A Case Study in the use of ROC Curves for Algorithm Design

P.A. Bromiley, P. Courtney and N.A. Thacker
Imaging Science and Biomedical Engineering, University of Manchester, Manchester, UK

THE UNIVERSITY of MANCHESTER

ISBE Imaging Science and Biomedical Engineering

Introduction

We describe the development of a vision system to detect natural events in a low-resolution image stream. The work involved the assessment of algorithmic design decisions to maximise detection reliability. This assessment was carried out by comparing measures and estimates made by the system under development with measures obtained independently.

We show that even when these independent measures are themselves noisy, their independence can serve to guide design decisions and allow performance estimates to be made. We believe that such an approach will be applicable to other situations when an image-based system is to be used in the analysis of natural scenes where a precise ground truth is not available.

The Human Fall Detector

The system described here was a human fall detector using a novel infrared sensor with limited spatial resolution. Its purpose was to monitor a natural scene in an environment such as sheltered housing, where falls amongst elderly residents often lead to aggravated injury, and may even result in death. The aim was to extract estimates of the velocity of human subjects from the infrared image stream and examine them for patterns of velocities typical of falls. A list of various types of falls was prepared, and simulations were performed by an actress. These fall simulations were recorded using both the infrared sensor and a colour CCD camera. Since different cameras and lens systems were used, we would expect them to suffer from noise, bias and distortion independently.

A frame of the colour video.

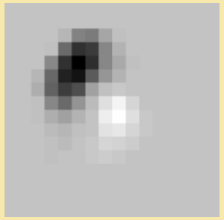

Corresponding frame from the infrared sensor.

System development was divided into stages:

- Velocity estimates extracted from the colour video data were used as a gold standard to demonstrate a correlation between the velocities extracted from the infrared data and the physical velocities present in the scene.
- An MLP neural network was trained to take temporal windows of velocity measurements from the infrared sensor and produce a fall/non-fall decision as output.

Velocity Correlation

The infrared sensor used was differential, and so any moving object warmer than the background created two regions on the sensor: positive values covering the pixels the object was moving into, and negative values covering the pixels the object was moving out of. A zero-crossing existed between the two regions. The average position of this zero-crossing was tracked over time to produce the velocity estimate.

Velocity estimates were extracted from the colour data using a colour segmentation algorithm, described elsewhere[1], to track the centroid of the actresses shirt.

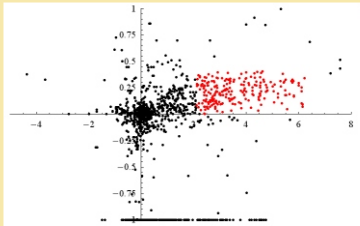

The infrared velocity estimate plotted against the colour video velocity estimate, in units of pixels per frame (at 30 f.p.s), after smoothing with the MRA (median rolling average) filter.

Various smoothing methods were applied to the data to reduce noise and improve the correlation between the velocity estimates. These included a five-point rolling average, a five-point median average, and the MRA filter, which took the average of the median three points of any five-point window. The best method was selected using ROC curves. Thresholds were applied to the colour and infrared velocity estimates to artificially classify the data points as "fast" or "slow". ROC curves were then produced for this artificial classification system by varying the threshold applied to the infrared velocity estimate.

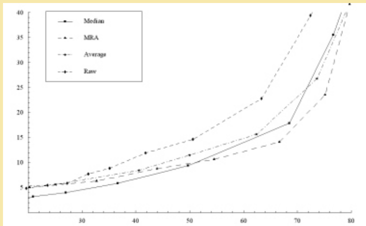

ROC curves (error rate against true acceptance rate) for the raw velocity data and the three smoothing methods (five-point median average, five-point rolling average and MRA filter).

The MRA smoothing filter clearly emerged as the best filter for this data set. Thus the independent (colour) measure was used to guide algorithmic development in the absence of a suitable ground truth.

1. P.A. Bromiley, N.A. Thacker and P. Courtney. Colour Image Segmentation using Non-Parametric Density Estimation in Colour Space, BMVC 2001.

Neural Network Training

In order to train a neural network to detect falls in the infrared data the positions of the falls were labelled by hand, and 20% of the data were extracted at random for training. Nine-point windows of velocity estimates were used as inputs to the network, and 120 networks were trained, varying all available parameters.

ROC curves were plotted showing the performance of each network on the data not used in training. The procedure applied is described in full in the proceedings. In order to guide the training process two additional classification systems were applied to the data:

- A nearest neighbour classifier, providing an upper bound on the network performance.
- A single data point decision system, providing a comparative measure of the performance improvement due to the neural network.

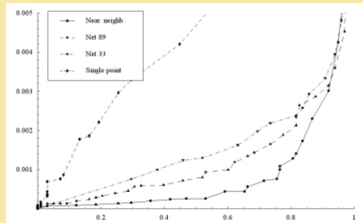

ROC curves for the best-performing networks trained with RPROP and CGM, for the best-performing nearest neighbour classifier, and for the single data point decision system.

Conclusions and Further Work

The study of performance evaluation is vital in placing machine vision on a solid scientific basis. We have described a case study giving two examples of the assessment of algorithmic design decisions. We have shown that:

- Independent measures can serve to guide rational design decisions in the absence of a genuine ground truth.
- Temporal identification of events can be subjected to a similar analysis: performance bounds obtained from independent classifiers can provide an estimate of the proximity of the system to optimal performance.
- With suitable consideration of the definitions of true and false detections, ROC curves can provide a unified, generic approach to performance evaluation.

All of our software is freely available from the TINA website:

http://www.niac.man.ac.uk/Tina