An Information Market Based Approach for Decision Fusion

Johan Perols, Kaushal Chari, Manish Agrawal
University of South Florida

Overview

In this project, we are designing and experimentally testing a novel multi-classifier combination combiner method for decision fusion. Our combiner method, referred to as Information Market based Fusion (IMF), is based on pari-mutual betting markets. Using computational experiments involving 17 datasets from the UCI Machine Learning Repository and 22 different base-classifiers from Weka, we have initial results that show that IMF significantly outperforms Majority (MAJ), Average (AVG) and Weighted Average (WAVG) combiner method schemes. We are currently performing sensitivity analysis.

Introduction

Multi-classifier combination (MCC) is a technique that can be used to improve the classification performance in various classification problems by combining the decisions of multiple individual classifiers (Kittler and Roli 2000). In MCC, individual classifiers, commonly referred to as base-classifiers, classify objects based on inputs consisting of object feature vectors (see Figure 1). These classifications or decisions are then combined using a combiner method into a single decision about an object's class label. The basic idea behind MCC is that different classifiers in an ensemble have different strengths and weaknesses, and therefore provide complementary information about the classification problem.

Different combiner methods have been proposed and examined in the literature, see (Suen and Lam 2000) for details on various combiner methods in the literature. Results from the literature indicate that MCC overall provides performance benefits, and that MAJ and AVG perform either at a similar level or significantly better than trained methods (Kittler, et al. 1998). Existing combiner methods that require training however assume that ensemble base-classifier composition remains constant and that training performance is a good proxy for subsequent actual performance. They furthermore are not designed to integrate with software agent architectures.

The development of software agent technology offers a new framework and infrastructure to support decision-making (Nissen and Sengupta 2006) where human experts driving software agents as well as autonomous software agents embodying classifiers and other intelligent algorithms can leverage their individual strengths to make collective decisions. In such decision-making endeavors it is important that humans as well as software agent experts are provided incentives to truthfully provide their actual classification decisions, and that the decision fusion mechanism integrates well with the coordination mechanism used. Existing combiner methods were not designed with agents-based decision-making platforms in mind and therefore do not provide these two features.

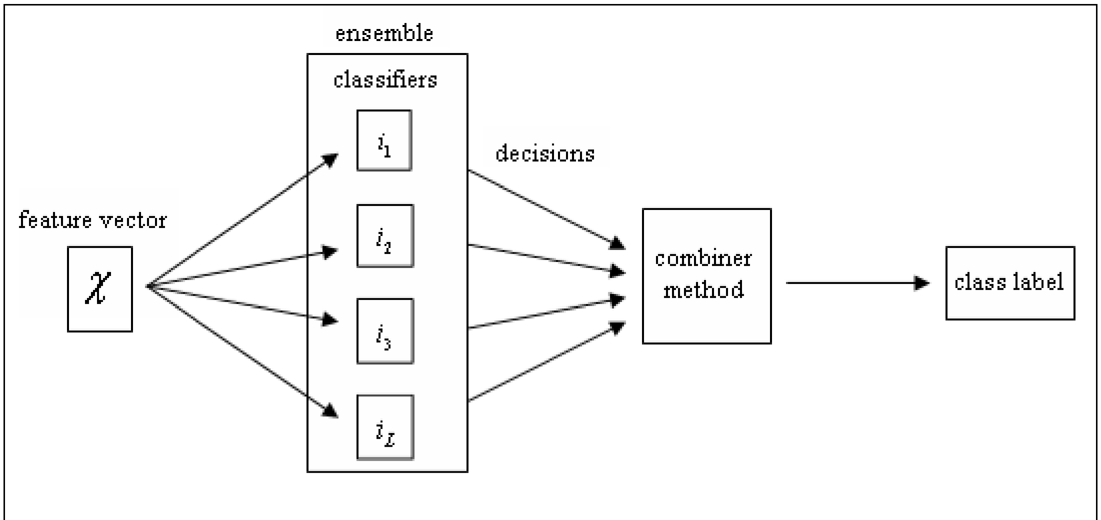

Figure 1. Generic Classifier Combiner Architecture

Our research objective is to design a relatively effective MCC combiner method that does not require training, can adapt to changes in ensemble composition and base-classifier accuracy, does not assume that the base-classifiers are cooperative agents and integrates with various MAS coordination mechanisms. To accomplish our research objective we propose a market based approach for MMC that we describe next.

Information Market Based Fusion

The combiner method proposed in this paper is theoretically grounded in information markets. Information markets are markets designed specifically for the purpose of information aggregation. Equilibrium prices, derived using conventional market mechanisms, provide information based on the private and public information maintained by the market participants about a specific situation, future event or object of interest (Hanson 2003). The aggregation mechanism used in IMF is based on Pari-mutuel betting, which is primarily used in horse race gambling. In Pari-mutuel betting the final track odd for a given horse is the total amount bet on all horses in the race divided by the total amount bet on the given horse. From a MCC perspective these odds are of great importance as they represent aggregated information where the inverted odd for a specific horse can be viewed as the probability estimate of that horse winning the race.

Assuming an ensemble E of m agents embodying different base-classifiers (agents) represented by indices i in the index set $D = \{1,..., m\}$ where any agent $i \in D$ uses the feature vector associated with t to determine the posterior probability estimate p_{itj} in [0, 1] that t belongs to class $j \in J,\ J = \{1, 2\}$. Agent i in E uses p_{itj} to determine the amount to bet q_{itj} on object class j and is paid following the pari-mutuel mechanism, based on q_{itj} and the total bets placed by all the agents Q_{tj}, and the true class of object t. Ensemble E's overall probability estimate that t belongs to $j \in J$ is given by $1/O_{tj}$ in [0,1], where O_{tj} is the odd that t belongs to j when O_{tj} is in equilibrium. O_{tj} is in equilibrium when the potential payouts $Q_{tj}O_{tj}$ for each $j \in J$ and the total amount bet on all events Q_t are equal, assuming house commission is zero.

Establishing equilibrium odds is however a nontrivial task, given that there is dependency between Q_{tj} and O_{tj}. In IMF, see figure 2, O_{tj} is found for each object t by iteratively updating the odds, and agents placing bets using these odds until the odds provided to the agents and their subsequent bets result in $Q_{tj}O_{tj} = Q_t$, at which time the market closes. A few observations are made here. First, bets placed in finding the equilibrium odds are only used for the purpose of updating the odds. Second, if agent bets are discontinuous over O_{tj} then the existence of equilibrium odds cannot be guaranteed (Carlsson, et al. 2001). To overcome oscillation of odds over multiple iterations, and find a good solution when no optimal solution exists, we use a combination of binary search and optimization.

The ensemble's overall probability estimate $1/O_{tj}$ is then compared to a cut-off value C_j to determine if t should be classified as belonging to class j. In the context of fraud detection application, if the ensemble classifies t as fraudulent, the agents are asked to place their final bets q_{itj} and after determining the true class of object t the winnings are distributed to the agents according to the pari-mutuel betting mechanism. To determine the amount to bet on classes j = 1, 2 and on the sure bet q_{it3} (not betting on either outcome) agent i solves an expected utility maximization problem given the current market odds O_{tj}, the agent's probability estimates p_{itj}, and the agent's current wealth w_{it}.

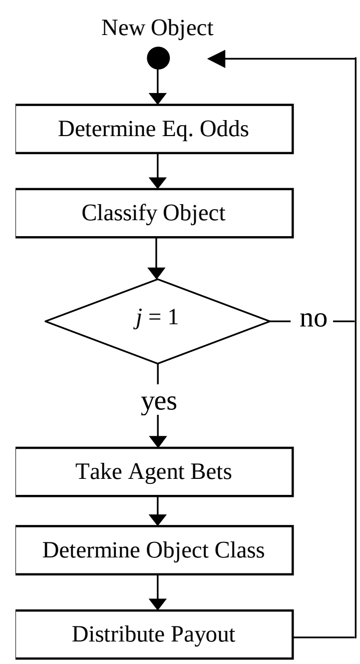

Figure 2. IMF Combiner Method

References Available Upon Request