Detecting Faces using Evolution

John Montgomery, MSc Advanced Computer Science
msc37jxm@cs.bham.ac.uk
Supervisor: Dr Jeremy Wyatt

Evolution of Ratio-Templates

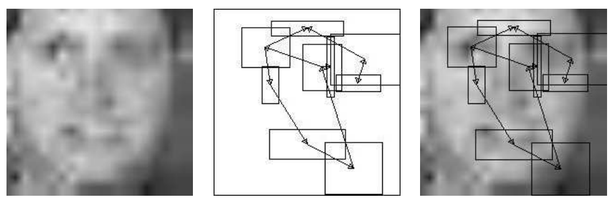

A Ratio-Templates[Sinha, 1996] is composed of several different "regions" and "ratios". Each region is used to find the average brightness in an area of an image. Each ratio represents a comparison between the average brightness of pairs of regions. Depending on these comparisons we are able to decide whether a "match" has occurs between the ratio-template and an image. Typically the regions and ratios in ratio-templates are specified by hand.

Genetic Algorithm

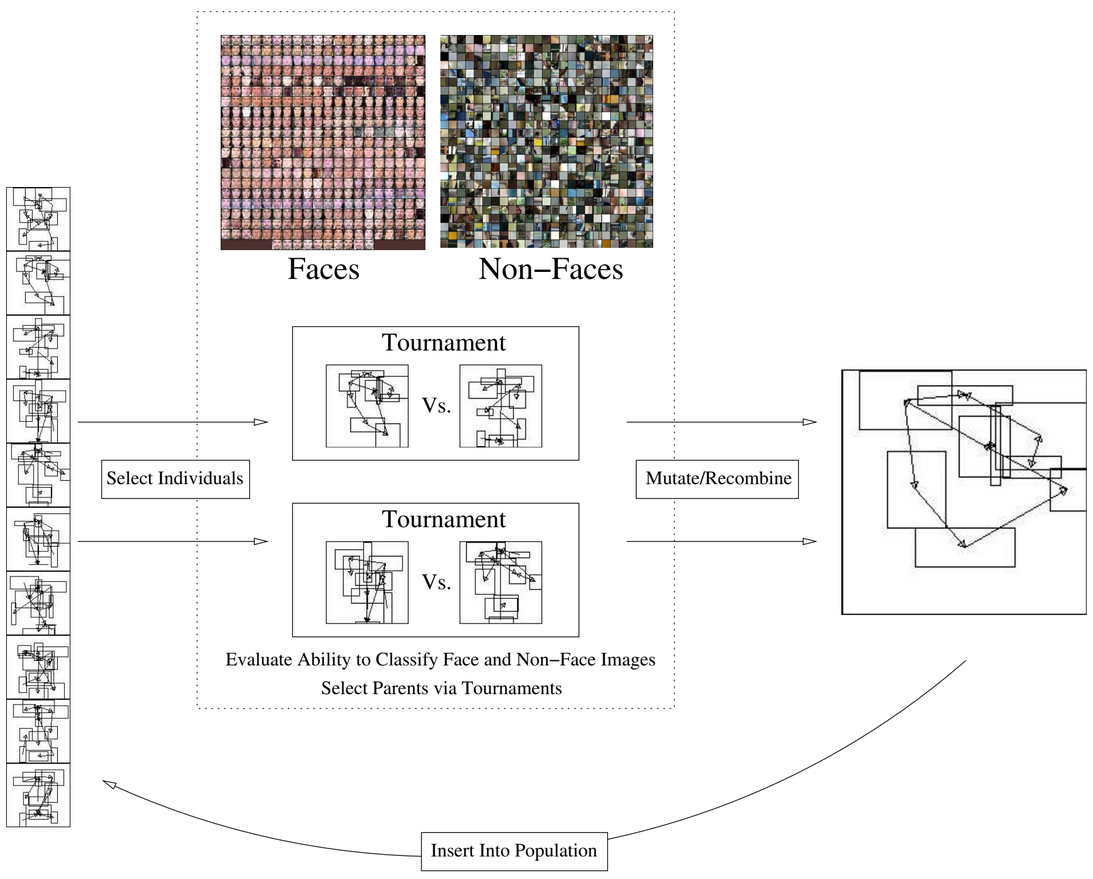

In this project evolution was used to create the ratio-templates. Evolution was used to control the locations, sizes and numbers of regions, as well as the nature and number of ratios used. The ability of candidate templates to correctly classify face and non-face images was used to guide this process.

Boot-Strapping Non-Face Images

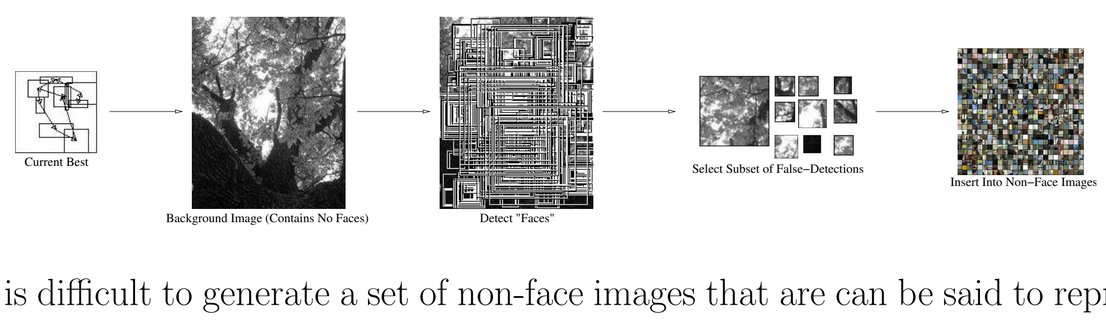

It is difficult to generate a set of non-face images that are can be said to represent all non-face images. So we "boot-strap" [Rowley et al., 1998] the non-face images once every 10 generations. We take the current best individual from the population and use it to detect faces on an image we know contains no faces. Any detections that occur on this image are implicitly wrong and can be used as examples of non-face images that are currently not correctly identified.

Detecting Faces

Once we have ratio-templates that are capable of discriminating between face and non-face images we need to harness them to detect faces in images.

Detection at Multiple Scales

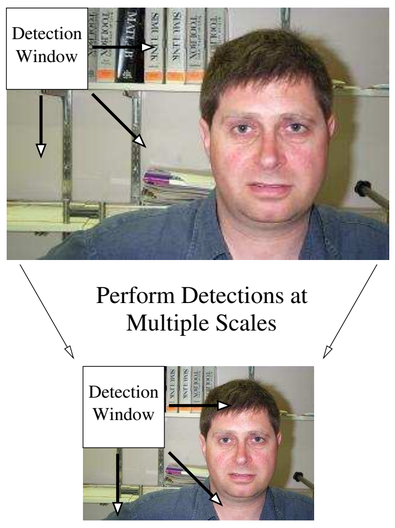

To detect faces of different sizes and at different positions in an image we repeatedly scan our detector over the image at several different scales.

This is a potentially very time consuming process, so it is important that a detector works quickly. As a particular piece of an image is much less likely to contain a face, than otherwise, it is often desirable for a detector to be able to quickly reject images that do not contain faces.

Reducing False-Positives

To reduce the number of false-positives that a detector finds we can combine the output of several ratio-templates.

The evolved ratio-templates make different mistakes, as they will have evolved in different ways. If every ratio-template has to agree before a detection occurs, then they will tend to agree about face images, but are less likely to agree about non-face images. This will lead to a lower false-positive rate.

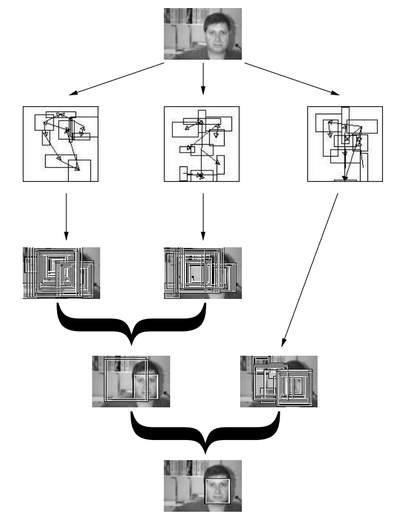

Figure 1: The Author's Face being Detected

References

[Rowley et al., 1998] Rowley, H., Baluja, S., and Kanade, T. (1998). Neural network-based face detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20:23–38.

[Sinha, 1996] Sinha, P. (1996). *Perceiving and Recognizing three-dimensional forms*. PhD thesis, Massachusetts Institue of Technology.