

Classification of Imbalanced Data with a Geometric Digraph Family

Artür Manukyan

*Graduate School of Sciences and Engineering
Koç University
Sarıyer, 34450, Istanbul, Turkey*

AMANUKYAN13@KU.EDU.TR

Elvan Ceyhan

*Department of Statistics
University of Pittsburgh
Pittsburgh, 15260, PA, USA*

ELVANCEYHAN@GMAIL.COM

Editor: William Cohen

Abstract

We use a geometric digraph family called class cover catch digraphs (CCCDs) to tackle the class imbalance problem in statistical classification. CCCDs provide graph theoretic solutions to the class cover problem and have been employed in classification. We assess the classification performance of CCCD classifiers by extensive Monte Carlo simulations, comparing them with other classifiers commonly used in the literature. In particular, we show that CCCD classifiers perform relatively well when one class is more frequent than the other in a two-class setting, an example of the *class imbalance problem*. We also point out the relationship between class imbalance and class overlapping problems, and their influence on the performance of CCCD classifiers and other classification methods as well as some state-of-the-art algorithms which are robust to class imbalance by construction. Experiments on both simulated and real data sets indicate that CCCD classifiers are robust to the class imbalance problem. CCCDs substantially undersample from the majority class while preserving the information on the discarded points during the undersampling process. Many state-of-the-art methods, however, keep this information by means of ensemble classifiers, but CCCDs yield only a single classifier with the same property, making it both appealing and fast.

Keywords: Class Cover Catch Digraphs, Class Cover Problem, Class Imbalance Problem, Class Overlapping Problem, Graph Domination, Prototype Selection, Support Estimation

1. Introduction

Class imbalance problem has recently become a topic of extensive research. In a two-class setting, imbalance in class(es) occurs when one class is represented by far more observations (points) than the other class in the data set (see, e.g., Chawla et al. (2004) and López et al. (2013)). Class imbalance problem is observed in many areas such as medicine, fraud detection and education. Some examples are clinical trials in which only 5% of patients in the data set have a certain disease, such as cancer (Mazurowski et al., 2008); detecting fraudulent customers where most individuals are law-abiding in insurance, credit card and telecommunications industries (Phua et al., 2004); and archives of college students where

mostly the ones who have fair results are kept (Thai-Nghe et al., 2009). In these and many other real life cases, majority class (i.e., the class with larger size) confounds the classifier performance by hindering the detection of subjects from the minority class (i.e., the class with fewer points).

The classification methods in machine learning usually suffer from the imbalance of class sizes in the data sets because most of these methods work on the assumption that class sizes are balanced (Japkowicz and Stephen, 2002). For example, the commonly used k -nearest neighbor (k -NN) classification algorithm is highly influenced by the class imbalance problem. In the k -NN approach, a new point is classified as the class of the most frequent one from its first k nearest neighbors (Fix and Hodges, 1989; Cover and Hart, 1967). As a result, in a two-class setting where one class substantially outnumbers the other, a point is more likely to be classified as the majority class by the k -NN classifier. In literature, sensitivity of k -NN classifier to the class imbalance problem and some solutions on choosing the appropriate k have been discussed in cases of imbalanced classes (see Mani and Zhang, 2003; García et al., 2008; Hand and Vinciotti, 2003). Decision trees and support vector machines (SVM) are also some of the well known classifiers that are sensitive to the class imbalance in a data set (Japkowicz and Stephen, 2002; Tang et al., 2009). SVMs are among the most commonly used algorithms in the machine learning literature due to their well understood theory and high performance among popular algorithms (Wu et al., 2008; Fernández-Delgado et al., 2014), but these methods have been demonstrated to be inefficient against highly imbalanced data sets, although SVMs are still robust to moderately imbalanced data sets (Akbari et al., 2004; Raskutti and Kowalczyk, 2004).

We approach the classification of imbalanced data sets with methods that solve class cover problem (CCP), where the goal is to find a region that encapsulates all the members of the class of interest (i.e., target class). This particular region can be viewed as a *cover*; hence the name *class cover* (Cannon and Cowen, 2004). This problem is closely related to another problem in statistics, namely *support estimation*: estimating the support of a particular random variable defined in a measurable space (Schölkopf et al., 2001). Here, each cover can be realized as estimates of its associated class support. Priebe et al. (2001) introduced the class cover catch digraphs (CCCD) to find graph theoretic solutions to the CCP problem, and provided some results on the minimum dominating sets and the distribution of the domination number of such digraphs for one dimensional data. Priebe et al. (2003a) applied CCCDs on classification and showed that approximate minimum dominating sets of CCCDs (which were obtained by a greedy algorithm) and radii of the covering balls can be used to establish efficient classifiers. Moreover, DeVinney et al. (2002) defined random walk CCCDs (RW-CCCDs) where balls of class covers are more relaxed compared to previously introduced so called pure-CCCDs (P-CCCDs). In P-CCCDs, no member of the non-target class is covered, but RW-CCCDs allow some points of non-target class to be covered by the cover of the target class. Some target class points may also be uncovered in the process. Hence, RW-CCCDs may potentially avoid overfitting. CCCDs have been applied in face detection (Socolinsky et al., 2003) and latent class discovery in gene expression data (Priebe et al., 2003b). There are several other approaches in the literature to solve the class cover problem, including covering the classes with a set of boxes (Bereg et al., 2012) or set of convex hulls (Takigawa et al., 2009).

In this article, we study the effects of class imbalance on two CCCD classifiers, P-CCCD and RW-CCCD. Moreover, we report on the effects of class overlapping problem (which is defined as deterioration of classification performance when class supports overlap) along with the class imbalance problem to further investigate the performance of CCCD classifiers when imbalance and overlapping between classes co-exist. Thus, we show that when there is a considerable amount of class imbalance, whether class supports overlap or not, the CCCD classifiers perform better than the k -NN classifier. We show the robustness of CCCD classifiers to the class imbalance by simulating cases having increasing levels of class imbalance. We also compare CCCD classifiers with SVM classifiers which are potentially robust to moderate levels of class imbalance but not to high levels. With respect to class imbalance problem, the k -NN, SVM and decision tree classifiers may be referred to as “weak” classifiers; that is, these methods perform weakly when there is imbalance in the data set. However, such classifiers can be modified to address the unequal priors in a data set, and hence, can be converted to “strong” classifiers which are potentially robust to the class imbalance problem. We show that CCCD classifiers are also inherently robust (i.e., robust to class imbalance without any modification), and we compare the CCCD classifiers against the state-of-the-art strong classification methods which are constructed to perform well when class imbalance occurs. We consider ensemble learning, cost sensitive learning and resampling schemes in conjunction with k -NN, SVM and decision tree classifiers, and show that RW-CCCDs and P-CCCDs perform comparable to those strong classifiers.

Among the two variations of CCCD classifiers, we show that the RW-CCCD is more appealing in many aspects. For both simulated and real life examples, RW-CCCDs perform better than P-CCCDs and weak classifiers, and perform comparable to strong classifiers when the classes of data sets are imbalanced and/or overlapping. Moreover, we report on the complexity of the two CCCD classifiers and demonstrate that RW-CCCDs reduce the data sets substantially more than the other classifiers, thus increasing the testing speed. But most importantly, while reducing the majority class to mitigate the effects of class imbalances, CCCDs preserve the information on the discarded points of the majority class. CCCDs provide a novel potential solution to the class imbalance problem; that is, they capture the density around prototype points (i.e., members of the dominating sets) as radii of the covering balls. Hence, CCCDs preserve the information while reducing the data set. In the literature, only the strong classifiers based on hybrids of ensembles and resampling schemes achieve a similar task which requires multiple classifiers to be employed, and thus, result in lengthy training and testing time. However, CCCDs define single classifiers that undersample the data set with, possibly, a slight loss of information.

We provide a short review of the existing methods for classifying data sets with class imbalance in Section 2, introduce P-CCCD and RW-CCCD classifiers in Section 3, discuss the balancing effect of CCCD classifiers in Section 4. Finally, in Section 5, we compare the CCCD classifiers with the classifiers that are both sensitive (weak) and non-sensitive (strong) classifiers to class imbalance by simulated and real data sets, and report on the computational complexity of all weak classifiers.

2. Methods for Handling Class Imbalance Problem

Solving the class imbalance problem received considerable attention in the machine learning literature (see Chawla et al., 2004; Kotsiantis et al., 2006; Longadge and Dongre, 2013). Almost all algorithms designed to mitigate the effects of class imbalance incorporate a “weak” classifier which is modified to show some level of robustness to the class imbalance problem. The weak algorithm is modified either (i) in data level which involves a pre-processing of the data set being used in training, or (ii) in algorithmic level such that a “strong” classifier is constructed with a decision rule suited for the imbalances in the data set. Many modern algorithms are hybrids of both types; but in particular, there are mainly three of them: resampling methods, cost-sensitive methods, and ensemble methods (He and Garcia, 2009).

Resampling methods are commonly employed to remove the effects of class imbalance in the classification process. Resampling methods provide solutions to the class imbalance problem by (i) downsizing the majority class (undersampling) or (ii) generating new (synthetic) points for the minority class (oversampling). Hence, such methods modify the classifiers only at the data level. It might be useful to clean or erase some points in the majority class to balance the data (Drummond et al., 2003; Liu et al., 2009). However, in some cases, all points from both classes may be valuable/important, and hence, should be kept despite the differences in the class sizes. Oversampling methods generate synthetic points similar to the minority class to mitigate the class imbalance problem while preserving the information (Han et al., 2005). On the other hand, Batista et al. (2004) suggest that the combination of both over and undersampling methods can further improve the classification performance. One such method is the SMOTE+ENN method where the oversampling method SMOTE of Chawla et al. (2002) and edited nearest neighbors (ENN) method of Wilson (1972) are applied to an imbalanced data set, consecutively. While SMOTE balances the classes of the data set by generating artificial points between members of the minority class, ENN cleans the data set to further increase the classification performance of the weak classifier. Here, ENN method is an undersampling method that primarily aims to remove noisy points from the data set but not to balance the classes.

Another family of methods, namely cost-sensitive learning methods, has originated from real life: the cost of misclassifying a minority and a majority class member is usually not the same (Elkan, 2001). Frequently, the minority class has higher misclassification cost than the majority class. Classification methods such as decision trees (e.g., C4.5), can be modified to take these costs into account (see Ling et al., 2004; Zadrozny et al., 2003). C5.0 is an extended version of C4.5 incorporating the cost of each class (Kuhn and Johnson, 2013). Most weak classifiers can be easily modified so as to recognizing misclassification costs. The constrained violation cost C of SVM classifiers can be adjusted to individual class costs (Chang and Lin, 2011). As for k -NN, one solution is to appoint weights to all points of the data set with respect to their classes. Hence, such weights are the costs of classes giving precedence to minority class points (Barandela et al., 2003). On the other hand, for those algorithms that costs are not inherently recognizable or available, meta-learning schemes can be used along with weak classifiers without modifying the classifiers. Such learning methods are similar to ensemble learning methods (Domingos, 1999).

A fast developing field called ensemble learning also contributes to the family of methods handling the class imbalance problem (see Galar et al., 2012). The idea is to combine several classifiers to create a new classifier which has significantly better performance than its constituents (Rokach, 2010). AdaBoost is a popular algorithm among this family of learning methods (Freund and Schapire, 1997; Wu et al., 2008). AdaBoost assigns weights to each of the points in the data set and updates these weights in accordance with how well the points are estimated by each classifier. Galar et al. (2012) provide a survey of the most important ensemble learning methods that solve the class imbalance problem. However, it has been observed in some studies that ensemble learning methods work best when used together with resampling methods (López et al., 2013). In fact, ensembling and resampling schemes compensate the shortcomings of each other. The EasyEnsemble is a classifier with two levels of ensembles. First, a random undersampled majority class and the original minority class are used to train an ensemble classifier, then another random sample is drawn in the same way to train a second ensemble. This process is repeated several times to mitigate the effects of information loss as each ensemble would be applied on a different random subset of the majority class.

3. Classification with Class Cover Catch Digraphs

Class Cover Catch Digraphs (CCCDs) offer graph theoretic solutions to CCP (Priebe et al., 2001, 2003a). The objective of CCP is to find a region that covers the members of a specific class. More specifically, let (Ω, M) be a measurable space and let $\mathcal{X}_n = \{x_1, x_2, \dots, x_n\} \subset \Omega$ and $\mathcal{Y}_m = \{y_1, y_2, \dots, y_m\} \subset \Omega$ be observations from two classes \mathcal{X} and \mathcal{Y} with class conditional distributions F_X, F_Y and a joint cdf $F_{X,Y}$, respectively. Let $\Omega = \mathbb{R}^d$ and, without loss of generality, assume that the target class (i.e., the class of interest) is \mathcal{X} . In a CCCD, for $x_i, x_j \in \mathcal{X}_n \subset \mathbb{R}^d$ is the center of a ball with radius $r_i = r(x_i)$. Each ball is represented by $B_i = B(x_i, r_i)$ and if $x_j \in B_i$ then x_i is said to cover (or catch) x_j . Here, $d(\cdot, \cdot)$ can be any dissimilarity measure but we use the Euclidean distance henceforth. A CCCD is a digraph $D = D(V, A)$ with vertex set $V(D) = \mathcal{X}_n$ and the arc set $A(D)$ where $(x_i, x_j) \in A(D)$ if and only if $x_j \in B_i$. The term “catch” refers to arc (x_i, x_j) of the digraph D where x_i is said to catch x_j . The binary relation $x_i \sim x_j$, which is defined as $x_j \in B_i$, is asymmetric, thus the adjacency of x_i and x_j is represented with directed edges or arcs which yield a digraph instead of a graph.

In CCCDs, the goal is to find a subset of balls $C_{\mathcal{X}} \subseteq B_{\mathcal{X}} = \{B_1, B_2, \dots, B_n\}$ such that $Q_{\mathcal{X}} \subseteq \cup_{B \in C_{\mathcal{X}}} B$ for $Q_{\mathcal{X}} \subseteq \mathcal{X}_n$ where the set $Q_{\mathcal{X}}$ is some desirable subset of the target class training set \mathcal{X}_n which we want to cover. Preferably, the goal is to find a set $C_{\mathcal{X}}$ such that $Q_{\mathcal{X}} = \mathcal{X}_n$, however it might be desirable that the class cover may ignore some target class points to avoid overfitting. If a class cover of a CCCD fails to cover some target class points, it is called an *improper* cover, otherwise it is a *proper* cover. For covering \mathcal{Y}_m , we reverse the roles of classes \mathcal{X} and \mathcal{Y} . The class \mathcal{Y} becomes the target class and \mathcal{X} becomes the non-target class. Finding an appropriate cover $C_{\mathcal{X}}$ is equivalent to finding the dominating set of the CCCD with $V(D) = \mathcal{X}_n$. Let $N(s) = \{t \in V(D) : (s, t) \in A(D)\}$ be the *open neighborhood* of a vertex $s \in V(D)$: the set of vertices that have an arc from the vertex s , or the neighbors of s . A dominating set of a digraph D is defined as a subset of vertices $S \subseteq V(D)$ such that union of the closed neighborhoods, defined by $\bar{N}(s) = N(s) \cup \{s\}$, of elements of S

is the vertex set of the digraph: $\cup_{s \in S} \bar{N}(s) = V(D)$. Among all dominating sets, usually the ones with minimum cardinality, called the minimum dominating sets, are preferable. The cardinality of the minimum dominating set(s) is referred to as the *domination number*, denoted as $\gamma(D)$. However, minimum dominating sets are often computationally intractable and finding them is, in general, an NP-hard optimization problem. Hence, greedy algorithms are often employed to find sets with approximately minimum cardinality (Chvatal, 1979; DeVinney, 2003).

CCCDs can easily be generalized to the multi-class case with k classes. To establish the set of covers $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$ associated with a set of classes $\mathcal{X} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_k\}$, we merge the classes into two classes as $\mathcal{X}_T = \mathcal{X}_i$ and $\mathcal{X}_{NT} = \cup_{i \neq j} \mathcal{X}_j$ for $i, j = 1, \dots, k$. We refer to classes \mathcal{X}_T and \mathcal{X}_{NT} as target class and non-target class, respectively. More specifically, target class is the class we want to find the cover of, and the non-target class is the union of the remaining classes. We transform the multi-class case into a two-class setting and find the cover of i -th class, C_i , for each $i = 1, \dots, k$.

We employ two families of CCCDs, pure-CCCDs (P-CCCDs) and random walk CCCDs (RW-CCCDs) that differ in the definition of the radius $r(x)$. In these two digraphs, the (approximate) minimum dominating set S and the classifier are defined in slightly different ways; with the main distinction between the two being the way the covers are defined. The covering balls of P-CCCDs do not contain any non-target class point (hence the name “pure”) whereas RW-CCCDs possibly allow some non-target class points inside of the class cover of the target class so as to avoid overfitting. Moreover, some target class points may also be excluded from the covers of RW-CCCDs. Therefore, P-CCCDs construct pure and proper covers but RW-CCCD covers are not necessarily pure or proper.

3.1 Classification with P-CCCDs

In P-CCCDs, the covering balls $B_x = B(x, r(x))$ exclude all non-target class points. Thus, for a target class point $x \in \mathcal{X}_n$, which is the center of a ball B_x , the radius $r(x)$ should be smaller than the distance between x and the closest non-target point $y \in \mathcal{Y}_m$: $r(x) < \min_{y \in \mathcal{Y}_m} d(x, y)$. Given $\tau \in (0, 1]$, the radius $r(x)$ is defined as follows (Marchette, 2010):

$$r(x) := (1 - \tau)d(x, l(x)) + \tau d(x, u(x)), \tag{1}$$

where

$$u(x) := \operatorname{argmin}_{y \in \mathcal{Y}_m} d(x, y)$$

and

$$l(x) := \operatorname{argmax}_{z \in \mathcal{X}_n} \{d(x, z) : d(x, z) < d(x, u(x))\}.$$

The effect of parameter τ on the radius $r(x)$ is illustrated in Figure 1 (DeVinney, 2003). The ball with radius $r(x)$ catches the neighboring target class points, and for any $\tau \in (0, 1]$, the ball B_x catches the same points as well. Hence, the choice of τ does not effect the structure of digraph but might affect the classification performance which will be shown later in Section 5. On the other hand, for all $x \in \mathcal{X}_n$, the definition of $r(x)$ in Equation (1) keeps any non-target point $y \in \mathcal{Y}_m$ out of the ball B_x , that is $\mathcal{Y}_m \cap B_x = \emptyset$ for all $B_x \in \mathcal{C}_{\mathcal{X}}$. Here, B_x is an open ball: $B_x = \{z \in \mathbb{R}^d : d(x, z) < r(x)\}$. The digraph D is “pure” since

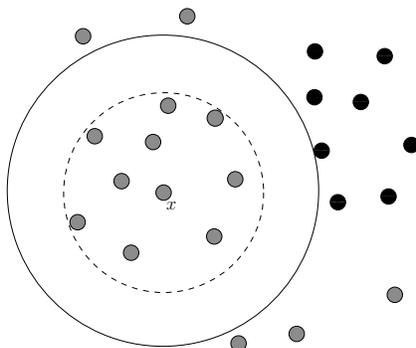


Figure 1: The effect of τ on the radius $r(x)$ of the target class point x in a two-class setting. Grey and black points represent the target and non-target classes, respectively. The solid circle centered at x is constructed with the radius associated with $\tau = 1$ and the dashed one with $\tau = 0.0001$ (DeVinney, 2003).

the balls contain only the target class points; hence, the name pure-CCCD. Once all balls are constructed, so is the digraph D . Therefore, we have to find a covering set $C_{\mathcal{X}}$ which is equivalent to finding a minimum dominating set $S \subseteq V(D)$. The greedy algorithm of finding an approximate minimum dominating set of a P-CCCD is given in Algorithm 1. At each iteration, the vertex which has the largest neighborhood (i.e., highest number of arcs) is removed from the graph together with its neighbors. Then, the process is repeated until all vertices of D are removed. The algorithm adds elements to the dominating set until all points are either dominated or dominate some other points. Hence, the covers established by P-CCCDs are proper covers: $Q_{\mathcal{X}} = \mathcal{X}_n$ and $Q_{\mathcal{Y}} = \mathcal{Y}_m$. The P-CCCD of one class, its associated class cover (constructed by the elements of the dominating set), and covers of both classes are illustrated in Figure 2.

Algorithm 1 The greedy algorithm for finding an approximate minimum dominating set of a digraph D . Here, $D[H]$ is the graph induced by the set of vertices $H \subseteq V(D)$ (see West, 2000).

Input: A digraph $D = D(V, A)$

Output: An approximate minimum dominating set, S

- 1: **set** $H = V(D)$ and $S = \emptyset$
 - 2: **while** $H \neq \emptyset$ **do**
 - 3: $v^* = \operatorname{argmax}_{v \in V(D)} |\bar{N}(v)|$
 - 4: $S = S \cup \{v^*\}$
 - 5: $H = V(D) \setminus \bar{N}(v^*)$
 - 6: $D = D[H]$
 - 7: **end while**
-

Before Algorithm 1 finds an approximate solution, we should first construct the digraph D . The P-CCCD cover $C_{\mathcal{X}}$ and the P-CCCD D depend on the distances between points of the target class \mathcal{X}_n , denoted by the matrix $\mathcal{M}_{\mathcal{X}}$, and the distances from all points of \mathcal{X}_n to all points of \mathcal{Y}_m , denoted by matrix $\mathcal{M}_{\mathcal{X}, \mathcal{Y}}$. Later, we construct the set of balls

$B_{\mathcal{X}} = \{B_1, B_2, \dots, B_n\}$, and get the set of arcs $A(D)$ where $V(D) = \mathcal{X}_n$. Hence, the minimum cardinality ball cover problem is reduced to a minimum dominating set problem. We find such a cover with Algorithm 2 which runs in quadratic time and, in addition, depends on the dimensionality of the training set $\mathcal{X}_n \cup \mathcal{Y}_m$.

Algorithm 2 The greedy algorithm for finding an approximate minimum cardinality ball cover $C_{\mathcal{X}}$ of the target class points \mathcal{X}_n given a set of non-target class points \mathcal{Y}_m .

Input: Points of the target class \mathcal{X}_n , the non-target class \mathcal{Y}_m and the P-CCCD parameter $\tau \in (0, 1]$

Output: An approximate minimum cardinality ball cover $C_{\mathcal{X}}$

- 1: $r(x) := (1 - \tau)d(x, l(x)) + \tau d(x, u(x))$ for all $x \in \mathcal{X}_n$
 - 2: Construct the digraph D with the set $B_{\mathcal{X}}$.
 - 3: Find the approximate minimum dominating set S of digraph D by Algorithm 1.
 - 4: $C_{\mathcal{X}} := \cup_{s \in S} B(s, r(s))$
-

Theorem 1 *Algorithm 2 is an $\mathcal{O}(\log n)$ -approximation algorithm and finds an approximate minimum cardinality ball cover $C_{\mathcal{X}}$ of the target class \mathcal{X} in $\mathcal{O}(n(n + m)d)$ time.*

Proof. The algorithm is polynomial time reducible to a greedy minimum set cover algorithm which finds an approximate solution with size at most $\mathcal{O}(\log n)$ times of the optimum solution (Chvatal, 1979; Cannon and Cowen, 2004). We first calculate the distance matrices $\mathcal{M}_{\mathcal{X}}$ and $\mathcal{M}_{\mathcal{X}\mathcal{Y}}$ which take $\mathcal{O}(n^2d)$ and $\mathcal{O}(nmd)$ time, respectively. Constructing the digraph D requires computing $l(x)$ and $u(x)$ in Equation (1) for all $x \in \mathcal{X}_n$, taking $\mathcal{O}(n^2 + nm)$ time in total. Then, we set the arc set $A(D)$ in $\mathcal{O}(n^2)$ time. Finally, the algorithm finds a solution for the digraph D in $\mathcal{O}(n^2)$ time, hence the total running time of the algorithm is $\mathcal{O}(n(n + m)d)$. ■

When \mathcal{Y} is the target class, observe that the time complexity is $\mathcal{O}(m(n + m)d)$, and an approximate solution is of size at most $\mathcal{O}(\log m)$ times the optimal solution by Theorem 1, since $m = |\mathcal{Y}_m|$. A P-CCCD classifier consists of the covers of all classes, hence the total training time of finding CCCDs of a data set with two-class setting is $\mathcal{O}((n + m)^2d)$.

After establishing both class covers $C_{\mathcal{X}}$ and $C_{\mathcal{Y}}$, any new data point can be classified in \mathbb{R}^d according to where it resides. Here, there are three cases according to the location of the given point, z , to be classified: z is (i) only in $C_{\mathcal{X}}$ or $C_{\mathcal{Y}}$, (ii) in both $C_{\mathcal{X}}$ and $C_{\mathcal{Y}}$ or (iii) in neither of $C_{\mathcal{X}}$ and $C_{\mathcal{Y}}$. The case (i) is straightforward: z belongs to class \mathcal{X} if $z \in C_{\mathcal{X}} \setminus C_{\mathcal{Y}}$ or to class \mathcal{Y} if $z \in C_{\mathcal{Y}} \setminus C_{\mathcal{X}}$. For cases (ii) and (iii), we need to find a way to decide the class of the point in a reasonable way. In fact, for all the cases, the estimated class of a given point z is determined by

$$\operatorname{argmin}_{C \in \{C_{\mathcal{X}}, C_{\mathcal{Y}}\}} \left[\min_{x: B(x, r) \in C} \rho(z, x) \right] \quad (2)$$

where $\rho(z, x) = d(z, x)/r(x)$ (Marchette, 2010). The dissimilarity measure $\rho(x, z)$ indicates whether or not the point z is in the ball of radius $r(x)$ with center x , since $\rho(x, z) \leq 1$ if z is inside the (closure of the) ball and > 1 otherwise. The measure $\rho : \Omega \times \Omega \rightarrow \mathbb{R}_+$ is simply a scaled dissimilarity measure, since Euclidean distance between two points, $d(x, y)$, is divided (or scaled) with the radius, $r(x)$ or $r(y)$. This measure violates the symmetry axiom among metric axioms since $\rho(x, y) \neq \rho(y, x)$ whenever $r(x) \neq r(y)$. However, Priebe et al. (2003a) showed that the dissimilarity measure ρ satisfies the continuity condition,

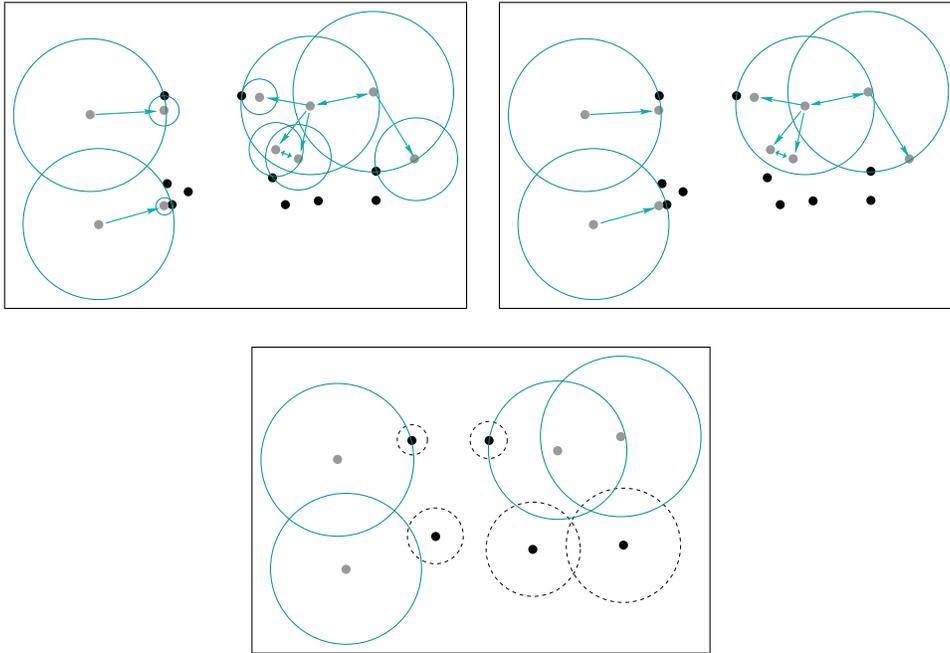


Figure 2: An illustration of the CCCDs (with the grey points representing the points from the target class) in a two-class setting. Presented in top left are all covering balls and the digraph $D = (V, A)$ and in top right are the balls that constitute a class cover for the target class and are centered at points which are the elements of the dominating set $S \subseteq V(D)$. In the bottom panel, we present the dominating sets of both classes and their associated balls which establish the class covers. The class cover of grey points is the union of solid circles, and that of black points is the union of dashed circles.

i.e., under the assumptions that both F_X and F_Y are continuous and strictly separable ($\inf_{x \in \mathcal{X}_n, y \in \mathcal{Y}_m} d(x, y) = \delta > 0$), P-CCCD classifiers are consistent; that is, their misclassification error approaches to the Bayes optimal classification error as $m, n \rightarrow \infty$. The measure ρ favors points with bigger radii; that is, for example, for a new point z equidistant to two points, the point with bigger radius is closer in terms of this scaled dissimilarity measure; for example, $\rho(x, z) < \rho(y, z)$ when $d(x, z) = d(y, z)$ and $r(x) > r(y)$. The radius $r(x)$ can be viewed as an indicator of the density around the point x . Thus, a point x with bigger radius might suggest that the point z is more likely to be drawn from the same distribution (or class) where x is drawn (i.e., from the denser class).

3.2 Classification with Random Walk CCCDs

For P-CCCDs, the class covers defined by CCCDs were “pure” of non-target class points; that is, no member of the non-target class was allowed inside the cover of the target class. As in Figure 1, the ball centered at the point x cannot expand any further since its radius is restricted by the distance to the closest non-target class point. This strategy may cause the

cover to overfit or be sensitive to noise or outliers in the non-target class. By allowing some neighboring non-target class points inside the cover and some target class points outside the cover, the random walk CCCDs (RW-CCCDs) catch as much target class points as possible with an adaptive strategy of choosing the radii (DeVinney et al., 2002). For $x \in \mathcal{X}_n$, $|\mathcal{X}_n| = n$ and $|\mathcal{Y}_m| = m$, RW-CCCDs define a function on radius of a ball given by

$$\begin{aligned} R_x(r) &= R_x(r; \mathcal{X}_n, \mathcal{Y}_m) \\ &:= \frac{m}{n} |\{z \in \mathcal{X}_n : d(x, z) \leq r\}| - |\{z \in \mathcal{Y}_m : d(x, z) \leq r\}|. \end{aligned} \quad (3)$$

where second and third arguments in $R_x(r; \mathcal{X}_n, \mathcal{Y}_m)$ are suppressed when there is no ambiguity. The function $R_x(r)$ can be viewed as a one-dimensional random walk. When the ball centered at $x \in \mathbb{R}^d$ expands, it hits either a target class point or a non-target class point which increases or decreases the random walk by one unit, respectively. The ratio m/n is included in the first term as to avoid the bias resulted by unequal sample sizes (i.e., class imbalance). An illustration is given in Figure 3 for the case $m = n$. The function $R_x(r)$ aims to find such radii that it contains a few non-target class points and sufficiently many target class points. In addition, we also want to avoid balls with large radii. Hence, the radius of x is the value maximizing $R_x(r)$ with an additional penalty function $P_x(r)$ which biases toward small radii:

$$r_x := \operatorname{argmax}_{r \in \{d(x, z) : z \in \mathcal{X}_n \cup \mathcal{Y}_m\}} R_x(r) - P_x(r). \quad (4)$$

Although a penalty function seems fit, DeVinney (2003) pointed out that the choice of $P_x(r) = 0$ usually works sufficiently well in practice. As in P-CCCDs, the radius of a ball represents the density of its center’s neighborhood. Maximizing $R_x(r)$ determines the best possible radius. Moreover, unlike P-CCCDs, the balls of RW-CCCDs are closed balls: $B_x = \{z \in \mathbb{R}^d : d(x, z) \leq r(x)\}$.

Similar to P-CCCDs, finding a cover, or a dominating set, of a RW-CCCD is an NP-hard problem. However, RW-CCCDs find the minimum dominating sets in a slightly different fashion. Instead of finding a set S such that $\cup_{s \in S} \bar{N}(s) = V(D)$ as in Algorithm 1, we first locate the vertex x^* (a target class point) which has maximum of some score, T_{x^*} , and remove all target and non-target class points covered with the ball of this vertex, B_{x^*} . In the next iteration, we recalculate the radii of remaining target class points, find the next point with the maximum score and continue until all target class points are covered. This greedy method of finding dominating set(s) S of RW-CCCDs is given in Algorithm 3. The resulting dominating set S has approximate minimum cardinality. For each target class point $x \in \mathcal{X}_n$, the score T_x is associated with $R_x(r_x)$ and is given by

$$T_x = R_x(r_x) - \frac{r_x n_u}{2d_m(x)} \quad (5)$$

where n_u is the number of uncovered target class points in the current iteration, and $d_m(x) = \max_{z \in \mathcal{X}_n} d(x, z)$. The term which is linear in r_x of the right hand side of Equation (5) is similar to $P(r)$ in Equation (4): it biases the scores toward choosing dominating points with smaller radii. On the other hand, Algorithm 3 is likely to choose dominating points with radius $r = 0$. These points only dominate themselves but they are thought of being

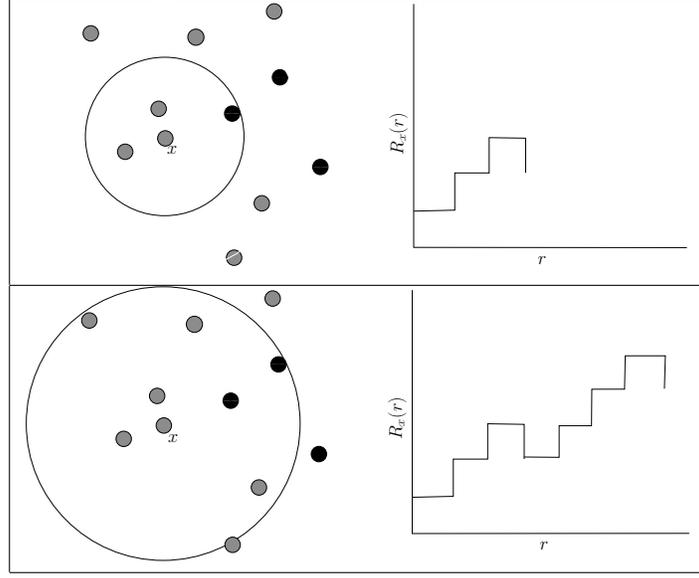


Figure 3: Two snapshots of $R_x(r)$ associated with the ball B_x centered at x for $m = n$.

Algorithm 3 The greedy algorithm for finding an approximate minimum dominating set for RW-CCCDs of points \mathcal{X}_n from the target class given non-target class points \mathcal{Y}_m .

Input: Target class points \mathcal{X}_n and non-target class points \mathcal{Y}_m

Output: Approximate dominating set S of \mathcal{X}_n

- 1: $H_0 = \mathcal{X}_n$, $H_1 = \mathcal{Y}_m$ and $S = \emptyset$
 - 2: $\forall x \in \mathcal{X}_n$, $d_m(x) = \max_{z \in \mathcal{X}_n} d(x, z)$
 - 3: **while** $H_0 \neq \emptyset$ **do**
 - 4: $n_u = |H_0|$
 - 5: **for all** $x \in \mathcal{X}_n$ **do**
 - 6: $r(x) = \operatorname{argmax}_r R_x(r; H_0, H_1)$ for $r \in \{d(x, z) : z \in H_0 \cup H_1\}$
 - 7: **end for**
 - 8: $x^* = \operatorname{argmax}_{x \in H_0} R_x(r(x); H_0, H_1) - \frac{r(x)n_u}{2d_m(x)}$
 - 9: $S = S \cup \{x^*\}$
 - 10: $H_0 = H_0 \setminus (\bar{N}(x^*) \cap \mathcal{X}_n)$ and $H_1 = H_1 \setminus (\bar{N}(x^*) \cap \mathcal{Y}_m)$
 - 11: **end while**
 - 12: $C_{\mathcal{X}} := \cup_{s \in S} B(s, r(s))$
-

not covered since their balls have radii $r = 0$. Hence, RW-CCCDs may establish improper covers.

Algorithm 3 is similar to Algorithm 2, however after each iteration, a point is added to the set S and the random walk $R_x(r)$ is recalculated for all uncovered $x \in A_0$. Hence, we need an additional sweep on the training set which makes Algorithm 3 run in cubic time.

Theorem 2 Algorithm 3 finds covers $C_{\mathcal{X}}$ of the target class \mathcal{X} and \mathcal{Y} in $\mathcal{O}((n + d + \log(n + m))(n + m)^2)$ time.

Proof. In Algorithm 3, the matrix of distances between points of training set $\mathcal{X}_n \cup \mathcal{Y}_m$ should be computed since, for all $x \in \mathcal{X}_n \cup \mathcal{Y}_m$, the entire data set is swept to maximize $R_x(r)$. This takes $\mathcal{O}((n + m)^2 d)$ time. The algorithm runs until all target class points

are covered, but for each iteration, the random walk $R_x(r)$ is recalculated. The maximum $R_x(r_x)$ could be found by sorting the distances for all $x \in A_0$ which could be done prior to the while loop. This sorting takes $\mathcal{O}((n+m)^2 \log(n+m))$ time. Since A_0 and A_1 are updated at each iteration, we can just erase the distances corresponding to points covered by $\bar{N}(x^*)$ which does not change the order of sorted list provided before the while loop. Hence, $\operatorname{argmax} R_x(r_x)$ is found and the covered points erased in $\mathcal{O}((n+m)^2)$ time. The while loop iterates n times in the worst case, and hence the algorithm runs in a total of $\mathcal{O}((n+d+\log(n+m))(n+m)^2)$ time. \blacksquare

Note that Algorithm 3 finds a cover of \mathcal{Y}_m in $\mathcal{O}((m+d+\log(n+m))(n+m)^2)$ time which makes a RW-CCCD classifier trained in $\mathcal{O}((n+m)^3)$ time for $d < n$ and $\log(n+m) < n$. RW-CCCD classifiers are much better classifiers that potentially avoid overfitting, but with a cost of being much slower compared to the P-CCCD classifiers.

Since P-CCCD covers are pure and proper covers, P-CCCD classifiers tend to overfit (DeVinney, 2003). In RW-CCCDs, covering balls allow some points of \mathcal{Y}_m inside $C_{\mathcal{X}}$ to increase average classification performance. In that case, Algorithm 3 cannot be reduced to a minimum set cover problem since the definition of sets change after adding a single point to the dominating set. Hence, the upper bound $\mathcal{O}(\log n)$ does not apply to RW-CCCDs. However, we expect to get bigger balls in RW-CCCDs compared the ones in P-CCCDs which intuitively suggests that the covers of RW-CCCDs are lower in cardinality. We conduct empirical studies to show that RW-CCCDs, in fact, produce dominating sets with lower size compared to P-CCCDs in some cases.

In RW-CCCD, once the class covers (or dominating sets) are determined, the scaled dissimilarity measure in Equation (2) is a good choice for estimating the class of a new point z . However, DeVinney (2003) incorporates the scores of each ball to produce better performing class covers in classification. Hence, the class of a new point z is determined by

$$\operatorname{argmin}_{C \in \{C_{\mathcal{X}}, C_{\mathcal{Y}}\}} \left[\min_{x: B(x,r) \in C} \rho(z,x)^{T_x^e} \right]$$

where $\rho(z,x)$ is defined as in Equation (2). Here, $e \in [0,1]$ controls at what level the score T_x is incorporated. We observe that for $d(z,x) < r(x)$, $\rho(z,x) = d(z,x)/r(x)$ decreases as T_x increases. Hence, if a new point z is in both covers, $z \in C_{\mathcal{X}} \cap C_{\mathcal{Y}}$, the score T_x is a good indicator to which class the new point z belongs since the bigger the T_x , the more likely the ball contains more target class points. For $e = 1$, we fully incorporate each score T_x of covering balls and with $e = 0$, we ignore the scores. By introducing a value for the parameter e in $(0,1)$, it is possible to further improve the performance of RW-CCCD classifiers.

4. Balancing the Class Sizes with CCCDs

The CCCD classifiers substantially reduce the number of majority class observations in a data set. The reason is that balls of majority class members are more likely to catch neighboring points of the same class. The greedy algorithm given in Algorithm 1 selects vertices with the largest closed neighborhood. Similarly, Algorithm 3 selects vertices so that their balls are as dense as possible (i.e., target class points are abundant in the balls) with some contaminating non-target class points. Both algorithms choose balls with a large

number of target class points, and hence substantially reduce the data set (in particular, majority class points). Points of the minimum dominating set correspond to the centers of balls that establish the class covers. Hence CCCD classifiers can also be viewed as *prototype selection* methods where the objective is finding a set of points, or *prototypes*, S ; from the training set to preserve or increase the classification performance while substantially reducing the sample size. However, the radii of dominating set(s) are also stored and used in the classification process.

In Figure 4, we illustrate the behavior of balls associated with P-CCCDs and RW-CCCDs. Note that in both families of digraphs, balls of the majority class tend to be larger and hence are more likely to catch more majority class points. Since the majority class has much more members than the minority class, balls of the majority points are more likely to catch the neighboring majority points. CCCD classifiers keep the information of ball centers and their associated radii. Larger cardinality of the majority class allows the construction of bigger balls and hence, larger values of radii are more likely to correspond to larger number of caught class members. As a result, CCCDs balance the data set and, at the same time, preserve the information of the local density by retaining the radii. The data set becomes balanced since the center of balls are the points of the new training data set which will be employed later in classification.

The loss of information in undersampling schemes are of course inevitable, however it is possible to preserve a portion of that discarded information by other means. EasyEnsemble is an ensemble classifier used for that very purpose; however, it needs multiple classifiers to be employed. Each classifier is trained on a different balanced subset of the original training data set, and hence the ensemble classifier preserves the information on the entire data set given by a collection of unbiased classifiers. On the other hand, CCCDs achieve the same goal by transforming the density around points into the radii. CCCDs resemble cluster based resampling methods in that regard. Instead of randomly sampling the data set, cluster based sampling schemes divide each class into clusters, and then, oversample the minority class or undersample the majority class proportional to each subclass. Covering balls of CCCDs have a similar purpose which has also been discussed in Priebe et al. (2003b). They use the covering balls of the minimum dominating sets to explore the latent subclasses of each class of gene expression data sets. In fact, the balls of CCCDs may correspond to clusters. Hence, sets of points associated with each cluster is undersampled to a single point (i.e., a prototype or a dominating point), and the information on the cluster is provided by the radius which represents the density of that cluster. The bigger the radius, the more influence a prototype has over the domain. In P-CCCDs, the radii may be sensitive to noise, but RW-CCCDs ignore noisy points to avoid overfitting. Moreover, in RW-CCCDs, we have an additional statistic provided by each cluster, the score given in Equation (5) based on the random walk. We use both the radii and these scores to define the RW-CCCD classifiers, and thus achieve better performing classifiers with more reduction and less information loss.

We approach the problem of class imbalance from the perspective of class overlapping problem as well. Several researchers on class imbalance revealed that overlap between the class supports degrade the classification performance of imbalanced data sets even more (see Prati et al., 2004; Batista et al., 2004, 2005; Galar et al., 2012). Let $E \subset \mathbb{R}^d$, and let $s(F_X)$ and $s(F_Y)$ be the supports of the classes \mathcal{X} and \mathcal{Y} , respectively. We define E as the overlapping region of these two class supports, $E := s(F_X) \cap s(F_Y)$. Moreover, let

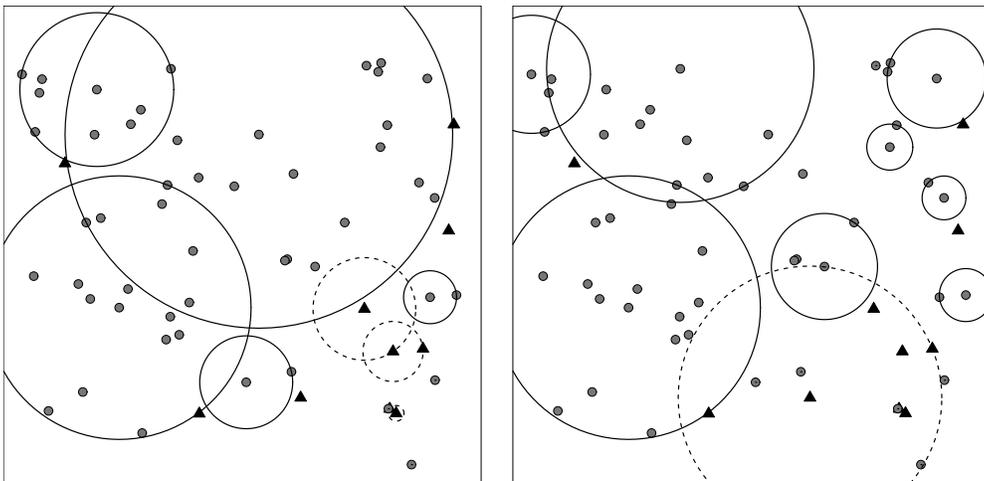


Figure 4: An illustration of the covering balls associated with majority and minority P-CCCD (left) and the corresponding RW-CCCDs ($\tau = 0.0001$) (right) of an imbalanced data set in a two-class setting where majority and minority class points are represented by grey dots and black triangles, respectively.

$q(E) := |\mathcal{Y}_m \cap E|/|\mathcal{X}_n \cap E|$ be ratio of class sizes restricted to the region $E \subset \mathbb{R}^d$. We say $q(E)$ is the “local” imbalance ratio with respect to E . Also, let the “global” imbalance ratio be $q = q(\mathbb{R}^d) = m/n$. Throughout this work, in both simulated and real data examples, we study and discuss the local imbalance ratio $q(E)$ restricted to the overlapping region E and the global imbalance ratio q . We specifically illustrate the performance of several classifiers for various levels of class imbalance (local or global) and class overlapping, and assess the performance of CCCD classifiers compared to weak and strong versions of k -NN, SVM and C4.5 classifiers.

5. Comparing CCCDs with Other Classifiers

We study the performance of CCCD classifiers in comparison with weak and strong classifiers in two separate sections. Recall that we call a classifier as “weak” when the method is inherently sensitive to class imbalance, and as “strong” when it is non-sensitive (or less sensitive). We use the area under curve (AUC) measure to evaluate the performance of the classifiers on the imbalanced data sets (López et al., 2013). AUC measure is often used on imbalanced real data classes. This measure has been shown to be better than the correct classification rate in general (Huang and Ling, 2005). We discuss the computational complexity of weak classifiers to emphasize the testing speed of CCCD classifiers when trained by imbalanced data sets. Finally, we compare both weak and strong classifiers with CCCDs on real data sets by considering the overlapping and imbalance ratios of all data sets.

5.1 Monte Carlo Simulation Study with Weak Classifiers

In this section, we compare the CCCD-based classifiers, namely P-CCCD and RW-CCCD, with k -NN, support vector machines (SVM) and C4.5, on simulated data sets. These classifiers are listed in Table 1. We employ the `cccd`, `e0171` and `RWeka` packages in R to classify test data sets with the P-CCCD, SVM (with Gaussian kernel) and C4.5 classifiers, respectively (Marchette, 2013; Meyer et al., 2014; R Core Team, 2015).

For each of four classification methods other than C4.5, we assign the optimum parameter values which are the best performing values among all considered parameters. For example, an optimum the P-CCCD parameter τ is found in a preliminary (pilot) Monte Carlo simulation study associated with the main simulation setting (i.e., the same setting of the main simulation). In the pilot study, we perform a Monte Carlo simulation with 200 replications and count how many times a τ value has the maximum AUC among $\tau = 0.0, 0.1, \dots, 1.0$ in 200 trials. Note that, since $\tau \in (0, 1]$, we denote $\tau = \epsilon$ (machine epsilon) as $\tau = 0$ for the sake of simplicity. For each replication of the pilot simulation, we (i) classify the test data set with all τ values, (ii) record the τ values with maximum AUC and (iii) update the count of the recorded τ values. Finally, we appoint the one that has the maximum count (the best performing τ) as the τ^* , the optimum τ . Then, we use τ^* as the parameter of P-CCCD classifier in our main simulation. The parameters of optimal k -NN, SVM and RW-CCCD classifiers are defined similarly. SVM methods often incorporate both a kernel parameter γ and a constrained violation cost C . We only optimize γ since the selection of an optimum C parameter will be more crucial for cost-sensitive SVM methods. Moreover, we consider two versions of the C4.5 classifier where both incorporate Laplace smoothing. The first tree classifier, C45-LP, prunes the decision tree with %25 confidence level but the second classifier, C45-LNP, does not use pruning at all.

We first consider a simulation setting similar to the one in DeVinney et al. (2002) where CCCD classifiers showed relatively good performance compared to the k -NN classifier. Here, we simulate a two-class setting where observations from both classes are drawn from separate multivariate uniform distributions: $F_X = U(0, 1)^d$ and $F_Y = U(0.3, 0.7)^d$ for $d = 2, 3, 5, 10$. Notice that $s(F_Y) \subset s(F_X)$; i.e., $E = s(F_Y)$. We perform Monte Carlo replications where on each replication, we train the data with equal sizes of observations ($m = n$) from each class for $n = 50, 100, 200, 500$. On each replication, we record the AUC measures of the classifiers on the test data set with 100 observations from each class, resulting a test data set of size 200. We simulate test data sets until AUCs of all classifiers achieve a standard error below 0.0005. Average of AUCs of all classifiers in Table 1 are given in Figure 5 for all (n, d) combinations. Additionally, in Figure 6, we report the τ values of best performing P-CCCD classifiers in our pilot simulation study for all (n, d) combinations. In Figure 6, there are separate histograms for each combination. Each histogram represents the number of times a τ value has the maximum AUC. Also in Figure 7, we report the e values of the best performing RW-CCCD classifiers of the same pilot simulation study for $e = 0, 0.1, \dots, 1.0$.

We start by investigating the effect of τ and e on CCCD classifiers. The relationship between τ , n and d can also be observed in Figure 6. The higher the τ value, the better the performance of P-CCCD classifier with increasing d and decreasing n . This may indicate that balls with $\tau = 0$ (i.e., $\tau = \epsilon$) represent the density around their centers better for low dimensional data sets. However, with increasing dimensionality and lower class sizes, the

Method	Description
P-CCCD	P-CCCD with the optimum τ (in the pilot study) among $\tau = 0, 0.1, \dots, 1.0$
RW-CCCD	RW-CCCD with the optimum e (in the pilot study) among $e = 0, 0.1, \dots, 1.0$
k -NN	k -NN with optimum k (in the pilot study) among $k = 1, 2, \dots, 30$
SVM	SVM with the radial basis function (Gaussian) kernel with the optimum γ (in the pilot study) among $\gamma = 0.1, 0.2, \dots, 3.9, 4.0$ (Joachims, 1999)
C45-LP	C4.5 with Laplace smoothing and reduced error pruning (%25 confidence)
C45-LNP	C4.5 with Laplace smoothing and no pruning

Table 1: The description of classifiers employed in the article.

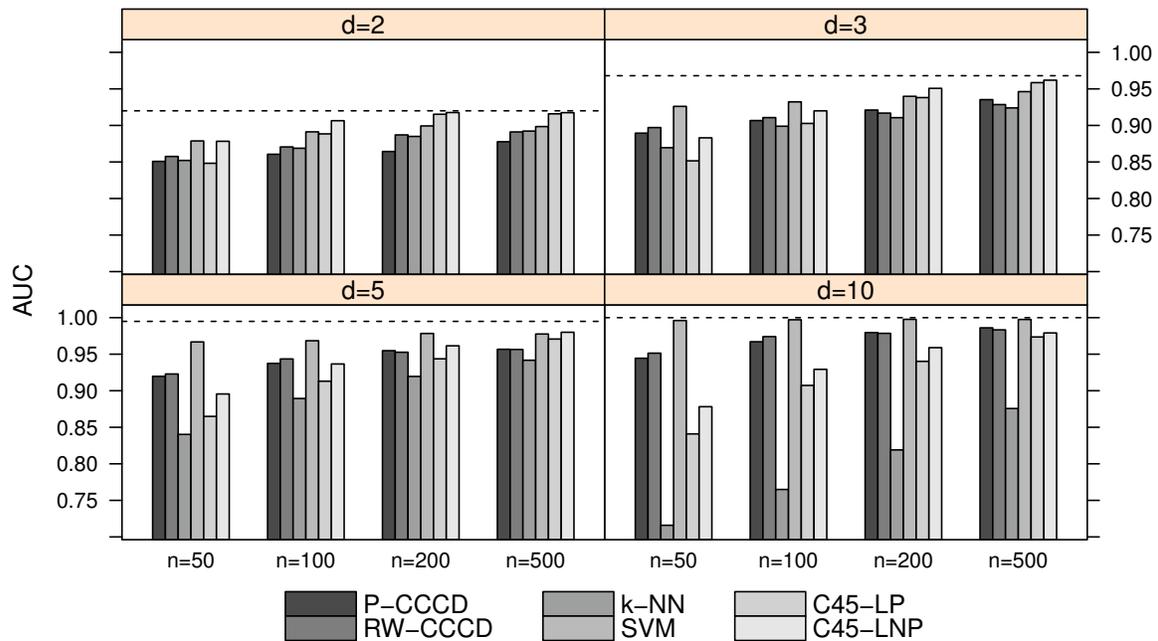


Figure 5: CCRs in the two-class setting, $F_X = U(0, 1)^d$ and $F_Y = U(0.3, 0.7)^d$ under various simulation settings, with $d = 2, 3, 5, 10$ and equal class sizes $m = n = 50, 100, 200, 500$.

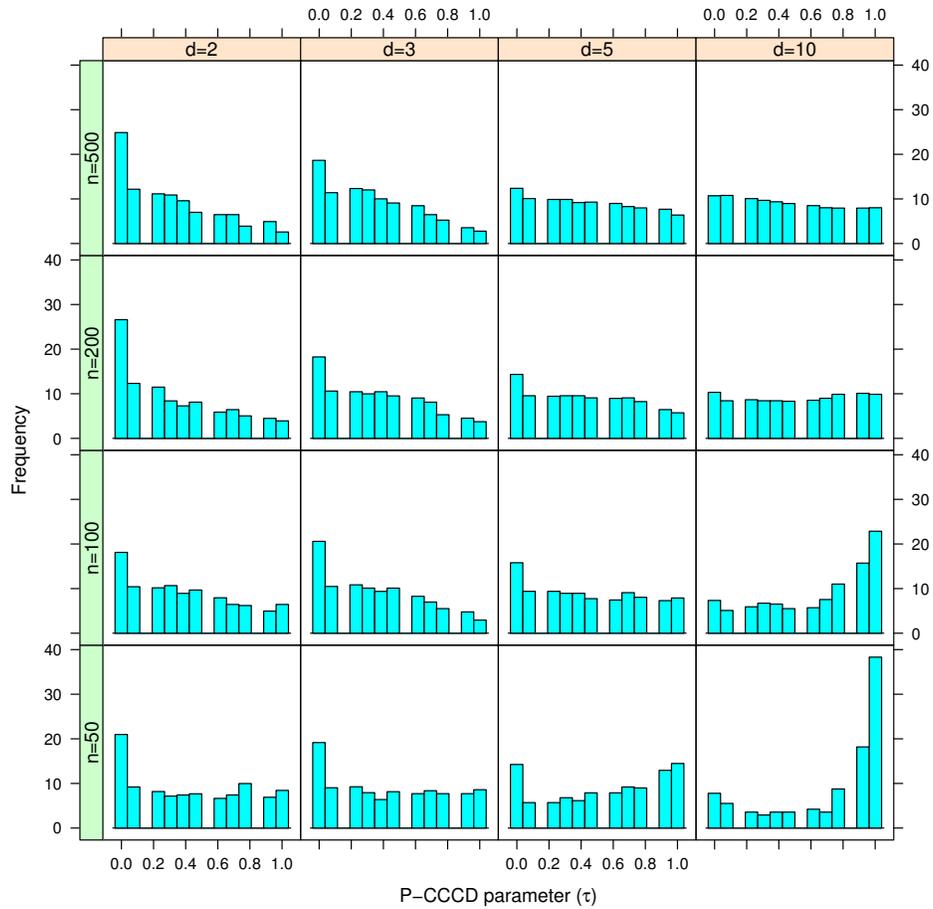


Figure 6: Frequencies of the best performing τ values among $\tau = 0.0, 0.1, \dots, 1.0$ in our pilot study (this is used to determine the optimal τ used in P-CCCD). The simulation setting is same as to the one presented in Figure 5.

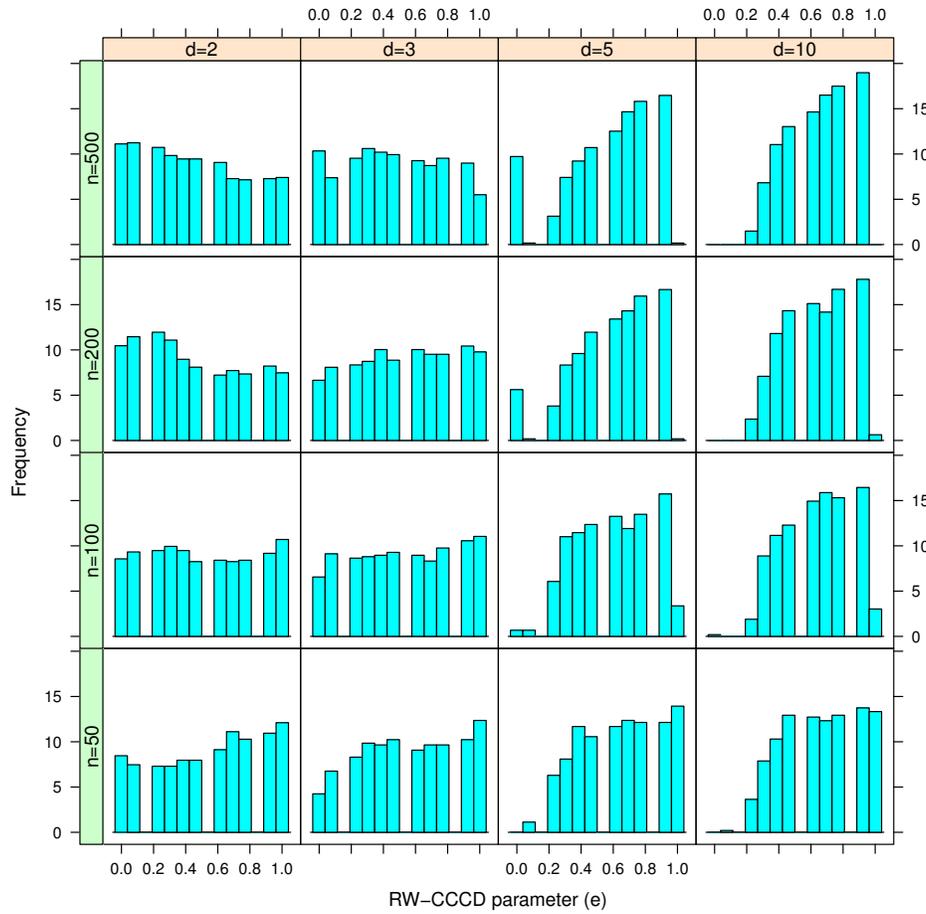


Figure 7: Frequencies of the best performing e values among $\tau = 0.0, 0.1, \dots, 1.0$ in our pilot study (this is used to determine the optimal e used in RW-CCCD). The simulation setting is same as to the one presented in Figure 5.

set of points gets sparser in \mathbb{R}^d . In the case of RW-CCCD, classifiers with high e values are either better or comparable to those with lower e values. The scores T_x of covering balls are definitely beneficial to the performance of the RW-CCCD classifiers, however with increasing n and decreasing d (especially for $n = 500$ and $d = 2$) RW-CCCD with lower e is better since the radii successfully represent the density around the prototype points due to the high number of observations in the data set.

Figure 5 illustrates the AUCs of all classifiers along with the Bayes optimal performance given with the dashed line. Comparing the performance of CCCD classifiers with other classification methods, we observe that RW-CCCD and P-CCCD classifiers outperform the k -NN classifier when the support of one class is entirely embedded inside that of the other class. These results are similar to the conclusions of DeVinney et al. (2002): with increasing dimensionality, the difference between k -NN and CCCD classifiers becomes more apparent, i.e., CCCD classifiers have nearly 0.20 AUC more than k -NN. On the other hand, the SVM classifier has about 0.05 more AUC than P-CCCD and RW-CCCD classifiers, especially for lower class sizes. Although, both versions of CCCD classifiers outperform the k -NN and C4.5 classifiers with increasing dimensionality, the gap between these two classifiers and CCCD classifier is getting narrower with increasing class sizes. The RW-CCCD classifier is slightly better than the P-CCCD classifier for lower n . In addition, C45-LNP achieves slightly better results than C45-LP.

In the setting presented in Figure 5, apparently, two classes overlap on the region $E = s(F_Y) = [0.3, 0.7]^d$ which is the entire support of the class \mathcal{Y} . For equal class sizes, $q = m/n = 1$ but $q(E) \approx (1/0.4)^d = \text{Vol}(s(F_Y))/\text{Vol}(s(F_X))$, where $\text{Vol}(\cdot)$ is the volume functional. The classes are clearly imbalanced in E , although $m = n$. Hence, class \mathcal{X} becomes the minority and class \mathcal{Y} becomes the majority class with respect to E . However, readjusting the class sizes m and n might change the performance of P-CCCD and RW-CCCD classifiers compared to the k -NN and C4.5 classifiers. Therefore, we conduct another simulation study with classes from the same uniform distributions, but we set $m = 50$ and $n = 200$ for $d = 2, 3$, and $m = 50$ and $n = 1000$ for $d = 5, 10$. In this experiment, we simulated 4 times more \mathcal{X} class members than \mathcal{Y} for $d = 2, 3$, and 20 times more for $d = 5, 10$. Results of this second experiment is given in Figure 8. k -NN and C4.5 classifiers outperform P-CCCD classifier in all d cases and has comparable AUC with SVM. However, only for $d = 2, 5$, RW-CCCD classifier achieves considerably more or comparable AUC compared to other classifiers. In this example, k -NN classifiers have nearly 0.05 more AUC than P-CCCDs, and also RW-CCCDs have, in general, 0.05 more AUC than k -NN classifiers.

Results from Figures 5 and 8 seem conflicting to each other, even though $E = s(F_Y)$. In the simulation setting of Figure 8, we draw more samples from the class \mathcal{X} to balance the class sizes with respect to E . In fact, the effect on the difference of AUCs between CCCD, k -NN and C4.5 classifiers depends heavily on the local class imbalance restricted to the overlapping region E . The classes in region E are less imbalanced in setting of Figure 8 than in the setting of Figure 5. Observe that $q(E) \approx (1/0.4)^d/4$ when $(m, n) = (50, 200)$, $q(E) \approx (1/0.4)^d/20$ when $(m, n) = (50, 1000)$, and $q(E) \approx (1/0.4)^d$ in $(m, n) = (50, 50)$. Hence, d does also affect the balance between classes. With increasing d , the region E gets smaller in volume compared to $s(F_X)$ and, as a result, fewer points of the class \mathcal{X} falls in E . Thus, we need to draw more samples from \mathcal{X} as dimensionality increases, in order to balance the classes with respect to E . These results suggest that, the more imbalanced the data set

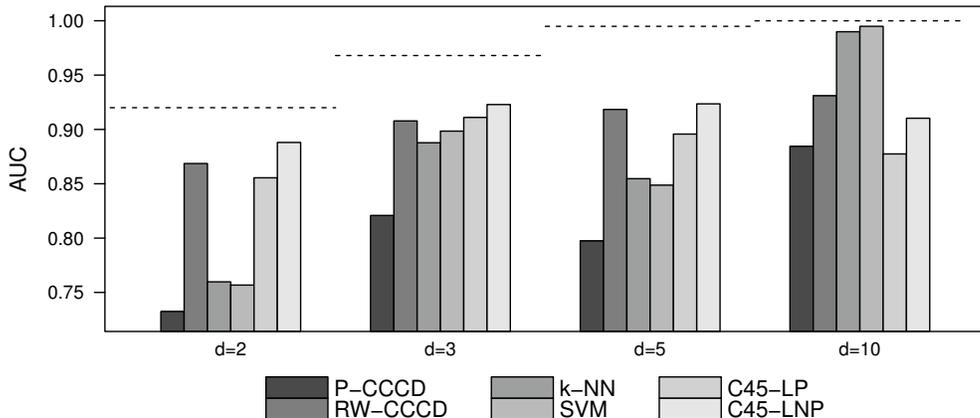


Figure 8: CCRs in a two-class setting, $F_X = U(0, 1)^d$ and $F_Y = U(0.3, 0.7)^d$ with fixed $n = 200$ and $m = 50$ in $d = 2, 3$, and with fixed $n = 1000$ and $m = 50$ in $d = 5, 10$.

in overlapping region E , the worse the performance of k -NN and C4.5 classifiers while CCCD classifiers preserve their classification performance. So, CCCD classifiers exhibit robustness (to the class imbalance problem). On the other hand, in Figure 5, we observe that the AUC of k -NN classifier approaches to the AUC of CCCD classifiers with increasing class sizes. Because, when q and $q(E)$ are fixed, the classification performance still depends on individual values of n or m . This result is in line with the results of Japkowicz and Stephen (2002) who reported that the effect of class imbalance on the classification performance diminishes if both class sizes are sufficiently large. Furthermore, SVM classifier performs better than all classifiers in Figure 5, and performs worse than RW-CCCD classifiers only for $d = 2, 5$ in Figure 8. This might be an indication that SVM classifier is also not affected by the local class imbalance with respect to E , and performs usually better than both P-CCCD and RW-CCCD classifiers if the support of one class is inside the other. For the C4.5 classifier, on the other hand, it is known for quite some time that the pruning is detrimental for classifying imbalanced data sets (Cieslak and Chawla, 2008). In any case, C45-LNP has more AUC than C45-LP in all simulation settings.

In a two-class setting with an overlapping region E , we should expect CCCD classifiers to outperform k -NN classifiers in cases of (global or local) class imbalance. Let $F_X = U(0, 1)^d$ and $F_Y = U(\delta, 1 + \delta)^d$ for $\delta, q = 0.05, 0.10, \dots, 0.95, 1.00$; $d = 2, 3, 5, 10$; $n = 400$ and $m = qn$. Here, the shifting parameter δ controls the level of overlap. The class supports get more overlapped with decreasing δ . Since $E = (\delta, 1)^d$ and the supports of both classes are unit boxes, observe that $q(E) \approx q$. The closer the value of q to 1, more balanced the classes are. We aim to address the relationship between the classifiers for various combinations of overlapping and global class imbalance ratios.

Figure 9 illustrates the difference between AUCs of CCCD and other classifiers (k -NN, SVM and C4.5) in separate heat maps for $d = 2, 3, 5, 10$. We use the unpruned C4.5 classifier

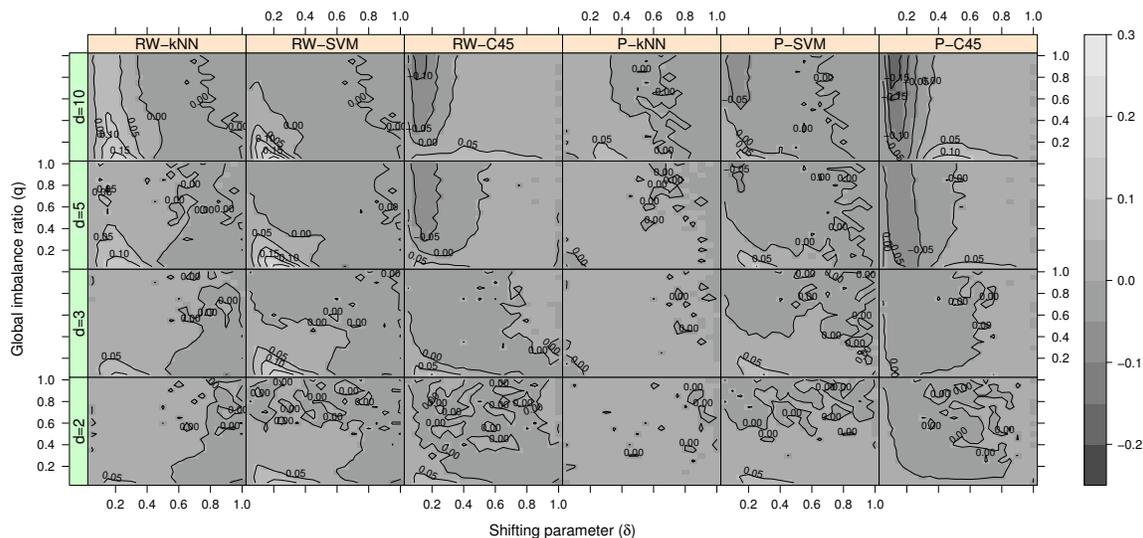


Figure 9: Differences between the AUCs of CCCD and other classifiers. For example, the panel titled with "RW-kNN" presents $AUC(\text{RW-CCCD}) - AUC(k\text{NN})$. In this two-class setting, classes are drawn from $F_X = U(0, 1)^d$ and $F_Y = U(\delta, 1 + \delta)^d$ with $d = 2, 3, 5, 10$. Each cell of the grey scale heat map corresponds to a single combination of simulation parameters $\delta, q = 0.05, 0.1, \dots, 0.95, 1.00$ with $n = 400$ and $m = qn$.

C45-LNP, since it tends to perform better for imbalanced data sets, and we refer to C45-LNP as C4.5 for simplicity. Each cell of a single heat map is associated with a combination of δ and q values. Lighter tone cells indicate that CCCD classifiers are better than the other classifiers in terms of AUC, and vice versa for the darker tones. When the classes are imbalanced and moderately overlapping, RW-CCCD classifier has at least 0.05 more AUC than all other non-CCCD classifiers but P-CCCD classifier is only better than all others provided that $d = 10$. If the classes are balanced or their supports are not considerably overlapping, there seem to be no visible difference between CCCD and the other classifiers. Thus, the other classifiers suffer from the imbalance of the data while CCCD classifiers show robustness to the class imbalance. But more importantly, this difference is getting more apparent with increasing dimensionality. When d is high, fewer points of the minority class fall in E although $q(E)$ is fixed. Even though the classes are imbalanced, if the minority class have substantially small size, the class imbalance problem becomes more detrimental (Japkowicz and Stephen, 2002). Under the conditions that the data set has substantial imbalance and overlapping, AUC of RW-CCCD classifier is followed, in order by, the AUC of C4.5, SVM and k -NN classifiers.

Unlike the comparison of CCCD and SVM classifiers in Figures 5 and 8, SVM classifier has less AUC than CCCD classifiers with low δ and low q values in Figure 9. In this setting, n is fixed to 400 and the lowest value of m is 20. Compared to our experiments in Figures 5 and 8, this setting produces highly imbalanced data sets (one class has far

more observations than the other, $m \ll n$). Akbani et al. (2004) conducted a detailed investigation and listed some reasons of SVM classifier being sensitive to highly imbalanced UCI data sets (Bache and Lichman, 2013). They did not, however, address the problem of overlapping class supports but offered a modification to SMOTE algorithm in order to improve the robustness of SVM. On the other hand, especially for $d = 5$ and $d = 10$, SVM, k -NN and C4.5 classifiers have more AUC than CCCD classifiers with increasing q and decreasing δ . This may indicate that other weak classifiers are better than CCCD classifiers for balanced classes.

The effects of class imbalance might also be observed when the class supports are well separated. If the class supports are disjoint, that is $s(F_X) \cap s(F_Y) = \emptyset$, the AUC is fairly high. However, it might still be affected by the global imbalance level, q . Therefore, we simulate a data set with two classes where $F_X = U(0, 1)^d$ and $F_Y = U(1 + \delta, 2 + \delta) \times U(0, 1)^{d-1}$. Figure 10 illustrates the results of this simulation study. Both class supports are d dimensional unit boxes as in the previous simulation setting, however they are now disjoint (separated along the first dimension). In addition, the parameter δ controls the smallest distance between the class supports where $\delta = 0.05, 0.10, \dots, 0.45, 0.50$. With increasing δ , the points of class \mathcal{Y} move further away from the points of \mathcal{X} . Figure 10 illustrates the difference between AUCs of CCCD and other classifiers under this simulation setting.

In Figure 10, unlike the performance of CCCD classifiers in Figure 9, P-CCCD classifiers have more AUC than RW-CCCD classifiers. When classes are imbalanced and supports are close, P-CCCD classifiers outperform both SVM and k -NN classifiers for all d values, but RW-CCCD classifiers have nearly 0.03 more AUC than these classifiers only in $d = 10$. However, this is not the case with C4.5 classifier since none of the classifiers outperform C4.5; that is, C4.5 yields over 0.04 more AUC than CCCD classifiers. A well separated data set is more likely to be classified better with C4.5 tree classifier because a single separating line exists between the two class supports. Hence, C4.5 locates such a line and efficiently classifies points regardless of the distance between class supports as long as the distance is positive. On the other hand, the balls of P-CCCD classifiers establish appealing covers for the class supports because the supports do not overlap. P-CCCD classifiers establish covering balls, big enough to catch substantial amount of points from the same class. Similarly, RW-CCCD classifiers establish pure covers, and this is the result of the separation between class supports. However, P-CCCD classifiers achieve better classification performance than RW-CCCD classifiers. When the classes are well separated, the radii of a ball in random walk, say from class \mathcal{X} , is likely $\max_{z \in \mathcal{X}_n} d(z, x)$ but in P-CCCD classifiers, it is $\min_{z \in \mathcal{Y}_n} d(z, x)$. In fact, the RW-CCCD classifiers are nearly equivalent to P-CCCD classifiers. Thus, when $\tau > 0$, P-CCCD classifiers are more likely to produce bigger balls than RW-CCCD classifiers, and potentially avoid overfitting.

In Figure 10, RW-CCCD classifiers have slightly or considerably less AUC than other classifiers when data sets are imbalanced and the supports are slightly far away from each other. The random walk contaminates the class cover with some non-target class points to improve the classification performance. However, since the classes are well separated and one class has substantially fewer points than the other, random walks are likely to yield balls to cover some points from the support of the non-target class, resulting in a degradation in the performance of RW-CCCD classifiers. On the other hand, P-CCCD classifiers outperform both k -NN and SVM classifiers for lower q and lower δ . The closer and more imbalanced the

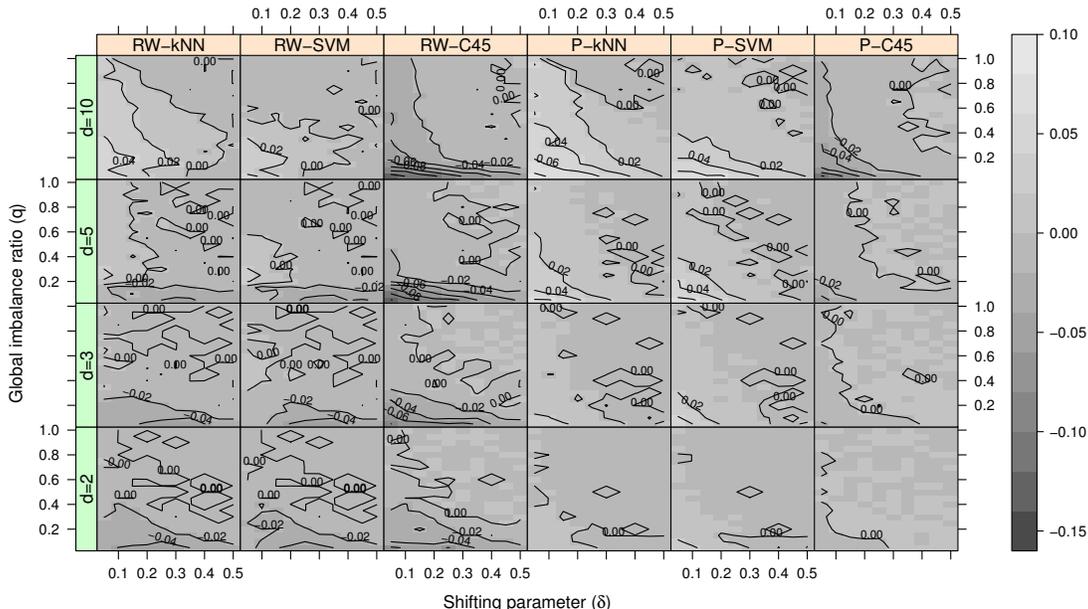


Figure 10: Differences between the AUCs of CCCD and other classifiers (see Figure 9 for details). In this two-class setting, classes are drawn from $F_X = U(0, 1)^d$ and $F_Y = U(1 + \delta, 2 + \delta) \times U(0, 1)^{d-1}$ where $d = 2, 3, 5, 10$, $\delta = 0.05, 0.1, \dots, 0.45, 0.50$ and $q = 0.05, 0.1, \dots, 0.95, 1.00$ with $n = 400$ and $m = qn$. AUCs of all classifiers are over 88% since the class supports are well separated.

data, the better the performance of P-CCCDs than other classifiers. Although the classes do not overlap, the effect of class imbalance is still observed when the supports are close. When there is mild imbalance between classes, CCCD classifiers have either comparable or less AUC. In addition, note that the performances of SVM and k -NN classifiers deteriorate but P-CCCD classifiers preserve their AUC with increasing d . Let $E \subset \mathbb{R}^d$ be some region that contains points of both classes which are sufficiently close to the decision boundary. With increasing d , fewer minority class points are in this region, and hence fewer members of this class fall in E . As a result, the performance of both SVM and k -NN classifiers suffer from local class imbalance with respect to E .

Finally, we investigate the effect of dimensionality when classes are balanced (i.e., $q = 1$) and their supports are overlapping. In this setting, $F_X = U(0, 1)^d$ and $F_Y = U(\delta, 1 + \delta)^d$. Here, let $q(E) \approx q = 1$, hence the classes are also locally balanced with respect to E as well as being globally balanced. Also, δ controls the level of overlap between two classes. However, we define δ in such a way that the overlapping ratio $\alpha \in [0, 1]$ is fixed for all dimensions. When α is 0, the supports are well separated, and when α is 1, the supports of classes are the same, i.e., $s(F_X) = s(F_Y)$. The closer α to 1, the more the supports overlap. Observe that $\delta \in [0, 1]$ can be expressed in terms of the overlapping ratio α and

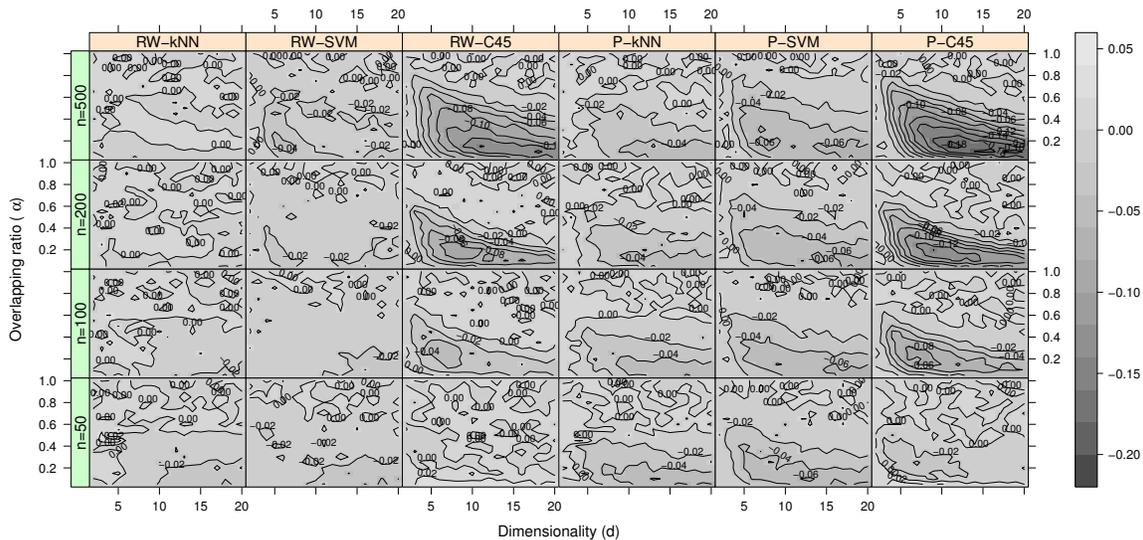


Figure 11: Differences between the AUCs of CCCD and other classifiers (see Figure 9 for details). In this two-class setting, classes are drawn from $F_X = U(0, 1)^d$ and $F_Y = U(\delta, 1 + \delta)^d$ where $n = 50, 100, 200, 500$, $\alpha = 0.05, 0.1, \dots, 0.45, 1.00$ and $d = 2, 3, 4, \dots, 20$.

dimensionality d :

$$\alpha = \frac{\text{Vol}(s(F_X) \cap s(F_Y))}{\text{Vol}(s(F_X) \cup s(F_Y))} = \frac{(1 - \delta)^d}{2 - (1 - \delta)^d} \iff \delta = 1 - \left(\frac{2\alpha}{1 + \alpha} \right)^{1/d}. \quad (6)$$

Hence, we calculate δ for each (d, α) combination by the Equation (6). In Figure 11, each cell of the grey scale heat map corresponds to a single combination of simulation parameters $\alpha = 0.05, 0.1, \dots, 0.95, 1.00$ and $d = 2, 3, 4, \dots, 20$. In Figure 11, the differences between the AUCs of CCCD classifiers and other classifiers are up to 0.20. The k -NN and SVM classifiers have comparable performance with CCCD classifiers, or outperform both CCCD classifiers. However, C4.5 has more AUC with increasing d . Employing CCCD classifiers do not considerably increase the classification performance over other classifiers when classes are balanced.

5.2 Empirical Comparison of CCCD-based and Strong Classifiers

In this section, we compare the CCCD-based classifiers with strong versions of k -NN, SVM and C4.5 classifiers on simulated data sets. Each classifier is modified in three different schemes, namely, resampling, ensemble and cost-sensitive schemes. We use SMOTE+ENN algorithm as the resampling scheme and EasyEnsemble algorithm as the ensembling scheme. As for the cost sensitive versions on weak classifier, we adjust the classifiers into recognizing class weights. For k -NN, we employ an algorithm giving more weight on neighboring minority class members; for SVM, we use two separate constrained violation costs for each

Method	Description
SMOTE+ENN	A combination of SMOTE ($t = 2$ and $k = 5$) and ENN ($k = 3$) (Batista et al., 2004).
EasyEnsemble	A combination of undersampling ($T = 4$) and Adaboost ($s_i = 10$) for $i = 1, 2, \dots, T$ (Liu et al., 2009)
C5.0	The cost sensitive version of C4.5 (Kuhn and Johnson, 2013).
CkNN	A cost sensitive version of k -NN (Barandela et al., 2003).
CSVM	A cost sensitive version of SVM (Chang and Lin, 2011).

Table 2: The description of classifiers employed in the article.

corresponding class; and for C4.5, we employ the C5.0. With three schemes and three weak classifiers, we get nine strong classifiers to study. We list and describe all these schemes in Table 2.

SMOTE+ENN algorithm, first, oversamples the entire training data set by generating artificial points in between a point and its neighbors. Specifically, for each point in the data set, t points among k neighbors are selected, and until the data set is balanced, new artificial points are generated in between these points and their selected neighbors. Later, ENN algorithm cleans the data set of noisy points by checking all points if the majority of their k neighbors are labeled as the class of the point. If not, the point is erased from the data set. Simply, SMOTE+ENN is a hybrid of over and undersampling methods. EasyEnsemble algorithm is a hybrid of undersampling and ensemble methods. An ensemble of weak classifiers is established by generating T many undersampled balanced data sets from the training data set. Then, each data set is used to train individual Adaboost classifiers with s_i many weak classifiers for $i = 1, 2, \dots, T$. Hence, EasyEnsemble is an ensemble of $\sum_{i=1}^T s_i$ many weak classifiers.

We choose one of the simulation settings conducted in Section 5.1. Since CCCD classifiers are observed to be better than other classifiers when both class imbalance and overlapping occurred, we only compare CCCD classifiers with strong classifiers on a single simulation setting. Hence we choose the setting presented in Figure 9, i.e., we let $F_X = U(0, 1)^d$ and $F_Y = U(\delta, 1 + \delta)^d$ for $\delta, q = 0.05, 0.10, \dots, 0.95, 1.00$, $n = 400$ and $m = qn$. We aim to highlight the differences between the strong classifiers and CCCD classifiers for various combinations of overlapping and class imbalance ratios. The results on average AUCs of each strong classifier is given in Figure 12. In general, RW-CCCDs seem to perform better than P-CCCDs. For $d > 2$, P-CCCDs have nearly 0.10 less AUC than RW-CCCDs when the classes are substantially overlapping and imbalanced, and it is observed that P-CCCDs are usually worse compared to the strong classifiers considered. However, the AUCs of RW-CCCD classifiers are either comparable or slightly less compared to others with the most difference being seen in the case of $d = 10$ when RW-CCCDs compared to EC4.5 and C5.0, ensemble and cost sensitive versions of the C4.5 classifier, respectively. However, with decreasing δ and q , the RW-CCCDs have only 0.05 less AUC than others. Also, RW-CCCDs seem to have 0.05 more AUC than C5.0 for moderately overlapping and imbalanced data sets, and seem to have 0.05 more AUC than ESVM, ensemble based SVMs, when the data set is both overlapping and imbalanced. This suggests that RW-CCCDs yield comparable results in comparison to the state-of-the-art robust methods when class imbalance and class overlapping co-exist. Additionally, we show in Section 5.3 that RW-CCCDs generate prototype sets that considerably reduce the training data set.

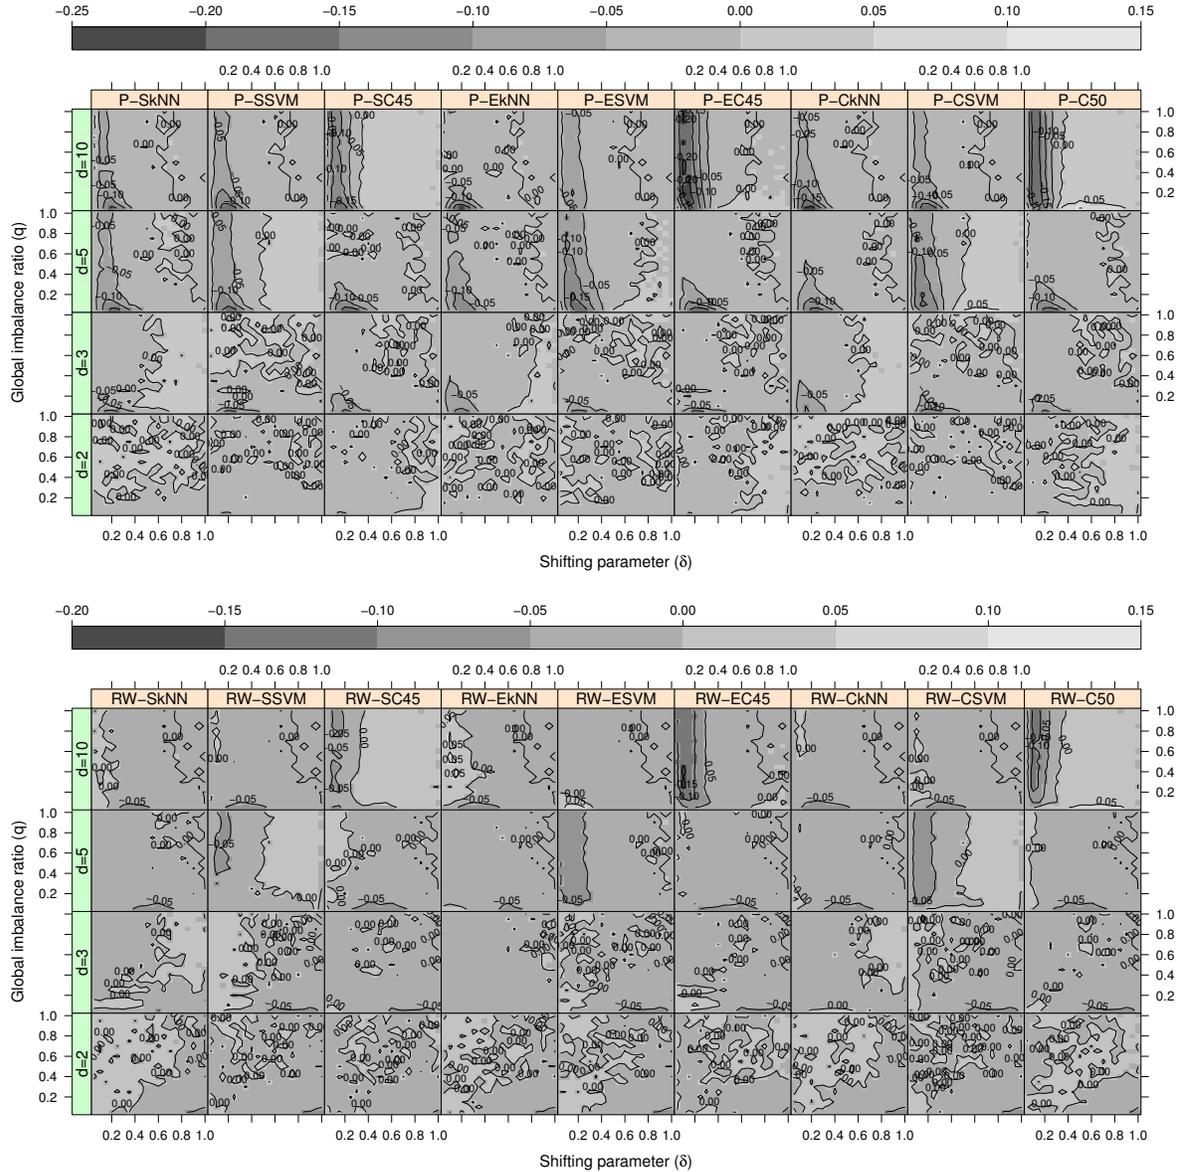


Figure 12: Differences between the AUCs of CCCD and other classifiers (see Figure 9 for details). P-CCCD in (top) and RW-CCCD in (bottom). Here, resampling scheme strong classifiers are coded with “S”, ensemble schemes with “E”, and cost sensitive schemes with “C”. For example, “SkNN” refers to the resampling schemed k -NN classifier. In this two-class setting, classes are drawn from $F_X = U(0, 1)^d$ and $F_Y = U(\delta, 1 + \delta)^d$ where $d = 2, 3, 5, 10$, $q, \delta = 0.05, 0.1, \dots, 0.95, 1.00$ with $n = 400$ and $m = qn$.

	Training		Testing	
	Time	Space	Time	Space
P-CCCD	$\mathcal{O}(N^2d)$	$\mathcal{O}(N^2)$	$\mathcal{O}(Nd)$	$\mathcal{O}(Nd)$
RW-CCCD	$\mathcal{O}(N^3 + N^2d)$	$\mathcal{O}(N^2)$	$\mathcal{O}(Nd)$	$\mathcal{O}(Nd)$
k -NN	\dots	\dots	$\mathcal{O}(Nd)$	$\mathcal{O}(Nd)$
SVM	$\mathcal{O}(N^3 + N^2d)$	$\mathcal{O}(N^2)$	$\mathcal{O}(Nd)$	$\mathcal{O}(Nd)$
C45-LNP	$\mathcal{O}(Nd^2)$	$\mathcal{O}(Nd)$	$\mathcal{O}(d)$	$\mathcal{O}(2^d)$

Table 3: Training and testing space and time complexities of the weak classifiers.

5.3 Complexity Analysis of the Classifiers

In Table 3, we compare training and testing time and space complexities of P-CCCDs, RW-CCCDs, k -NN, SVM and C4.5 classifiers. Let $N = n + m$ be the size of training data set. C4.5 is the fastest among all classifiers and requires the least space. However, unpruned C4.5 constitute a tree with its space complexity increasing exponentially on d since the data set is divided into at most two for all dimensions. The remaining classifiers are all instance based learning methods which depend on a matrix of distances between the points of training data set. Hence, their space complexity is at least $\mathcal{O}(N^2)$ and they run in at least $\mathcal{O}(N^2d)$ time. Both SVM and RW-CCCD classifiers run in $\mathcal{O}(N^3)$ time for $d < N$, and P-CCCD runs in $\mathcal{O}(N^2d)$ time. Minimum dominating set problem of P-CCCDs are polynomial time reducible to minimum set cover problems, and hence they run in $\mathcal{O}(N^2)$ time in the worst case but they require the computation of the distance matrix which takes the most time. However, in RW-CCCDs, covering balls are re-defined each time a new point is added to the prototype set. As a result, this operation requires an additional sweep on the training set on each iteration which makes RW-CCCD run in $\mathcal{O}(N^3)$ time, for $d < n$. For SVM, the training time of usual optimization algorithms is $\mathcal{O}(N^3)$ for $d < n$. However, it is possible to reduce the complexity to $\mathcal{O}(N^{2.3})$ with sequential minimal optimization (SMO) method (Chang and Lin, 2011).

Note that, k -NN does not require any training time or space, and should use the entire training data set to classify the test data set. However, CCCD and SVM classifiers reduce the training data set by means of prototype sets (minimum dominating sets in CCCD and support vectors in SVMs) even though their worst case testing space complexity is $\mathcal{O}(Nd)$. The entire training data set could be chosen as the prototype set for some cases, but we show that the data set is substantially reduced when the classes are imbalanced. In Figure 13, we compare the sizes of the set of prototypes in RW-CCCDs and the set of support vectors in SVM and CSVM classifiers. We consider the simulation settings with two classes for $F_X = U(0, 1)^d$ and $F_Y = U(\delta, 1 + \delta)^d$, $\delta, q = 0.1, 0.4, 0.7, 1.0$, $d = 2, 3, 5, 10$, $n = 400$ and $m = qn$.

In Figure 13, the number of both support vectors and prototypes decrease with increasing δ . The prototype set heavily depends on the overlapping ratio between class supports. Undoubtedly, when points of either class are further away from each other, covering balls get bigger for CCCDs, and the separating hyperplane requires less support vectors. On the other hand, observe that the number of support vectors are much higher than the number of prototypes of RW-CCCDs. The number of support vectors decreases with decreasing q .

The more imbalanced the data set, the fewer support vectors are generated. But in any case, RW-CCCDs still reduce the training data set more than SVMs. In Figure 14, we compare the number of prototypes in both CCCD classifier and the size of the C4.5 and C5.0 classifier trees for the same simulation setting.

The number of prototypes in P-CCCDs are, in general, much higher than that of other classifiers. Also, notice that the less imbalanced the classes are, the less the data reduction in P-CCCDs. However, there is not much change in the number in RW-CCCDs, C4.5 and C5.0, and since the size of trees grows exponentially on d , the size of trees get bigger than the size of CCCDs for some substantially high δ and d . Moreover, the size of trees in C5.0 is considerably less than that in C4.5 (Kuhn and Johnson, 2013). Although the number of prototypes are much higher than the size of trees in highly overlapped and imbalanced cases, RW-CCCDs reduce the training set substantially more than C4.5 and C5.0 in moderately imbalanced and moderately overlapped higher dimensional settings.

5.4 Real Data Examples

In this section, we compare the performance of CCCD classifiers and all other weak and strong classifiers on several data sets from UC Irvine (UCI) Machine Learning and KEEL repositories (Bache and Lichman, 2013; Alcalá-Fdez et al., 2011). To test the difference between the AUC of classifiers, we employ the 5x2 cross validation (CV) paired t -test (see Dietterich, 1998) and the combined 5x2 CV F -test (see Alpaydm, 1999). The 5x2 CV test has been devised by Dietterich (1998) and found to be the most powerful test among those with acceptable type-I error. However, the test statistics of 5x2 t -tests depend on which one of the ten folds is used. Hence, Alpaydm (1999) offered a combined 5x2 CV F -test which works as an omnibus test for all ten possible 5x2 t -tests (for each five repetitions there are two folds, hence ten folds in total). Basically, if a majority of ten 5x2 t -tests suggests that two classifiers are significantly different in terms of performance, the F -test also suggests a significant difference. Hence, an F -test with high p -value suggests that some of the ten t -tests fail to have low p -values.

We also provide the overlapping ratios and imbalance levels of these data sets. In a simulation study such as the one in Section 5.1, we have control on the overlapping region of two classes since we can choose the supports of the classes, hence their overlapping region is exactly known. However, in real data sets where the support of classes are neither defined nor available, we need methods to estimate the supports and hence estimate the overlapping ratios for the two classes. We employ the support vector data description (SVDD) method of Tax and Duin (2004) for this purpose. The method finds a description (or a region) of a data set, which covers a desired percentage of the points. SVDDs are also used in novelty or outlier detection. It has been inspired by the SVM classifiers and is based on defining a sphere around the data set. Similar to SVM, kernel functions can be employed to define more relaxed regions. SVDD is also a one-class learning method where the goal is to decide if a new point belongs to this particular class or not (Juszczak et al., 2002). By using SVDD approach, Xiong et al. (2010) found the SVDD regions of each class and its overlapping region. We also use SVDD to find the overlapping region E of each pair and report on the imbalance ratio with respect to E . The overlapping ratio is the percentage of points from both classes that reside in A . We use the Ddtools toolbox (Tax, 2014) of

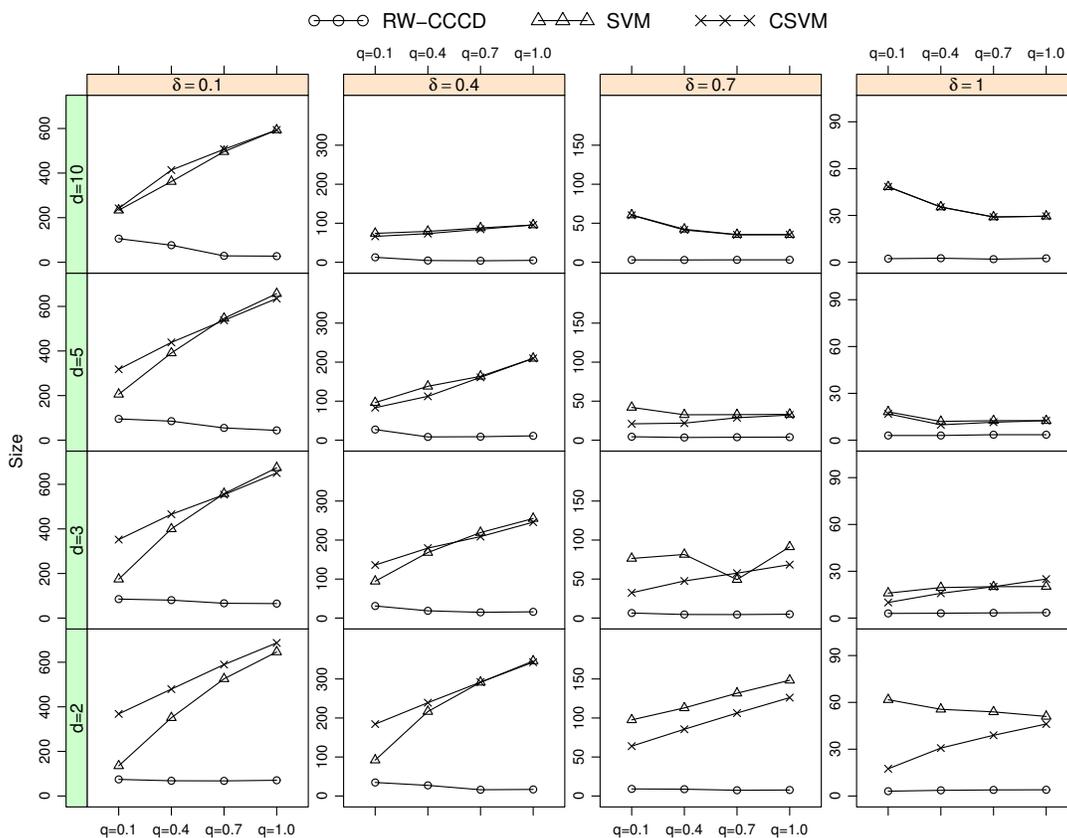


Figure 13: Comparison of the sizes of reduced data sets in RW-CCCDs, SVM and CSVM classifiers. Here “size” refers to the number of covering balls in RW-CCCD or the number of support vectors in SVM classifiers. In this two-class setting, classes are drawn from $F_X = U(0, 1)^d$ and $F_Y = U(\delta, 1 + \delta)^d$ where $\delta, q = 0.1, 0.4, 0.7, 1.0$ with $n = 400$ and $m = qn$.

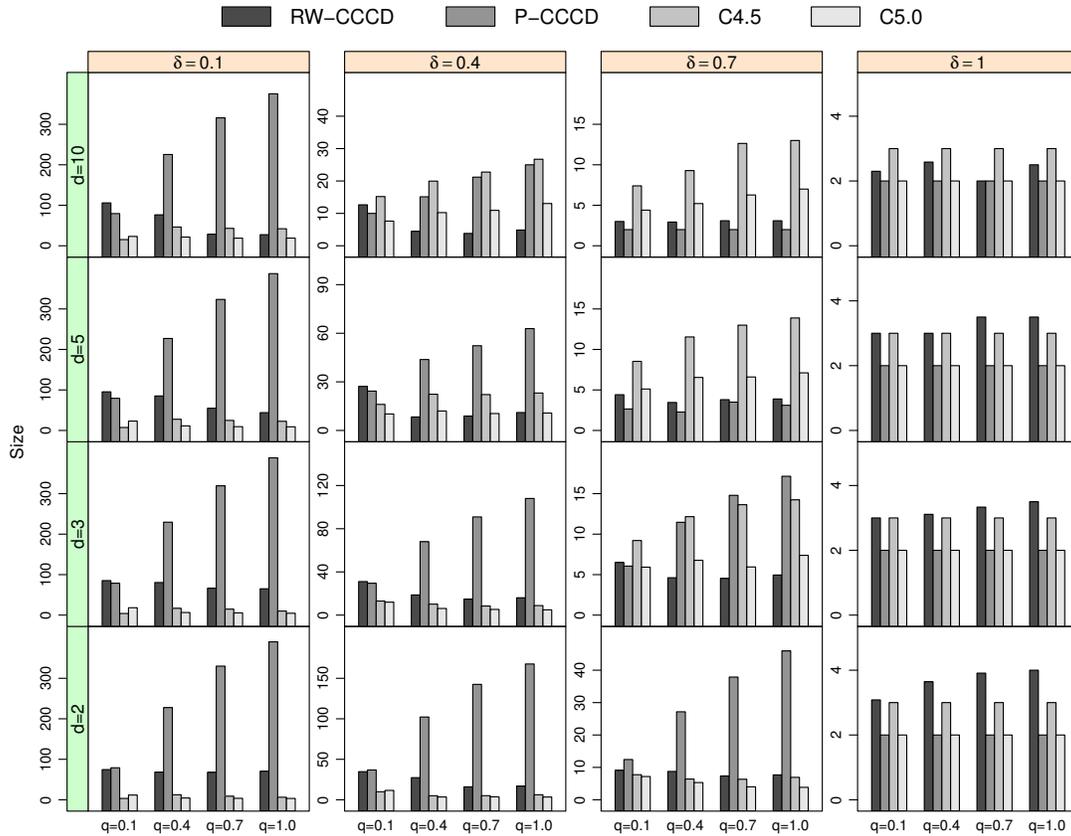


Figure 14: Comparison of the sizes of reduced data sets in CCCDs and C4.5. Here “size” refers to the number of covering balls in CCCDs or the number of nodes in the decision tree of C4.5 classifiers. The simulation setting is same as in Figure 13.

Data	$q = m/n$	N	d		$\sigma = 2$	$\sigma = 3$	$\sigma = 4$	$\sigma = 5$	$\sigma = 6$	$\sigma = 7$	$\sigma = 8$	$\sigma = 9$	$\sigma = 10$
Sonar	1.14	208	61	OR	4%	19%	23%	25%	26%	26%	27%	28%	28%
				IR	1.22	1.04	0.96	0.93	0.96	0.97	0.97	0.90	0.90
Ionosphere	1.78	351	35	OR	25%	36%	66%	69%	66%	79%	61%	76%	81%
				IR	90.00	62.50	8.70	6.20	5.44	3.67	5.02	3.98	3.31
Segment0	6.02	2308	20	OR	0%	0%	0%	0%	0%	0%	0%	0%	0%
				IR	NA								
Page-Blocks0	8.79	5472	11	OR	0.6%	0.6%	0.6%	0.8%	0.9%	1%	1%	1%	1%
				IR	0.47	0.22	0.22	0.25	0.40	0.39	0.43	0.54	0.63
Vowel0	9.98	988	14	OR	0%	0%	0%	0%	0%	0%	0%	0%	0%
				IR	NA								
Shuttle0vs4	13.87	1829	10	OR	0%	0%	0%	0%	0%	0%	0%	0%	0%
				IR	NA								
Yeast4	28.10	1484	9	OR	45%	27%	39%	37%	26%	26%	26%	26%	31%
				IR	18.97	99.75	18.50	18.24	392.00	390.00	391.00	390.00	41.18
Yeast1289vs7	30.70	947	9	OR	45%	44%	69%	43%	30%	29%	29%	29%	29%
				IR	24.47	23.76	24.34	23.00	47.33	45.83	45.83	45.66	45.50
Yeast5	32.70	1484	9	OR	6%	3%	0%	0%	0%	0%	6%	7%	0.1%
				IR	NA	1.15	NA						
Yeast6	41.40	1484	9	OR	30%	46%	42%	31%	30%	30%	10%	13%	13%
				IR	64.14	21.03	27.59	76.33	73.66	73.66	38.25	63.00	63.00
Abalone19	129.40	4174	9	OR	25%	20%	15%	14%	13%	12%	12%	11%	11%
				IR	104.30	104.30	163.75	197.33	278.00	262.50	253.00	244.00	234.00

Table 4: Overlapping ratios and (local) imbalance ratios in the overlapping region of data sets. “IR” stands for the imbalance ratio in the overlapping region and “OR” stands for the overlapping ratio which is the percentage of points from both classes residing in the overlapping region. IR=“NA” indicates that one of the classes has no members in the intersections of SVDD regions of classes.

MATLAB environment to produce the SVDDs of classes. Our choice of the kernel is the same as we have used with SVM classifiers in this study, the radial basis (i.e., Gaussian) kernel; for consistency. However, the selection of σ in the kernel is crucial for the SVDD region.

In Table 4, we present the overlapping ratios and the imbalance in the overlapping areas of all data sets for $\sigma = 2, 3, \dots, 10$. Although the value of σ produces different overlapping ratios, it is apparent that classes of data sets Ionosphere, Abalone19, Yeast4, Yeast6 and Yeast1289vs7 have more overlap than others, and these overlapping data sets have substantial local class imbalance in their respective overlapping regions. Other data sets have almost no overlapping nor imbalance in the overlapping regions even though their classes are globally imbalanced. One of these data sets is Yeast5 which has a imbalance ratio of $q = 32.70$ but has no imbalance in the small overlapping region.

In Table 5, we give the average AUC measures and their standard deviations of all CCCD-based and other classifiers according to the 5x2 CV scheme for the data sets. All other classifiers, weak or strong, have been two-way tested with 5x2 CV F -test against both RW-CCCD and P-CCCD classifiers. Their p -values are also provided in Table 5. For each of five repetitions, we divide the data into two folds. The AUC of fold 1 is given by using fold 1 as a training set and fold 2 as the test set. For fold 2, the process is similar. We repeated these experiments five times for all three classifiers. Looking at results from 11 data sets, RW-CCCD usually performs better than P-CCCD classifiers, and in addition,

ensemble based classifiers perform the best in general. Moreover, ensemble classifiers seem to perform better than RW-CCCDs but this difference is usually not significant, meaning RW-CCCDs perform comparable to ensemble classifiers in more than few folds of all ten folds. For example, compared to ensemble methods, RW-CCCD has nearly 0.07 less AUC in Yeast5, 0.02 less AUC in Yeast6, and 0.1 less AUC in Yeast1289vs7 data set. The difference is significant, however, with the data set Abalone19 with a level of < 0.03 . Although RW-CCCD achieves an average AUC value 0.6, ensemble classifiers achieve over 0.7. On the other hand, there is no significant difference between AUCs of RW-CCCD and ensembles in other highly overlapped and locally imbalanced data sets. On these data sets, RW-CCCD have significantly more AUC than weak classifiers and have AUC comparable to strong classifiers. Thus, these results from real data sets seem to resonate with the results from our simulations and further support the robustness of CCCD classifiers to the class imbalance problem.

6. Summary and Discussion

We assess the classification performance of various classifiers such as RW-CCCD, pure-CCCD, k -NN, SVM and C4.5 classifiers and their variants when class imbalance occurs, and we illustrate the robustness of CCCD classifiers to the class imbalance in data sets. This imbalance often occurs in real life data sets where, in two-class settings, minority class (the class with fewer number of observations) is usually dwarfed by the majority class. Class imbalances hinder the performance of many classification algorithms. We studied the performance of CCCD classifiers under class imbalance problem by first simulating a two-class setting similar to the one used in DeVinney (2003). In this setting, the support of one class is entirely embedded in the support of the other. Drawing equal number of observations from both class supports results in an imbalance between two classes with respect to their overlapping region, called *local (or restricted) class imbalance*. This difference in the class sizes was also the case in the example of DeVinney (2003), and it is the reason that CCCD classifiers show better results than the k -NN classifier. We show that P-CCCD classifiers with lower τ values tend to perform better than the ones with higher τ values. This is merely a result of balls with $\tau = 0$ representing the local density of the target class points better. Similarly, the RW-CCCD classifiers with lower e values are better when the dimensionality is low and the class sizes are high. This might indicate that the denser the data set in \mathbb{R}^d , the less useful the scores T_x . However, fully utilizing the scores usually increases the classification performance.

Analysis of both simulated and real data sets indicate that both CCCD classifiers show robustness to the class imbalance problem. We demonstrated this by studying the effects of the class overlapping problem together with the class imbalance problem. In fact, there are studies in the literature focusing on the performance of classification methods when class overlapping and class imbalance problems occur simultaneously (Prati et al., 2004; Denil and Trappenberg, 2010). Overlapping of classes is an important factor in the classification of imbalanced data sets; that is, it drastically affects the classification performance of most algorithms. When classes are both imbalanced and overlapping, performance of k -NN, SVM and C4.5 classifiers deteriorate whereas CCCD classifiers are not affected as severely as these methods. We use two alternatives of C4.5 classifiers where we prune the decision

		Ionosphere	Sonar	Yeast6	Yeast5	Yeast4	Yeast1280vs7	Vowel0	Shuttle0vs4	Abalone19	Segment0	Page-Blocks0
RW-CCCD	AUC	0.917±0.023	0.722±0.050	0.866±0.051	0.898±0.063	0.807±0.048	0.643±0.057	0.877±0.046	0.996±0.003	0.603±0.065	0.895±0.011	0.875±0.019
	AUC	0.934±0.032	0.805±0.045	0.755±0.053	0.793±0.094	0.605±0.051	0.555±0.088	0.958±0.025	0.988±0.016	0.506±0.019	0.957±0.010	0.869±0.009
P-CCCD	AUC	0.803±0.019	0.804±0.027	0.786±0.082	0.839±0.063	0.619±0.088	0.562±0.04	0.971±0.031	0.996±0.004	0.512±0.015	0.988±0.004	0.863±0.009
	<i>p</i> -value (vs RW)	0.000	0.107	0.072	0.573	0.009	0.087	0.037	0.693	0.104	0.000	0.604
	<i>p</i> -value (vs P)	0.005	0.473	0.552	0.209	0.650	0.177	0.504	0.534	0.735	0.026	0.600
SVM	AUC	0.949±0.010	0.719±0.057	0.710±0.045	0.741±0.069	0.527±0.024	0.507±0.014	0.956±0.039	0.984±0.012	0.500±0.000	0.564±0.006	0.901±0.013
	<i>p</i> -value (vs RW)	0.088	0.763	0.043	0.096	0.002	0.047	0.095	0.459	0.101	0.000	0.249
	<i>p</i> -value (vs P)	0.582	0.091	0.086	0.467	0.129	0.084	0.700	0.632	0.535	0.000	0.016
C4.5	AUC	0.851±0.024	0.712±0.049	0.742±0.048	0.803±0.075	0.613±0.057	0.570±0.043	0.948±0.032	0.999±0.001	0.503±0.009	0.982±0.004	0.917±0.012
	<i>p</i> -value (vs RW)	0.008	0.747	0.015	0.429	0.084	0.270	0.090	0.242	0.088	0.000	0.229
	<i>p</i> -value (vs P)	0.039	0.358	0.634	0.734	0.819	0.508	0.137	0.533	0.535	0.100	0.020
Sk-NN	AUC	0.836±0.022	0.735±0.073	0.877±0.042	0.969±0.015	0.807±0.05	0.667±0.04	0.942±0.04	0.998±0.001	0.561±0.024	0.98±0.005	0.907±0.009
	<i>p</i> -value (vs RW)	0.008	0.172	0.350	0.402	0.673	0.735	0.330	0.458	0.424	0.000	0.056
	<i>p</i> -value (vs P)	0.039	0.374	0.165	0.063	0.036	0.006	0.602	0.571	0.167	0.121	0.015
SSVM	AUC	0.949±0.014	0.632±0.054	0.865±0.042	0.946±0.025	0.800±0.031	0.687±0.047	0.984±0.023	0.998±0.001	0.525±0.017	0.987±0.004	0.905±0.011
	<i>p</i> -value (vs RW)	0.106	0.008	0.664	0.477	0.667	0.312	0.012	0.324	0.164	0.000	0.176
	<i>p</i> -value (vs P)	0.592	0.031	0.273	0.120	0.017	0.055	0.160	0.567	0.391	0.031	0.002
SC4.5	AUC	0.867±0.016	0.680±0.050	0.760±0.075	0.872±0.070	0.688±0.073	0.605±0.060	0.942±0.025	0.997±0.002	0.500±0.009	0.982±0.005	0.922±0.019
	<i>p</i> -value (vs RW)	0.066	0.450	0.074	0.572	0.216	0.229	0.203	0.590	0.086	0.000	0.027
	<i>p</i> -value (vs P)	0.141	0.165	0.559	0.116	0.472	0.334	0.616	0.642	0.509	0.088	0.005
Ek-NN	AUC	0.856±0.025	0.813±0.028	0.889±0.043	0.964±0.004	0.857±0.024	0.755±0.040	0.900±0.008	0.996±0.004	0.731±0.041	0.955±0.006	0.913±0.006
	<i>p</i> -value (vs RW)	0.019	0.090	0.494	0.379	0.267	0.101	0.092	0.693	0.007	0.000	0.060
	<i>p</i> -value (vs P)	0.069	0.320	0.108	0.077	0.005	0.000	0.638	0.534	0.006	0.787	0.001
ESVM	AUC	0.948±0.008	0.783±0.041	0.896±0.037	0.970±0.004	0.862±0.023	0.749±0.047	0.973±0.019	0.998±0.002	0.744±0.026	0.735±0.028	0.953±0.004
	<i>p</i> -value (vs RW)	0.138	0.027	0.376	0.336	0.175	0.209	0.023	0.347	0.028	0.001	0.011
	<i>p</i> -value (vs P)	0.669	0.454	0.060	0.074	0.004	0.004	0.312	0.546	0.000	0.000	0.001
EC4.5	AUC	0.909±0.021	0.800±0.031	0.858±0.043	0.937±0.008	0.835±0.041	0.662±0.066	0.963±0.012	1.000±0.000	0.694±0.054	0.990±0.006	0.959±0.005
	<i>p</i> -value (vs RW)	0.397	0.068	0.430	0.397	0.494	0.794	0.101	0.112	0.004	0.000	0.008
	<i>p</i> -value (vs P)	0.415	0.691	0.153	0.084	0.018	0.083	0.553	0.529	0.014	0.017	0.001
Ck-NN	AUC	0.834±0.024	0.809±0.023	0.874±0.039	0.972±0.008	0.819±0.022	0.692±0.036	0.986±0.023	0.998±0.001	0.583±0.039	0.987±0.002	0.889±0.010
	<i>p</i> -value (vs RW)	0.008	0.110	0.660	0.330	0.716	0.459	0.007	0.366	0.549	0.000	0.552
	<i>p</i> -value (vs P)	0.042	0.382	0.138	0.067	0.006	0.035	0.030	0.573	0.013	0.032	0.008
CSVM	AUC	0.951±0.010	0.735±0.062	0.835±0.042	0.948±0.025	0.760±0.045	0.690±0.062	0.983±0.027	0.984±0.012	0.685±0.057	0.590±0.011	0.953±0.006
	<i>p</i> -value (vs RW)	0.068	0.725	0.569	0.370	0.207	0.638	0.009	0.459	0.257	0.000	0.016
	<i>p</i> -value (vs P)	0.532	0.164	0.408	0.164	0.054	0.140	0.109	0.632	0.043	0.000	0.002
C5.0	AUC	0.865±0.038	0.709±0.052	0.802±0.020	0.905±0.085	0.692±0.066	0.610±0.068	0.939±0.022	1.000±0.000	0.521±0.037	0.987±0.006	0.935±0.015
	<i>p</i> -value (vs RW)	0.020	0.683	0.395	0.694	0.219	0.614	0.265	0.112	0.132	0.000	0.027
	<i>p</i> -value (vs P)	0.005	0.347	0.514	0.034	0.439	0.109	0.436	0.529	0.342	0.144	0.002

Table 5: Average of AUC values of ten folds, and standard deviations, of CCCD, weak and strong classifiers for data sets. The *p*-values of 5x2 CV *F*-tests show the results of two-way tests comparing both CCCDs with other classifiers. Some of best performers are given in bold.

tree in one and do not in the other. It is known for some time that pruning deteriorates the performance of tree classifiers under class imbalance. Moreover, SVM is robust to moderately imbalanced class sizes but demonstrates no robustness in highly imbalanced cases. However, whether the data set is highly or moderately imbalanced, CCCD classifiers seem to preserve their AUC compared to k -NN, SVM and C4.5 classifiers. Hence, our study suggests that CCCD classifiers are appealing alternatives when data have class imbalance. In addition, we mention the effect of the individual class sizes on the class imbalance problem (Japkowicz and Stephen, 2002). Whatever the ratio between class sizes is, if the minority class has a substantially high number of points, the effect of imbalances between classes tend to diminish.

The classifiers k -NN, SVM and C4.5 are referred to as weak classifiers since, by construction, they are sensitive to imbalances between classes in data sets. In addition, we consider three distinct families of methods to establish strong classifiers based on weak classifiers, and compare them with CCCD classifiers. We conduct simulation studies to determine how the classification performance jointly depends on both (global) class imbalance and class overlapping, parameterized as q and δ , respectively. Finally, we apply all these classifiers on several UCI and KEEL data sets. By using the SVDD method of Tax and Duin (2004), we estimated the overlapping ratios of all these data sets. We show that CCCD classifiers outperform or perform comparable to k -NN, SVM and C4.5 classifiers for some overlapping and imbalance ratios in both simulated and real data sets. In particular, CCCDs are better than SVM classifiers in highly imbalanced cases. The effect of high class imbalance on SVM classifier is also studied in Akbani et al. (2004) and Raskutti and Kowalczyk (2004). However, when no imbalance occurs between classes, CCCD classifiers usually show either comparable or slightly worse performance than the other classifiers. As for strong classifiers, we employ the most successful methods from three families of schemes where EasyEnsemble and SMOTE+ENN methods are among them. In our simulation studies, we demonstrated that CCCD classifiers, especially RW-CCCDs, work well compared to these strong classifiers when there are considerable overlap and the high (local) imbalance between classes. However, these methods are slightly better than RW-CCCDs as these strong classifiers are the best performing ones among their respective families (Batista et al., 2004; López et al., 2013). Nevertheless, RW-CCCDs have still high performance compared to these classifiers with additional increase in testing speed.

We also investigate the performance of CCCD classifiers under different conditions. Specifically, in two different experiments, we simulate two classes where (i) classes are imbalanced but supports are not overlapping (well separated) and (ii) classes are balanced and supports are overlapping with increasing dimensionality. P-CCCD classifiers are better than RW-CCCD classifiers in experiment (i). Both CCCD classifiers mostly outperform k -NN and SVM classifiers when classes are imbalanced and not overlapping, however RW-CCCD classifiers outperform these classifiers only when dimensionality is sufficiently high. In experiment (ii), the classification performance of CCCD classifiers slightly degrade compared to k -NN and SVM classifiers, especially with increasing d . Among CCCD classifiers, random walk covers appear to be better when classes are both overlapping and imbalanced, however our results suggest the use of pure covers when classes are imbalanced and well separated (i.e., not overlapping). In fact, class supports are often overlapping in real life data sets, hence RW-CCCD classifiers seem to be more appealing in practice.

In practice, classifiers based on CCCD classifiers resemble prototype selection methods. CCCDs balance the class sizes by defining balls that catch surrounding points of the same class, and discard these points from the training set. The resulting data set is composed of the centers of these balls and associated radii which are used in scaled dissimilarity measures. Although, CCCD classifiers remove substantial amount of observations from the majority class, they preserve (most of) the information with the radii. The bigger the radius, the more likely that the balls of CCCD classifiers contain more points. The radii could be considered as an indicator of the local density of the target class. The real advantage of CCCD classifiers are these prototype sets which are of (approximately) minimum cardinality, although training time and space of P-CCCDs and RW-CCCDs may be considerably high. However, the number of points in the prototype set is substantially low, and hence testing speed is increased. In some cases, RW-CCCDs provide classifiers with the least testing space complexity. Only the decision tree based classifiers, C4.5 and C5.0, achieve comparable or slightly more reduction to RW-CCCDs. However, with increasing dimensionality, sizes of these trees grow exponentially, making them less appealing than RW-CCCDs in the sense of classification space complexity. Hence, CCCDs preserve important information regarding the data sets while substantially increasing the testing speed. In literature, many classifiers have been devised to preserve the information on the deleted majority class points, however they are all ensemble based classifiers which substantially increase both training and testing time complexities. In that regard, CCCDs offer a novel approach to this particular problem.

Eveland et al. (2005) modified RW-CCCD classifiers as to increase the speed of the face detection in which imbalances between classes occur naturally. They did only refer to the real life applications which consist of class imbalances. They did not, however, investigate the relationship between class imbalance and overlapping problems as thoroughly as our study does. On the other hand, establishing class covers with Euclidean balls raise the possibility of using different regions (the regions are Euclidean hyperballs around target class points in CCCD) to balance the data and, thus, construct non-parametric classifiers with more classification performance. Along this line, CCCDs can be generalized using *proximity maps* (Jaromczyk and Toussaint, 1992). For example, Ceyhan (2005) defined *proximity catch digraphs* (PCDs) that are generalized versions of CCCDs. Ceyhan (2005) has introduced three families of PCDs and used them to test spatial data patterns of segregation and association (see Ceyhan and Priebe, 2005; Ceyhan et al., 2006, 2007). PCDs can also be used to derive new graph-based classifiers which are potentially robust to the class imbalance problem. The study of their properties and performance is a topic of ongoing research by the authors.

Acknowledgments

Most of the Monte Carlo simulations presented in this article were executed at Koç University High Performance Computing Laboratory. This research was supported by the European Commission under the Marie Curie International Outgoing Fellowship Programme via Project # 329370 titled PRinHDD.

References

- R. Akbani, S. Kwek, and N. Japkowicz. Applying support vector machines to imbalanced datasets. In *Proceedings of 15th European Conference on Machine Learning*, pages 39–50, Pisa, Italy, 2004.
- J. Alcalá-Fdez, A. Fernández, J. Luengo, J. Derrac, and S. Garc´a. KEEL data-mining software tool: Data set repository, integration of algorithms and experimental analysis framework. *Multiple-Valued Logic and Soft Computing*, 17(2-3):255–287, 2011.
- E. Alpaydm. Combined 5×2 cv F test for comparing supervised classification learning algorithms. *Neural Computation*, 11(8):1885–1892, 1999.
- K. Bache and M. Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- R. Barandela, J. S. Sánchez, V. Garcia, and E. Rangel. Strategies for learning in class imbalance problems. *Pattern Recognition*, 36(3):849–851, 2003.
- G. E. Batista, R. C. Prati, and M. C. Monard. A study of the behavior of several methods for balancing machine learning training data. *ACM SIGKDD Explorations Newsletter*, 6(1):20–29, 2004.
- G. E. Batista, R. C. Prati, and M. C. Monard. Balancing strategies and class overlapping. In *Proceedings of 6th International Symposium on Intelligent Data Analysis: Advances in Intelligent Data Analysis VI*, pages 24–35, Madrid, Spain, 2005.
- S. Bereg, S. Cabello, J. M. Díaz-Báñez, P. Pérez-Lantero, C. Seara, and I. Ventura. The class cover problem with boxes. *Computational Geometry*, 45(7):294–304, 2012.
- A. H. Cannon and L. J. Cowen. Approximation algorithms for the class cover problem. *Annals of Mathematics and Artificial Intelligence*, 40(3):215–223, 2004.
- E. Ceyhan. *An investigation of proximity catch digraphs in Delaunay tessellations*. PhD thesis, Johns Hopkins University, Baltimore, MD, USA, 2005.
- E. Ceyhan and C. E. Priebe. The use of domination number of a random proximity catch digraph for testing spatial patterns of segregation and association. *Statistics & Probability Letters*, 73(1):37–50, 2005.
- E. Ceyhan, C. E. Priebe, and J. C. Wierman. Relative density of the random r -factor proximity catch digraph for testing spatial patterns of segregation and association. *Computational Statistics & Data Analysis*, 50(8):1925–1964, 2006.
- E. Ceyhan, C. E. Priebe, and D. J. Marchette. A new family of random graphs for testing spatial segregation. *Canadian Journal of Statistics*, 35(1):27–50, 2007.
- C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.

- N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. SMOTE: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16(1):321–357, 2002.
- N. V. Chawla, N. Japkowicz, and A. Kotcz. Editorial: Special issue on learning from imbalanced data sets. *ACM SIGKDD Explorations Newsletter*, 6(1):1–6, 2004.
- V. Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
- D. A. Cieslak and N. V. Chawla. Learning decision trees for unbalanced data. In *Proceedings of the ECML PKDD 2008 Machine Learning and Knowledge Discovery in Databases: European Conference*, pages 241–256, Antwerp, Belgium, 2008.
- T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1):21–27, 1967.
- M. Denil and T. Trappenberg. Overlap versus imbalance. In *Proceedings of the 23rd Canadian Conference on Advances in Artificial Intelligence*, pages 220–231, Berlin, Heidelberg, 2010.
- J. DeVinney, C. Priebe, D. Marchette, and D. Socolinsky. Random walks and catch digraphs in classification. In *Proceedings of the 34th Symposium on the Interface, Volume 34: Computing Science and Statistics*, Montreal, Quebec, Canada, 2002.
- J. G. DeVinney. *The class cover problem and its application in pattern recognition*. PhD thesis, Johns Hopkins University, Baltimore, MD, USA, 2003.
- T. G. Dietterich. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10(7):1895–1923, 1998.
- P. Domingos. MetaCost: A general method for making classifiers cost-sensitive. In *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '99, pages 155–164, New York, NY, USA, 1999.
- C. Drummond, R. C. Holte, et al. C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In *Workshop on Learning from Imbalanced Datasets (II)*, Washington DC, USA, 2003.
- C. Elkan. The foundations of cost-sensitive learning. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 973–978, Melbourne, Australia, 2001.
- C. K. Eveland, D. A. Socolinsky, C. E. Priebe, and D. J. Marchette. A hierarchical methodology for class detection problems with skewed priors. *Journal of Classification*, 22(1):17–48, 2005.
- M. Fernández-Delgado, E. Cernadas, S. Barro, and D. Amorim. Do we need hundreds of classifiers to solve real world classification problems. *Journal of Machine Learning Research*, 15(1):3133–3181, 2014.

- E. Fix and J. L. Hodges. Discriminatory analysis. nonparametric discrimination: Consistency properties. *International Statistical Review*, 57(3):238–247, 1989.
- Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- M. Galar, A. Fernandez, E. Barrenechea, H. Bustince, and F. Herrera. A review on ensembles for the class imbalance problem: Bagging-, boosting-, and hybrid-based approaches. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 42(4):463–484, 2012.
- V. García, R. A. Mollineda, and J. S. Sánchez. On the k -NN performance in a challenging scenario of imbalance and overlapping. *Pattern Analysis and Applications*, 11(3-4):269–280, 2008.
- H. Han, W.-Y. Wang, and B.-H. Mao. Borderline-SMOTE: A new over-sampling method in imbalanced data sets learning. In *Proceedings of the 2005 International Conference on Advances in Intelligent Computing - Volume Part I*, pages 878–887, Berlin, Heidelberg, 2005.
- D. J. Hand and V. Vinciotti. Choosing k for two-class nearest neighbour classifiers with unbalanced classes. *Pattern Recognition Letters*, 24(9):1555–1562, 2003.
- H. He and E. A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284, 2009.
- J. Huang and C. X. Ling. Using AUC and accuracy in evaluating learning algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 17(3):299–310, 2005.
- N. Japkowicz and S. Stephen. The class imbalance problem: A systematic study. *Intelligent Data Analysis*, 6(5):429–449, 2002.
- J. W. Jaromczyk and G. T. Toussaint. Relative neighborhood graphs and their relatives. *Proceedings of the IEEE*, 80(9):1502–1517, 1992.
- T. Joachims. Making large-scale SVM learning practical. In B. Schölkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, pages 169–184. MIT Press, Cambridge, MA, 1999.
- P. Juszczak, D. M. Tax, and R. Duin. Feature scaling in support vector data description. In *Proceedings of 8th Annual Conference of the Advanced School for Computing and Imaging*, pages 95–102, Delft, Netherlands, 2002.
- S. Kotsiantis, D. Kanellopoulos, P. Pintelas, et al. Handling imbalanced datasets: A review. *GESTS International Transactions on Computer Science and Engineering*, 30(1):25–36, 2006.
- M. Kuhn and K. Johnson. *Applied Predictive Modeling*. Springer, New York, USA, 2013.

- C. X. Ling, Q. Yang, J. Wang, and S. Zhang. Decision trees with minimal costs. In *Proceedings of the 21th International Conference on Machine Learning*, page 69, Banff, Alberta, Canada, 2004.
- X.-Y. Liu, J. Wu, and Z.-H. Zhou. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 39(2):539–550, 2009.
- R. Longadge and S. Dongre. Class imbalance problem in data mining: Review. *International Journal of Computer Science and Network*, 2(1):83–87, 2013.
- V. López, A. Fernández, S. García, V. Palade, and F. Herrera. An insight into classification with imbalanced data: Empirical results and current trends on using data intrinsic characteristics. *Information Sciences*, 250:113–141, 2013.
- I. Mani and I. Zhang. kNN approach to unbalanced data distributions: A case study involving information extraction. In *Proceedings of ICML'2003 Workshop on Learning from Imbalanced Datasets II*, Washington, DC, USA, 2003.
- D. J. Marchette. Class cover catch digraphs. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(2):171–177, 2010.
- D. J. Marchette. *cccd: Class Cover Catch Digraphs*, 2013. URL <http://CRAN.R-project.org/package=cccd>. R package version 1.04.
- M. A. Mazurowski, P. A. Habas, J. M. Zurada, J. Y. Lo, J. A. Baker, and G. D. Tourassi. Training neural network classifiers for medical decision making: The effects of imbalanced datasets on classification performance. *Neural Networks*, 21(2):427–436, 2008.
- D. Meyer, E. Dimitriadou, K. Hornik, A. Weingessel, and F. Leisch. *e1071: Misc Functions of the Department of Statistics (e1071), TU Wien*, 2014. URL <http://CRAN.R-project.org/package=e1071>. R package version 1.6-4.
- C. Phua, D. Alahakoon, and V. Lee. Minority report in fraud detection: Classification of skewed data. *ACM SIGKDD Explorations Newsletter*, 6(1):50–59, 2004.
- R. C. Prati, G. E. Batista, and M. C. Monard. Class imbalances versus class overlapping: An analysis of a learning system behavior. In *Proceedings of 3rd Mexican International Conference on Artificial Intelligence*, pages 312–321, Mexico City, Mexico, 2004.
- C. E. Priebe, J. G. DeVinney, and D. J. Marchette. On the distribution of the domination number for random class cover catch digraphs. *Statistics & Probability Letters*, 55(3):239–246, 2001.
- C. E. Priebe, D. J. Marchette, J. G. DeVinney, and D. A. Socolinsky. Classification using class cover catch digraphs. *Journal of Classification*, 20(1):3–23, 2003a.
- C. E. Priebe, J. L. Solka, D. J. Marchette, and B. T. Clark. Class cover catch digraphs for latent class discovery in gene expression monitoring by DNA microarrays. *Computational Statistics & Data Analysis*, 43(4):621–632, 2003b.

- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2015. URL <http://www.R-project.org/>.
- B. Raskutti and A. Kowalczyk. Extreme re-balancing for SVMs: A case study. *ACM SIGKDD Explorations Newsletter*, 6(1):60–69, 2004.
- L. Rokach. Ensemble-based classifiers. *Artificial Intelligence Review*, 33(1):1–39, 2010.
- B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471, 2001.
- D. A. Socolinsky, J. D. Neuheisel, C. E. Priebe, J. G. DeVinney, and D. J. Marchette. Fast face detection with a boosted CCCD classifier. In *Proceedings of the 35th Symposium on the Interface, Volume 34: Computing Science and Statistics*, Salt Lake City, Utah, USA, 2003.
- I. Takigawa, M. Kudo, and A. Nakamura. Convex sets as prototypes for classifying patterns. *Engineering Applications of Artificial Intelligence*, 22(1):101–108, 2009.
- Y. Tang, Y.-Q. Zhang, N. V. Chawla, and S. Krasser. SVMs modeling for highly imbalanced classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 39(1):281–288, 2009.
- D. M. Tax. Ddtools, the data description toolbox for MATLAB, July 2014. version 2.1.1.
- D. M. Tax and R. P. Duin. Support vector data description. *Machine Learning*, 54(1):45–66, 2004.
- N. Thai-Nghe, A. Busche, and L. Schmidt-Thieme. Improving academic performance prediction by dealing with class imbalance. In *Proceedings of 19th International Conference on Intelligent Systems Design and Applications*, pages 878–883, Pisa, Italy, 2009.
- D. B. West. *Introduction to Graph Theory*. Prentice Hall, New Jersey, USA, 2 edition, 2000.
- D. L. Wilson. Asymptotic properties of nearest neighbor rules using edited data. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-2(3):408–421, July 1972.
- X. Wu, V. Kumar, J. R. Quinlan, J. Ghosh, Q. Yang, H. Motoda, G. J. McLachlan, A. Ng, B. Liu, S. Y. Philip, et al. Top 10 algorithms in data mining. *Knowledge and Information Systems*, 14(1):1–37, 2008.
- H. Xiong, J. Wu, and L. Liu. Classification with class overlapping: A systematic study. In *Proceedings of the 1st International Conference on e-Business Intelligence*, Shanghai, China, 2010.
- B. Zadrozny, J. Langford, and N. Abe. Cost-sensitive learning by cost-proportionate example weighting. In *Proceedings of 3rd IEEE International Conference on Data Mining*, pages 435–442, Melbourne, Florida, USA, 2003.