Effective Supra-Classifiers for Knowledge Base Construction ⋆

Kurt D. Bollacker and Joydeep Ghosh

Dept of Electrical and Computer Engineering
University of Texas at Austin
Austin, TX 78712
{kdb,ghosh}@lans.ece.utexas.edu

Abstract

In this work, we explore the use of the *supra-classifier framework* in the construction of a classifier knowledge base. Previously, we introduced this framework within which labels produced by old classifiers are used to improve the generalization performance of a new classifier for a different but related classification task (Bollacker & Ghosh, 1998). We showed empirically that a simple *Hamming nearest neighbor* is superior to other techniques (e.g. MLP, decision trees, Naive Bayes, Combiners) as a supra-classifier. Here, we describe theoretically how the probability that the Hamming nearest neighbor supra-classifier will predict the true target class approaches certainty at an exponential rate as more classifiers are reused. The scalability of the Hamming nearest neighbor with large numbers of previously created classifiers makes it a good choice as a supra-classifier in the application of building a repository of domain knowledge organized as a *classifier knowledge base.*

Key words: Knowledge Reuse, Evolving Knowledge Base, Hamming Nearest Neighbor

⋆ The authors would like to thank Ross Baldick, Gustavo De Veciana, and Ray Mooney for useful discussions and the Army Research Office contracts DAAH-94-G-0417 and DAAH 049510494, Texas ATP grant #442, as well as the National Science Foundation for sustaining support.

1 Introduction

1.1 Knowledge Reuse via Supra-Classifiers

Knowledge Reuse is the process of applying knowledge embodied in previously learned tasks to help solve a new, related task.

We recently introduced the *supra-classifier* knowledge reuse framework (Bollacker & Ghosh, 1997), which can apply the results of existing classifiers to help solve a new classification task. A schematic of this framework can be seen in Figure 1. The supra-classifier knowledge reuse process is to present

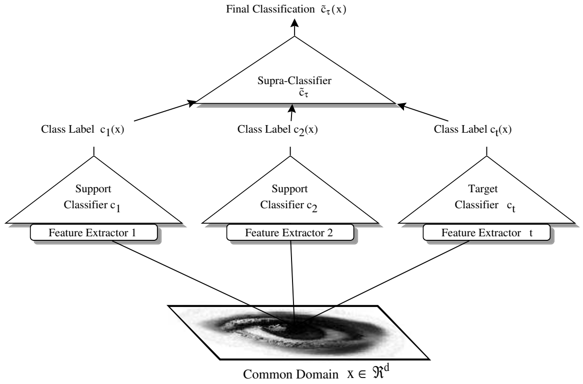

Fig. 1. A *Supra-Classifier* based knowledge reuse framework.

the training samples for a new, current target classification task to all available previously trained classifiers and then use the resulting output vector of classification labels as the input for a second stage supra-classifier. Only label information is used from support classifiers, thus allowing the reuse of any type of classifier architecture. The supra-classifier then makes the final classification decision for the current target classification task τ. Previously trained classifiers, termed *support classifiers*, are generally (but not always) designed for tasks other than the current target classification task of interest.

1.2 Supra-Classifier Design

Supra-classifiers make classification decisions on a vector of categorical (support class label) input values. Just like any ordinary classifier, they use the training samples for the target classification task (and perhaps *a priori* information) to make a decision. For an ordinary classifier, the feature set is static, and so to improve classification performance, more and/or better training samples are needed. In contrast to this, the premise of the supra-classifier framework is that knowledge can *also* be added by increasing the number of relevant support classifiers (input features). Designing a good supra-classifier is simply designing a classifier for a large number of discrete valued features that works well even with very few training samples. Although several architectures are possible, the Hamming nearest neighbor (HNN) classifier has been shown empirically to outperform decision trees, multilayer perceptron neural networks, naive Bayes, and other types of classifiers in the supra-classifier role. Furthermore, we have theoretically demonstrated HNN supra-classifier scalability to large numbers of support classifiers (Bollacker, 1998), with a brief reiteration here.

Consider a space $\Omega = \Re^d \times \mathcal{S}$ where $\mathcal{S}$ is a set of all possible discrete labels. Let $\Omega_\tau = \Re^d \times \mathcal{S}_\tau$ be a space defined for a classification task τ where $\mathcal{S}_\tau \subset \mathcal{S}$ and $|\mathcal{S}_\tau|$ is finite. $\Re^d$ is the input space where values to be classified exist, and $\mathcal{S}_\tau$ is the set of possible labels for that data for task τ. Let $X_\tau : X_\tau \in \Re^d$ and $Y_\tau : Y_\tau \in \mathcal{S}_\tau$ be the random variables for the prior distribution on Ω_τ associated with task τ. It is assumed that $\forall(x, y) \in (\Re^d, \mathcal{S}_\tau)$, $P(y|x)$ is defined. The most likely value of a class label y, given an input value x is defined as the maximum posterior probability function $t_\tau(x) = \arg\max_y P(Y_\tau = y|X_\tau = x)$. Thus, $t_\tau(\cdot) : \Re^d \to \mathcal{S}_\tau$ is the target function to be approximated by a target classifier. Let $\mathcal{B}$ be a set of support classification tasks which have the same input domain space $\Re^d$ as task τ. For each classification task $b \in \mathcal{B}$, there is an associated set $\mathcal{S}_b$ and classifier $c_b(\cdot) : \Re^d \to \mathcal{S}_b$.[1] Let the function $\vec{c}(\cdot) = (c_b(\cdot))_{b\in\mathcal{B}}$. For all b, let $C_b : C_b = c_b(X_\tau)$ be the random variables resulting from the application of X_τ to support classifier b and let $\vec{C} = \vec{c}(X_\tau)$. Let $T_\tau : T_\tau = t_\tau(X_\tau)$ be defined as the random variable associated with the target function on X_τ. Assumed in this context is a support classifier set $\mathcal{B}$, which has n members; i.e. $n = |\mathcal{B}|$, and for convenience it is assumed that $\mathcal{B} = \{1 \dots n\}$ to allow b to index entries of $\mathcal{B}$.

Using these definitions, the Hamming Nearest Neighbor (HNN) can be defined as a simple classifier for discrete input features (e.g. support classifier labels) similar to traditional nearest neighbor classifiers, which operate in a Euclidean space. If $1(\cdot)$ is the indicator function, then the Hamming distance between two sample input values x_r and x_s can be calculated as

$$D_n(x_r, x_s) = \sum_{b=1}^{n} 1(c_b(x_r) \neq c_b(x_s)).$$

where $c_b(\cdot)$ is the labelling bth support For each test sample s , the Hamming Nearest Neighbor (HNN) supra-classifier will choose the class label of the training sample with the smallest Hamming distance from it. If the previous classifiers are conditionally independent given the current classification

[1] Although some of the support classifiers may have been trained for task τ directly, in general $b \neq \tau$ and $\mathcal{S}_\tau \neq \mathcal{S}_b$, as the tasks are different.

task but are not independent of that current task, then the HNN approaches a perfect classification test rate with an exponential lower bound for an increasing number of previous classifiers. This is more precisely stated in Theorem 1.

Theorem 1:
If the random variables $C_b : b \in \mathcal{B}$ are independent of each other conditionally on the target class Y_τ, $\mu_\infty \geq 0$, the μ_n are equal for all values of n, making μ_n a constant, the training samples α, β, and test sample γ are value pairs that are randomly and independently drawn trials from the random variable pairs (X_α,Y_α), (X_β,Y_β) and (X_γ,Y_γ) respectively, X_α, X_β, X_γ, and X_τ are IID, Y_α, Y_β, Y_γ, and Y_τ are IID,samples α and β and test sample γ, and $\forall j, k \in \mathcal{S}_\tau, p_j = p_k$), then

$$P(D_n(X_\beta, X_\gamma) > D_n(X_\alpha, X_\gamma)|Y_\alpha = Y_\gamma, Y_\beta \neq Y_\gamma) \geq 1 - (1 - \mu_n^2)^{\frac{n}{2}}. \qquad (1)$$

Note that the only information this bound uses is the mean of the $E[Z_b]$ values and the boundedness of Z_b. As discussed in (Hofri, 1995), a better bound may be achieved by considering higher moments, but in general it is often difficult to do so. Since the Z_b random variables can only take the values $\{-1, 0, 1\}$, they are completely described by their first and second moments, potentially making this more tractable.

Theorem 1 states that as the number of conditionally independent and (at least barely useful) support classifiers increases, the probability that the (weighted) HNN classifier will predict the true target class approaches certainty bounded by an exponential rate. It should be noted that Theorem 1 holds even if there is only one training sample of each target class. This result leads to the observation that *under certain conditions, a wealth of features may substitute for a dearth of samples.* This is counter to the conventional wisdom that classifiers require more training samples when more features are present(Jimenez & Landgrebe, 1998). The trick is in avoiding a Euclidean feature space and the resulting typical curse of dimensionality problems.

2 A Classifier Knowledge Base

A classifier knowledge base is a set of classifiers for a domain of interest. For a sample in this domain, each classifier can provide a label. The resulting set of labels is a description of the sample, and if (at least some of) the classifiers "say something useful" about the sample, the labels can be used to aid in new characterizations of that sample. The new labels that become available when a new classifier is built can be contributed back to the classifier knowledge base. In this manner, each new classifier both takes existing knowledge from the knowledge base and gives some new knowledge back. As the number of classifiers grows, the knowledge base learns its domain better, and can thus provide increasing knowledge to help build new classifiers.

2.1 Experiments in Prediction of Personal Preference

In order to investigate the supra-classifier framework as applied to a classifier knowledge base, an experimental knowledge base on a small, fixed set of images was created. This experimental knowledge base was used to predict whether a user would "like" an image or not, based on previous classifications made by that user about those images. Although this target class is similar to those used in collaborative recommender systems (e.g. (Balabanović, 1997; Menczer, 1997)), the difference is that in many collaborative filtering systems, only like/dislike classifications are allowed, while in a classifier knowledge base, the previous classifications may be output for any set of labels. Like most collaborative filtering systems, our setup uses humans as support classifiers, making generalization to new images difficult. However, in this experiment we are primarily considering the situation where the data set is important and slow to change. We are trying to predict the user's behavior, not simply an image classification.

A dataset was created consisting of 30 photographic images, mostly acquired from the author's personal collection and a commercial CD-ROM. Some of these images can be seen in Figure 2. Seventy-three classifications were de-

Fig. 2. Nine of the 30 images in the image dataset.

signed for this set of images. Most of class labels sets were high level, such as "Big vs. Small", "Clean vs. Dirty", "Busy vs. Calm", and "Solid vs. Liquid vs. Gas". A Web site was created to present the 30 images to six human users (graduate students), who were asked to classify the images in each of 72 ways. These 72 manual (support) classifications can be thought of as a "personal profile" of knowledge about the images for each user. For the 73rd classification, users were asked to decide whether they "liked" each image more than average. This "like/dislike" label was used as the target classification for experiments in supra-classifier testing. One hundred experimental trials each of several combinations of training set and support classifier sizes were performed for each user and the results were averaged over all the trials and users. For each trial, the 30 images were randomly split into training and test sets, where the training sets ranged in size from from 3 to 27. The order of inclusion of the 72 support classifiers was also random for each trial.

2.1.1 Comparative Results

Although it is useful for a supra-classifier to perform better when there are more training samples, the most important role of a classifier knowledge base is to be able to perform well when there are only few training samples available for a new task but a large number of previously created classifiers. We compared the HNN supra-classifier with other architectures, including a Naive Bayes (NBAY) classifier, a C4.5 decision tree, a multilayer perceptron neural network (MLP), and a combiner based classifier (VOTE) where each component was a naive Bayes classifier that saw only one of each of the 72 previous classifications. A most common class (MCC) classifier always chose the most frequent target class in the training set, and was given as a baseline. The test rate of using each type of supra-classifier versus the number of support classifiers in the case of six training samples is given in Figure 3. The combiner (VOTE) supra-classifier has the best performance when the number of support classifiers was very small, but as the number of support classifiers increased, its performance actually was reduced. Many of the supra-classifiers had increasing performance with more support classifiers up to about 32, but only the HNN continued to have increased performance up to the full set of 72 support classifiers.

2.1.2 Results in a Knowledge Space

In order to understand the effect of knowledge reuse, consider a 2D *knowledge space* in which classifier performance is measured for a specific number of training samples and support classifiers. The average target classification test rate for the HNN supra-classifier given a range of number of training samples and support classifiers can be seen in the knowledge space represented in Figure 4. To achieve a given classification test rate, the HNN supra-classifier needed less training samples as the number of support classifiers grew. This

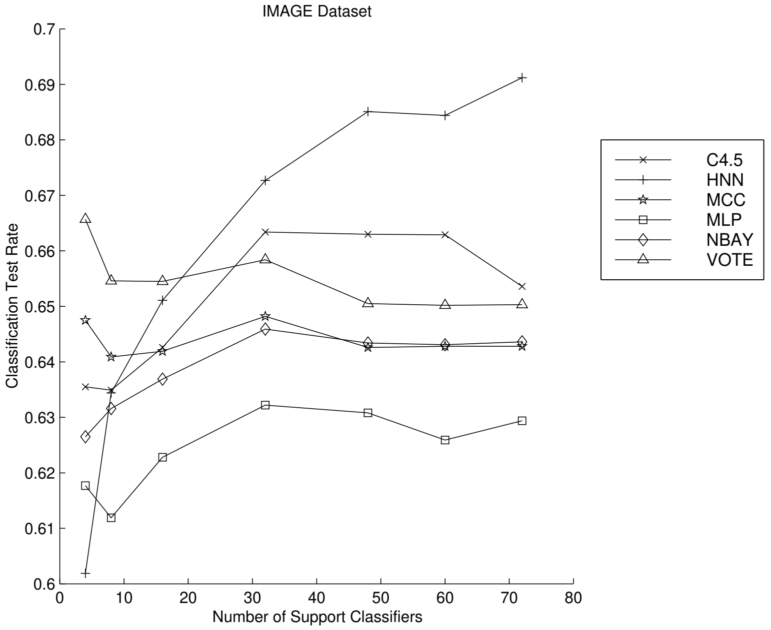

Fig. 3. Average classification test rate versus number of support classifiers for several supra-classifiers for the image dataset

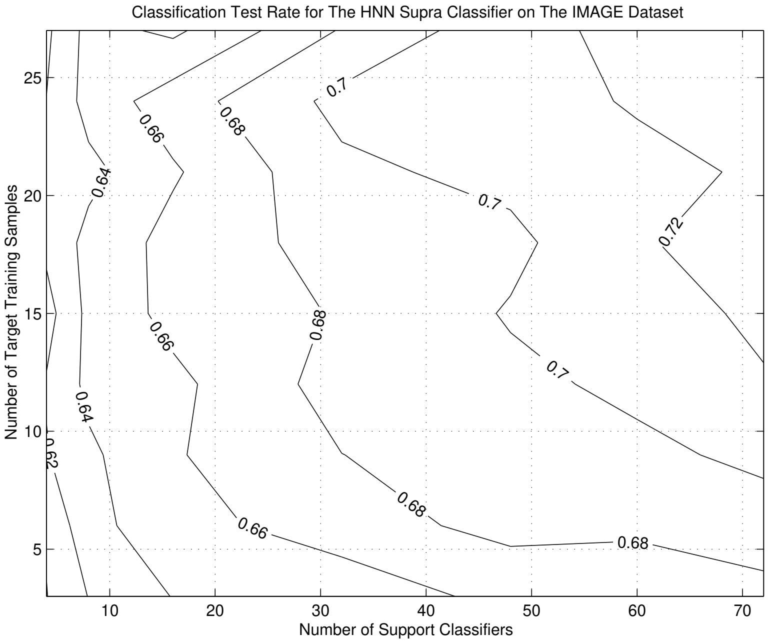

Fig. 4. Classification test rate contours in the knowledge space for the HNN supra-classifier on a classifier knowledge base for the image dataset.

is evidence that in a knowledge base, the HNN would perform better as the knowledge base grows. Contrast this to the poor performance from a Naive Bayes classifier as a supra-classifier in Figure 5. Here, the Naive Bayes classifier does not take advantage of support classifiers as they are added, probably because of increasing disparity with the required independence assumption.

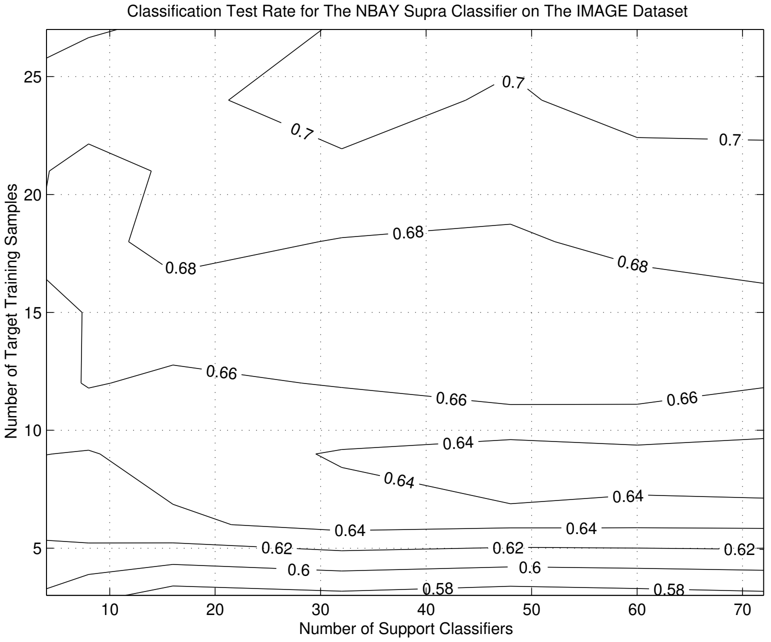

Fig. 5. Classification test rate contours in the knowledge space for the naive Bayes supra-classifier on experimental knowledge base

3 Conclusions and Future Work

The supra-classifier framework (and the HNN in particular) is discussed as an approach to the creation of a classifier knowledge base, and empirical evidence indicates its effectiveness. A prototype classifier knowledge base has been created and is available on the World Wide Web for testing and public use at `http://www.lans.ece.utexas.edu/cgi-bin/cgiwrap/kdb/knowledgebase`.

The most important direction for future development in the application of a practical classifier knowledge base is to actually create a knowledge base that provides some utility to a large number of users. If created with an easily accessible interface like the prototype shown here, a large classifier knowledge base could be built, and further investigations could be performed in the context of a truly practical system. Possible examples of such a system would be an example-queried Web page recommender or a predictor of legislative votes that uses post voting records as a knowledge base [2].

[2] Databases of such records are readily available on the Web at such sites as `http://thomas.loc.gov`

References

Balabanović, M. (1997). An adaptive Web page recommendation service. In *Proceedings of the First International Conference on Autonomous Agents*.

Bollacker, K. D. (1998). *The Supra-Classifier Framework for Knowledge Reuse*. Ph.D. thesis, The University of Texas at Austin, Austin, TX.

Bollacker, K. D., & Ghosh, J. (1997). Knowledge reuse in multiple classifier systems. *Pattern Recognition Letters*, *18(11-13)*, 1385–1390.

Bollacker, K. D., & Ghosh, J. (1998). A supra-classifier architecture for scalable knowledge reuse. In *The 1998 International Conference on Machine Learning*, pp. 64–72.

Hofri, M. (1995). *Analysis of Algorithms*. Oxford University Press, New York.

Jimenez, L. O., & Landgrebe, D. A. (1998). Supervised classification in high-dimensional space: Geometric, statistical, and asymptotical properties of multivariate data. *IEEE Transactions on Systems, Man, and Cybernetics, PartC: Applications and Reviews*, *28*, 39–54.

Menczer, F. (1997). ARACHNID: Adaptive retrieval agents choosing heuristic neighborhoods for information discovery. In *Machine Learning: Proceedings of the fourteenth International Conference*, pp. 227–235.