Recognising the Bad Eggs

Mike Lopez

Manukau Institute of Technology
Auckland, New Zealand

mike.lopez@manukau.ac.nz

Proceedings of the 15th Annual NACCQ, Hamilton New Zealand July, 2002 www.naccq.ac.nz

ABSTRACT

The variation in natural products poses substantial challenges to recognition systems. In the examined automated processing environment, software is used to manage the grading, cleaning and packing of free-range eggs. The software needs to distinguish between eggs that are clean or dirty and also between those that are cracked or not.

Conventional approaches to this problem typically involve sophisticated algorithms with a matrix of parameters that are “trained” on a subset of data and then used to classify the wider dataset. There is a broad variation in natural products and, in consequence, the accuracy of these classifications falls short of that required for production systems. Such algorithms also typically require substantial processing power and are difficult to integrate into real-time systems that must classify products within a short fixed time interval.

This paper describes an alternative approach in which a high level abstraction of the problem is used to guide the selective application of the algorithms.

1. INTRODUCTION

The problem examined here relates to the implementation of a real-time system that classifies organic free-range eggs into categories such as clean, cracked, dirty as well as shape and weight. The natural variation in eggs eliminates the possibility of using a standard reference for comparison. The business costs of misclassification are skewed; classifying a dirty egg as clean involves a major cost to the business, whereas classifying a clean egg as dirty costs only the cost of a wasted egg. The organic status of the business precludes the use of commercial chemical cleaners and cleaning each egg by hand would be uneconomic.

To be commercially viable, the algorithms used must be able to classify eggs reliably (with target error rates of the order of 10-6 for the major costs referred to above) and rapidly (with target processing time less than 200ms per egg).

Algorithms for the classification of natural things have been investigated by many researchers worldwide and have immediate applications in fields as diverse as Precision Agriculture and computer vision for robotics. The general approach taken has been to emulate the assumed visual processing of humans and animals. This approach typically involves preprocessing a large volume of raw data to extract a feature set, and then applying algorithms to this feature set to achieve the classification.

Often, these algorithms are parameterised general algorithms with the parameters discovered by applying the algorithms to a "training set" of data with known classifications.

Using these techniques, Dror *et al.* (2001) have achieved an accuracy of 94% in classifying ellipsoid surfaces and Bennedsen and Rasmussen (2001) achieved 95.1% accuracy in classification of soil/ wheat/rape using spectrographic images. Although these techniques offer significant promise, the accuracy achieved is not yet high enough for widespread commercial application and the substantial computer processing time involved causes problems for real-time systems.

From a computer science viewpoint, the key inefficiency is preprocessing in a sequential fashion the large volume of raw data that is handled by parallel processing in natural vision. Much of this data has little effect on the outcome of the classification and may even be discarded by the later parts of the algorithms.

The question addressed in this paper is whether this preprocessing could be applied selectively. Could the algorithm give feedback to the preprocessing on where processing time could be best invested?

2. APPROACH

The approach suggested is to replace the linear model (feature extraction, algorithm) with a cyclical model (fig1).

After an initial sample is chosen, the cycle (feature extraction, classification and sample selection) is repeated until a "good enough" classification is reached. The classification updates the current interpretation of the data, and this interpretation is used to guide each of the processes in the cycle.

The current interpretation is simply an abstraction of the problem that becomes progressively constrained as the iterations proceed.

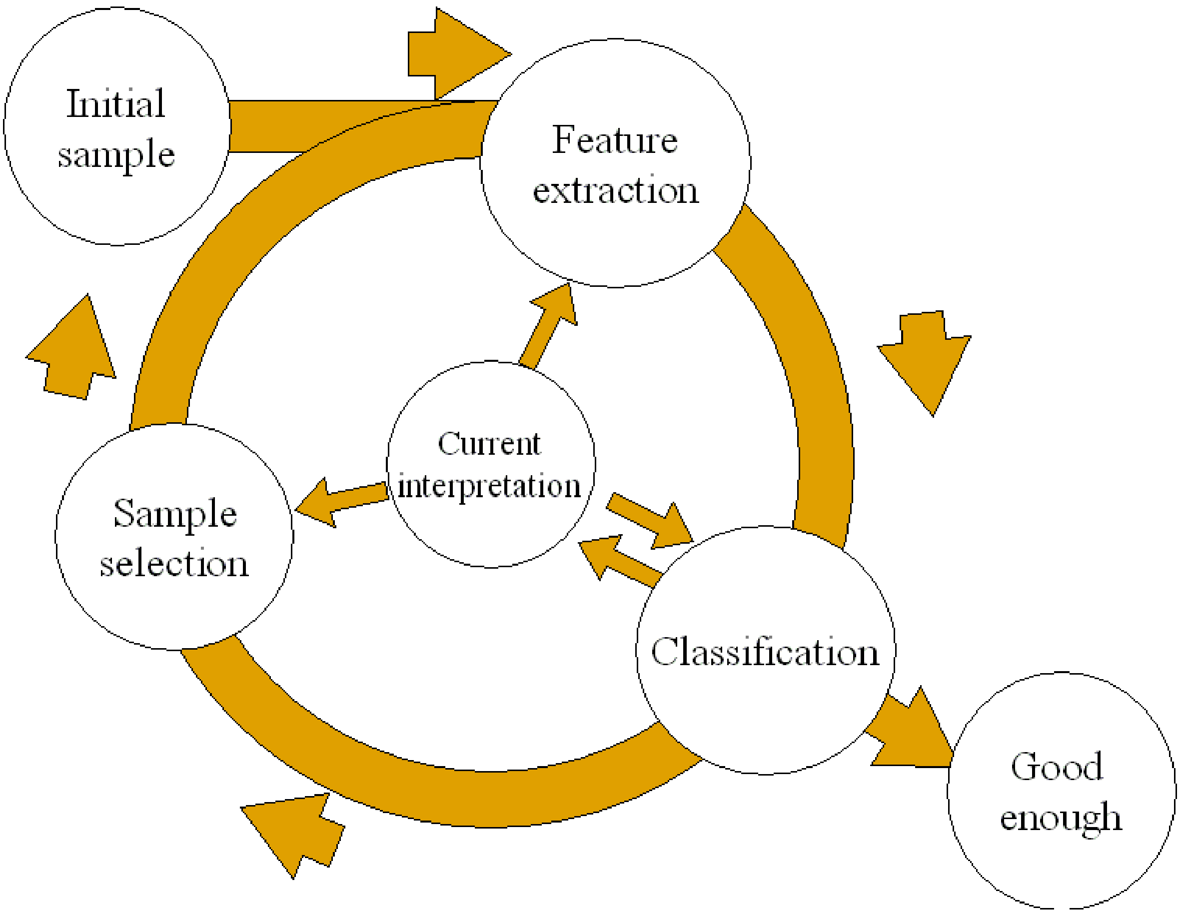

Figure 1

Two additional elements are needed in order to apply this model in a real-time system; the number of iterations must be limited in order to provide an upper bound to processing time and, in consequence, provision must be made for an additional category "unable to classify".

The commercial viability then rests on the ability to classify an acceptable proportion of the items within the desired confidence level and processing time constraints.

3. INITIAL FINDINGS

To test this approach, prototype software was developed using Visual Basic 6. The program used a library of patterns and a standard pattern-matching algorithm. The algorithm produced a "best guess" classification together with a measure of its confidence in the classification. The images that were analysed were monochrome images of 160 x 120 pixels.

The program used the same algorithm for both the linear and cyclic methods. When using the linear method, the whole image was submitted to the algorithm. With the cyclic method, a sample of the image was submitted and the cycle repeated up to 10 times until a confidence level of 99.9999 % was reached.

The precision multi-media timer in the Windows API (timeGetTime) was used to measure the time taken for a series of classifications and the average time for each method was calculated.

Both methods achieved 100% accuracy in classification on the data supplied. Over a series of tests, the average classification time for the linear model varied between 195 and 200 ms; that for the cyclic model varied between 0 and 1.25 ms. The prototype was run on a Pentium III 700 workstation running under Windows 2000. In this context, the resolution of the multi-media timer can vary between 1 and 5 milliseconds, so a larger series of tests on more complex images would be needed to get more accurate timings.

4. CONCLUSIONS

The sampling approach shows considerable promise in producing accurate classification within the constraints of limited processing time availability. The results so far suggest a ratio of around 200:1 in processing time, but the ratio may well widen as the complexity of the problem increases. This will be the subject of further investigation.

Further work will also be carried out on optimizing the feedback mechanisms from the abstraction to the sampling, feature extraction and classification elements of the algorithms.

REFERENCES

Dror, R.O., Adelson, E.H., Willsky, A.S. (2001). "Recognition of Surface Reflectance Properties from a Single Image under Unknown Real-World Illumination", Proceedings of the IEEE Workshop on Identifying Objects Across Variations in Lighting: Psychophysics & Computation, Kauai, Hawaii, December 2001.

Bennedsen, B.S., Rasmussen, P. (2001). "Plant Spectra Classification by Neural Networks". Proceedings of the third European conference on Precision Agriculture.