

Adaptive Committees of Feature-Specific Classifiers for Image Classification

by Fabrizio Falchi, Tiziano Fagni, and Fabrizio Sebastiani

Researchers from ISTI-CNR, Pisa, are working on the effective and efficient classification of images through a combination of adaptive image classifier committees and metric data structures explicitly devised for nearest neighbour searches.

An automated classification system is normally specified by defining two essential components. The first is a scheme for internally representing the data items to be classified; this representation scheme, which is usually vectorial in nature, must be such that a suitable notion of similarity (or closeness) between the representations of two data items can be defined. Here, 'suitable' means that similar representations must be attributed to data items that are perceived to be similar. If this is the case, a classifier may identify, within the space of all representations of data items, a limited region of space where lie the objects belonging to a given class; here, the assumption of course is that data items that belong to the same class are 'similar'. The second component is a learning device that takes as input the representations of training data items and generates a classifier from them.

In this work, we address single-label image classification, ie the problem of setting up an automated system that classifies an image into exactly one from a predefined set of classes. Image classification has a long history, with most existing systems conforming to the pattern described above.

We take a detour from this tradition by designing an image classification system that uses not one but five different representations for the same data item. These representations are based on five different descriptors or 'features' from the MPEG-7 standard, each analysing an image from a different point of view. As a learning device we use a 'committee': an appropriately combined set of five feature-specific classifiers, each based on the representation of the image specific to a single MPEG-7 feature. The committees that we use are adaptive; to classify each image, they dynamically decide which of the five classifiers should be entrusted with the classification decision, or determine whose decisions should be trusted more. We study experimentally four dif-

ferent techniques for combining the decisions of the five individual classifiers, namely dynamic classifier selection and weighted majority voting, each of which is realized in standard and in confidence-rated forms.

As a technique for generating the individual members of the classifier committee, we use distance-weighted k nearest neighbours, a well-known example-based learning technique.

implemented a system that makes use of metric data structures explicitly devised for nearest neighbour searching.

We have run experiments on these techniques using a dataset consisting of photographs of stone slabs classified into different types of stone. All four committee-based methods largely outperform, in terms of accuracy, a baseline consisting of the same distance-weighted k nearest neighbours fed with




Figure 1: Photographs of stone slabs classified into different types of stone, used to experiment automated classification.

Technically, this method does not require a vectorial representation of data items to be defined, since it simply requires that the distance between any two data items is defined. This allows us to abstract away from the details of the representation specified by the MPEG-7 standard, and simply specify our methods in terms of distance functions between data items. This is not problematic, since distance functions both for the individual MPEG-7 features and for the image as a whole have already been studied and defined in the literature.

Since distance computation is so fundamental to our methods, we have also studied how to compute distances between data items efficiently, and have

a 'global' distance measure obtained by linearly combining the five feature-specific distance measures. Among the four committee-based methods, the confidence-rated methods are not uniformly superior to those that do not use confidence values, while dynamic classifier selection methods are found to be definitely superior to weighted majority voting methods.

Link:

<http://www.isti.cnr.it/People/F.Sebastiani/Publications/IMTA09.pdf>

Please contact:

Fabrizio Sebastiani
ISTI-CNR, Italy

Tel: +39 050 3152 892

E-mail: fabrizio.sebastiani@isti.cnr.it