

Unknown malcode detection and the imbalance problem

Robert Moskovitch · Dima Stopel · Clint Feher ·
Nir Nissim · Nathalie Japkowicz · Yuval Elovici

Received: 11 September 2008 / Accepted: 6 May 2009 / Published online: 11 July 2009
© Springer-Verlag France 2009

Abstract The recent growth in network usage has motivated the creation of new malicious code for various purposes. Today's signature-based antiviruses are very accurate for known malicious code, but can not detect new malicious code. Recently, classification algorithms were used successfully for the detection of unknown malicious code. But, these studies involved a test collection with a limited size and the same malicious: benign file ratio in both the training and test sets, a situation which does not reflect real-life conditions. We present a methodology for the detection of unknown malicious code, which examines concepts from text categorization, based on n -grams extraction from the binary code and feature selection. We performed an extensive evaluation, consisting of a test collection of more than 30,000 files, in which we investigated the class imbalance problem. In real-life scenarios, the malicious file content is expected to be low, about 10% of the total files. For practical purposes, it is unclear as to what the corresponding percentage in the training set

should be. Our results indicate that greater than 95% accuracy can be achieved through the use of a training set that has a malicious file content of less than 33.3%.

1 Introduction

The term *malicious code* (malcode) commonly refers to pieces of code, not necessarily executable files, which are intended to harm, generally or in particular, the specific owner of the host. Malcodes are classified, mainly based on their transport mechanism, into four main categories: *worms*, *viruses*, *Trojans* and new group that is becoming more common, which is comprised of *remote access Trojans* and *backdoors*. The recent growth in high-speed internet connections and in internet network services has led to an increase in the creation of new malicious codes for various purposes, based on economic, political, criminal or terrorist motives (among others). Some of these codes have been used to gather information, such as passwords and credit card numbers, as well as behavior monitoring.

Current anti-virus technology is primarily based on two approaches: *signature-based* methods, which rely on the identification of unique strings in the binary code; while being very precise, it is useless against unknown malicious code [1]. Moreover, these can be overcome by a variant after checking it with a black box anti-virus check [2]. The second approach involves *heuristic-based* methods, which are based on rules defined by experts, which define a malicious behavior, or a benign behavior, in order to enable the detection of unknown malcodes [3]. Other proposed methods include *behavior blockers*, which attempt to detect sequences of events in operating systems, and *integrity checkers*, which periodically check for changes in files and disks. However, besides the fact that these methods can be bypassed

R. Moskovitch (✉) · D. Stopel · C. Feher · N. Nissim · Y. Elovici
Deutsche Telekom Laboratories,
Department of Information Systems Engineering,
Ben Gurion University, 84105 Be'er Sheva, Israel
e-mail: robertmo@bgu.ac.il

D. Stopel
e-mail: stopel@bgu.ac.il

C. Feher
e-mail: clint@bgu.ac.il

N. Nissim
e-mail: nirni@bgu.ac.il

Y. Elovici
e-mail: elovici@bgu.ac.il

N. Japkowicz
School of Information Technology and Engineering,
University of Ottawa, Ottawa, ON K1N 6N5, Canada
e-mail: nat@site.uottawa.ca

by viruses, their main drawback is that, by definition, they can only detect the presence of a malcode after it has been executed.

Therefore, generalizing the detection methods to be able to detect unknown malcodes is crucial. Recently, classification algorithms were employed to automate and extend the idea of heuristic-based methods. As we will describe in more detail shortly, the binary code of a file is represented by n -grams and *classifiers* are applied to learn patterns in the code and classify large amounts of data. A classifier is a rule set which is learnt from a given *training-set*, including examples of classes, both malicious and benign files in our case. Recent studies, which we survey in the next section [4–8], have shown that this is a very successful strategy. However, these studies present evaluations based on test collections, having similar proportion of malicious versus benign files in the test collections (50% of malicious files). This proportion has two potential drawbacks. These proportions do not reflect real life situation, in which malicious code is commonly significantly less than 50% and additionally these studies, as will be shown later, might report optimistic results. Recent survey¹ made by McAfee indicates that about 4% of search results from the major search engines on the web contain malicious code. Additionally, it was found that above 15% of the files in the KaZaA network contained malicious code.² Thus, we assume that the percentage of malicious files in real life is about or less than 10%, but we also consider other possible percentages.

In this study, we present a methodology for *malcode categorization* based on concepts from *text categorization*. We present an extensive and rigorous evaluation of many factors in the methodology, based on eight types of classifiers. The evaluation is based on a test collection 10 times larger than any previously reported collection, containing more than 30,000 files. We introduce the class imbalance problem, which refers to domains in which the proportions of each class instances is not equal, in the context of our task, in which we evaluate the classifiers for five levels of malcode content (percentages) in the training-set and 17 (percentages) levels of malcode content in the test-set. We start with a survey of previous relevant studies. We describe the methods we used, including: concepts from *text categorization*, data preparation, and classifiers. We present our results and finally discuss them.

2 Background

2.1 Detecting unknown malcode via machine learning

Over the past five years, several studies have investigated the direction of detecting unknown malcode based on its binary code. Schultz et al. [4] were the first to introduce the idea of applying machine learning (ML) methods for the detection of different malcodes based on their respective binary codes. They used three different feature extraction (FE) approaches: *program header*, *string features* and *byte sequence features*, in which they applied four classifiers: a *signature-based method* (anti-virus), *Ripper*—a rule-based learner, *Naïve Bayes* and *Multi-Naïve Bayes*. This study found that all of the ML methods were more accurate than the signature-based algorithm. The ML methods were more than twice as accurate when the out-performing method was Naïve Bayes, using strings, or Multi-Naïve Bayes using byte sequences. Abou-Assaleh et al. [5] introduced a framework that used the common n -gram (CNG) method and the k nearest neighbor (KNN) classifier for the detection of malcodes. For each class, malicious and benign, a representative profile was constructed and assigned a new executable file. This executable file was compared with the profiles and matched to the most similar. Two different data sets were used: the *I-worm collection*, which consisted of 292 Windows internet worms and the *win32 collection*, which consisted of 493 Windows viruses. The best results were achieved by using 3–6 n -grams and a profile of 500–5,000 features. Kolter and Maloof [6] presented a collection that included 1971 benign and 1651 malicious executables files. N -grams were extracted and 500 features were selected using the *information gain* measure [8]. The vector of n -gram features was binary, presenting the presence or absence of a feature in the file and ignoring the frequency of feature appearances (in the file). In their experiment, they trained several classifiers: IBK (KNN), a similarity based classifier called TFIDF classifier, Naïve Bayes, Support Vector Machines (SVM) (SMO) and Decision tree (J48). The last three of these were also boosted. Two main experiments were conducted on two different data sets, a small collection and a large collection. The small collection included 476 malicious and 561 benign executables and the larger collection included 1651 malicious and 1971 benign executables. In both experiments, the four best-performing classifiers were Boosted J48, SVM, boosted SVM and IBK. Boosted J48 out-performed the others. The authors indicated that the results of their n -gram study were better than those presented by Schultz et al. [4]. Recently, [6] reported an extension of their work, in which they classified malcodes into families (classes) based on the functions in their respective payloads. In the categorization task of multiple classifications, the best results were achieved for the classes' *mass mailer*, *backdoor* and *virus* (no benign classes). In attempts

¹ McAfee Study Finds 4% of Search Results Malicious, By Frederick Lane, 4th June 2007 [http://www.newsfactor.com/story.xhtml?story_id=010000CEUEQO].

² S. Shin, J. Jung, H. Balakrishnan, Malware Prevalence in the KaZaA File-Sharing Network, Internet Measurement Conference (IMC), Brazil, October 2006.

to estimate the ability to detect malicious codes based on their issue dates, these techniques were trained on files issued before July 2003, and then tested on 291 files issued from that point in time through August 2004. The results were, as expected, lower than those of previous experiments. Those results indicate the importance of maintaining the training set by acquisition of new executables, in order to cope with unknown new executables. Henchiri and Japkowicz [8] presented a hierarchical feature selection approach which enables the selection of n -gram features that appear at rates above a specified threshold in a specific virus family, as well as in more than a minimal amount of virus classes (families). They applied several classifiers: ID3, J48 Naïve Bayes, SVM- and SMO to the data set used by Schultz et al. [4] and obtained results that were better than those obtained through traditional feature selection, as presented in [4], which mainly focused on 5-grams. Additionally, [9] presented to use the frequency of the n -grams in the files to select them as alternative to information gain based selection criterion.

2.2 The imbalance problem

In machine learning the data is, often, presented as a list of labeled examples, in which an example is described by a vector of features and an additional special feature which is the class (e.g., malicious/benign). Thus, the data is actually a matrix, in which each example is a row having n features and a class, which are the columns. Often there are equal numbers of examples for each class. These general proportions are important since most of the classifiers are probabilistic and thus they induce the general proportions of the classes in the dataset. For evaluation purposes the dataset is divided into two datasets: training set, which is used to train a classifier and which actually represents the world to the learner, and a test set which represents the real life scenario. Whenever there is a significant difference in the proportions of the numbers of examples for the classes, which happens often as a result of less available examples of a specific class, it might affect the accuracy of the classifier. This case is called the class imbalance problem.

The class imbalance problem was first noticed by the machine learning research community a little over a decade ago (e.g., [10–12]). As just discussed, it typically occurs when there are significantly more instances from one class relative to other classes. In such cases most standard classifiers tend to misclassify the instances of the low represented classes. In certain cases of extreme imbalances, the classifier may go as far as to classify all the data with the label of the large class, thus, completely ignoring the data from the small class. More and more researchers realized that the performance of their classifiers may be suboptimal due to the fact that the datasets are not balanced. This problem is even more important in fields where the natural datasets are

highly imbalanced in the first place [13], like the problem we describe.

Over the years, the machine learning community has addressed the issue of class imbalances following two general strategies. The first one, which is classifier-independent, consists of balancing the original data set, using different kinds of undersampling or oversampling approaches. In particular, researchers have experimented with random sampling, where instances from the training set are either duplicated or eliminated at random (e.g., [14]); directed sampling, where specific instances are targeted for undersampling or oversampling with the idea of strengthening the most relevant data and weakening the least relevant ones (e.g., [14, 10]); and artificial sampling, where the smaller class is oversampled with artificially generated data designed to augment the minority class without creating the risk of overfitting [15]. The second way involves modifying the classifiers in order to adapt them to the data sets. In particular, these approaches look for ways of incorporating misclassification costs into the classification process and assigning higher misclassification costs to the minority class so as to compensate for its small size. This was done for a variety of different classifiers such as Neural networks [16], Random Forests [17], and SVM [18].

However, in our problem unlike in other problems, the data is not imbalanced in the training set, but rather in real life conditions, which we reflect by the test set. Thus, we don't need an algorithm to overcome the imbalanced data, but rather to understand the optimal construction of a training set to achieve the best performance in real life conditions. Our research is, thus, more in line with the work of [19], which considers the question of what proportion of examples of each class is most appropriate for learning if a only a limited number of training instances can be used altogether. Their work, considers the case of decision tree induction on twenty-six different data sets. We, on the other hand, focus on the single problem of interest here—malware detection—but consider eight different classifiers.

Another way in which our work relates to the research emanating from the class imbalance community concerns the choice of an evaluation metric, as discussed in Sect. 4.2.

3 Methods

3.1 Text categorization

For the detection and acquisition of unknown malicious code, we suggest the use of well-studied concepts from *information retrieval* (IR) and more specific *text categorization*. In our problem, binary files (executables) are parsed and n -gram terms are extracted. Each n -gram term in our problem is analogous to a word in the textual domain. We hereby describe IR concepts which we used in this study.

Salton presented the *vector space model* [20] to represent a textual file as a *bag of words*. After parsing the text and extracting the words, a vocabulary, of the entire collection of words is constructed. Each of these words may appear zero to multiple times in a document and at least in a single document. The *vocabulary* is the vector of terms which was extracted from the entire set of documents. Each term in the vocabulary can be described by its frequency in the entire collection, often called *document frequency*, which is later used for the term weighting. For each document a vector of terms in the size of the vocabulary is created, such that each index in the vector represents the *term frequency (TF)* in the document. Equation 1 shows the definition of a normalized TF, in which the term frequency is divided by the maximal appearing term in the document with values in the range of [0–1]. An extended representation is the *TF Inverse Document Frequency (TFIDF)*, which combines the frequency of a term in the document (TF) and its frequency in the documents collection, denoted by *Document Frequency (DF)*, as shown in Eq. 2, in which the term's (normalized) TF value is multiplied by the $IDF = \log(N/DF)$, where N is the number of documents in the entire file collection and DF is the number of files in which it appears.

$$TF = \frac{\text{term frequency}}{\max(\text{term frequency in document})} \quad (1)$$

$$TFIDF = TF * \log\left(\frac{N}{DF}\right) \quad (2)$$

The TF representation is actually the representation which was used in previous papers in our domain of malicious code classification. However, in the textual domain it was shown that the *tfidf* is a richer and more successful representation for terms for retrieval and categorization purposes [20], thus, we expected that using the *tfidf* weighting will lead to better performance than the *tf*. In the textual domain often the *stop words*, which are words that appear often, such as *the*, *to*, etc, are removed. These terms can be characterized by having high DF value.

3.2 Data set creation

We created a data set of malicious and benign executables for the Windows operating system, as this is the system most commonly used and most commonly attacked. To the best of our knowledge and according to a search of the literature in this field, this collection is the largest one ever assembled and used for research. We acquired the malicious files from the VX Heaven website.³ The dataset contains 7,688 malicious files. To identify the files, we used the Kaspersky⁴ anti-virus and the Windows version of the Unix 'file' command for file

type identification. The malicious files included: The files in the benign set, including executable and DLL (Dynamic Linked Library) files, were gathered from machines running Windows XP operating system on our campus. More specifically the set included applications, such as messenger, visual studio executables, anti-virus applications, zipping applications, as well as windows inner and driver dlls, service packs installers and other installation files and executables related to varying applications which were installed on the machines. The benign set contained 22,735 files. The Kaspersky anti-virus program was used to verify that these files do not contain any malicious code.

3.3 Data preparation and feature selection

N-grams extraction We parsed the files using several *n-gram* sequence lengths, denoted by n . Vocabularies of 16,777,216, 1,084,793,035, 1,575,804,954 and 1,936,342,220, for 3-gram, 4-gram, 5-gram and 6-gram respectively were extracted. Later TF and TFIDF representations were calculated for each *n-gram* in each file.

In machine learning applications, the large number of features (many of which do not contribute to the accuracy and may even decrease it) in many domains presents a significant problem. Moreover, in our problem, the reduction of the number of features is crucial, but must be performed while maintaining a high level of accuracy. This is due to the fact that, as shown earlier, the vocabulary size may exceed billions of features, far more than can be processed by any feature selection tool within a reasonable period of time. Additionally, it is important to identify those terms that appear in most of the files, in order to avoid vectors that contain many zeros. Thus, we first extracted the features having the top 5,500 *document frequency* (Eq. 2) values as a preliminary aggressive feature selection, on which later three feature selection methods were applied. In order to check whether the *stop words* phenomena happens in our problem domain we selected the 5,500 top features and 1,000–6,500 top features from the entire list ranked by the DF. The features selected from the top 1,000–6,500, in which the top 1,000 features were removed, represented the idea of stop-words in the textual domain, which we examined their potential effect here.

Feature selection We used a *filters* approach, in which a measure is used to quantify the correlation of each feature to the class (malicious or benign) and estimate its expected contribution to the classification task. After applying the filter each feature gets a rank which quantifies its expected contribution in the classification task, from which later the features with the top ranks are used. Note that the filter is applied on the dataset and the measure is independent of any classification algorithm, which enables to compare the performances of the different classification algorithms on the same subset

³ <http://vx.netlux.org>.

⁴ <http://www.kaspersky.com>.

of features. We used three feature selection measures. As a baseline, we used the *document frequency* measure DF (the amount of files in which the term appeared in), *Gain Ratio* (GR) [8] and *Fisher Score* (FS) [21].

3.3.1 Gain ratio

Gain Ratio was originally presented by Quinlan in the context of *Decision Trees* [8], which was designed to overcome a bias in the *Information Gain* (IG) measure, and which measures the expected reduction of entropy caused by partitioning the examples according to a chosen feature. Given entropy $E(S)$ as a measure of the impurity in a collection of items, it is possible to quantify the effectiveness of a feature in classifying the training data. Equation 4 presents the formula of the entropy of a set of items S , based on C subsets of S (for example, classes of the items), presented by S_c . *Information Gain* measures the expected reduction of entropy caused by partitioning the examples according to attribute A , in which V is the set of possible values of A , as shown in Eq. 3. These formulas refer to discrete values; however, it is possible to extend them to continuous values attribute.

$$IG(S, A) = E(S) - \sum_{v \in V(A)} \frac{|S_v|}{|S|} \cdot E(S_v) \quad (3)$$

$$E(S) = - \sum_{c \in C} \frac{|S_c|}{|S|} \cdot \log_2 \frac{|S_c|}{|S|}. \quad (4)$$

The IG measure favors features having a high variety of values over those with only a few. GR overcomes this problem by considering how the feature splits the data (Eqs. 5 and 6). S_i are d subsets of examples resulting from partitioning S by the d -valued feature A .

$$GR(S, A) = \frac{IG(S, A)}{SI(S, A)} \quad (5)$$

$$SI(S, A) = - \sum_{i=1}^d \frac{|S_i|}{|S|} \cdot \log_2 \frac{|S_i|}{|S|} \quad (6)$$

3.3.2 Fisher score

The Fisher score ranking technique calculates the difference, described in terms of mean and standard deviation, between the positive and negative examples relative to a certain feature. Equation 7 defines the Fisher score, in which R_i is the rank of feature i , describing the proportion of the substitution of the mean of the feature i values in the positive examples (p) and the negative examples (n), and the sum of the standard deviation. The bigger the R_i , the bigger the difference between the values of positive and negative examples relative to feature i ; thus, this feature is more *important* for separating the positive and negative examples. This technique is described in details in [21].




Fig. 1 A typical architecture of a feed forward ANN, having five hidden neurons and a single output neuron. The number of the hidden neurons and the number of the output neurons may vary according to the analyzed data

$$R_i = \frac{|\mu_{i,p} - \mu_{i,n}|}{\sigma_{i,p} + \sigma_{i,n}} \quad (7)$$

Based on each feature selection measure we selected the top 50, 100, 200 and 300 features.

3.4 Classification algorithms

We employed four commonly used classification algorithms: *Artificial Neural Networks* (ANN), *Decision Trees* (DT), *Naïve Bayes* (NB), and their boosted versions, BDT and BNB respectively, as well as SVM with three kernel functions. We briefly describe the classification algorithms we used in this study.

3.4.1 Artificial neural networks

An Artificial Neural Network (ANN) [22] is an information processing paradigm inspired by the way biological nervous systems, such as the brain, process information. The key element is the structure of the information processing system, which is a network composed of a large number of highly interconnected neurons working together in order to approximate a specific function, as shown in Fig. 1. An ANN is configured for a specific application, such as pattern recognition or data classification, through a *learning process* during which the individual weights of different neuron inputs are updated by a *training algorithm*, such as back-propagation. The weights are updated according to the examples the network receives, which reduces the *error function*. Equation 8 presents the output computation of a two-layered ANN, where x is the input vector, v_i is a weight in the output neuron, g is the activation function, w_{ij} is the weight of a hidden neuron and $b_{i,o}$ is a bias. All the ANN manipulations

were performed within the MATLAB(r) environment using the Neural Network Toolbox.

$$f(x) = g \left[\sum_i v_i g \left(\sum_j w_{ij} x_j + b_i \right) + b_o \right] \quad (8)$$

3.4.2 Decision trees

Decision tree learners [24] are a well-established family of learning algorithms. Classifiers are represented as trees whose internal nodes are tests of individual features and whose leaves are classification decisions (classes). Typically, a greedy heuristic search method is used to find a small decision tree, which is induced from the data set by splitting the variables based on the *expected information gain*. This method correctly classifies the training data. Modern implementations include pruning, which avoids the problem of over-fitting. In this study, we used J48, the Weka [24] version of the C4.5 algorithm [23]. An important characteristic of decision trees is the explicit form of their knowledge, which can be easily represented as rules.

3.4.3 Naïve bayes

The Naïve Bayes classifier is based on the *Bayes theorem*, which in the context of classification states that the posterior probability of a class is proportional to its prior probability, as well as to the conditional likelihood of the features, given this class. If no independent assumptions are made, a Bayesian algorithm must estimate conditional probabilities for an exponential number of feature combinations. Naive Bayes simplifies this process by assuming that features are *conditionally independent*, given the class, and requires that only a linear number of parameters be estimated. The prior probability of each class and the probability of each feature, given each class is easily estimated from the training data and used to determine the posterior probability of each class, given a set of features. Empirically, Naive Bayes has been shown to accurately classify data across a variety of problem domains [25].

3.4.4 Adaboost.M1 (BDT and BNB)

Boosting is a method for combining multiple classifiers. Adaboost was introduced by [26] and among its many variants is the Adaboost.M1 that is implemented in Weka. Given a set of examples and a base classifier, it generates a set of hypotheses combined by weighted majority voting. Learning is achieved in iterations. In each iteration a new set of instances is selected by favoring misclassified instances of previous iterations. This is done using an iteratively updated distribution that includes a probability for each instance to




Fig. 2 An SVM that separates the training set into two classes, having maximal margin in a two-dimensional space

be selected in the next iteration. We used the Adaboost.M1 to boost J48 decision trees and Naïve Bayes.

3.4.5 Support vector machines

Support Vector Machines is a binary classifier, which finds a linear hyperplane that separates the given examples of two classes known to handle large amounts of features. Given a training set of labeled examples in a vector format, the SVM attempts to specify a linear hyperplane that has the maximal margin, defined by the maximal (perpendicular) distance between the examples of the two classes. The examples lying closest to the hyperplane are known as the supporting vectors. The normal vector of the hyperplane (denoted as w in Eq. 9, in which n is the number of the training example) is a linear combination of the supporting vectors multiplied by LaGrange multipliers (alphas). Figure 2 illustrates a two-dimensional space, in which the examples (vectors) are located according to their features values in two groups based on their labels (classes +1 and -1) and the hyperplane which is derived to separate them linearly according to their label.

Often the data set cannot be linearly separated, so a kernel function K is used. The SVM actually projects the examples into a higher dimensional space to create a linear separation of the examples. Note that when the kernel function satisfies Mercer's condition, as was explained by Burges [27], For the general case, the SVM classifier will be in the form shown in Eq. 9, while n is the number of examples in training set, and w is normal of the hyperplane. We examined three commonly used kernels: *Linear* (SVM-LIN), *Polynomial* (SVM-POL) and *RBF* (SVM-RBF). We used the Lib-SVM implementation.⁵

$$f(x) = \text{sign}(w \cdot \Phi(x)) = \text{sign} \left(\sum_{i=1}^n \alpha_i y_i K(x_i x) \right) \quad (9)$$

⁵ <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.

Fig. 3 The Polynomial (*Left*) and Linear (*Right*) kernels applied to the same training set. The Polynomial has successfully separated the training-set hyperplane whereas the Linear hasn't





Fig. 4 The Polynomial (*Left*) and RBF kernels (*Right*) are applied to the same training set. The RBF has successfully separated the data-set whereas the Polynomial hasn't, it can be observed that the RBF kernel is very sophisticated and powerful also towards complicated training-set



As it derived from the theoretic basis of SVM and was also empirically shown by Joachim [28], one should select the appropriate kernel function with the appropriate configurations of the parameters, while, usually, the more sophisticated the kernel is, the better the performance and the results are. The ranks of the sophistication are: linear, Polynomial, RBF so that linear is the simplest one, and RBF is more sophisticated than the two others. Note that with the sophistication of the kernel, the training time and the computational resources requirements are larger.

Figure 3 (which were produced by the applet which is available from the LIBSVM website [29]) illustrate the use of the kernels given the same training-set in an SVM with

Linear and Polynomial kernels. While the SVM with Linear kernel (right-side) is not sophisticated enough to determine a hyperplane that separates the training-set optimally (thus, each class vectors are located separately), the SVM with the Polynomial kernel (left-side) has successfully determined such one.

Figure 4 illustrates the use of the RBF and Polynomial kernels in an SVM applied to the same complicated dataset. While the SVM with Polynomial kernel (left-side) is not sophisticated enough to determine hyperplane that separates the training-set, the SVM with the RBF kernel (right-side) has successfully determined such one.

4 Evaluation

4.1 Research questions

We wanted to evaluate the proposed methodology for the detection of unknown malicious codes through two main experiments. The first experiment was designed to determine the best conditions, including four aspects:

1. Which *term representation* is better, *TF* or *TFIDF*?
2. Which *n*-gram is the best: 3, 4, 5 or 6?
3. What is the better range for *global* feature selection, top 5,500 or 1,000–6,500?
4. Which *top-selection* is the best: 50, 100, 200 or 300 and which features selection: *DF*, *FS* and *GR*?
5. Which Malicious Code Percentage in the training set will be the best for any Malicious Code Percentage in the test set, and for real life conditions?

To answer the listed questions we first performed a wide set of experiments to identify the best term representation, *n*-gram, global feature selection and top selection and feature selection measure. Using the best settings we performed additional set of experiments focusing on the imbalance problem.

4.2 Evaluation measures

For evaluation purposes, we wanted to measure the accuracy of the classification algorithms, as well as the false positive and true positive which are often very important, in order to be able to tune the classifier for the best needs. For that we used the common set of measures, which included the *True Positive Rate (TPR)* measure, which is the number of *positive* instances classified correctly, as shown in Eq. 10, *False Positive Rate (FPR)*, which is the number of *negative* instances misclassified (Eq. 10), and the *Total Accuracy*, which measures the number of absolutely correctly classified instances, either positive or negative, divided by the entire number of instances shown in Eq. 11.

$$TPR = \frac{|TP|}{|TP| + |FN|}; \quad FPR = \frac{|FP|}{|FP| + |TN|} \quad (10)$$

$$\text{Total Accuracy} = \frac{|TP| + |TN|}{|TP| + |FP| + |TN| + |FN|} \quad (11)$$

Total Accuracy is the most intuitive evaluation measure and is very often used in Machine Learning. It simply returns the percentage of right choices made by the classifier. One thing it does not do, however, is indicate whether the classifier is more adept at classifying positive or negative examples. This is often important information, like in our problem where we are interested in finding out what proportion of the malcodes present in the data are actually

detected by the classifier (TPR) and what proportion of the virus-free data is wrongly classified as virus data (FPR). Even if a classifier is good at detecting viruses, it might be discarded from consideration because of the large number of false alarms (high FPR) it generates. Information of this kind could not be obtained from the Total Accuracy alone, and this is why analyses of TPR and FPR results were also included.

In fact, for the imbalance analysis, *Total Accuracy* is not only misinformed, but it is, often, simply an inappropriate measure of performance. Indeed, in such circumstances, a trivial classifier that predicts every case as the majority class could achieve very high accuracy in extremely skewed domains. Several proposals have been made to address this issue including the decomposition of accuracy into its basic components (TPR and FPR) [14], the use of ROC Analysis [30] or the G-Mean [31]. In this paper, we selected to decompose accuracy into its basic components along with the use of the G-mean. This approach is conceptually simpler than using ROC Analysis and sheds sufficient light on our results.

The g-means measure (Eq. 13) which is often used in imbalance datasets evaluation studies, is based on the sensitivity and specificity measures (Eq. 12).

$$\text{Sensitivity} = \frac{|TP|}{|TP| + |FN|}; \quad \text{Specificity} = \frac{|TN|}{|TN| + |FP|} \quad (12)$$

$$G - \text{means} = \sqrt{\text{Sensitivity} * \text{Specificity}} \quad (13)$$

Sensitivity is exactly the same thing as the TPR introduced earlier. The difference in name simply stems from the fact that various fields of study came up with the same measures of success, but named them differently. The term “Sensitivity” was coined in the medical domain while TPR is the name used in the Machine Learning community. Specificity is the opposite of FPR. It measures the proportion of negative data rightly labeled as negative. In our problem, this corresponds to the proportion of uninfected data rightly labeled as such. Sensitivity and Specificity, thus, give us the same information as TPR and FPR. The G-mean, however, combines this information in a way different from the way in which Total Accuracy does. By multiplying the components together, indeed, the G-mean sheds light on whether the classifier is lacking on one or the other aspect of classification (detection of positive examples and recognition of a negative example). This is information that is not provided by Total Accuracy and which is critical, as previously discussed, in the case of class imbalances. This is why G-mean results were also provided in the class imbalance study.




Fig. 5 While their mean accuracies were quite similar, the top 5,500 features out-performed the top 1,000–6,500, TF out-performed TFIDF and the 5-gram out-performed the other n -gram sizes

5 Experiments and results

5.1 Experiment 1

To answer the four questions, presented earlier, we designed a wide and comprehensive set of evaluation runs, including all the combinations of the optional settings for each of the aspects, amounting in 1,536 runs in a 5-fold cross validation format for all eight classifiers. Note that the files in the test-set were not in the training-set presenting unknown files to the classifier.

5.1.1 Global feature selection versus n -grams

First we wanted to find the best terms representation, tf vs tfidf, and the global feature selection. Figure 5 presents the mean accuracy of the combinations of the term representations and n -grams. While the mean accuracies are quite similar, the top 5,500 features performed better, as did the TF representation and the 5-gram. Having the TF out-performing has meaningful computational advantages; we will elaborate on these advantages in the Discussion. Additionally, the 5-grams outperformed the other n -gram sizes.

5.1.2 Feature selections and top selections

To identify the best feature selection method and the top amount of features we calculated the mean accuracy of each option, as shown in Fig. 6. Generally, the Fisher score was the superior method, starting with high accuracy, even with 50 features. Unlike the other methods, in which the 300 features out-performed, the DF's accuracy decreased after the selection of more than 200 features, while the GR accuracy significantly increased as more features were added.




Fig. 6 The Fisher score was very accurate when used with just 50 features, and maintained this high performance level as the number of features increased. When more than 200 features were added, the accuracy of GR increased significantly and the accuracy of DF decreased

Table 1 Classifiers performance: the BDT, DT and ANN out-performed (bold), while maintaining low levels of false positives

Classifier	Accuracy	FP	FN
ANN	0.941	0.033	0.134
DT	0.943	0.039	0.099
NB	0.697	0.382	0.069
BDT	0.949	0.040	0.110
BNB	0.697	0.382	0.069
SVM-lin	0.921	0.033	0.214
SVM-poly	0.852	0.014	0.544
SVM-rbf	0.939	0.029	0.154

5.1.3 Classifiers

The results of each classifier under the best settings identified before for all the classifiers (Sect. 5.1.2), including the top 300 Fisher score-selected 5-gram terms represented by TF from the top 5,500 features are presented in Table 1. The BDT, DT and ANN outperform and demonstrated low false positive rates, while the SVM classifiers also performed very well. We suggest that the poor performance of the Naïve Bayes, may be explained by the independence assumption of the NB classifier.

5.2 Experiment 2: the imbalance problem

In the second experiment, we present our main contribution in this study. In this experiment we investigated rigorously the imbalance problem in our domain and to actually answer the fifth research question. The fifth research question presents the question of what are the optimal proportions of the benign

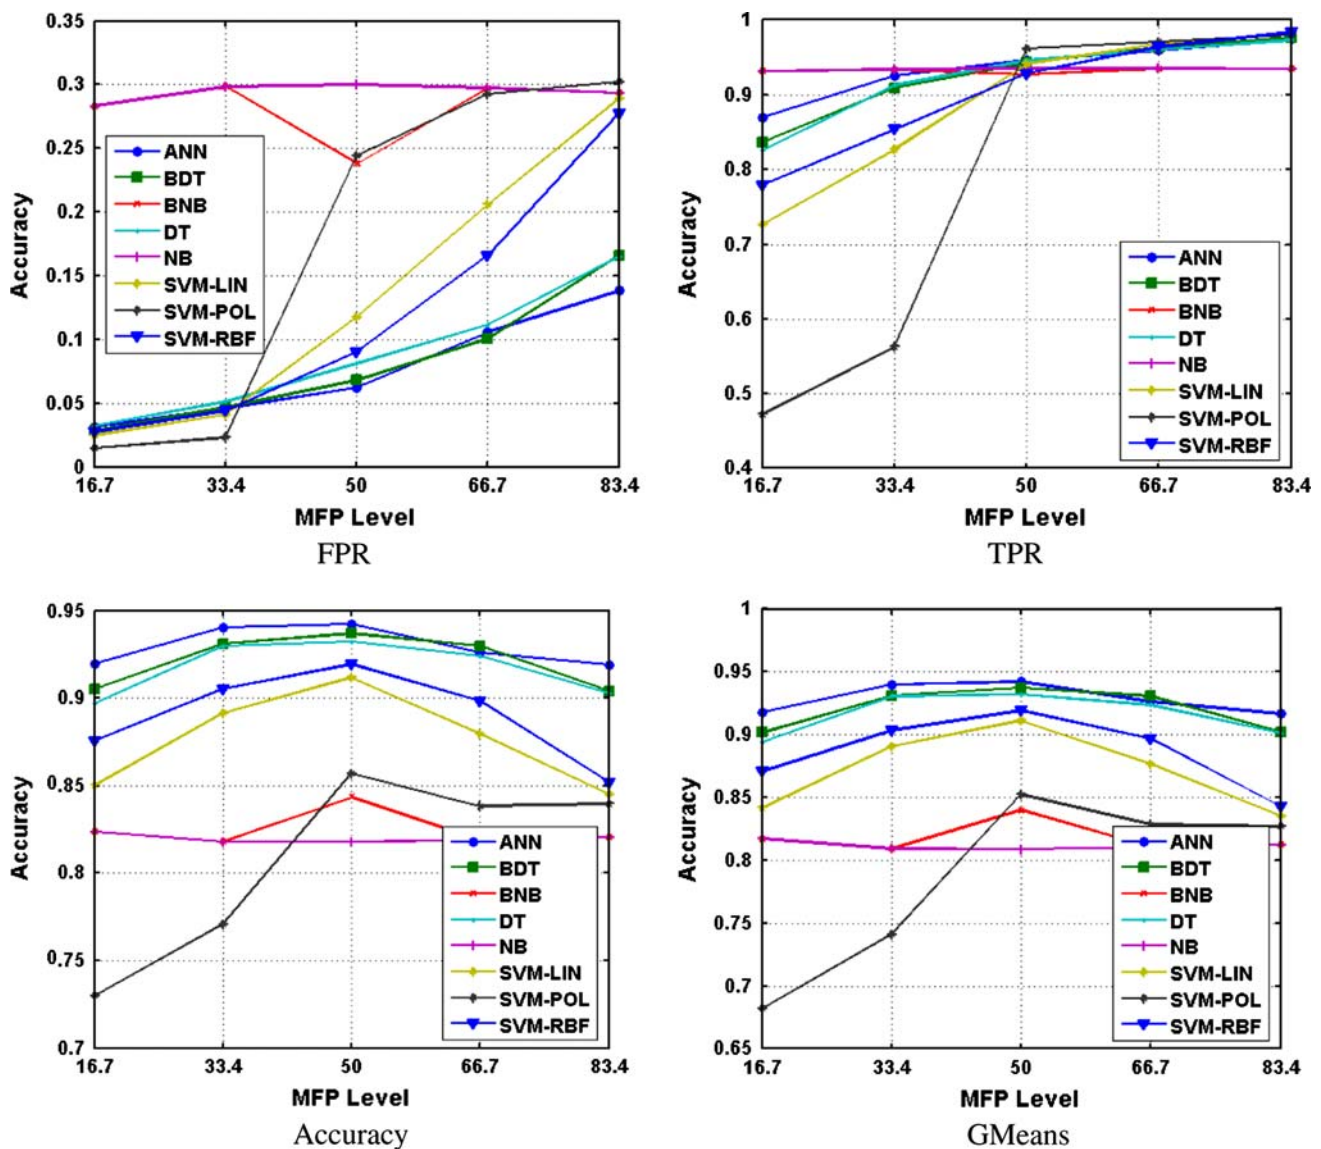


Fig. 7 ANN, BNB and DT out-performed, with consistent accuracy, across the different malcode content levels

and malicious contents in the training set for varying levels of proportions in the test set, which reflect the situation in real life. This is a very useful investigation for practical purposes, since when applying this technique the proportions in the training set should be considered according to the expected proportions in the stream of the file, which was represented in this experiment by the test set.

We used the best configuration and the top 300 Fisher Score-selected 5-gram terms represented by TF from the top 5,500 features. We created five levels of Malicious Files Percentage (MFP) in the training set (16.7, 33.4, 50, 66.7 and 83.4%), which represent the proportions which can be controlled when applying this technique. For example, when

referring to 16.7%, we mean that 16.7% of the files in the training set were malicious and 83.4% were benign. The test set represents the 'real-life' situation, while the training set represents the set-up of the classifier, which is controlled. While we assume that a MFP above 30% (of the total files) is not a realistic proportion in the stream of real networks, but we used test set that included high percentages of malicious files in order to gain insights into the behavior of the classifiers in these situations. Our study examined 17 levels of MFP (5, 7.5, 10, 12.5, 15, 20, 30, 40, 50, 60, 70, 80, 85, 87.5, 90, 92.5 and 95%) in the test sets. Eventually, we ran all the product combinations of five proportions in the training sets and 17 test sets, for a total of 85 runs for each

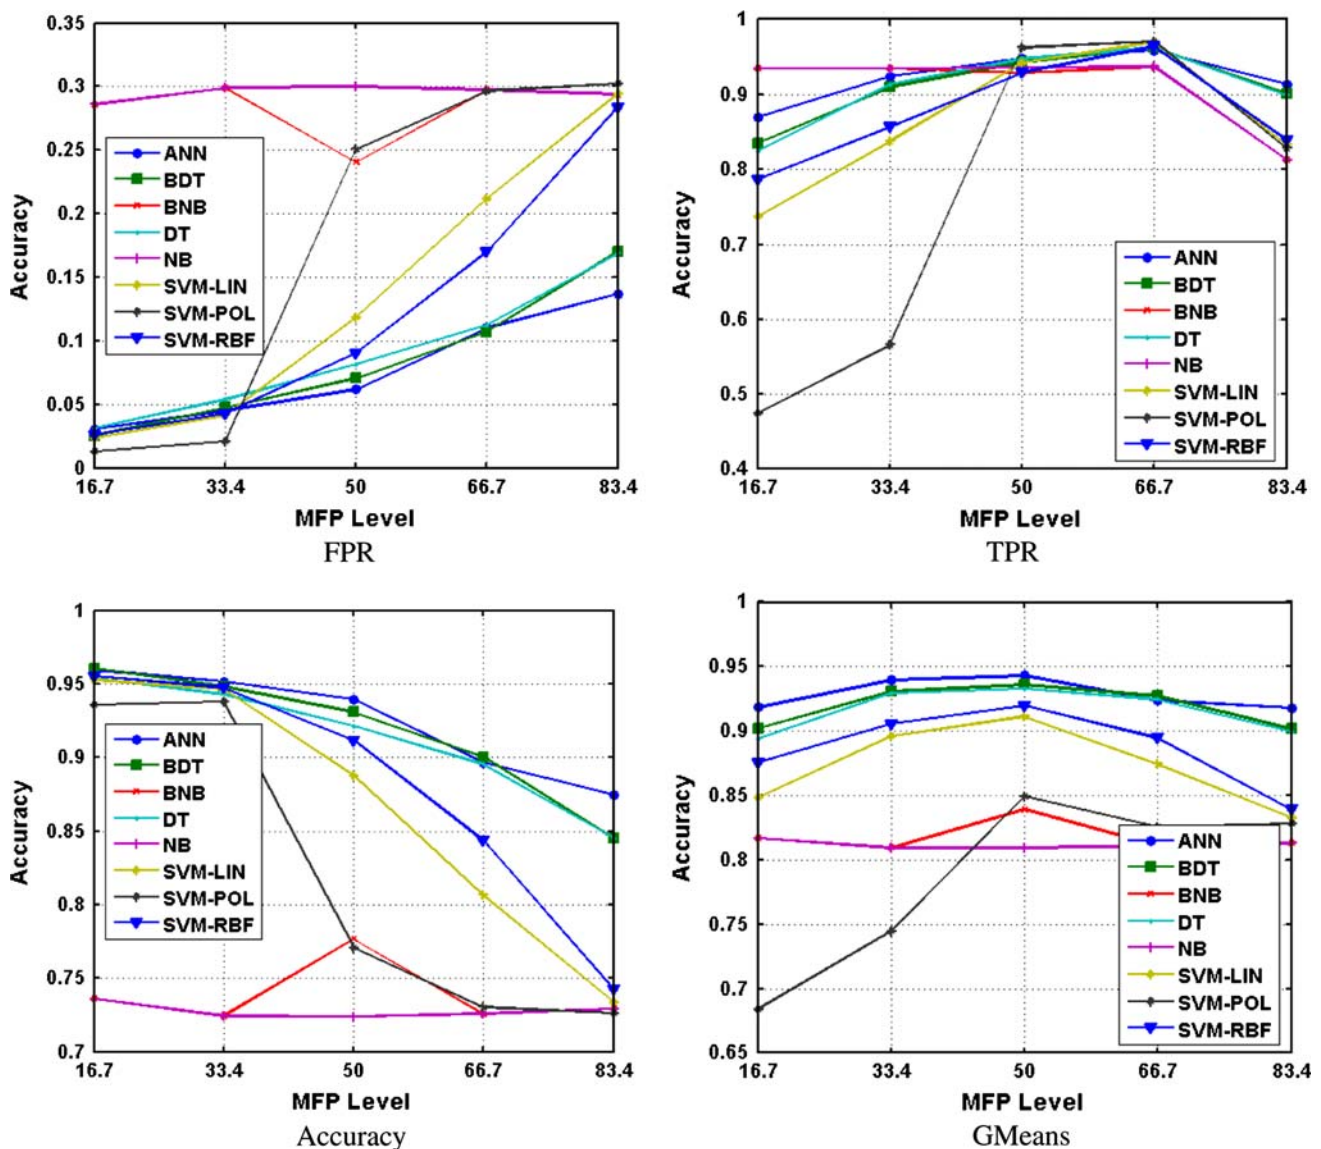


Fig. 8 Greater than 95% accuracy was achieved for the scenario involving a low level (16.7%) of malicious file content in the training set

classifier. We created two sets of data sets in order to perform a 2-fold cross validation-like evaluation to make the results more significant. We analyze all the results using four evaluation measures: the *True Positive Rate*, the *False Positive Rate*, the combination of them by the *Accuracy* and the *g-means* measure which is often used in imbalance problems.

5.2.1 Training-set malware percentage

In this analysis we evaluated the performance of each training-set MFP settings against the varying levels of MFP in the test-sets. Thus, for each MFP in the training set the mean values of the measures are presented. Figure 7 presents the mean

accuracy (averaged over all the MFP levels in the test-sets) of each classifier for each MFP level in the training set.

All the classifiers, beside NB, demonstrated an increased FPR and TPR with the increase of the MFP in the training-set. According to the Accuracy and g-means measures the classifiers behaved similarly. ANN, BDT and DT demonstrated the highest accuracy, and relatively stable, performance across the different MFP levels, while BNB, NB and SVM-POL generally performed poorly. SVM-RBF and SVM-LIN performed well, but not consistently. They were both most accurate at the 50% level, while the accuracy of SVM-POL increased as more malicious examples were presented.

5.2.2 10% malcode percentage in the test set

We consider the 10% MFP level in the test set as a realistic scenario, which reflects real life conditions, in which there are 10% of malicious contents.

Figure 8 presents the mean accuracy in the 2-fold cross validation of each classifier for each MFP level in the training set, with a fixed level of 10% MFP in the test set. Thus, each point in the curve is the average of all the runs with the varying MFPs in the training sets. Accuracy levels above 95% were achieved when the training set had a MFP of 16.7%, while a rapid decrease in accuracy was observed when the MFP in the training set was increased. Thus, the optimal proportion in a training set for practical purposes should be in the range of 10–40% malicious files. This is in line with [19] who concluded, from their study, that when accuracy is used, the optimal class distribution in the training set tends to be near the natural class distribution.

5.2.3 Relations among MFPs in training and test sets

In Sects. 5.2.1 and 5.2.2 we presented the mean results of varying MFP levels of the test set (5.2.1) and for 10% MFP in the test (5.2.2) for the each MFP in the training set. Here we present specifically the accuracy for each experiment of a MFP level in the training vs a MFP level in the test. Thus, in the following figures a three-dimension results presentation, in which the horizontal axes are the training set MFP and the test set MFP and the vertical axis is the accuracy, is given for each classifier. This presentation gives a more detailed guideline for setting the MFP in the training set for each expected MFP in the stream, reflected by the test set.

Most of the classifiers behaved optimally when the MFP levels in the training-set and test-set were similar, except for the NB and BDT, which showed low performance levels earlier. This indicates that when configuring a classifier for a real-life application, the MFP in the training-set has to be set accordingly.

6 Discussion and conclusions

We presented a methodology for the representation of malicious and benign executables for the task of unknown malicious code detection. This methodology enables the highly accurate detection of unknown malicious code, based on previously seen examples, while maintaining low levels of false alarms. In the first experiment, we found that the TFIDF representation has no added value over the TF, which is not the case in information retrieval applications. This is very important, since using the TFIDF representation introduces some computational challenges in the maintenance of the collection when it is updated. In order to reduce the number of




Fig. 9 ANN




Fig. 10 DT




Fig. 11 NB

n -gram features, which ranges from millions to billions, we first used the DF measure to select the top 5,500 features. The Fisher Score feature selection outperformed the other




Fig. 12 BNB




Fig. 15 SVM-POL




Fig. 13 BDT




Fig. 16 SVM-RBF




Fig. 14 SVM-LIN

methods and using the top 300 features resulted the best performance. Generally, the ANN and DT achieved high *mean* accuracies, exceeding 94%, with low levels of false alarms.

In the second experiment, we examined the relationship between the MFP in the test set, which represents real-life scenario, and in the training-set, which being used for training the classifier. In this experiment, we found that there are classifiers which are relatively inert to changes in the MFP level of the test set. In general, the best mean performance (across all levels of MFP) was associated with a 50% MFP in the training set (Fig. 7). However, when setting a level of 10% MFP in the test-set, as a real-life situation, we looked at the performance of each level of MFP in the training set. A high level of accuracy (above 95%) was achieved when less than 33% of the files in the training set were malicious, while for specific classifiers, the accuracy was poor at all MFP levels (Fig. 8). Finally, we presented a three-dimensional representation of the results at all the MFP levels for each classifier (Figs. 9, 10, 11, 12, 13, 14, 15 and 16). In General, the best performance was on the diagonal, where the MFP levels in the training-set and the test-set were equal. We found a decreased accuracy as the MFP of the training

set and test set differs, while NB did not seem to be affected by the level of MFP in the training-set and was influenced only by the MFP level in the test-set. In NB the accuracy increased as the MFP in the test-set increased.

Based on our extensive and rigorous experiments, we conclude that when one sets up a classifier for use in a real-life situation, he should consider the expected proportion of malicious files in the stream of data. Since we assume that, in most real-life scenarios, low levels of malicious files are present, training sets should be designed accordingly. In [32] a more general version of n -grams, called n -perms is proposed in which the order of the sequence ignored in order to detect similar permutations of code. As future work we would like to examine the use of n -perms as a more general representation of the code.

References

- Filiol, E., Josse, S.: A statistical model for undecidable viral detection. *J. Comput. Virol.* **3**, 65–74 (2007)
- Filiol, E.: Malware pattern scanning schemes secure against black-box analysis. *J. Comput. Virol.* **2**, 35–50 (2006)
- Gryaznov, D.: Scanners of the year 2000: Heuristics. In: *Proceedings of the 5th International Virus Bulletin* (1999)
- Schultz, M., Eskin, E., Zadok, E., Stolfo, S.: Data mining methods for detection of new malicious executables. In: *Proceedings of the IEEE Symposium on Security and Privacy*, 178–184 (2001)
- Abou-Assaleh, T., Cercone, N., Keselj, V., Sweidan, R.: N-gram based detection of new malicious code. In: *Proceedings of the 28th Annual International Computer Software and Applications Conference (COMPSAC'04)* (2004)
- Kolter, J.Z., Maloof, M.A.: Learning to detect malicious executables in the wild. In: *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 470–478. ACM Press, New York (2004)
- Mitchell, T.: *Machine Learning*. McGraw-Hill, New York (1997)
- Henchiri, O., Japkowicz, N.: A feature selection and evaluation scheme for computer virus detection. In: *Proceedings of ICDM-2006*, pp. 891–895. Hong Kong (2006)
- Reddy, D., Pujari, A.: N-gram analysis for computer virus detection. *J. Comput. Virol.* **2**, 231–239 (2006)
- Kubat, M., Matwin, S.: Addressing the curse of imbalanced data sets: one-sided sampling. In: *Proceedings of the Fourteenth International Conference on Machine Learning*, pp. 179–186 (1997)
- Fawcett, T.E., Provost, F.: Adaptive fraud detection. *Data Min. Knowl. Discov.* **1**(3), 291–316 (1997)
- Ling, C.X., Li, C.: Data mining for direct marketing: problems and solutions. In: *Proceedings of the Fourth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 73–79 (1998)
- Chawla, N.V., Japkowicz, N., Kotcz, A.: Editorial: special issue on learning from imbalanced data sets. *SIGKDD Explor. Newsl.* **6**(1), 1–6 (2004)
- Japkowicz, N., Stephen, S.: The class imbalance problem: a systematic study. *Intel. Data Anal. J.* **6**, 5 (2002)
- Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: SMOTE: synthetic minority over-sampling technique. *J. Artif. Intel. Res. (JAIR)* **16**, 321–357 (2002)
- Lawrence, S., Burns, I., Back, A.D., Tsoi, A.C., Giles, C.L.: Neural network classification and unequal prior class probabilities. In: Orr, G., Muller, R.-R., Caruana, R. (eds.) *Tricks of the Trade. Lecture Notes in Computer Science State-of-the-Art Surveys*, pp. 299–314. Springer, Heidelberg (1998)
- Chen, C., Liaw, A., Breiman, L.: Using random forest to learn unbalanced data. Technical Report 666, Statistics Department, University of California at Berkeley (2004)
- Morik, K., Brockhausen, P., Joachims, T.: Combining statistical learning with a knowledge-based approach—a case study in intensive care monitoring. In: *Proceedings of the International Conference of Machine Learning*, pp. 268–277 (1999)
- Weiss, G., Provost, F.: Learning when training data are costly: the effect of class distribution on tree induction. *J. Artif. Intel. Res.* **19**, 315–354 (2003)
- Salton, G., Wong, A., Yang, C.S.: A vector space model for automatic indexing. *Commun. ACM* **18**, 613–620 (1975)
- Golub, T., Slonim, D., Tamaya, P., Huard, C., Gaasenbeek, M., Mesirov, J., Coller, H., Loh, M., Downing, J., Caligiuri, M., Bloomfield, C., Lander, E.: Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* **286**, 531–537 (1999)
- Bishop, C.: *Neural Networks for Pattern Recognition*. Clarendon Press, Oxford (1995)
- Quinlan, J.R.: *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers, Inc., San Francisco (1993)
- Witten, I.H., Frank, E.: *Data Mining: Practical Machine Learning Tools and Techniques*, 2nd edn. Morgan Kaufmann Publishers, Inc., San Francisco (2005)
- Domingos, P., Pazzani, M.: On the optimality of simple Bayesian classifier under zero-one loss. *Mach. Learn.* **29**, 103–130 (1997)
- Freund, Y., Schapire, R.: A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.* **55**, 119–139 (1997)
- Burges, C.J.C.: A tutorial on support vector machines for pattern recognition. *Data Min. Knowl. Discov.* **2**(2), 955–974 (1998)
- Joachims, T.: Making large-scale support vector machine learning practical. Schölkopf, B., Burges, C., Smola, A. (eds.) *Advances in Kernel Methods: Support Vector Machines*. MIT Press, Cambridge (1998)
- Chang, C., Lin, C.: LIBSVM: a library for support vector machines (2001)
- Provost, F., Fawcett, T.: Robust classification systems for imprecise environments. In: *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)* (1998)
- Kubat, M., Holte, R., Matwin, S.: Machine learning for the detection of oil spills in satellite radar images. *Mach. Learn.* **30**, 195–215 (1998)
- Karim, Md., Walenstein, A., Lakhota, A., Parida, L.: Malware phylogeny generation using permutations of code. *J. Comput. Virol.* **1**, 13–23 (2005)