8. We see that the window that seems to most influence the cover type is the first and fifth quintile of the features Elevation and Horizontal distance to hydrology. Indeed, all the top ten windows contain those two features, and either their first or fifth quintile. As we can see in Table 7, the influence of the fifth quintile of $\chi = (E,HH,HF)$, the top window, is given by the fact that no terrain of the types 3, 4, 5 and 6 is in this quintile. Note that, again, our approach allows a better interpretation of the selected features by the analysis of the local dependence between the features and the cover type.

Finally, applying Algorithm 3 we choose the profiles displayed in Table 9 for $y \in \{1, 2, 3, 4, 5, 6, 7\}$. We see, for example, that the profile most prone to be of type 1 is (E,HH,HF) = Quintile 5 and of type 3 is (E,HH,HR,HF) = Quintile 1. Note that it does not mean that most of the terrains with these profiles are of type 1 and 3, but rather that the probability of a terrain with these profiles being of types 1 and 3, respectively, is 87% and 396% greater than the probability of a terrain for which we do not know the profile. Therefore, we see again the difference between the LLDS approach and the classification problem. In the LLDS approach, given a profile is greater than the marginal probability, while in the classification problem, given a profile, we are interested in determining the type for which the conditional probability given the profile is greater than the marginal probability, while in the classification problem, given a profile, we are interested in determining the type for which the conditional probability given the profile is greater than the marginal probability, while in the classification problem, given a profile, we are interested in determining the type for which the conditional probability given the profile is greater.

Features (χ)	Window (W)	$\eta_{\chi}(Y \mid W)$
(E,HH,HF)	Quintile 5	0.38
(E,A,HH,HF)	Quintile 5	0.38
(E,HH)	Quintile 5	0.37
(E,A,HH)	Quintile 5	0.37
(E,HH,VH,HF)	Quintile 5	0.37
(E,A,HH,VH,HF)	Quintile 5	0.36
(E,HH,HF)	Quintiles 1 & 5	0.36
(E,A,HH,VH)	Quintile 5	0.36
(E,HH,VH)	Quintile 5	0.36
(E,HH)	Quintiles 1 & 5	0.36

Table 8. Features selected applying Algorithm 2 to the Covertype dataset.

Table 9. Profiles selected applying Algorithm 3 to the Covertype dataset for $y \in \{1, 2, 3, 4, 5, 6, 7\}$.

Cover Type Features		LF Maximum	Profile			
1	(E,HH,HF)	1.87	Quintile 5			
2	(E,HH,HR)	1.63	Quintile 2			
3	(E,HH,HR,HF)	4.96	Quintile 1			
4 (E,HH,HF) ¹		5	Quintile 1			
5 (E,HR,HF)		3.31	Quintile 2			
6 (E,HH,HF)		4.90	Quintile 1			
7 (E,HH)		4.89	Quintile 5			
¹ Among other profiles.						

As an example, suppose the joint distribution that generated the LF of Table 7 and the profile Quintile 1. We have that the maximum conditional probability given this profile is the probability of type 2 (54, 473/116, 203 = 0.47), while the maximum lift is that of type 4, although its conditional probability is only 2,747/116,203 = 0.02. However, the conditional probability of type 4 given the profile, even though absolutely small, is relatively great: it is 5 times the marginal probability 0.004. Therefore, on the one hand, if there is a new terrain whose profile is (E,HH,HF) = Quintile 1, we classify it as being of type 2. On the other hand, if we want to sample terrains from a population and are interested in maximizing the number of terrains of type 4, we may sample from the population with profile (E,HH,HF) = Quintile 1 instead of the whole population, expecting to sample four times more terrains of type 4.

5. Final Remarks

The feature selection algorithms based on the LLDS extend the classical approach to feature selection to a higher resolution one, as they take into account the local dependence between the features and the variable of interest. Indeed, classical feature selection may be performed by walking through a tree in which each node is a vector of features, i.e., a BLFS, while feature selection based on the LLDS is established by walking through an extended tree, i.e., a CBLOP, in which inside each node there is another tree, that represents the windows of the features, as displayed in the example in Figure 1. Therefore, feature selection based on the LLDS increases the reach of feature selection algorithms to a new variety of applications.

The LLDS may treat a problem that is the opposite of that of classification, i.e., when we are interested in, given a class y, finding the profile x of which we may sample from its population in order to maximize the number of instances of class y. Indeed, in the classification problem we want to do the exact opposite: classify a instance with known profile x into a class of Y. Therefore, although LLDS tools may also be applied to the classification problem (as they are in the literature), they are of great importance in problems that we may call the *reverse engineering* of the classification one. Thus, our approach broadens the application of feature selection algorithms to a new set o problems by the extension of their search spaces from BLFs to CBLOPs.

The algorithms proposed in this paper may be optimized in order to not walk through the entire CBLOP, as its size increases exponentially with the number of features, so that the algorithm may not be computable for a great number of features. Moreover, the algorithms may be subjected to *overfitting* if the sample size is relatively small, so that their search space may be restricted. The methods of [1-7,9,10,22-29], for example, may be adapted to the multi-resolution algorithms in order to optimize them. Furthermore, the properties of the η coefficients and the LF must be studied in a theoretical framework, in order to establish their variances, sample distributions and develop statistical methods to estimate and test hypothesis about them.

The LLDS adapts classical measures, such as the MI and the Kullback-Leibler Divergence, into coherent dependence coefficients that assess the dependence between random variables in multiple resolutions, presenting a wide view of it. As it does not make any assumption about the dependence kind, the LLDS measures the raw dependence between the variables and, therefore, may be relevant for numerous purposes, being feature selection just one of them. We believe that the algorithms proposed in this paper, and the LLDS in general, bring advances to the state-of-art in dependence measuring and feature selection, and may be useful in various frameworks.

Supplementary Materials: The following are available online at www.mdpi.com/1099-4300/20/2/97/s1: an **R** [34] package called *localift* that performs the algorithms proposed by this paper; an **R** object that contains the results of the algorithms used to analyse the performances dataset; and an **R** code that apply the algorithms to the Congressional Voting Records and Covertype datasets.

Acknowledgments: We would like to thank Antonio Carlos Hernandes who kindly provided the performances dataset used in the application section. The Covertype dataset is Copyrighted 1998 by Jock A. Blackard and Colorado State University.

Author Contributions: Diego Marcondes wrote the paper and analysed the data of the application section. All authors substantially contributed to the theoretical content of the paper. All authors have read and approved the final manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

- BLFS Boolean lattice of feature sets
- CBLOP Collection of Boolean lattices of ordered pairs
- CE Conditional Entropy
- LF Lift Function
- LLDS Local Lift Dependence Scale
- MI Mutual Information
- SFFS Sequential Forward Floating Selection
- SFS Sequential Forward Selection

References

- 1. Guyon, I.; Elisseeff, A. An introduction to variable and feature selection. *J. Mach. Learn. Res.* 2003, 3, 1157–1182.
- 2. Kohavi, R.; John, G.H. Wrappers for feature subset selection. Artif. Intell. 1997, 97, 273–324.
- Lal, T.; Chapelle, O.; Weston, J.; Elisseeff, A. Embedded methods. In *Feature Extraction. Studies in Fuzziness and Soft Computing*; Guyon, I., Nikravesh, M., Gunn, S., Zadeh, L.A., Eds.; Springer: Berlin/Heidelberg, Germany, 2006; Volume 207, pp. 137–165.
- John, G.H.; Kohavi, R.; Pfleger, K. Irrelevant features and the subset selection problem. In Proceedings of the Eleventh International Conference on International Conference on Machine Learning, New Brunswick, NJ, USA, 10–13 July 1994; pp. 121–129.
- Hall, M.A. Correlation-based Feature Selection for Discrete and Numeric Class Machine Learning. In Proceedings of the Seventeenth International Conference on Machine Learning, Stanford, CA, USA, 29 June 29–2 July 2000; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 2000; pp. 359–366.

- Das, S. Filters, Wrappers and a Boosting-Based Hybrid for Feature Selection. In Proceedings of the Eighteenth International Conference on Machine Learning, Williamstown, MA, USA, 28 June–1 July 2001; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 2001; pp. 74–81.
- Yu, L.; Liu, H. Feature selection for high-dimensional data: A fast correlation-based filter solution. In Proceedings of the 20th International Conference on Machine Learning (ICML-03), Washington, DC, USA, 21–24 August 2003; pp. 856–863.
- 8. Martins, D.C.; Cesar, R.M.; Barrera, J. W-operator window design by minimization of mean conditional entropy. *Pattern Anal. Appl.* **2006**, *9*, 139–153.
- 9. Ris, M.; Barrera, J.; Martins, D.C. U-curve: A branch-and-bound optimization algorithm for U-shaped cost functions on Boolean lattices applied to the feature selection problem. *Pattern Recognit.* **2010**, *43*, 557–568.
- 10. Atashpaz-Gargari, E.; Reis, M.S.; Braga-Neto, U.M.; Barrera, J.; Dougherty, E.R. A fast Branch-and-Bound algorithm for U-curve feature selection. *Pattern Recognit.* **2018**, *73*, 172–188.
- 11. Mitra, P.; Murthy, C.; Pal, S.K. Unsupervised feature selection using feature similarity. *IEEE Trans. Pattern Anal. Mach. Intell.* **2002**, *24*, 301–312.
- 12. Lichman, M. UCI Machine Learning Repository, 2013. Available online: http://archive.ics.uci.edu/ml/ index.php (accessed on 26 January 2018).
- 13. Coppock, D.S. Why lift? Data modelling and mining. 2002. Available online: https://www.informationmanagement.com/news/why-lift (accessed on 26 January 2018).
- 14. Tufféry, S.; Riesco, R. Data Mining and Statistics for Decision Making; Wiley: Hoboken, NJ, USA, 2011.
- 15. Shannon, C.E.; Weaver, W. *The Mathematical Theory of Communication*; University of Illinois Press: Urbana, IL, USA, 1949; Volume 29.
- 16. Bouma, G. Normalized (Pointwise) Mutual Information in Collocation Extraction. 2009. Available online: https://svn.spraakdata.gu.se/repos/gerlof/pub/www/Docs/npmi-pfd.pdf (accessed on 26 January 2018).
- Role, F.; Nadif, M. Handling the Impact of Low Frequency Events on Co-occurrence based Measures of Word Similarity. In Proceedings of the International Conference on Knowledge Discovery and Information Retrieval, Paris, France, 26–29 October 2011; pp. 226–231.
- 18. Kullback, S.; Leibler, R.A. On information and sufficiency. Ann. Math. Stat. 1951, 22, 79–86.
- 19. Deng, L.Y. The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation, and Machine Learning; Taylor & Francis: Boca Raton, FL, USA, 2006.
- 20. Neto, U.M.B.; Dougherty, E.R. Error Estimation for Pattern Recognition; Wiley: Hoboken, NJ, USA, 2015.
- 21. Amaldi, E.; Kann, V. On the approximability of minimizing nonzero variables or unsatisfied relations in linear systems. *Theor. Comput. Sci.* **1998**, 209, 237–260.
- 22. Marill, T.; Green, D. On the effectiveness of receptors in recognition systems. *IEEE Trans. Inf. Theory* **1963**, *9*, 11–17.
- 23. Whitney, A.W. A direct method of nonparametric measurement selection. *IEEE Trans. Inf. Theory* **1971**, *100*, 1100–1103.
- 24. Pudil, P.; Novovičová, J.; Kittler, J. Floating search methods in feature selection. *Pattern Recognit. Lett.* **1994**, 15, 1119–1125.
- 25. Somol, P.; Pudil, P.; Novovičová, J.; Paclık, P. Adaptive floating search methods in feature selection. *Pattern Recognit. Lett.* **1999**, *20*, 1157–1163.
- Somol, P.; Novovičová, J.; Pudil, P. Flexible-hybrid sequential floating search in statistical feature selection. *Structural, Syntactic, and Statistical Pattern Recognition*; Springer: Berlin/Heidelberg, Germany, 2006; pp. 632–639.
- 27. Nakariyakul, S.; Casasent, D.P. An improvement on floating search algorithms for feature subset selection. *Pattern Recognit.* **2009**, *42*, 1932–1940.
- 28. Narendra, P.M.; Fukunaga, K. A branch and bound algorithm for feature subset selection. *IEEE Trans. Comput.* **1977**, *9*, 917–922.
- Ris, M.S. Minimization of Decomposable in U-Shaped Curves Functions Defined on Poset Chains—Algorithms and Applications. Ph.D. Thesis, Institute of Mathematics and Statistics, University of Sao Paulo, Sao Paulo, Brazil, 2012. (In Portuguese)
- Peng, H.; Long, F.; Ding, C. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans. Pattern Anal. Mach. Intell.* 2005, 27, 1226–1238.

- 31. Śmieja, M.; Warszycki, D. Average information content maximization—A new approach for fingerprint hybridization and reduction. *PLoS ONE* **2016**, *11*, e0146666.
- 32. Kwak, N.; Choi, C.H. Input feature selection by mutual information based on Parzen window. *IEEE Trans. Pattern Anal. Mach. Intell.* **2002**, *24*, 1667–1671.
- 33. Mahalanobis, P.C. On the generalized distance in statistics. Proc. Natl. Inst. Sci. (Calcutta) 1936, 2, 49–55.
- 34. R Core Team. *R: A Language and Environment for Statistical Computing;* R Foundation for Statistical Computing: Vienna, Austria, 2016.



 \odot 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0/).