Extending the Classification Paradigm to Temporal Domains

Mohammed Waleed Kadous

School of Computer Science & Engineering
University of New South Wales
Sydney, NSW 2052, Australia
waleed@cse.unsw.edu.au

Introduction

One of the primary areas of machine learning research has been supervised concept learning – given some information about examples whose class is known, the goal is to produce a classifier which can classify examples whose class is not known. In general, research in this area has focused on situations where an object's attributes do not change in the short term. However, in many real-world domains, such as speech, sign language, robotics and medicine, many of the classification tasks involve dynamic attributes. Furthermore, temporal properties are critical to classification.

The current work involves developing a temporal classification learner that works in a variety of domains, does not require excessive amounts of data and is able to produce comprehensible descriptions of the concepts, while still having high predictive accuracy.

Model

Each domain has a specific set of channels, which are analogous to attributes in static classification. These channels are "sampled" at regular time intervals. The set of values of all channels at a particular instant is called a frame. A stream is a sequence of frames. The inputs to the learner are thus: a description of the channels; the set of classes; and a set of training streams labelled with their class. The output is a classifier that is capable of taking an unlabelled stream and predicting its class.

Testbed Applications

Currently work is progressing on three testbed applications:

Auslan Sign Recognition. Using instrumented gloves, isolated signs from Australian sign language are sampled. The goal is to produce a system which can recognise isolated signs from one of 95 classes.

ECG Analysis. Doctors can diagnose patients by examining their ECG's (electrocardiographs). It would also be possible to apply a temporal classification algorithm to the data. It might prove interesting to compare machine-generated rules to those used by doctors.

Robotics & Sensor Fusion. Much of what is interesting in a robot's domain may be the temporal variation and the correlation between sensors. A robot could be taught by example to deal with particular events, e.g. "door opening", "enemy robot firing weapon" etc.

Technique

Several techniques have been developed so far. There are three main families of techniques:

"Data-space" approaches. These deal with the channels directly. Techniques of this family so far have included treating each frame-channel pair as an attribute, then using instance-based learning; using entropy measures for each frame-channel pair as a bias to instance-based learning; extracting templates and using a χ^2 test to determine usefulness to classification.

"Segment-space" approaches. An algorithm for converting a channel to a sequence of line segments using a minimum description length heuristic has been developed. Once converted to a sequence of lines, the data is simpler to manipulate. For example, this information could be provided to an ILP system to extract rules from, or segments could be clustered, labelled and then a grammar induced.

A "hybrid" approach. Current research is based on the hypothesis that it is not important to recognise every part of the stream; but rather to isolate "markers" that are distinctive of a particular class. To the author's knowledge, this is a novel approach.

The approach being investigated uses an entropy-based data-space method to isolate markers, then applies a segment-space generalisation algorithm to generate a readable description. These are then used to build simple one-class classifiers which output confidence factors, which can then be combined to form a general classifier.

Acknowledgements

Thanks be to the Creator for giving me the ability to do this research. Thanks, secondarily, to Claude Sammut for many suggestions. This work is funded by an Australian Postgraduate Award.