Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT

PERGAMON

Pattern Recognition 37 (2004) 1245–1265

PATTERN RECOGNITION
THE JOURNAL OF THE PATTERN RECOGNITION SOCIETY

www.elsevier.com/locate/patcog

Analysis of error-reject trade-off in linearly combined multiple classifiers

Giorgio Fumera, Fabio Roli*

Department of Electrical and Electronic Engineering, University of Cagliari, Piazza d'Armi, Cagliari 09123, Italy

Received 21 January 2003; accepted 15 December 2003

Abstract

In this paper, a theoretical and experimental analysis of the error-reject trade-off achievable by linearly combining the outputs of an ensemble of classifiers is presented. To this aim, the theoretical framework previously developed by Tumer and Ghosh for the analysis of the simple average rule without the reject option has been extended. Analytical results that allow to evaluate the improvement of the error-reject trade-off achievable by simple averaging their outputs under different assumptions about the distributions of the estimation errors affecting a posteriori probabilities, are provided. The conditions under which the weighted average can provide a better error-reject trade-off than the simple average are then determined. From the theoretical results obtained under the assumption of unbiased and uncorrelated estimation errors, simple guidelines for the design of multiple classifier systems using linear combiners are given. Finally, an experimental evaluation and comparison of the error-reject trade-off of the simple and weighted averages is reported for five real data sets. The results show the practical relevance of the proposed guidelines in the design of linear combiners.

Keywords: Multiple classifier systems; Classifier fusion; Linear combiners; Reject option; Error-reject trade-off

1. Introduction

During the past decade, many research communities, among others the pattern recognition community and the machine learning community, have shown a growing interest in the so-called Multiple Classifier Systems (MCSs) [1]. It is now widely accepted that a combination of multiple classifiers can provide advantages over the traditional monolithic approach to classifier design. Besides the many experimental works showing the improvement in performance that can be achieved by MCSs in several applications, a few works have also provided theoretical analyses of the simplest combining techniques proposed in the literature. For instance, Tumer and Ghosh [2,3] developed a theoretical framework for analysing the performance improvement achievable by the simple average of classifier outputs. A theoretical analysis of the majority voting rule was provided by Lam and Suen [4]. Kittler et al. [5] developed a theoretical framework for the combination of classifiers that use distinct pattern representations. Kuncheva [6] compared the classification error at a given point in the feature space, for majority voting, simple average, and order statistics rules. Kittler and Alkoot [7] compared the sum and majority vote rules theoretically and by experiments. Despite these important works, a general theoretical framework for classifier combination is currently beyond the state of the art [1]. Consequently, many important topics lack theoretical explanations, and a comparison of various fusion rules is only possible by experiments. Further theoretical analyses aimed at investigating these topics and comparing a limited set of fusion rules, though under strict assumptions, are necessary steps towards a general framework of classifier fusion.

* Corresponding author. Tel.: +39-070-675-5779; fax: +39-070-675-5782.
E-mail addresses: fumera@diee.unica.it (G. Fumera), roli@diee.unica.it (F. Roli).

0031-3203/$30.00 © 2003 Pattern Recognition Society. Published by Elsevier Ltd. All rights reserved.
doi:10.1016/j.patcog.2003.12.005

In this paper, we address a topic that has not been considered from a theoretical viewpoint, i.e. the improvement of the error-reject trade-off achievable by classifier combination. Theoretical works, like the ones quoted above, have analysed the performance of MCSs only in terms of error probability, without taking into account the reject option. Only few experimental works have evaluated the performance of MCSs with the reject option. For instance, Battiti and Colla [8] have investigated the error-reject trade-off provided by MCSs using majority voting rule or linear combiners by experiments. Lam and Suen's experiments [9] analysed MCSs using the Bayesian and the weighted majority voting rules. The experimental work by Foggia et al. [10] dealt with the Bayesian rule.

In this work we focus on linear combiners, one of the simplest and most widely used combining techniques, and analyse the theoretical and experimental improvement of the error-reject trade-off that can be achieved by classifier combination. The main purpose of this paper is to provide an analytical evaluation of the improvement in error-reject trade-off achievable by the linear combination of multiple classifiers. To this end, we have extended the analytical framework developed by Tumer and Ghosh [2,3]. This framework allows to evaluate the error probability, without the reject option, achievable by simple averaging the outputs of classifiers that provide estimates of the class posterior probabilities. In previous works [11,12], we used this framework to compare the performance of the simple and weighted average rules. In this work, we extend the framework to the evaluation of the expected risk of individual classifiers, and of the linear combination of multiple classifiers. This allows us to assess and compare the performances of individual and linearly combined classifiers when the reject option is used. A preliminary analysis has been presented by the authors in Ref. [13]. In this paper, we extend both the theoretical analysis and the experimental investigation. Furthermore, we address the problem of obtaining practical guidelines for the design of linear combiners with the reject option, based on the results obtained from the analysis of our theoretical framework.

The paper is organised as follows. In Section 2, the theoretical background of statistical classification with the reject option is reviewed briefly. In Section 3, the framework by Tumer and Ghosh is summarised, and our extension to classification with the reject option described. The quantitative analysis of the error-reject trade-off achievable by the simple average and weighted average combining rules is presented in Section 4. This analysis allows to compare the error-reject trade-off of these two rules. In Section 5, we show that the theoretical framework suggests some practical guidelines for the design of a linear combiner with the reject option. The results of an experimental comparison, guided by the analysis in Sections 4 and 5, are reported in Section 6. Conclusions are drawn in Section 7.

2. Classification with the reject option in statistical pattern recognition: basic concepts

It is well known that the reject option is useful in improving classification reliability in pattern recognition applications for which the cost of rejecting certain patterns and handling them with different procedures (e.g., manual classification) is lower than that of wrong classifications. In the framework of the minimum risk theory, Chow [14,15] defined the optimal classification rule with the reject option described in what follows.

Consider a classification problem with a loss function defined by costs c_C, c_R, and c_E, which indicate the costs of correct classification, rejection, and misclassification, respectively. It is easy to see that the meaningful values of the costs are the ones that satisfy the constraints $c_C < c_R < c_E$. For any given classification rule, the expected risk (i.e., the expected value of the classification cost of any given pattern) is [15]:

$$c_C P(correct) + c_R P(reject) + c_E P(error),$$

where $P(correct)$, $P(reject)$, and $P(error)$ denote the probability of correct classification, rejection, and misclassification, respectively. Since they sum up to 1, it follows that the added risk can be rewritten as

$$(c_R - c_C)P(reject) + (c_E - c_C)P(error) \tag{1}$$

Since the error probability can be reduced only by increasing the reject probability, it follows from (1) that to minimise the expected risk, a trade-off between the error and reject probabilities must be found, as a function of classification costs. Chow proved that the optimal classification rule with the reject option (that is, the rule which provides the best error-reject trade-off) is the following. Given an input pattern $\mathbf{x}$, consider the maximum of its class posterior probabilities $P_i(\mathbf{x}) = \max_j P_j(\mathbf{x})$ (the subscript indicates the class). Then, assign $\mathbf{x}$ to the ith class, if $P_i(\mathbf{x}) \geqslant T$; otherwise, reject $\mathbf{x}$. The term T is the so-called reject threshold, whose value is given by

$$T = \frac{c_E - c_R}{c_E - c_C} \tag{2}$$

(see Ref. [15]). It is easy to see that the above rule is a generalisation of the standard Bayes rule. In particular, for the extreme case in which the cost c_R of rejections equals the cost c_E of misclassifications, the reject threshold equals zero; in this case, no pattern is rejected, and Chow's rule reduces to the standard Bayes rule.

In principle, the error-reject trade-off of a classifier must be evaluated in terms of the expected risk (1) for varying values of the classification costs. However, it has been shown that minimising the expected risk for any value of the classification costs is equivalent to maximising classification accuracy for any value of the reject probability [15].