Expert Systems with Applications 30 (2006) 746–760

Expert Systems with Applications

www.elsevier.com/locate/eswa

Response modeling with support vector machines

HyunJung Shin[a], Sungzoon Cho[b,*]

[a]Friedrich Miescher Laboratory, Max Planck Society, Spemannstr. 37, 72076 Tübingen, Germany
[b]Department of Industrial Engineering, College of Engineering, Seoul National University, San 56-1, Shillim-Dong, Kwanak-Gu, 151-744 Seoul, South Korea

Abstract

Support Vector Machine (SVM) employs Structural Risk Minimization (SRM) principle to generalize better than conventional machine learning methods employing the traditional Empirical Risk Minimization (ERM) principle. When applying SVM to response modeling in direct marketing, however, one has to deal with the practical difficulties: large training data, class imbalance and scoring from binary SVM output. For the first difficulty, we propose a way to alleviate or solve it through a novel informative sampling. For the latter two difficulties, we provide guidelines within SVM framework so that one can readily use the paper as a quick reference for SVM response modeling: use of different costs for different classes and use of distance to decision boundary, respectively. This paper also provides various evaluation measures for response models in terms of accuracies, lift chart analysis, and computational efficiency.

Keywords: Response modeling; Direct marketing; Support vector machines (SVMs); Pattern selection; Class imbalance; Scoring

1. Introduction

Direct marketing is concerned with identifying likely buyers of certain products or services and promoting them to the potential buyers through various channels. A response model predicts a probability that a customer is going to respond to a promotion or offer. Using the model, one can identify a subset of customers who are more likely to respond than others. A more accurate response model will have more respondents and fewer non-respondents in the subset. By doing so, one can significantly reduce the overall marketing cost without sacrificing opportunities.

Various statistical and machine learning methods have been proposed for response modeling. These researches will be reviewed in Section 2. Most recent is Support Vector Machine (SVM) that has been spotlighted in the machine learning community thanks to its theoretical soundness and practical performance. First, it is quite satisfying from a theoretical point of view. SVM can be characterized by three statements (Vapnik, 1999). SVM attempts to position a decision boundary so that the margin between the two classes is maximized. The major parameters of SVM are taken from the training patterns. Non-linear SVM is based on the use of kernels to deal with high dimensional feature space without directly working in it. Conventional neural networks tend to overfit the training dataset, resulting in poor generalization since parameter selection is based on Empirical Risk Minimization (ERM) principle, which minimizes the error on the training set. On the contrary, the SVM formulation embodies the Structural Risk Minimization (SRM) principle, which minimizes the error on the training set with the lowest capacity. The difference allows SVM to generalize better, which is the goal in statistical learning. Theoretically, SVM includes a large class of neural networks (including radial basis functions networks), yet it is simple enough to be analyzed mathematically. Second, SVM achieved great success in practical applications as diverse as face detection and recognition, handwritten character and digit recognition, text detection and categorization, etc. (Dumais, 1998; Heisele, Poggio, & Pontil, 2000; Moghaddam & Yang, 2000; Osuna, Freund, & Girosi, 1997). In particular, Dumais (1998) and Joachims (1998) applied a number of learning methods to text categorization, such as SVMs, nearest neighbor classifiers, probabilistic Bayesian models, decision trees, and neural networks. Among them, SVMs achieved most substantial improvements over the currently best

* Corresponding author.
E-mail addresses: hyunjung.shin@tuebingen.mpg.de (H.J. Shin), zoon@snu.ac.kr (S. Cho).

0957-4174/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.eswa.2005.07.037

performing methods and they behaved robustly over a variety of different learning tasks. Aside from the aforementioned research efforts, Byun and Lee (2002) gave a comprehensive up-to-date survey on SVM applications.

However, there are some difficulties one would face when SVM is attempted to be applied to response modeling. First, SVM training can become computationally intractable. Generally, retailers keep huge amounts of customer data. Moreover, a new customer's record will be added on top of it on and on. Unfortunately, in SVM quadratic programming (QP) formulation, the dimension of kernel matrix ($M \times M$) is equal to the number of training patterns (M) (Vapnik, 1999). Thus, when the dataset is huge, training cannot be finished in a reasonable time even if we manage to load the kernel matrix on the memory. Most standard SVM QP solvers have time complexity of O(M3): MINOS, CPLEX, LOQO and MATLAB QP routines. And the solvers using decomposition methods have time complexity of $I \cdot O(Mq+q3)$ where I is the number of iterations and q is the size of the working set: Chunking, SMO, SVMlight and SOR (Hearst, Schölkopf, Dumais, Osuna, & Platt, 1997; Platt, 1999). Needless to say, I increase as M increases. Second, response modeling is likely to have a severe class imbalance problem since the customers' response rates are typically very low. Most of customers belong to the non-respondents' group (class 1), while only a few customers belong to the respondents' group (class 2). Under such a circumstance, most classifiers do not behave well, and neither does SVM. Third, one has to find a way to estimate scores or likelihoods from SVM. Given a limited amount of marketing expenses, a marketer wants to maximize the return or total revenue. Thus, one would like to know who is more likely to purchase than others. Response models compute each customer's likelihood or propensity to respond to a particular offer of a product or a service. These likelihood values or scores are then used to sort the customers in a descending order. Now, the marketer simply applies a cut-off value based on the marketing expenses and only those customers whose scores are larger than the value are identified. However, an SVM classifier returns a binary output, not a continuous output, which can be interpreted as a score.

In this paper, we provide solutions for the obstacles mentioned above. For the intractability problem of SVM training, we present a pattern selection algorithm that reduces the training set without accuracy loss. The algorithm selects only the patterns near the decision boundary based on neighborhood properties. Its performance was previously validated for various problems in (Shin & Cho, 2003a). For the remaining two problems, we provide guidelines within SVM framework so that one can readily use the paper as a quick reference for SVM response modeling. For the class imbalance problem, we show how to incorporate different misclassification costs into the objective function, and for the scoring problem, we show theory-backup of how the distance from a pattern to the decision hyperplane in the feature space can be used as a score. In addition, we provide various measures for evaluating the response models in both accuracy and profit.

The remaining part of this paper is organized as follows. Section 2 presents related work on various statistical or machine learning methods applied to direct marketing domain. The section also includes literature reviews on pattern selection, class balancing, and scoring. Section 3 briefly explains the SVM theory, in particular, the patterns critically affecting the training. Section 4 addresses the obstacles in applying SVM to response modeling. The section proposes ways to reduce the training set, to handle the class imbalance problem, and to obtain the customer scores from an SVM classifier. Section 5 provides the experimental results on a direct marketing dataset. The section includes the data set description, experimental design, and performance measurements. We conclude this paper with some future works in Section 6.

2. Related work

2.1. Various methods applied to direct marketing

Traditionally, statistical methods, mostly regression techniques, have been applied to response modeling. Most textbooks cover logistic regression as the de facto method due to its simplicity, explainability and availability (Hosmer & Lemeshow, 1989; Sen & Srivastava, 1990). Malthouse (1999) compared ridge regression with stepwise regression on the Direct Marketing Educational Foundation data set 2 (DMEF2). In his study, both methods were used for determining the moderate number of variables in response modeling. Empirically, he showed that ridge regression is a more stable and less risky method than dropping variables. In his recent report, a similar approach, which additively considered the dollars spent in response to an offer, was proposed (Malthouse, 2002). Colombo and Jiang (1999) proposed a simple Recency–Frequency–Monetary (RFM) stochastic model for ranking (or scoring) customers. The RFM stochastic model derived from the response distribution of the past was used to estimate the likelihood of future responses. A customer mailing list obtained from a tele-markeing company was used for comparing the performance of the stochastic model with that of regression and cross-tabulation model. They reported that the stochastic model provided a more insightful alternative to ranking customers.

Recently, machine-learning methods have been proposed. They include decision trees and neural networks, etc. Haughton and Oulabi (1997) compared the response lifts of two mostly common decision tree algorithms: Classification and Regression Tree (CART) and Chi-Square Automatic Interaction Detector (CHAID). Although the two models are different in their tree-generating mechanism, there was no significant difference in the response-lift perspective. Ling and Li (1998) compared a Naive Bayes response model and a C4.5 response model. Applying the ada-boost algorithm (Freund & Schapire, 1996) to each base model for better performance, they conducted experiments on three direct marketing problems such as loan product promotion, life insurance product campaign, and bonus program. All experiments were designed to discuss the difficulties, which can arise during the response modeling process, such as class imbalance and justifiability of evaluating measures. Coenen, Swinnen, Vanhoof, and Wets (2000) proposed to combine C5, a decision tree algorithm, and case-based reasoning (CBR). In this approach, the C5 based response modeling was conducted in the first step. Then, the respondents classified by the initial model were ranked by a CBR similarity measure. They improved the classification quality by accommodating a better ranking rather than the accuracy of the base response model itself. Chiu (2002) integrated genetic algorithm (GA) into a CBR based response model. For better case identification accuracy, the fittest weighting values on the cases were searched by GA. On the application of an insurance product purchase dataset, the base response model, CBR, achieved better classification accuracy. Deichmann, Eshghi, Haughton, Sayek, and Teebagy (2002) investigated the use of Multiple Adaptive Regression Splines (MARS) as a response model. MARS is an advanced decision tree technique enabling piecewise linear regression. The MARS response model outperformed the logistic regression model on the DMEF2.

There have also been many reports on neural networks. Moutinho, Curry, Davies, and Rita (1994) predicted bank customers' responses using neural networks, and Bounds and Ross (1997) showed that neural network based response models improved the response rate from 1 or 2% up to 95%. Zahavi and Levin (1997a) addressed unique merits and demerits of neural networks for response modeling. Viaene, Baesens, Van den Poel, Dedene, and Vanthienen (2001a) proposed to select relevant variables for neural network based response models. Ha, Cho, and MacLachlan (2005) proposed a response model using bagging neural networks. The experiments over a publicly available DMEF4 dataset showed that bagging neural networks give more improved and stabilized prediction accuracies than single neural networks and logistic regression. Performance comparison of the methods has been one of the controversial issues in direct marketing domain. Suh, Noh, and Suh (1999) and Zahavi and Levin (1997a,1997b) found that neural network did not outperform other statistical methods. They suggested to combine the neural network response model and the statistical method. On the other hand, Bentz and Merunkay (2000) reported that neural networks outperformed multinomial logistic regression. Potharst, Kaymak, and Pijls (2001) applied neural networks to direct mailing campaigns of a large Dutch charity organization. According to their results, the performance of neural networks surpassed that of CHAID or logistic regression.

Although SVM is applied to a wide variety of application domains, there have been only a couple of SVM application reports in response modeling. Cheung, Kwok, Law, and Tsui (2003) used SVM for content-based recommender systems. Web retailers implement a content-based system to provide recommendations to a customer. The system automatically matches his/her interests with product-contents through web pages, newsgroup messages, and new items. It is definitely a form of direct marketing that has emerged by virtue of recent advances in the World Wide Web, e-business, and on-line companies. They compared Naive Bayes, C4.5 and 1-nearest neighbor rule with SVM. The SVM yielded the best results among them. More specific, SVM application to response modeling was attempted by Viaene et al. (2001b). They proposed a Least Square SVM (LS-SVM) based wrapper approach. Wrapper indicates an input variable selection procedure working together with a learning algorithm, and it is frequently compared with alternative procedure, filter, that performs variable selection independently from a learning algorithm. In their study, the input variable pool was composed of RFM and non-RFM variables from the customer dataset provided by a major Belgian mail-order company. Then, the wrapper approach was performed in a sequential backward fashion, guided by a best-first variable selection strategy. Their approach, a wrapper around the LS-SVM response model, could gain significant reduction of model complexity without degrading predictive performance.

2.2. Pattern selection

Now, let us focus on the researches related to the difficulties we addressed in this paper. First, the most straightforward method to reduce a large training set is random sampling. In SVM, however, the patterns near the decision boundary are critical to learning. The training set reduced by random sampling may omit those, thus would lead to significantly poorer prediction results. Some SVM researchers thus have attempted to identify those training patterns near the decision boundaries. Lyhyaoui et al. (1999) implemented RBF classifiers which somewhat resemble SVMs, to clear the difference between both methods. RBF classifiers were built on the patterns near the decision boundary. To find them, they proposed 1-nearest neighbor method in the opposite class after class-wise clustering. But this method makes an impractical assumption that the training set is clean. An approach focusing more on SVM was proposed by Almeida, Braga, and Braga (2000) who conducted k-means clustering on the entire training set. All patterns were selected for heterogeneous clusters (whose patterns disagree in their class membership) while only the centroids were selected for homogeneous clusters. The drawbacks of this research are that it is not clear how to determine the number of clusters, and that the clustering performance is generally unstable (Liu & Nakagawa, 2001). More recently, Shin and Cho (2003a) proposed a neighborhood properties based pattern selection algorithm (NPPS), which will be introduced in Section 4.1.

2.3. Class-balancing

Regarding class imbalance, many researchers have recognized this problem and suggested several methods: enlarging the small class dataset by random sampling, reducing the large class dataset by random sampling, and ignoring the small class dataset and using only the large class dataset to build a one-class recognizer. Japkowicz (2000) compared the three commonly used methods above on the degree of concept complexity using a standard neural network classifier. All the methods generally improved the performance of the learning algorithm. In particular, the first two methods were very effective especially as the concept complexity increases while the last one was relatively less accurate. Ling and Li (1998) addressed the specificity of the class imbalance problem, which resides in marketing datasets. They did not attempt to balance the imbalanced class ratio for better predictive accuracy. Instead, to circumvent the class imbalance problem, a marketing specific evaluation measure, lift index, was suggested. Lift index provides the customer's rank (score) by reflecting the confidence of classification result. They argued that even if all of the patterns were predicted as one class, as long as the learning algorithm produces suitable ranking of the patterns, the imbalanced class distribution in the training set would no longer be a problem. However, in their experiments the entire best lift index was obtained when the sizes of the classes were equal. Thus, they recommended to reduce the large class dataset so that its size becomes equal to that of the small class. Alternatively, different misclassification costs can be incorporated into classes, which avoid direct artificial manipulation on the training set (Lee, Gunn, Harris, & Reed, 2001).

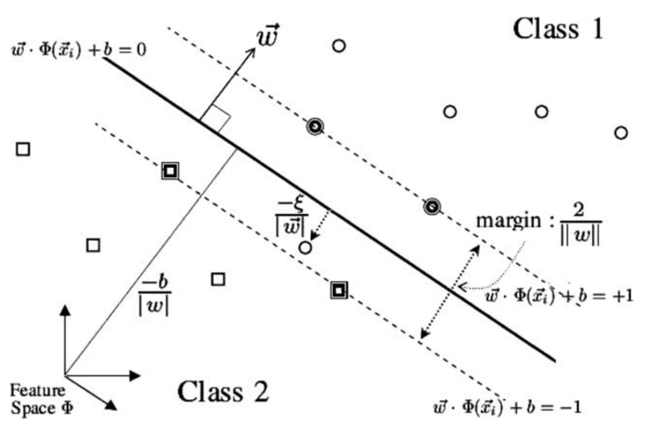
$\vec{w}\cdot\Phi(\vec{x}_i)+b=0$
$\vec{w}$
Class 1
margin : $\frac{2}{\|w\|}$
$\frac{-\xi}{|\vec{w}|}$
$\frac{-b}{|w|}$
$\vec{w}\cdot\Phi(\vec{x}_i)+b=+1$
Feature Space Φ
Class 2
$\vec{w}\cdot\Phi(\vec{x}_i)+b=-1$

Fig. 1. SVM classification problem: Through a mapping function $\Phi(\cdot)$, the class patterns are linearly separated in a feature space. The patterns determining both margin hyperplanes are outlined. The decision boundary is the half-way hyperplane between margins.

2.4. Scoring

Getting scores from a logistic regression model or a neural network model with sigmoidal output function is well known. The output gives a value of probability belonging to the class that is ranged from 0 to 1. Thus the output value is used as a score for sorting the customers. Ling and Li (1998) made use of the ada-boost algorithm (Freund & Schapire, 1996), an ensemble approach, to get the customers' scores. Basically, ada-boost maintains a sampling probability distribution on the training set, and modifies the probability distribution after each classifier is built. The probability of patterns with an incorrect prediction by the previous classifier is increased. So these patterns will be sampled more likely in the next round of boosting, to be learnt correctly. A pattern's probability to be incorrectly predicted allowed a corresponding rank. Sometimes, scores could be directly estimated by regression model having continuous target value, i.e. the dollars spent or the amount of orders. To do that, however, one needs to diagnose the problems the target variable has and conduct suitable remedies to cure them. Malthouse (2001) built a regression model to estimate the dollars spent on DMEF4. There were a large number of extreme values and the distribution was highly skewed. The extreme values could have a large influence on estimate values under least squares. And the variance of target variable most likely increased with its mean (heteroscedasticity). Thus, he performed log transformation to alleviate skewness and heterocedasticity, and used winsorization to exclude some extreme values of target. The predicted value of the dollars spent was used as a score in lift chart analysis. The lift result by means of regression based scoring will be briefly compared with that by means of classification based scoring in Section 5.4. Generally speaking, regression problem requires more information from input variables than classification problem does. In other words, binary classification is the simplest subproblem of regression. Producing good scores from marketing regression model is difficult at the present time. In addition, since SVM theory stemmed from classification context (Schölkopf, Burges, & Smola, 1999), it is natural to get scores from an SVM classifier.

3. Support vector machines and critical training patterns

Support Vector Machines (SVMs) are a general class of statistical learning architectures that perform structural risk minimization on a nested set structure of separating hyperplanes (Cristianini & Shawe-Taylor 2000; Schölkopf & Smola 2002, and Vapnik 1999). Consider a binary classification problem with M patterns $(\vec{x}_i, y_i)$,

$i=1,\ldots,M$ where $\vec{x}_i \in \Re^d$ and $y_i \in \{-1, 1\}$. Let us assume that patterns with $y_i=1$ belongs to class 1 while those with $y_i=-1$ belong to class 2. SVM training involves solving the following quadratic programming problem, which yields the largest margin ($2/\|w\|$) between classes,

$$\min \Theta(\vec{w},\xi) = \frac{1}{2}\|\vec{w}\|^2 + C\sum_{i}^{M}\xi_i, \qquad (1)$$
$$\text{s.t.} y_i(\vec{w}\cdot\Phi(\vec{x}_i)+b) \geq 1-\xi_i, \quad \xi_i \geq 0, \quad i=1,\ldots,M,$$

where $\vec{w} \in \Re^{\mathrm{d}}$, $b \in \Re$ (see Fig. 1). Eq. (1) is the most general SVM formulation allowing both non-separable and nonlinear cases. The ξs are nonnegative slack variables for a non-separable case, which play a role of allowing a certain level of misclassification. The $\Phi(\cdot)$ is a mapping function for a nonlinear case that projects patterns from the input space into a feature space. This nonlinear mapping is performed implicitly by employing a kernel function, $K(\vec{x},\vec{x}')$, to avoid the costly calculation of inner products, $\Phi(\vec{x})\cdot\Phi(\vec{x}')$. There are three typical kernel functions, RBF, polynomial, and tansig in due order,

$$K(\vec{x},\vec{x}') = \exp(-\|\vec{x}-\vec{x}'\|^2/2\sigma^2),$$
$$K(\vec{x},\vec{x}') = (\vec{x}\cdot\vec{x}'+1)^p, \qquad (2)$$
$$K(\vec{x},\vec{x}') = \tan h(\rho(\vec{x}\cdot\vec{x}')+\delta).$$

The optimal solution of Eq. (1) yields a decision function of the following form,

$$\begin{aligned} f(\vec{x}) &= \mathrm{sign}(\vec{w}\cdot\Phi(\vec{x})+b) \\ &= \mathrm{sign}\left(\sum_{i=1}^{M} y_i\alpha_i\Phi(\vec{x}_i)\cdot\Phi(\vec{x})+b\right) \\ &= \mathrm{sign}\left(\sum_{i=1}^{M} y_i\alpha_i K(\vec{x}_i,\vec{x})+b\right), \end{aligned} \qquad (3)$$

where α_is are nonnegative Lagrange multipliers associated with training patterns, respectively. The solutions, α_is, are obtained from the dual problem of Eq. (1), which minimizes a convex quadratic objective function under constraints

$$\min_{0\leq\alpha_i\leq C} W(\alpha_i,b) = \frac{1}{2}\sum_{i,j=1}^{M}\alpha_i\alpha_j y_i y_j K(\vec{x}_i\cdot\vec{x}_j) - \sum_{i=1}^{M}\alpha_i + b\sum_{i=1}^{M} y_i\alpha_i.$$

The first-order conditions on $W(\alpha_i, b)$ are reduced to the Karush–Kuhn–Tucker (KKT) conditions,

$$\frac{\partial W(\alpha_i,b)}{\partial\alpha_i} = \sum_{j=1}^{M} y_i y_j K(\vec{x}_i,\vec{x}_j)\alpha_j + y_i b - 1 = y_i f(\vec{x}_i) - 1 = g_i,$$
$$\frac{\partial W(\alpha_i,b)}{\partial b} = \sum_{j=1}^{M} y_j\alpha_j = 0, \qquad (4)$$

where $\bar{f}(\cdot)$ is the function inside the parentheses of sign in Eq. (3). The KKT complementarity condition, Eq. (4), partitions the training pattern set into three categories according to the corresponding α_is.

(a) $g_i>0 \rightarrow \alpha_i=0$: irrelevant patterns
(b) $g_i=0 \rightarrow 0<\alpha_I<C$: margin support vectors
(c) $g_i<0 \rightarrow \alpha_i=C$: error support vectors

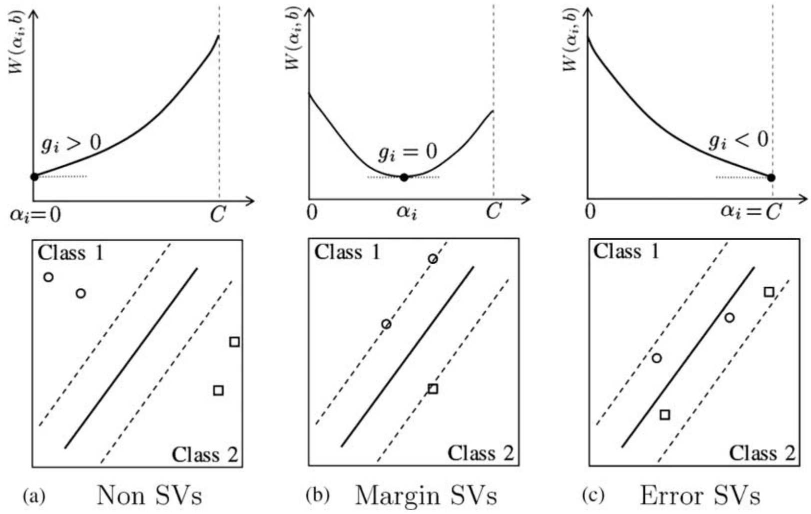

Fig. 2. Three categories of training patterns.

Fig. 2 illustrates those categories (Cauwenberghs & Poggio, 2001; Pontil & Verri, 1998). The patterns belonging to (a) are out of the margins, thus irrelevant to training, while the patterns belonging to (b) and (c) are critical ones directly affecting training. They are called support vectors (SVs). The patterns of (b) are strictly on the margin, hence called margin SVs. On the other hand, the patterns of (c) lie between two margins, hence called error SVs but are not necessarily misclassified. Going back to Eq. (3), we can now see that the decision function is a linear combination of kernels on only those critical training patterns (denoted as SVs) because the patterns corresponding to $\alpha_i = 0$ have no influence on the decision result.

$$f(\vec{x}) = \operatorname{sign}\left(\sum_{i=1}^{M} y_i \alpha_i K(\vec{x}_i, \vec{x}) + b\right) = \operatorname{sign}\left(\sum_{i \in \mathbf{SVs}} y_i \alpha_i K(\vec{x}_i, \vec{x}) + b\right). \tag{5}$$

To sum up, it is clear in SVM theory, which patterns are of importance to training. Those are distributed near the decision boundary, and fully and succinctly define the classification task at hand. And the SVMs trained with different kernel functions (RBF, polynomial, tansig) on the same training set have been founded to select almost identical subset as support vectors (Schölkopf et al., 1995). Therefore, it is worth finding such would-be support vectors in advance.

4. Support vector machines for response modeling

4.1. Large training dataset

We propose to use a neighborhood property based pattern selection algorithm (NPPS) (Shin & Cho, 2002, 2003a). The idea is to select only those patterns located around decision boundary since they are the ones that contain most information. Contrary to a usually employed 'random sampling,' this approach can be viewed as 'informative or intelligent sampling'. Fig. 3 conceptually shows the difference between NPPS and random sampling in selecting a subset of the training data. NPPS selects the patterns in the region around the decision boundary, while random sampling selects those from the whole input space. Obviously, the decision boundary is not known until a classifier is built. Thus, the algorithm utilizes neighborhood properties to infer the proximity of a pattern to the decision boundary. The first neighborhood property is that 'a pattern located near the decision boundary tends to have more heterogeneous neighbors in their class membership'. Thus, the proximity of pattern $\vec{x}$'s to the decision boundary was estimated by 'Neighbors_Entropy ($\vec{x}$, k)', which is defined as the entropy of the pattern $\vec{x}$s k-nearest neighbors' class labels,

$$\text{Neighbors\_Entropy}(\vec{x}, k) = \sum_{j=1}^{J} P_j \log_J \frac{1}{P_j},$$

where P_j is defined as k_j/k where k_j is the number of neighbors belonging to class j among the k nearest neighbors of $\vec{x}$ in J class classification problem. A pattern with a positive Neighbors_Entropy ($\vec{x}$,k) value is assumed to be close to the decision boundary, thus selected for training. Those patterns are likely to be SVs, which correspond to the margin SVs in Fig. 2(b) or the error SVs in Fig. 2(c). Among the patterns having a positive value of Neighbors_Entropy ($\vec{x}$,k), however, overlapping or noisy patterns are also present. These patterns have to be identified and removed as much as possible since they are more likely to be the error SVs misclassified in Fig. 2(c). The second neighborhood property thus dictates that 'an overlap or a noisy pattern tends to belong to a different class from its neighbors'. If a pattern's own label is different from the majority label of its neighbors, it is likely to be incorrectly labeled. The measure 'Neighbors_Match ($\vec{x}$, k)' is defined as the ratio of $\vec{x}$'s

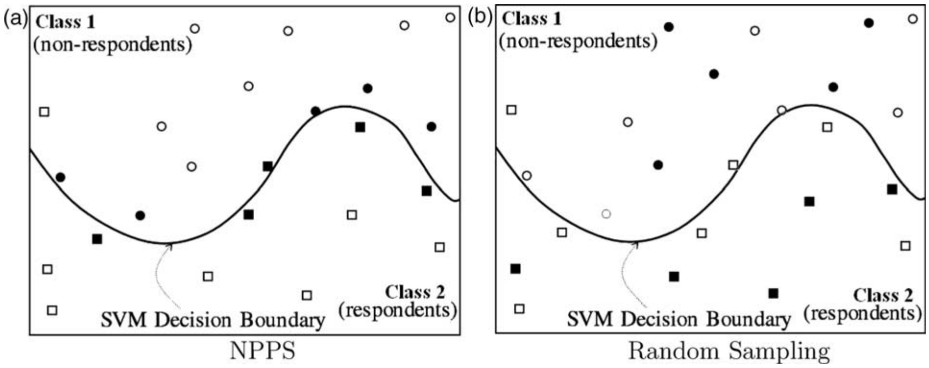

Fig. 3. NPPS and random sampling select different subsets: outlined circles and squares are the patterns belonging to class 1 (non-respondents' group) and class 2 (respondents' group), respectively. Black solid circles and squares are the selected patterns.

```
NPPS (D, k) {
    [0] Initialize $\mathbf{D}_\mathbf{e}^0$ with randomly chosen patterns from $\mathbf{D}$.
        Constants $k$ and $J$ are given. Initialize $i$ and various sets as follows:
        $i \leftarrow 0, \quad \mathbf{S}_\mathbf{o}^0 \leftarrow \emptyset, \quad \mathbf{S}_\mathbf{x}^0 \leftarrow \emptyset, \quad \mathbf{S}^0 \leftarrow \emptyset.$

    while $\mathbf{D}_\mathbf{e}^i \neq \emptyset$ do
      [1] Choose $\vec{x}$ satisfying [Expanding Criteria].
            $\mathbf{D}_\mathbf{o}^i \leftarrow \{\vec{x} \mid Neighbors\_Entropy(\vec{x},k) > 0,\ \vec{x} \in \mathbf{D}_\mathbf{e}^i\}.$
            $\mathbf{D}_\mathbf{x}^i \leftarrow \mathbf{D}_\mathbf{e}^i - \mathbf{D}_\mathbf{o}^i.$

      [2] Select $\vec{x}$ satisfying [Selecting Criteria].
            $\mathbf{D}_\mathbf{s}^i \leftarrow \{\vec{x} \mid Neighbors\_Match(\vec{x},k) \geq \beta/J,\ \vec{x} \in \mathbf{D}_\mathbf{o}^i\}.$

      [3] Update the pattern sets.
            $\mathbf{S}_\mathbf{o}^{i+1} \leftarrow \mathbf{S}_\mathbf{o}^i \cup \mathbf{D}_\mathbf{o}^i$ : the expanded,
            $\mathbf{S}_\mathbf{x}^{i+1} \leftarrow \mathbf{S}_\mathbf{x}^i \cup \mathbf{D}_\mathbf{x}^i$ : the non-expanded,
            $\mathbf{S}^{i+1} \leftarrow \mathbf{S}^i \cup \mathbf{D}_\mathbf{s}^i$ : the selected.

      [4] Compute the next evaluation set $\mathbf{D}_\mathbf{e}^{i+1}$.
            $\mathbf{D}_\mathbf{e}^{i+1} \leftarrow \bigcup_{\vec{x} \in \mathbf{D}_\mathbf{o}^i} k\mathbf{NN}(\vec{x}) - (\mathbf{S}_\mathbf{o}^{i+1} \cup \mathbf{S}_\mathbf{x}^{i+1}).$

      [5] $i \leftarrow i + 1.$
    end
    return $\mathbf{S}^i$
}
```

Fig. 4. NPPS.

Table 1
Notation

Symbol	Meaning
$\mathbf{D}$	the original training set whose cardinality is M
$\mathbf{D}_\mathbf{e}^i$	the evaluation set at ith step
$\mathbf{D}_\mathbf{o}^i$	A subset of $\mathbf{D}_\mathbf{e}^i$, the set of patterns to be 'expanded' from $\mathbf{D}_\mathbf{e}^i$ each element of which will compute its k nearest neighbors to constitute the next evaluation set, $\mathbf{D}_\mathbf{e}^{i+1}$
$\mathbf{D}_\mathbf{x}^i$	A subset of $\mathbf{D}_\mathbf{e}^i$, the set of patterns 'not to be expanded' from $\mathbf{D}_\mathbf{e}^i$, or $\mathbf{D}_\mathbf{x}^i = \mathbf{D}_\mathbf{e}^i - \mathbf{D}_\mathbf{o}^i$
$\mathbf{D}_\mathbf{s}^i$	the set of 'selected' patterns from $\mathbf{D}_\mathbf{o}^i$ at ith step
$\mathbf{S}_\mathbf{o}^i$	the accumulated set of expanded patterns, $\cup_{j=0}^{i-1}\mathbf{D}_\mathbf{o}^j$
$\mathbf{S}_\mathbf{x}^i$	the accumulated set of non-expanded patterns, $\cup_{j=0}^{i-1}\mathbf{D}_\mathbf{x}^j$
S^i	the accumulated set of selected patterns, $\cup_{j=0}^{i-1}\mathbf{D}_\mathbf{s}^j$ the last of which SN is the reduced training pattern set
k$\mathbf{NN}$ $(\vec{x})$	the set of k nearest neighbors of $\vec{x}$

neighbors whose label matches that of $\vec{x}$,

$$\text{Neighbors\_Match}(\vec{x},k) = \frac{|\{\vec{x}'|\text{label}(\vec{x}') = \text{label}(\vec{x}), \vec{x}' \in kNN(\vec{x})\}|}{k},$$

where $kNN(\vec{x})$ is the set of k nearest neighbors of $\vec{x}$. The patterns with a small Neighbors_Match $(\vec{x}, k)$ value is likely to be the ones incorrectly labeled. Only the patterns satisfying the two conditions, Neighbors_Entropy $(\vec{x},k) > 0$ and Neighbors_Match $(\vec{x},k) \geq \beta \cdot (1/J)(0 < \beta \leq 1)$, are selected. However, the NPPS evaluating kNNs for M patterns have time complexity of $O(M^2)$, so the pattern selection process itself can be time-consuming. To accelerate the pattern selection procedure, let us consider the third neighborhood property, 'the neighbors of a pattern located near the decision boundary tend to be located near the decision boundary as well'. Assuming the property, one may compute only the neighbors' label entropy for the patterns near the decision boundary instead of all the training patterns. Only the neighbors of a pattern satisfying Neighbors_Entropy $(\vec{x},k) > 0$, are evaluated in the next step. This lazy evaluation reduce the time complexity from $O(M^2)$ to $O(vM)$, where v is the number of patterns in the overlap region. In most practical problems, $v < M$ holds. We provided the time complexity analysis for the fast NPPS in (Shin & Cho, 2003b), and a systematic procedure for determining the value of k in (Shin & Cho, 2003c). The algorithm and related notations are shown in Fig. 4 and Table 1.

Fig. 5 depicts one of the experimental results previously reported (Shin & Cho, 2003a). The training set consisted of 600 patterns generated from four Gaussian densities

$$N_{1A} = \left\{ \vec{x} | N\left(\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0.5^2 & 0 \\ 0 & 0.5^2 \end{bmatrix} \right) \right\},$$

$$N_{1B} = \left\{ \vec{x} | N\left(\begin{bmatrix} -1 \\ -1 \end{bmatrix}, \begin{bmatrix} 0.5^2 & 0 \\ 0 & 0.5^2 \end{bmatrix} \right) \right\},$$

$$N_{2A} = \left\{ \vec{x} | N\left(\begin{bmatrix} -1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0.5^2 & 0 \\ 0 & 0.5^2 \end{bmatrix} \right) \right\},$$

$$N_{2B} = \left\{ \vec{x} | N\left(\begin{bmatrix} 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 0.5^2 & 0 \\ 0 & 0.5^2 \end{bmatrix} \right) \right\}.$$

And the classes, C_1 and C_2, were defined as

$$C_1 = \left\{ \vec{x} | \vec{x} \in N_{1A} \cup N_{1B}, \begin{bmatrix} -3 \\ -3 \end{bmatrix} \leq \vec{x} \leq \begin{bmatrix} 3 \\ 3 \end{bmatrix} \right\},$$

$$C_2 = \left\{ \vec{x} | \vec{x} \in N_{2A} \cup N_{2B}, \begin{bmatrix} -3 \\ -3 \end{bmatrix} \leq \vec{x} \leq \begin{bmatrix} 3 \\ 3 \end{bmatrix} \right\}$$

Fig. 5(a) is a typical result of SVM trained with all patterns. Fig. 5(b) shows the result of SVM trained with the 180 selected patterns by the naive or the fast NPPS (30.0% of the training set). The decision boundaries in both figures look quite similar, thus, generalization performance is similar. Furthermore, the NPPS reduced SVM training time 98%, almost two orders of magnitude.

4.2. Class imbalance

Usually there are many more non-respondents than respondents in training datasets. Thus, sub-sampling of non-respondent class data is the most widely used method to balance the datasets. However, random sampling allows 'important' patterns near the decision boundary to be missed. Those patterns are likely to become support vectors.

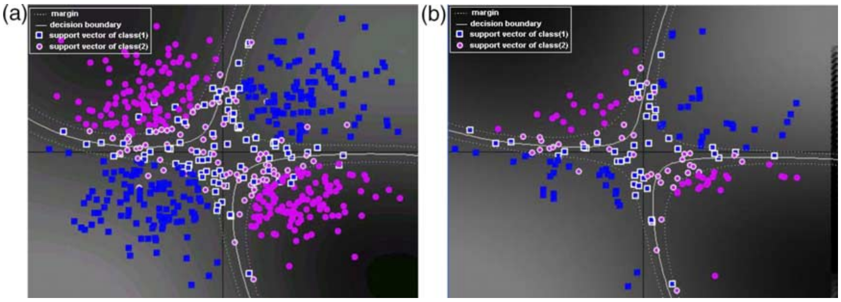

Fig. 5. Patterns and SVM decision boundaries of Continuous XOR Problem: decision boundary is depicted as a solid line and the margins are defined by the dotted lines in both sides of it. Support vectors are outlined.

Loss of those patterns could result in a poor generalization performance of SVM. Thus, instead, we propose to employ different misclassification costs to different class errors in the objective function, which is naturally allowed in SVM. This approach is not only safer, but also more principled.

Let m_1 and m_2 denote the size of class 1 and class 2 data sets, respectively, with $m_1 \gg m_2$ and $M = m_1 + m_2$. One way to alleviate data imbalance problem is to assign to a large class a smaller cost while assign to a small class a larger cost, which assures that a small class is not 'neglected.' In response modeling, there are many more non-respondents than respondents, thus the size of non-respondents is m1 while that of respondents is m_2. One way to accomplish it is to define and assign C_1 and C_2 to each class as below

$$\Theta(\vec{w}, \xi) = \frac{1}{2}||\vec{w}||^2 + C_1 \sum_{i \in \text{ nonrespondents}} \xi_i + C_2 \sum_{i \in \text{ respondents}} \xi_i, \tag{6}$$

where C_1 and C_2 are defined, respectively, as

$$C_1 = \frac{m_2}{M} \cdot C, \quad C_2 = \frac{m_1}{M} \cdot C. \tag{7}$$

In order to emphasize small respondent data set, a larger cost C_2 is assigned to its error term. Constant C is the original cost term used in Eq. (1).

4.3. Getting scores from an SVM classifier

The objective of response modeling is to compute the likelihood or propensity of each customer to respond to a particular offer so that the mailing response or profit is maximized. Lift chart is commonly used for this purpose, which sorts the customers by the descending order of their estimated value (score), and then the customers in the first several deciles are finally decided to be mailed. Although an SVM classifier returns a binary output (-1 or 1) as shown in Eq. (5), one can still estimate a score based on the distance between a pattern and the decision boundary. In other words, we assume that a pattern located further from the decision boundary has a higher probability of belonging to that class. The decision boundary hyperplane $\bar{f}(\vec{x})$ in a feature space Φ is represented as

$$\bar{f}(\vec{x}) = \sum_{i \in \mathbf{SVs}} y_i \alpha_i \Phi(\vec{x}_i) \cdot \Phi(\vec{x}) + b = 0 \tag{8}$$

from Eq. (5). It should be noted that the decision boundary is a hyperplane in the feature space Φ even though it is a nonlinear hyper-surface in the input space. In the feature space, hence, the distance from a pattern $\Phi(\vec{x})$ to the decision boundary hyperplane $\bar{f}(\vec{x})$ can be calculated by

$$\text{dist}(\Phi(\vec{x}), \bar{f}(\vec{x})) = \frac{|\bar{f}(\vec{x})|}{|\sum_{i \in \mathbf{SVs}} y_i \alpha_i \Phi(\vec{x}_i)|^2}. \tag{9}$$

The exact value of the distance is possible to obtain from Eq. (9) by using kernel trick (Vapnik, 1999) even though the actual mapping function $\Phi(\cdot)$ is not known: a kernel function $K(\vec{x}, \vec{x}/)$ replaces $\Phi(\vec{x}) \cdot \Phi(\vec{x}/)$ particularly during the denominator calculation. However, one does not need to know the exact value of the distance, since only a relative score or rank is all that is required in lift chart analysis. The denominator in Eq. (9) is common for all patterns, thus the signed function value in the numerator, $\bar{f}(\vec{x})$, can be used in computing ranks. The larger the value of $\bar{f}(\vec{x})$, the lower the rank of that particular customer's likelihood becomes.

5. Experiments

This section provides the empirical results of SVM based response modeling with the proposed approach. In particular, the performance evaluation measures pertinent to response modeling are also proposed and measured.

5.1. Dataset

In machine learning literature, so-called standard and public datasets are used. But, in response modeling, or in direct marketing for that matter, such datasets do not seem to exist. Many papers use a unique dataset, which is not available for other researchers. The only exception seems to be datasets from the Direct Marketing Educational Foundation (DMEF) (The Direct Marketing Association). The DMEF makes marketing datasets available to researchers. Dataset DMEF4, was used in various researches (Ha et al., 2005; Malthouse, 2001, 2002). It is concerned with an up-scale gift business that mails general and specialized catalogs to its customer base several times each year. The problem is to estimate how much each customer will spend during the test period, 09/1992–12/1992, based on the training period, 12/1971–06/1992. There are 101, 532 patterns in the dataset, each of which represents the purchase history information of a customer. Each customer is described by 91 input variables. A subset of 17 input variables, some original and some derived, were employed just as in (Malthouse, 2001) (see Table 2). The dataset has two target variables, target mailing dollars (TARGDOL) and target mailing orders (TARGORD). The former indicates the purchase dollar amount during the test period, and the latter indicates the number of orders during the test period. The TARGDOL or the TARGORD could be directly estimated by building a regression model. Malthouse (2001) built a regression model to estimate the value of TARGDOL. But due to the problems of regression (Section 2), we formulated the problem into a classification one. A new target variable, RESPONSE, was defined as follows: 1 if TARGDOL (TARGORD) >0, 0 otherwise. Ha et al. (2005) used the same derivation to fit a neural network classifier. Thus, all the customers were categorized into either a non-respondent (class 1) or a respondent (class 2). The response rate is 9.4%, which means the class distribution of the dataset is highly imbalanced.

Table 2
Input Variables

Variable	Formula	Description
Original variables		
Purseas		Number of seasons with a purchase
Falord		Life-to-date (LTD) fall orders
Ordtyr		Number of orders this year
Puryear		Number of years with a purchase
Sprord		LTD spring orders
Derived variables		
Recency		Order days since 10/1992
tran38	1/recency	
tran51	$0\le \text{recency}<90$	
tran52	$90\le \text{recency}<180$	
tran53	$180\le \text{recency}<270$	Five dummy variables (tran51–55) having the value 1, if the condition is satisfied, otherwise the value 0
tran54	$270\le \text{recency}<366$	
tran55	$366\le \text{recency}<730$	
comb2	$\sum_{i=1}^{14} \text{prodgrp } i$	Number of product groups purchased from this year
tran25	$1/(1+\text{lorditm})$	Inverse of latest-season items
tran42	$\log(1+\text{ordtyr}\times \text{falord})$	Interaction between the number of orders
tran44	$\sqrt{\text{ordhist}\times \text{sprord}}$	Interaction between LTD orders and LTD spring orders
tran46	$\sqrt{\text{comb 2}}$	

5.2. SVM models

To verify the effectiveness of NPPS described in Section 4.1, we considered seven SVMs trained with randomly selected patterns. They are denoted as R*-SVM where '*' indicates the ratio of random samples drawn without replacement. S-SVM denotes the SVM trained with the patterns selected by NPPS (see Table 3). Each model was trained and evaluated using five-fold cross-validation. The number of neighbors (k) of NPPS, was set to 4 according to guidelines suggested in Shin and Cho (2003c). All the SVM models in Table 3 use the same hyper-parameter values to equalize their effects. The RBF kernel in Eq. (2) was used with parameter σ set to 0.5, and the misclassification tolerance parameter C in Eq. (1) set to 10. These parameter settings were determined through a trial-error approach over the combination of C and σ, ($\{0.1, 1, 10, 100, 1000\}\times\{0.25, 0.5, 1, 2, 3\}$), using ten fold cross-validation performance. The class imbalance problem addressed in Section 4.2 appeared in all the eight datasets. The sets selected by random sampling showed the common class ratio of m_1:m_2 = 90.6:9.4%. That is also the same ratio as the original training set since we conducted a stratified random sampling by the target variable. The training set reduced by NPPS, however, showed a different class ratio, m_1:m_2 = 65.5:34.5% (=5810:3061) on average. Even though NPPS improved the ratio of the smaller class from 9.4% up to 34.5%, the imbalance problem still remained. Thus, the different misclassification costs, C_1 and C_2 were set on every dataset as they were defined in Eq. (7). C_1 and C_2 of R*-SVM were 0.94 ($=0.094\times 10$) and 9.06($=0.906\times 10$), respectively. On the other hands, those of S-SVM were 3.45($=0.345\times 10$) and 6.55($=0.655\times 10$).

5.3. Performance measurements

The performances of the eight SVM response models were compared in terms of three criteria: accuracies lift chart and computational efficiency.

Table 3
SVM models: the number of patterns selected from NPPS slightly varies with the given set of each fold, thus it is represented as an average over the five reduced training sets

Model	No. of training data	Training data
R05-SVM	4060	5% random samples
R10-SVM	8121	10% random samples
R20-SVM	16244	20% random samples
R40-SVM	32490	40% random samples
R60-SVM	48734	60% random samples
R80-SVM	64980	80% random samples
R100-SVM	81226	100% random samples
S-SVM	Avg. 8871	The patterns selected by NPPS

Table 4
Confusion Matrix: FP, FN, TP and TN means false positive, false negative, true positive, and true negative in due order where TP and TN are the correct classification

		Classified		
		class 1 (non-respondent)	class 2 (respondent)	
Actual	class 1(non-respondent)	M_{11} (TN)	m_{12} (FP)	m_1
	class 2 (respondent)	M_{21} (FN)	m_{22} (TP)	m_2

5.3.1. Accuracies

The accuracy of a classifier can be described by a confusion matrix (see Table 4). Let m$_{ij}$ denote the number of patterns, which were classified as class j but whose actual class label is class i. A most widely used accuracy measurement is an Average Correct-classification Rate (ACR), which is defined as:

$$\text{Average Correct classification Rate (ACR)} = \frac{\text{TN}+\text{TP}}{M}$$

$$= \frac{m_{11}+m_{22}}{M}.$$

But, the average correct-classification rate can be misleading in an imbalanced dataset where the heavily represented class is given more weight. Receiver Operating Characteristic (ROC) analysis is usually performed as well (Provost & Fawcett, 1997), which measures the classifier's accuracy over the whole range of thresholds in terms of Specificity (Sp) and Sensitivity (Se) (SAS Institute Inc., 1998). They are defined as

$$\text{Specificity (Sp)} = \frac{\text{TN}}{\text{TN}+\text{FP}} = \frac{m_{11}}{m_{11}+m_{12}} = \frac{m_{11}}{m_1},$$

$$\text{Sensitivity (Se)} = \frac{\text{TP}}{\text{FN}+\text{TP}} = \frac{m_{22}}{m_{21}+m_{22}} = \frac{m_{22}}{m_2}.$$

Since we fixed the classification threshold at 0 in the SVM decision function Eq. (5), however, only one pair of Sp and Se per model was available. Thus, here the ROC plot has the eight pairs of (1-Sp, Se) scattered for their comparison. Another accuracy measure, Balanced Correct-classification Rate (BCR), was defined so as to incorporate Sp and Se into one term. BCR enforces balance in the correct classification rate between two classes. It is defined as

$$\text{Balanced Correct} - \text{classification Rate (BCR)} = \text{Sp}\cdot\text{Se}$$

$$= \left(\frac{m_{11}}{m_1}\right)\cdot\left(\frac{m_{22}}{m_2}\right).$$

5.3.2. Lift chart analysis of response rate and profit

Once the test patterns were sorted in a descending order according to $f(\vec{x})$, two kinds of lift charts were investigated. One is for response rate, and the other for profit. From the business point of view, the ultimate goal of direct mailing is to maximize the profit rather than the response rate itself (Malthouse, 1999). Thus we evaluated the eight competing

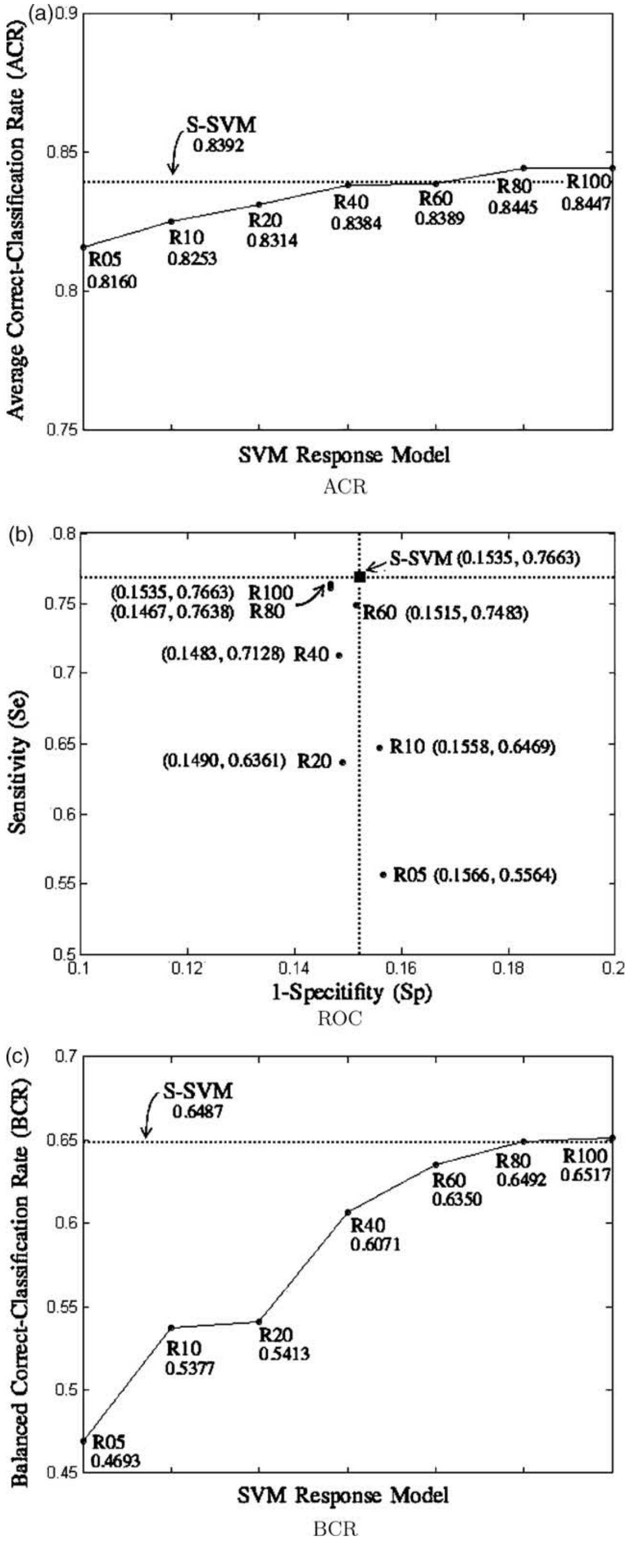

ACR

ROC

BCR

Fig. 6. Accuracies: accuracy of R*-SVM is depicted as a solid circle while that of S-SVM is represented as a dotted reference line.

SVM models from a profit aspect as well. For profit lift chart analysis, another target variable of DMEF4 dataset, *TARGDOL* (*target mailing dollar*), was associated with the rank of $f(\vec{x})$, which indicates the purchase dollar amount during the test period. Two measurements were used in evaluating lift charts. One is the average response rate or profit in the top decile, 'Top-Decile'. This measures how well two models identifies a small number of highly likely respondents. The other is 'Weighted-Decile' defined as

$$\text{Weighted}-\text{Decile} = \frac{\{1.0\times d_1+0.9\times d_2+0.8\times d_3\cdots+0.1\times d_{10}\}}{1.0+0.9+0.8+\cdots+0.1},$$

where d_i, ($i=1,\ldots,10$) is a cumulative average response rate or profit till *i*th decile in the lift table. This measures how well the model identifies a larger number of likely respondents in a larger rollout. A similar evaluation by two measurements has been adopted in data mining competitions (Ling & Li, 1998).

5.3.3. Computational efficiency

The evaluation was done in several measures: the number of training patterns, training time, and the number of support vectors, and recall time. The number of patterns directly influences the time complexity. The training time of SVM increases in proportion to the cube of the number of training patterns (in case of standard QP solver). The recall time increases linearly to the number of support vectors. Training time is of important concern to a direct marketer who is in charge of SVM modeling with a huge amount of data, while recall time is critical when the model is deployed to work in a real-time application such as fraud detection. Although recall time is not a primary issue in response modeling, we measured it for potential use to another application.

5.4. Results

We now give the experimental results of the eight SVM response models in the order of accuracies, lift chart analysis, and computational efficiency.

Fig. 6 shows how the eight SVM response models performed in terms of ACR, ROC, and BCR. First, Fig. 6(a) indicates a mean ACR over five-fold cross-validation of each SVM model. For the sake of convenience, *R**-SVM is briefly denoted as '*R**' in the figure. Sampling more patterns results in higher ACR, but the increasing rate is not very high. From R05 to R100, only about 3.52% ($=\{0.8447-0.8160\}/0.8160\times100\%$) of accuracy was gained from 1.900% ($=\{100-5\}/5\times100\%$) data increase. The S-SVM achieved ACR in the range of those from R60-R80. However, we could not make good evaluation of the model comparison using ACR because of class imbalance. In Fig. 6(b), the eight pairs of (1-Sp, Se) were plotted in ROC chart. A point located upper left corresponds to a better performance. The ACR is effectively broken down into two classwise accuracies, Sp for non-respondents (class 1) and Se for respondents (class 2). The Sps of the eight SVM models are similar, while the Ses shows significant differences. It should be noted that it is Se, accuracy for respondents' group, that is of greater importance to direct marketers, since their primary goal is to identify the respondents, not the non-respondents. S-SVM achieved a best Se, better than that of even R100-SVM. Fig. 6(c) shows the BCRs of the eight SVM response models. BCR clearly distinguished the accuracies of the eight SVM models. Sampling more data results in a larger BCR also. The BCR of S-SVM is almost same as that of R100-SVM.

Fig. 7 illustrates the lift chart of the cumulative average response rate. The base average response rate of DMEF4 dataset was 9.4%, which is represented as a solid horizon at the bottom of the chart. Two observations can be made. First, all the SVM response models did better than the base response rate. Second, more training patterns lead to a better

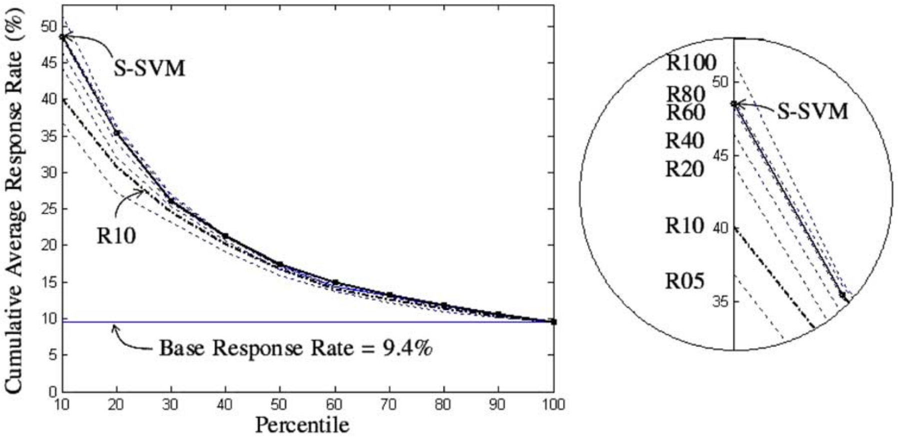

Fig. 7. Lift chart of cumulative average response rate: *R**-SVMs are depicted dotted lines but among them R10-SVM is represented as a dash-dot line. S-SVM is represented as a solid-dot line.

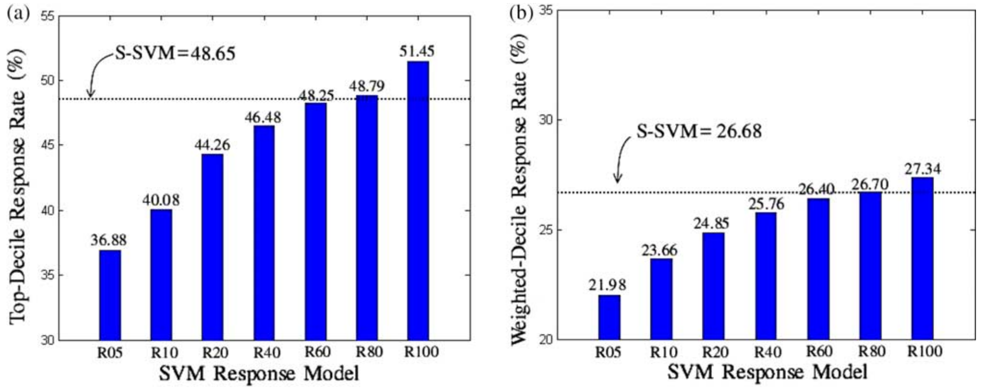

Fig. 8. Top-Decile response rate and Weighted-Decile response rate: R^*-SVM is depicted as a bar while S-SVM is represented as a dotted reference line.

lift of the response rate. R100-SVM showed the best performance while the R05-SVM showed the worst. Models trained with more patterns showed a steeper lift in the first several deciles. The lift curve of S-SVM was almost identical to that of R80-SVM. It is illuminating to compare the curve shape of S-SVM with that of R10-SVM represented as a dash-dot line. Although the two models had almost the same number of training patterns, they were significantly different in the lift performance. Fig. 8 shows the results of the lift measures described in Section 5.3.2: Top-Decile and Weighted-Decile. From the top 10 percentile of customers, R100-SVM obtained 51.45% response rate (see, Fig. 8(a)). The Top-Decile response rate of S-SVM was 48.65%, which is almost equal to that of R80-SVM, 48.79%. Fig. 8(b) shows the results of Weighted-Decile response rates. R100-SVM still did best, and S-SVM and R80-SVM came second. But the gap between the first and the second was not so big as in the Top-Decile response rate.

Now, Figs. 9 and 10 describe the lift chart results in terms of the profit. The average purchase dollar amount of DMEF4 was $48 when averaged over the respondents' group, but $4.5 when averaged over all customers. The horizon line in the lift chart of Fig. 9 represents the $4.5 base average profit. All the models did better than the base average profit and an SVM with more training pattern produced a higher profit in the first several deciles. But in terms of the profit lift, S-SVM showed a performance comparable to that of R100-SVM. It is also remarkable that the profit lifts of R100-SVM or S-SVM outperformed those of Malthouse who got the scores by fitting the problem as a regression one (Malthouse, 2001). For the cumulative average profit (dollars) of the second decile, Malthouse' regression model recorded $12–$15 while the SVM classification model recorded $17–$18. Fig. 10 illustrates the Top-Decile profit and the Weighted-Decile profit. The Top-Decile profit and the Weighted-Decile profit of R100-SVM were $23.78 and $12.99, respectively, and those of R80-SVM were $22.25 and $12.56. S-SVM was $23.53 in the Top-Decile profit and $12.77 in the Weighted-Decile profit, which were slightly less than those of R100-SVM but more than those of R80-SVM.

Finally, Table 5 shows the results of computational efficiency measures in columns: the number of training patterns, training time, the number of support vectors, its proportion to training patterns, and recall time. We used OSU SVM Classifier Matlab Toolbox, which is a hybrid algorithm of SMO and SVM^{light}, and is known as one of

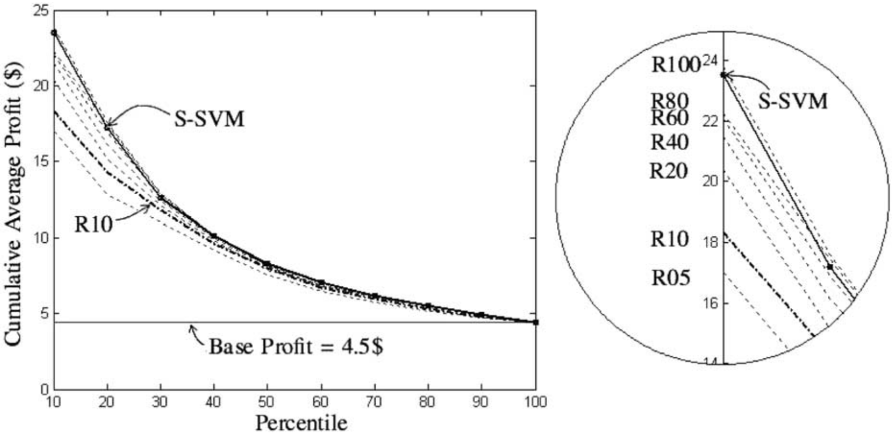

Fig. 9. Lift chart of cumulative average profit: R^*-SVMs are depicted dotted lines but among them R10-SVM is represented as a dash-dot line. S-SVM is represented as a solid-dot line.

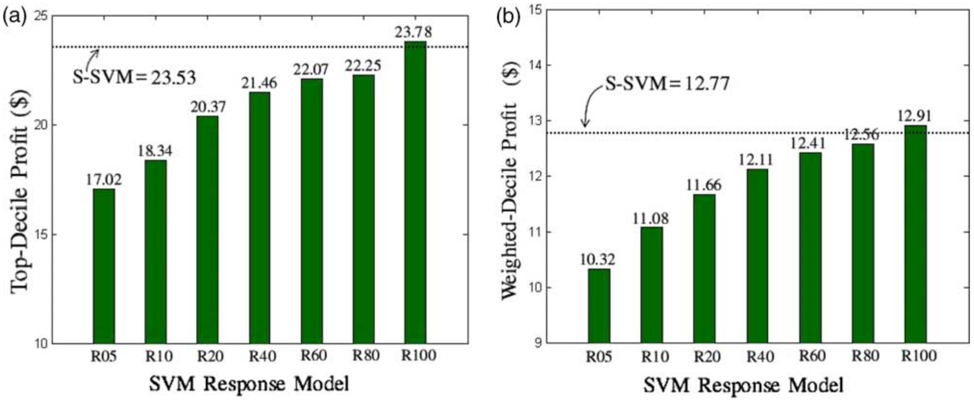

Fig. 10. Top-Decile profit and Weighted-Decile profit: *R**-SVM is depicted as a bar while S-SVM is represented as a dotted reference line.

Table 5
Computational efficiency of SVM response models

	Num. of training patterns	Training time (s)	Num. of SVs (proportion)	Recall time (s)
R05	4060	13.72	1975 (48.65%)	17.22
R10	8121	56.67	4194 (51.64%)	31.39
R20	16,244	149.42	7463 (45.94%)	56.17
R40	32,490	652.11	14,967 (46.07%)	112.08
R60	48,734	1,622.06	22,193 (45.54%)	166.11
R80	64,980	2,906.97	28,968 (44.58%)	237.31
R100	81,226	4,820.06	35,529 (43.74%)	381.31
S	8871	68.29	6,624 (74.67%)	45.13

the fastest solvers (Hearst, Schölkopf, Dumais, Osuna, & Platt, 1997). Training time increased proportionally to the number of training patterns with the peak of 4820 (sec) for R100-SVM. On the other hand, S-SVM took only 68 (sec). The total time of S-SVM was 129 (sec), when the NPPS running time, 61 (sec), was included. Note that SVM training is usually performed several times to find a set of optimal parameters, but the pattern selection is performed only once. In the fourth column, the number of support vectors is represented. At most, half of the random sampling training patterns were support vectors while 74% of the NPPS selected training patterns were support vectors. The result confirms that the NPPS' selection of training patterns was more efficient. Recall time was proportional to the number of support vectors as shown in the last column. Overall, the computational efficiency of S-SVM was comparable to that of R10-SVM or R20-SVM.

6. Conclusions and discussions

In this paper, we introduced SVM, a powerful classification model, and practical difficulties when applying it to response modeling in direct marketing: large training data, class imbalance and binary SVM output. We then proposed how to alleviate and solve those difficulties: informative sampling, different costs for different classes, and use of distance to decision boundary. In the experiments, we showed that the proposed solutions worked quite well. In particular, several models were trained and evaluated in terms of accuracies, lift chart analysis and computational efficiency. The SVM trained with the patterns selected by proposed NPPS (S-SVM) were compared with the ones trained with random samples (*R**-SVMs where '*' indicates the sampling percentage). Fig. 11 summarizes the results in terms of various measures. The horizontal bars in the figure shows the performance of S-SVM relative to those of R*-SVMs in various measures. S-SVM achieved the accuracies and uplifts comparable to those of R80-SVM and R100-SVM with a computational cost comparable to those of R10-SVM and R20-SVM.

Here, we would like to address some future research works. First, in lift chart analysis, we used two measures, Top-Decile and Weighted-Decile. The former is for specifying a small number of customers in the top decile, while the latter is for covering a larger number of customers in all deciles. If the mailing depth is optimized through a break-even analysis between revenue and cost, then more accurate and practical evaluation measure needs to be created. Second, the proposed pattern selection algorithm, NPPS, can also be utilized to reduce the lengthy training time of neural network classifiers. But it is necessary to add extra correct patterns to the selected pattern set in order to enhance the overlap region near the decision boundary (Choi & Rockett, 2002; Hara & Nakayama, 2000).

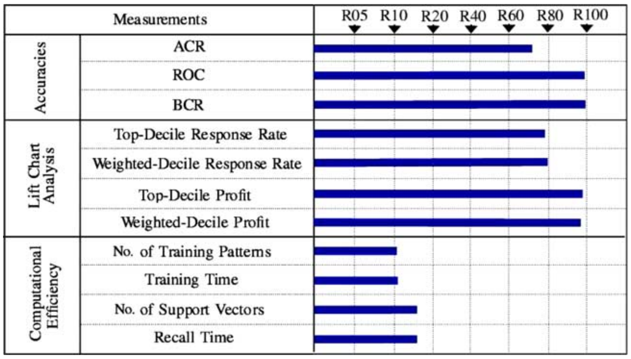

Fig. 11. How well S-SVM performed relative to *R**-SVMs.

The rationale is that ‘overlap patterns’ located on the ‘wrong’ side of the decision boundary cause the MLP training to take a longer time. Since the derivatives of the back-propagated errors are evaluated at those patterns, the derivatives are very small if they are grouped in a narrow region on either side of the decision boundary. By means of adding extra correct patterns, however, the network training converged faster. Third, the current version of NPPS works for classification problems only, thus is not applicable to regression problems. In a regression problem, the patterns located away from others, such as outliers, are less important to learning. Thus, a straightforward idea would be to use the mean (μ) and variance (Σ) of k nearest neighbors’ outputs. A pattern having a small value of Σ can be replaced by μ of its neighbors and itself. That is, $k+1$ patterns can be replaced by one pattern. On the contrary, a pattern having a large value of Σ can be totally eliminated, and its neighbors will be used for the next pattern searching. A similar research was conducted in (Shin & Cho, 2001) based on ensemble neural network, but more extended study based on k nearest neighbors is still under consideration. Regression NPPS will also be helpful for direct marketing problems with large datasets.

References

Almeida, M. B., Braga, A., & Braga, J. P. (2000). *SVM-KM: Speeding SVMs learning with a priori cluster selection and k-means Proceedings of the sixth brazilian symposium on neural networks* pp. 162–167.

Bentz, Y., & Merunkay, D. (2000). Neural networks and the multinomial logit for brand choice modeling: A hybrid approach. *Journal of Forecasting*, *19*(3), 177–200.

Bounds, D., & Ross, D. (1997). Forecasting customer response with neural networks. *Handbook of neural computation*, *G6.2*, 1–7.

Byun, H., & Lee, S. (2002). *Applications of support vector machines for pattern recognition: A survey International workshop on pattern recognition with support vector machines (SVM2002), lecture notes in computer science (LNCS 2388), Niagara Falls, Canada* pp. 213–236.

Cauwenberghs, G., & Poggio, T. (2001). *Incremental and decremental support vector machine learning Advances in neural information processing systems*, Vol. 13. Cambridge MA: MIT Press pp. 409–415.

Cheung, K.-W., Kwok, J. K., Law, M. H., & Tsui, K.-C. (2003). Mining customer product rating for personalized marketing. *Decision Support Systems*, *35*, 231–243.

Chiu, C. (2002). A case-based customer classification approach for direct marketing. *Expert Systems with Applications*, *22*, 163–168.

Choi, S. H., & Rockett, P. (2002). The training of neural classifiers with condensed dataset. *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics*, *32*(2), 202–207.

Coenen, F., Swinnen, G., Vanhoof, K., & Wets, G. (2000). The improvement of response modeling: Combining rule-induction and case-based reasoning. *Expert Systems with Applications*, *18*, 307–313.

Colombo, R., & Jiang, W. (1999). A stochastic RFM model. *Journal of Interactive Marketing*, *13*(3), 1–12.

Cristianini, N., & Shawe-Taylor, J. (2000). *An introduction to support vector machines and other kernel-based learning methods*. Cambridge: Cambridge University Press.

Deichmann, J., Eshghi, A., Haughton, D., Sayek, S., & Teebagy, N. (2002). Application of multiple adaptive splines (MARS) in direct response modeling. *Journal of Interactive Marketing*, *16*(4), 15–27.

Dumais, S. (1998). Using SVMs for text categorization. *IEEE Intelligent Systems* , 21–23.

Freund, Y., & Schapire, R. (1996). *Experiments with a new boosting algorithm Proceedings of the thirteenth international conference on machine learning* pp. 148–156.

Ha, K., Cho, S., MacLachlan, D. (2005) *Response models based on bagging neural networks*. Submitted for publication. Journal of Interactive Marketing 19(1), 17–30.

Hara, K., & Nakayama, K. (2000). *A training method with small computation for classification Proceedings of the IEEE-INNS-ENNS international joint conference*, Vol. 3 pp. 543–548.

Haughton, D., & Oulabi, S. (1997). Direct marketing modeling with CART and CHAID. *Journal of Direct Marketing*, *11*(4), 42–52.

Hearst, M. A., Schölkopf, B., Dumais, S., Osuna, E., & Platt, J. (1997). Trends and controversies—support vector machines. *IEEE Intelligent Systems*, *13*, 18–28.

Heisele, B., Poggio, T., Pontil, M. (2000). Face detection in still gray images, *Technical report AI memo 1687*, MIT AI Lab.

Hosmer, D. W., & Lemeshow, S. (1989). *Applied logistic regression*. London: Wiley.

Japkowicz, N. (2000). *Learning from imbalanced data sets: A comparison of various strategies AAAI workshop on learning from imbalanced data sets*. Menlo Park, CA: AAAI Press.

Joachims, T. (1998). *Text categorization with support vector machines: Learning with many relevant features Proceedings of 10th European conference on machine learning* pp. 137–142.

Lee, K. K., Gunn, S. R., Harris, C. J., & Reed, P. A. S. (2001). *Classification of imbalanced data with transparent kernels Proceedings of INNS-IEEE international joint conference on neural networks* pp. 2410–2415.

Ling, C. X., & Li, C. (1998). *Data mining for direct marketing: Problems and solutions Proceedings of ACM SIGKDD international conference on knowledge discovery and data mining (KDD-98)* pp. 73–79.

Liu, C. L., & Nakagawa, M. (2001). Evaluation of prototype learning algorithms for nearest-neighbor classifier in application to handwritten character recognition. *Pattern Recognition*, *34*, 601–615.

Lyhyaoui, A., Martinez, M., Mora, I., Vazquez, M., Sancho, J., & Figueiras-Vaidal, A. R. (1999). Sample selection via clustering to construct support vector-like classifiers. *IEEE Transactions on Neural Networks*, *10*(6), 1474–1481.

Malthouse, E. C. (1999). Ridge regression and direct marketing scoring models. *Journal of Interactive Marketing*, *13*(4), 10–23.

Malthouse, E. C. (2001). Assessing the performance of direct marketing models. *Journal of Interactive Marketing*, *15*(1), 49–62.

Malthouse, E. C. (2002). Performance-based variable selection for scoring models. *Journal of Interactive Marketing*, *16*(4), 10–23.

Moghaddam, B., & Yang, M. H. (2000). *Gender classification with support vector machines Proceedings of international conference on pattern recognition, Barcelona, Spain, 2000, and also appeared in proceedings of 4th IEEE international conference on automatic face and gesture recognition, Grenoble, France* pp. 306–311.

Moutinho, L., Curry, B., Davies, F., & Rita, P. (1994). *Neural network in marketing*. New York: Routledge.

Osuna, E., Freund, R., & Girosi, F. (1997). *Training support vector machines: An application to face detection Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* pp. 130–136.

Platt, J. C. (1999). *Fast training of support vector machines using sequential minimal optimization Advances in kernel methods: Support vector machines*. Cambridge, MA: MIT press pp. 185–208.

Pontil, M., & Verri, A. (1998). Properties of support vector machines. *Neural Computation*, *10*, 955–974.

Potharst, R., Kaymak, U., Pijls W. (2001). Neural networks for target selection in direct marketing, provided by Erasmus Research Institute of Management (ERIM), Erasmus University Rotterdam in its series Discussion Paper with number 77, http://ideas.repec.org/s/dgr/eureri.html.

Provost, F., & Fawcett, T. (1997). *Analysis and visualization of classifier performance: Comparison under imprecise class and cost distributions Proceedings of third international conference on knowledge discovery and data mining*. Menlo Park, CA: AAAI press pp. 43–48.

SAS Institute Inc. (1998). Enterprise mining premier.

Sen, A., & Srivastava, M. (1990). *Regression analysis: Theory, method, and applications, (Springer texts in statistics)*. New York: Springer.

Schölkopf, B., Burges, C. J. C., & Smola, A. J. (1999). *Advances in kernel methods: Support vector learning*. Cambridge, MA: MIT press.

Schölkopf, B., Burges, D., & Vapnik, V. (1995). *Extracting support data for a given task Proceedings of first international conference on knowledge discovery and data mining*. Menlo Park, CA: AAAI press pp. 252–257.

Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels: Support vector machines, regularization, optimization, and beyond*. Cambridge, MA: MIT press.

Shin, H. J., & Cho, S. (2001). Pattern selection using the bias and variance of ensemble. *Journal of the Korean Institute of Industrial Engineers*, *28*(1), 112–127.

Shin, H. J., Cho, S. (2002). Pattern selection for support vector classifiers, The third international conference on intelligent data engineering and automated learning (IDEAL), lecture notes in computer science (LNCS 2412), Manchester, UK, pp. 469–474.

Shin, H. J., & Cho, S. (2003a). *Fast pattern selection for support vector classifiers Proceedings of the seventh Pacific-Asia conference on knowledge discovery and data mining (PAKDD), Lecture notes in artificial intelligence (LNAI 2637), Seoul, Korea* pp. 376–387.

Shin, H.J., Cho, S. (2003). Fast pattern selection algorithm for support vector classifiers: Time complexity analysis, The fourth international conference on intelligent data engineering and automated learning (IDEAL), lecture notes in computer science (LNCS 2690), Hong Kong, China, pp. 1008–1015.

Shin, H. J., & Cho, S. (2003c). *How many neighbors to consider in pattern pre-selection for support vector classifier? Proceedings of the international joint conference on neural networks (IJCNN), Portland, U.S.A.* pp. 565–570.

Suh, E. H., Noh, K. C., & Suh, C. K. (1999). Customer list segmentation using the combined response model. *Expert Systems with Applications*, *17*(2), 89–97.

The Direct Marketing Association. http://www.the-dma.org/dmef/dmefdset.shtml.

Vapnik, V. (1999). *The nature of statistical learning theory* (2nd ed.). Berlin: Springer.

Viaene, S., Baesens, B., Van den Poel, D., Dedene, G., & Vanthienen, J. (2001a). Wrapped input selection using multilayer perceptrons for repeat-purchase modeling in direct marketing. *International Journal of Intelligent Systems in Accounting, Finance & Management*, *10*, 115–126.

Viaene, S., Baesens, B., Van Gestel, T., Suykens, J. A. K., Van den Poel, D., Vanthienen, J., et al. (2001b). Knowledge discovery in a direct marketing case using least squares support vector machines. *International Journal of Intelligent Systems*, *16*, 1023–1036.

Zahavi, J., & Levin, N. (1997a). Issues and problems in applying neural computing to target marketing. *Journal of Direct Marketing*, *11*(4), 63–75.

Zahavi, J., & Levin, N. (1997b). Applying neural computing to target marketing. *Journal of Direct Marketing*, *11*(4), 76–93.