

Emanuele Trucco · Marco Razeto

Robust iris location in close-up images of the eye

Received: 30 October 2004 / Accepted: 20 May 2005 / Published online: 12 November 2005
© Springer-Verlag London Limited 2005

Abstract This paper presents a robust, reliable iris location system for close-up, grey scale images of a single eye. The system is meant as a bootstrap or recovery module for automated iris tracking within medical applications. We model the iris contour with an active ellipse, sensitive to intensity gradients across its perimeter. In this way, we avoid modelling the noisy appearance of the iris (e.g. corneal reflections). The iris–sclera intensity transition is modelled at two spatial scales with Petrou–Kittler optimal ramp filters. The optimal ellipse is identified by a simulated annealing algorithm tuned to the problem characteristics. The system performed accurately and robustly with 327 real images against substantial occlusion levels and varying image quality, subject, eye shape and skin colour.

Existing iris and pupil location techniques can be divided into *invasive* and *non-invasive*. Invasive techniques involve the use of one or more devices to be applied to the subject, in the form of electrodes, contact lenses and even head-mounted photodiodes or cameras [2]. Non-invasive techniques, instead, do not make use of external devices, but often rely on special illumination to highlight relevant eye characteristics. Classical examples are the so-called Purkinje images that exploit the reflections of infrared light off the cornea-lens boundaries [2, 5, 22, 25]. Neither invasive devices nor structured illumination are admissible in our case.

The requirement of robustness in uncontrolled conditions of illumination, image quality and iris appearance prevents us from using well-established techniques introduced to achieve iris recognition for people identification. Ma et al. [17] use only high-quality, unoccluded frontal images and detect the *limbus* (the boundary between the iris and the white of the eye, also known as sclera) with a combination of edge detection and Hough transform. The integro-differential operator used by Daugman [7, 11] in the iris recognition context proves very robust, but is restricted to frontal images where the limbus appears circular. We cannot make similar assumptions, and we adopt a detection technique based on active ellipses.

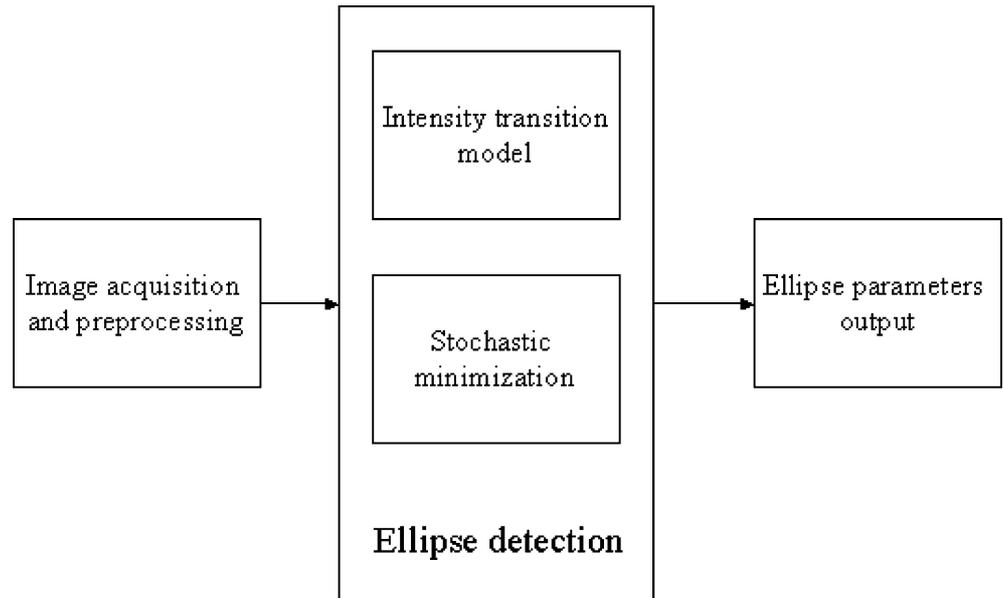
Active ellipses and, in general, *active contours* have been widely studied in computer vision [10, 15, 19, 23] and used, among others, in medical applications [8, 12, 14]. The main problems in our case are iris occlusion, unwanted features (e.g. corneal reflections) and varying image quality (e.g. blur, skin type, race). Occlusion is by far the biggest problem, since part of the iris is almost invariably covered by the eyelid or, occasionally, by other facial features. This is efficiently combated by a robust active ellipse location algorithm that recovers the iris shape even from limited parts of the limbus. Our system maximizes a criterion comparing intensity variations across the ellipse perimeter (Fig. 3)

1 Introduction and related work

This paper presents a robust iris location system, especially designed for close-up images of the eye (Figs. 1, 2). The system, in its present state, is meant as a single-image bootstrapping or failure recovery module for an eye tracker. The eye is assumed to cover a significant portion of the image acquired by a standard video camera. The system must be robust to occlusion (e.g. eyelids), skin and eye shape variations, image quality, and perform reliably in unstructured, general illumination conditions. Gaze and iris position within the image are unconstrained. The primary target is the determination of the iris centre with an accuracy of 5 pixels in a 360×280 image.

E. Trucco (✉) · M. Razeto
Vision, Signal and Image Processing Research Group School
of Engineering and Physical Sciences, Heriot Watt University,
Riccarton, Edinburgh, EH14 4AS, UK
E-mail: E.Trucco@hw.ac.uk

Fig. 1 Essential architecture of the active-ellipse iris detector



with a model derived from observations. When the ellipse overlaps the limbus, the criterion achieves a global maximum.

We use simulated annealing (SA henceforth) given the extremely craggy energy landscape. Non-stochastic optimization algorithms tried failed to reach the correct minimum with comparable consistency in the presence of occlusions and poor-quality images. In addition, a stochastic minimization algorithm can be initialized from significantly far-from-target states and still reach the correct minimum. This is a desirable property in our case, as no constraints are assumed on the iris position in the image. We detect the iris independent of other eye features, instead of using a complete eye template [1, 3, 6], because our images do not guarantee that other eye parts (pupil, eye contours) can be detected reliably. Moreover, contours occur at different spatial scales (rather sharp

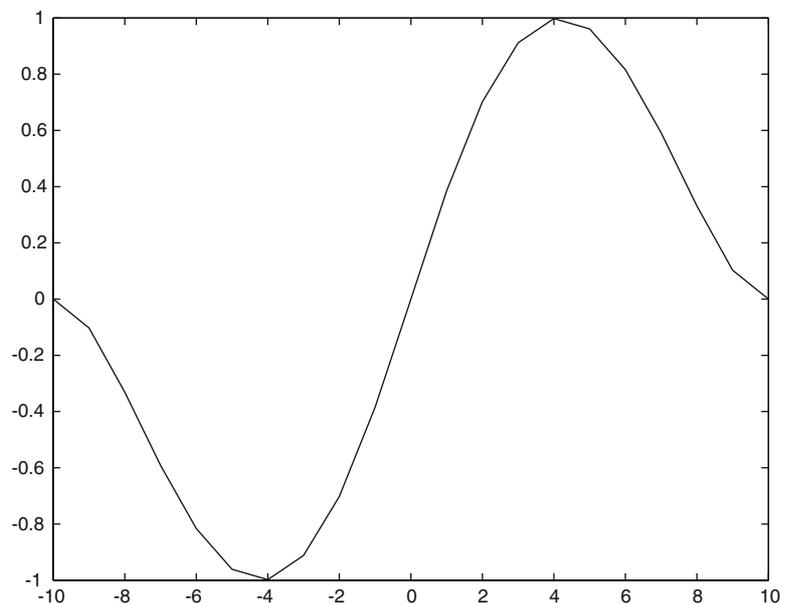
intensity transition between iris and sclera, wide and noisy one between sclera and eyelid) and some of the features appear blurred, making location especially difficult.

The rest of this paper is organised as follows. Section 2 presents the modules of our system (pre-processing, iris modelling, active ellipse location), Section 3 a performance evaluation, and Section 4 our conclusions and thoughts for future work.

2 System overview

The iris is located by an active ellipse search returning the estimated position and shape of the iris. The system incorporates a pre-processing stage designed to suppress reflections and similar distracting artefacts. Figure 1 sketches the essential system architecture.

Fig. 2 Example of a Petrou–Kittler optimal ramp detector (ramp size is 10). Pixel coordinates on x axis, signal values on y axis



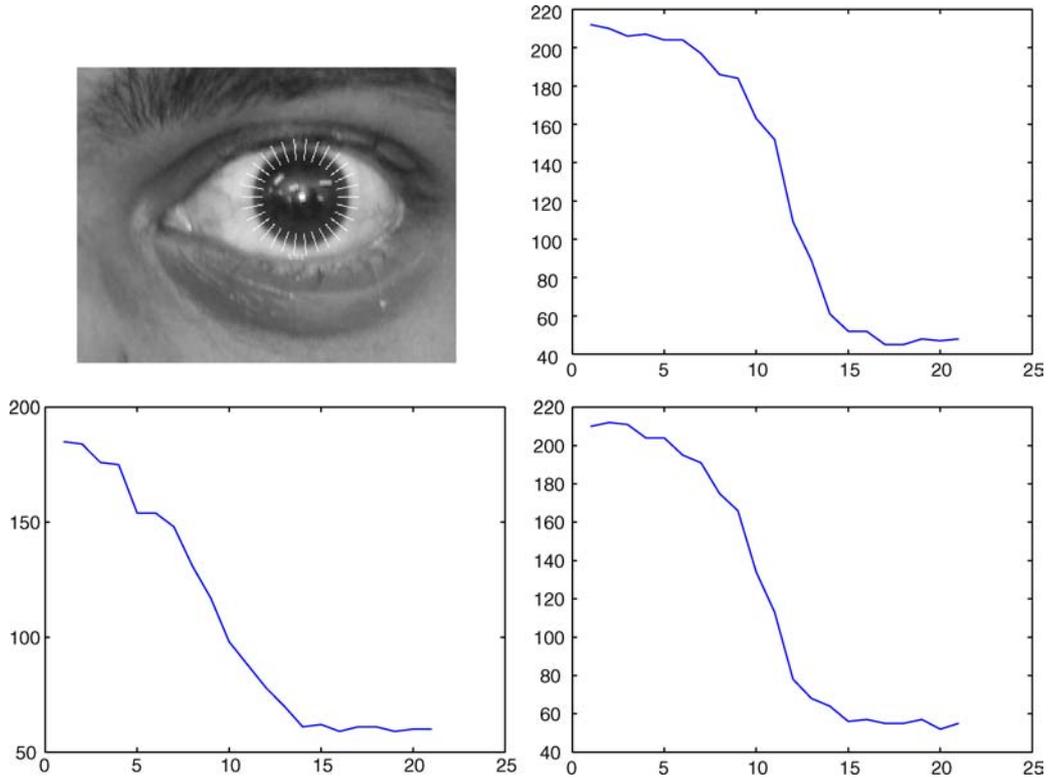


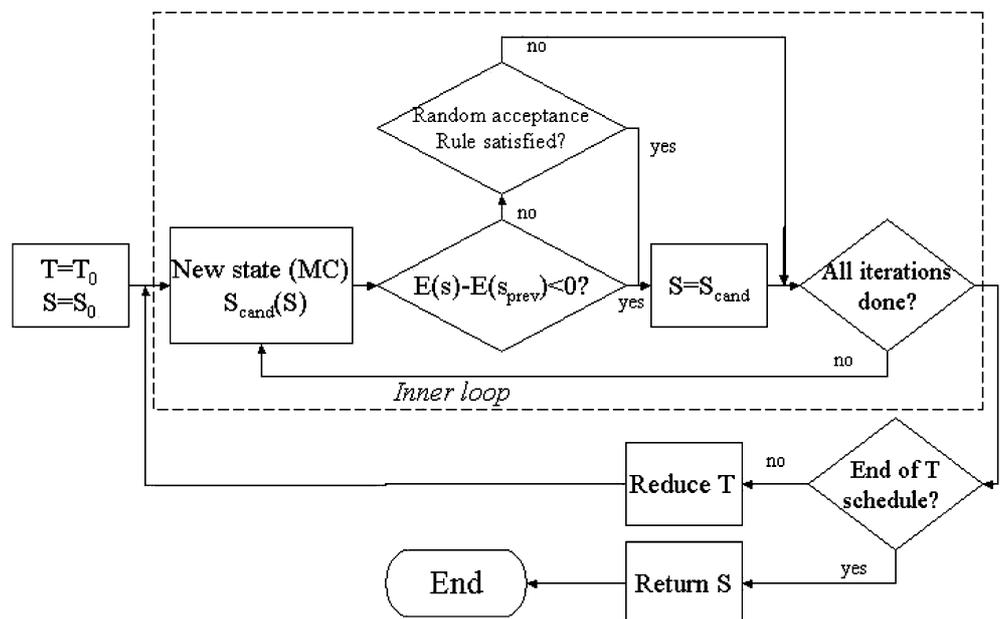
Fig. 3 The segments normal to the ellipse, along with intensities, are analysed by the Petrou–Kittler ramp detectors. The diagrams show measured intensity profiles along some of the segments (pixel co-ordinates on x axis, intensity values on y axis)

2.1 Input and preprocessing

Reflections and other unwanted artefacts, such as motion blur, can affect iris location seriously. Corneal reflections of room lights, for instance, prove particu-

larly disruptive as they introduce distracting, strong contours. A 10×10 median filter run on the whole image prior to minimization proved a simple and very effective solution to combat the effects of artefacts. We can afford such a relatively large mask (the average image size is

Fig. 4 SA block diagram. T is temperature, S is current state, $E(S)$ is the energy state of S . MC stands for “move class”. All other symbols are explained in the text



350×270), compromising texture details, as the algorithm uses only iris boundary information, which is well preserved by median filtering.

2.2 Iris location via active ellipses

We pose the problem of the location of the elliptical iris–sclera boundary as an optimization in the parameter space of an active ellipse model. The result identifies the elliptical image contour (location and shape) which best follows a predicted model of intensity changes. This block has two key components, the cost function and the minimization method. They will be described separately.

2.2.1 Intensity transition model: cost function

The sclera–iris boundary is characterized by a light (sclera) to dark (iris) transition. The spatial extent of such a transition depends on many factors, primarily image resolution and quality, and the particular eye shape. In images typical for our application and acquisition set-up, it can be between 3 and 12 pixels approximately. We model this transition with two Petrou–Kittler ramp edges [18] at two different spatial scales.

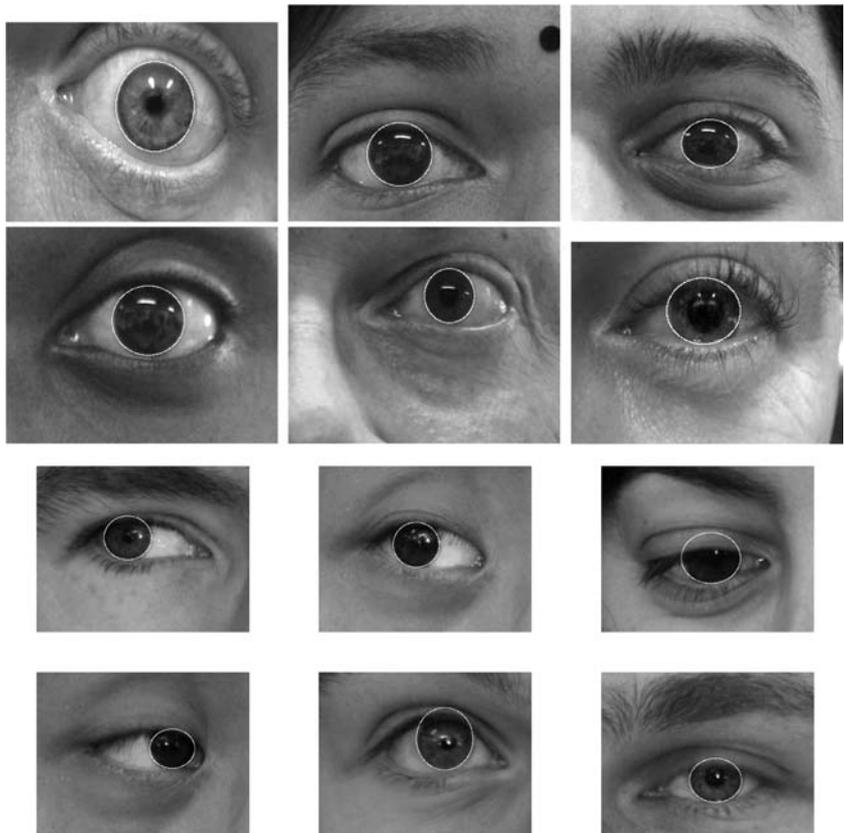
The ellipse is parametrized by its centre and axes. The latter are assumed aligned with the image axes, as head orientation is constrained in our application. The cost

function selects a fixed number n of control points (usually 30) along the ellipse perimeter, and calculates the normal to the ellipse at each control point (Fig. 3). Intensity profiles are extracted along these normals on both sides of the perimeter, yielding n numerical profiles. We convolve each of these profiles with two optimal ramp detection filter masks, introduced by Petrou and Kittler [18], at two different spatial scales. The graph of such filters is illustrated in Fig. 2, and their analytical model is

$$h(x) = e^{Ax}[L_1 \sin(Ax) + L_2 \cos(Ax)] \\ + e^{-Ax}[L_3 \sin(Ax) + L_4 \cos(Ax)] + L_5x + L_6e^{sx} \\ + L_7,$$

for $-w \leq x \leq 0$, where w is the filter half-width. The values of A and $L_1 \dots L_7$ are tabulated for the two target spatial scales. The parameter s represents the inverse of the transition length. We can easily tune the filter for different values of s by rescaling A and w . These filters are antisymmetric and therefore unaffected by uniform changes of illumination. The two-scale filtering scheme adopted, involving processing signals at separate scales and combining results, is inspired to [13]. Extensive testing suggested that two masks optimal for transitions of 4 and 10 pixels respond well to typical limbus edges encountered in our target imagery, and poorly to transitions related to non-iris features such as eyebrows or eyelashes. As we are interested only in the value of the filter output in the centre of the normal segments (i.e. on

Fig. 5 Examples of detection on low-difficulty images: moderate iris occlusions, acceptable image quality, subject looking mostly towards the camera



the ellipse perimeter), we compute only *one* filtered value per segment. Filtered values are summed over all control points and over both filter sizes to obtain the criterion to optimize, c , given below:

$$c = - \sum_{i=1}^N \left(\int_{-w}^w S_i(x) f_1(x) dx + \int_{-w}^w S_i(x) f_2(x) dx \right), \quad (1)$$

where N is the number of control points, S_i is the intensity profile extracted at the control point i and f_1 and f_2 are the filters at the two different scales.

2.2.2 Optimization scheme

Simulated annealing [21] has been variously used in computer vision [4, 9, 20]. SA is well-suited to tormented energy landscapes with several local minima as it can explore efficiently wide search regions of parameter space, homing in on promising regions. Figure 4 shows a sketchy flowchart of our version. A characteristic of SA is that all its blocks must be specialized for the specific problem being addressed [21]. A problem-specific move

class ensures sufficiently wide search regions at high temperatures (initial stages), and focuses search at lower temperatures (final stages). The number of new ellipses tested is also progressively reduced with temperature. The active ellipse is parameterized by a , b (semiaxes), O_x , O_y (centre co-ordinates), forming a 4D state vector S . This is initialized at the image centre with a default size. We do not consider the ellipse orientation in the image plane as the system is designed for applications in which the patient's head is reasonably vertical. Small inclinations are possible (and indeed included in our test), as the head is not constrained, but do not upset the fit enough to introduce orientation as a fifth parameter to optimize. Notice that the values of key parameters (number of iterations per loop, temperature reduction schedule) have been selected through extensive experimental analysis of the algorithm performance in varying, controlled conditions [24].

The algorithm is organized in two nested loops, an inner and an outer loop.

Inner loop A *move class* (MC) (see below) generates a new candidate state S_{cand} , for which an *energy difference* $\Delta = E(S_{\text{cand}}) - E(S)$ is computed. If $\Delta < 0$ (energy

Fig. 6 Examples of detection for challenging images, including heavily occluded iris, lower image quality (e.g. blur, motion)



decrease), the new state is accepted and used as the starting point for the next iteration; else the state is accepted or rejected using a random *acceptance rule* (AR). Both MC and AR depend on the temperature T , controlling the width of the search region. The procedure is repeated for a number of iterations depending on temperature (see below), and the lowest-energy state encountered is recorded.

Outer loop The temperature is lowered (affecting both MC and AR) and the inner loop repeated with a reduced number of iterations, starting from the best state so far. The algorithm ends when a stipulated minimum temperature is reached (details below). Typical values in our experiments are $T_{\text{start}} = 500$ and $T_{\text{end}} = 1$, which have been decided by sampling the cost function over several images and calculating the relative acceptance ratio, whose desirable value at high temperature is around 50%.

Move class The new state is generated stochastically, subject to constraints imposed by expectations on the iris size. For every new parameter a new candidate value is generated from a Gaussian distribution centred in the previous state's parameter value, and with standard deviation $\sigma_{\text{new}} = R\sigma_{\text{old}}$, much in the spirit of Fast Annealing by Ingber [16]. R controls the search range, starting from 2 pixels for ellipse centre and 1 pixel for axes lengths and decreasing with an independent annealing schedule (see below). Starting standard deviations are 3 pixels for the axes and 5 pixels for the centre co-ordinates. These values provide an adequately large step at high temperature and allows rapid exploration of large search areas. New states are tested only if falling within an allowed range reflecting the likely appearance of the iris in the images. Typical ranges used are 30–60 pixels for axis lengths and 0.8–1.2 for the axes' ratio, reflecting the appearance of the iris in our images.

Acceptance rule The standard (Metropolis) acceptance rule proves a simple and effective solution for our case: if the candidate state brings an energy increase, it is accepted with a probability $e^{-\Delta/T}$ (notice the dependence on T).

Annealing schedule the annealing schedule is numerically tuned to our problem. The temperature is reduced according to $T_{\text{new}} = T_{\text{old}}\alpha^{t+1}$, where t is the outer loop index (number of T values so far), and $\alpha = 0.98$. The number n of new states tested before decreasing T is also decreased, as annealing proceeds, following

$$n = \frac{N}{\sqrt{t+1}} + 100,$$

where N is the starting value (500). The number n ranges from 600 (high T) to little over 100 (low T).

The annealing schedule affects the move class via the range parameter R :

$$R_{\text{new}} = \left(\frac{1}{\sqrt{t+1}} + 0.3 \right) R_{\text{old}}$$

with all the symbols as above.

3 Experimental results and performance analysis

The current implementation runs on a Pentium 4 PC workstation, running Matlab 6.5.0 release 13. We used a database of 327 selected test images of various degrees of difficulty. Roughly, half of the images have a high degree of difficulty, including heavily occluded iris, blurred or unfocused images, images of eyes facing away from the camera, images of people wearing glasses or a combination of the above. The images are greyscales of various sizes, roughly averaging 350×270 pixel, captured by a digital camera or a digital camcorder under uncontrolled room lighting (Figs. 5 and 6).

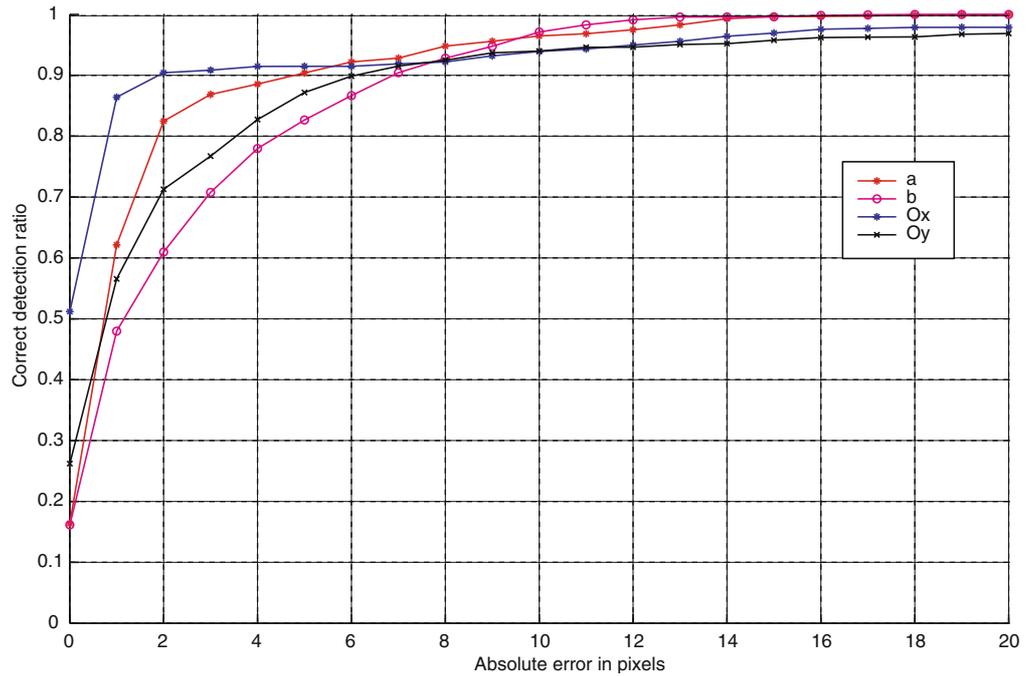
The unoptimized and uncompiled Matlab code processes one image in about 5 seconds. Previous experiences with porting MATLAB 6 code to target application platforms with machine-level programming led to execution speeds up to 20 times higher than MATLAB prototypes. In our case, this would mean approximately 10 frames per second.

Ground truth for quantitative tests was established manually by tracing ellipses following the limbus. We performed 50 runs on each image, totalling 50×327 = 16,350 runs. The initial ellipse for all tests is at the centre of the image, with semiaxes of 40 pixels each. We compute the difference between estimates of ellipse parameters and the corresponding ground truth values, showing the results as error distributions (histograms, Fig. 9). Good performance is indicated by distributions centred on zero and vanishing rapidly. An error of under 5 pixels on all parameters is considered a correct detection. Notice that, as the location of the pupil centre is



Fig. 7 Three examples of misdetection.

Fig. 8 Error probability of absolute error (*in pixels*) for the four ellipse parameters. *X* axis: error value in pixel. *Y* axis: fraction of experiments leading to correct detection, i.e. within set tolerances, for each ellipse parameter



the target, correct location of the centre is more important than the precise measurements of the iris size (axes).

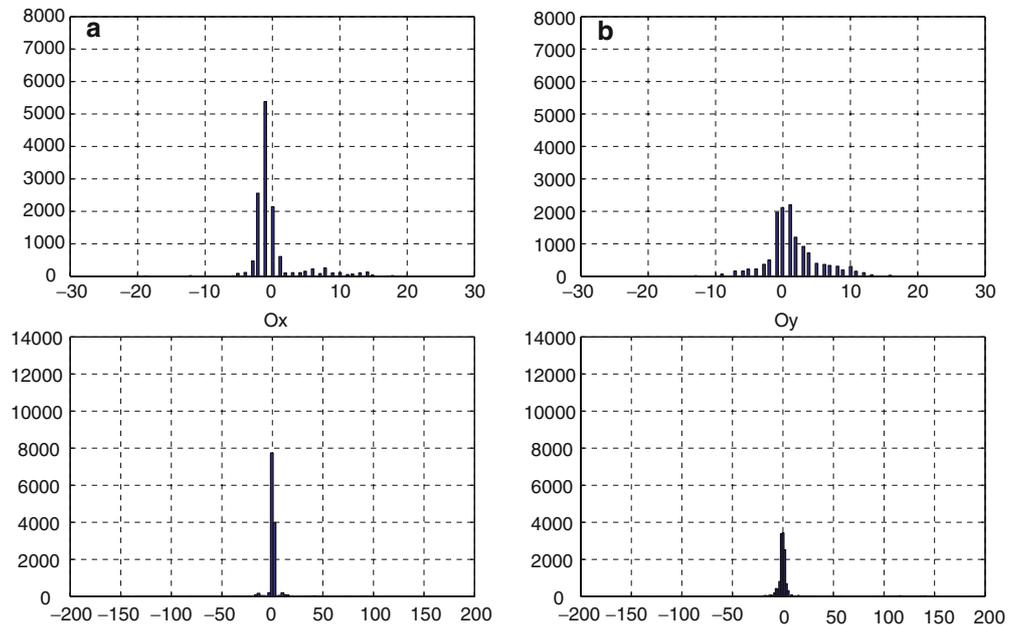
In addition to the quantitative analysis below, we stress that the results indicate excellent stability (Fig. 7).

Figure 9 shows the error distributions for all the images. Errors are quite limited: the standard deviation of the error on *b* is approximately 4 pixels, and the mean approximately 3 pixels. The standard deviation of *O_y* is less than 10 pixels, and the mean approximately 0. The errors on *b* and *O_y* in both image groups are on average

larger than those on *a* and *O_x*, as expected as iris occlusions occur mostly along the vertical axis (the limbus is indeed very rarely visible completely).

Figure 8 shows, for each ellipse parameter, the *estimated* a posteriori cumulative probability of a given error value in pixels. For instance, we can see that 91.5% of the *O_x* histogram falls within the 5-pixel tolerance interval, suggesting an *approximate* probability of 91.5% for correct detection (in our definition) of the horizontal component. For *O_y*, this figure is 88%. We notice that, given a 91.5% probability of correct location in one attempt, the probability of correct location

Fig. 9 Error distributions (*histograms*) of the image database for parameters *a* (*top left*), *b* (*top right*), *O_x* (*bottom left*), and *O_y* (*bottom right*). Throughout, *X* axis is error value in pixels.



within two independent attempts is 0.993 (probability of correct detection at the first attempt plus the probability of correct detection at second attempt given the first failed). Although two runs on the same image will not be completely independent, this suggests that running the algorithm twice provides a high probability of correct detection, assuming of course there is enough time for two runs *and* a criterion to decide which answer is likely to be the correct one. Possible decision criteria include the fact that wrong ellipses tend to place iris candidates outside the eye boundaries or generate candidates smaller than the average ellipses detected previously.

4 Conclusions and future work

We have presented an iris location system that proves robust, reliable and accurate with close-up, grey-level images of an eye. Performance is good in almost all of the images containing various levels of occlusion, distracting artefacts, image quality, and different eye shapes and skin colours. The robustness, reliability and precision of our system seem comparable with or better than those reported for other iris/pupil location and tracking systems, for instance [23] (pupil tracking using statistical snakes and infrared sequences) or [12] (iris location using deformable iris models).

It seems appropriate to point out a crucial difference between iris recognition and the class of applications we target. Iris recognition assumes that the subject opens his/her eye to expose the iris to facilitate iris location which is sensible given the operational context of iris recognition systems. We do not assume, instead, that the subject is asked to expose the iris; moreover, we do not control either illumination, or eye, or head motion. The appearance of the iris then becomes much less predictable, and iris detection much more complicated. Attempts to deploy edge-based algorithms followed by search for groups of edge chains forming ellipses did not meet our reliability targets, and this is the reason why we turned to active ellipses and stochastic optimisation.

Future work directions include developing a real-time C version, to be incorporated into an iris tracker, and incorporating learning models of the iris appearance as an alternative to active-ellipse location.

5 Originality and contributions

This paper has presented a robust, reliable iris location system for close-up, grey scale images of a single eye. A key contribution is that detection is achieved with very good reliability even with significant occlusions and uncontrolled illumination. This entirely original work has excellent potential for applications requiring reliable estimation of the position of the iris, including medical, transport (e.g. driver monitoring) and biometrics.

6 About the authors



Emanuele Trucco received his Laurea (MEng) and PhD degrees in Electronic Engineering from the University of Genoa, Italy, in 1984 and 1990, respectively. He has been active in computer vision research and applications ever since, publishing more than 120 papers and one well-known textbook on the subject. He is currently a Reader (Associate Professor US) at Heriot Watt University, Edinburgh, UK. Dr Trucco serves regularly on scientific or organising committees of most of the major conferences in the field. He is an Editor for the IEEE Transactions on Systems, Man and Cybernetics Part C, and the IEE Proceedings on Vision, Signal and Image Processing.



Marco Razeto graduated from the University of Genoa, Italy in Physics in 2000, and obtained his PhD degree from Heriot Watt University, UK, in 2005. His research interests are in tracking, recognition, detection, and medical imaging. He is currently a researcher at Voxar Barco, Edinburgh.

References

1. Blake A, Yuille AL (1992) Active vision. MIT, Cambridge
2. Duchowsky AT (2003) Eye tracking methodology. Springer, Berlin Heidelberg New York
3. Bettinger F, Cootes TF, Taylor CJ (2002) Modelling facial behaviours. In: Proceedings. BMVC'02, vol 2, pp 797–806
4. Luo B, Hancock ER (2001) Structural graph matching using the EM algorithm and simulated annealing. IEEE Trans PAMI 23(10)

5. Morimoto CH, Koons D, Amir A, Flickner M (2000) Pupil detection and tracking using multiple light sources. *Image Vis Comput*
6. Cristinacce D, Cootes TF (2003) Facial feature detection using ADABOOST with shape constraints. In: *Proceedings of BMVC'03*, 1:231–240
7. Daugman J (2004) How iris recognition works. *IEEE Trans Circuits Syst Video Technol* 14(1)
8. Davis RH, Twining C, Cootes T, Taylor C (2002) A minimum description length approach to statistical shape modelling. *IEEE Trans Medical Imaging* 21:535–527
9. Tan HL, Gelfand SB, Delp EJ (1992) A cost minimization approach to edge detection using simulated annealing. *IEEE Trans PAMI* 14:3–18
10. Isard M, Blake A (1998) *Active contours*. Springer, Berlin Heidelberg New York
11. Daugman J (1993) High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans PAMI* 15(11)
12. Ivins JP, Porrill J, Frisby JP (1997) A deformable model of the human iris driven by non-linear least-squares minimisation. In: *Proceedings of IPA'97*, 1:234–238
13. Weber J, Malik J (1994) Robust computation of optical flow in a multi-scale differential framework. *Int J Comput Vis*
14. Kerrigan S, McKenna S, Ricketts IW, Widgerowitz C (2003) Analysis of total hip replacements using active ellipses. In: *Proceedings of Medical Images Understanding and Analysis (MIUA)*
15. Leroy B, Medioni GG, Johnson E, Matthies L (2001) Crater detection for autonomous landing on asteroids. *Image Vis Comput* 19(11):787–792
16. Ingber L (1989) Very fast simulated re-annealing. *Mathematical and computer modelling* 12:967–973
17. Ma L, Tan T, Wang Y, Zhang D (2003) Personal identification based on iris texture analysis. *IEEE Trans PAMI* 25(12)
18. Petrou M, Kittler J (1991) Optimal edge detectors for ramp edges. *IEEE Trans. PAMI* 13
19. Vincze M, Ayromlou M, Zillich M (2000) Fast tracking of ellipses using edge-projected integration of cues. In: *International Conference on Pattern Recognition (ICPR'00)*
20. Ritter N, Owens R, Cooper J, Eikelboom R, van Saarloos PP (1999) Registration of stereo and temporal images of the retina. *IEEE Trans Med Imaging* 18(5)
21. Salamon P, Sibani P, Frost R (2002) *Facts, Conjectures and Improvements for Simulated Annealing*. SIAM
22. Shih S-W, Liu J (2004) A novel approach to 3-d gaze tracking using stereo cameras. *IEEE SMC (Part B)* 34(1):234–245
23. Ramadan S, Abd-Elmageed W, Smith CE (2002) Eye tracking using active deformable models. In: *Proceedings of ICVGIP'02*
24. Richard S (2004) *Comparative experimental assessment of four optimisation algorithms applied to iris location*. Master's Thesis, ECE/EPS, Heriot-Watt University
25. Cornsweet TN, Crane HD (1973) Accurate two-dimensional eye tracker using first and fourth Purkinje images. *J Opt Soc Am* 63(8):921–928