

Daugman's Iris Scanning Algorithm

W. A. Barrett

Biometrics Test Center, San Jose State University

2/4/00

DaugmanAlg1.doc, vs. 2.0

Abstract

This is a review of the patent and certain iris scanning papers by John Daugman [1, 2, 3] concerning the recognition of persons through their highly distinctive iris patterns. These papers reveal the basic technology of the recognition process, but do not reveal several critical details and parameters required by any practical recognition algorithm. We speculate through some educated guesses on what is likely an optimal choice of parameters for an accurate recognition of individuals.

Our aim in this is to develop our own approach to iris scanning for academic research purposes. Certain of our undergraduate engineering students have approached this problem, but much work remains.

Outline of Daugman's Algorithm

Daugman's algorithm (hereafter referred to as *DA*) requires these steps in order to reduce an image of the eye to a reduced encoded form. The image is from a gray-scale video camera, typically 640 pixels wide x 480 pixels high.

- Locate the center of the pupil, and determine its diameter. This part of the algorithm is reasonably well described in the papers, except for various optimizations designed to improve resolution and performance.
- Determine the diameter of the outer edge of the iris, where the iris meets the white sclera. The outer diameter will typically be approximately 320 pixels. The inner diameter will vary (for the same person), depending on the eye's accommodation to the average light level, from approximately 60 pixels to 200 pixels.
- Find the useful region of the iris, unobstructed by an upper or lower eyelid, and not obliterated (in the image) by light reflections. A region in the upper portion is always excluded since this is typically covered by the upper eyelid. A fan-shaped 45 degree in the lower portion is always removed since this is typically obscured by a strong reflection from the illuminating light source.
- Partition the useful region into 8 concentric rings of equal radial width, such that the rings cover the useful region of the iris. These eight rings can be seen in Figure 1 below.
- Apply a series of Gabor filters to the pixels found in these 8 rings, using polar coordinates. Each Gabor convolution is at some characteristic radius r , angle θ , with analysis scales α , β , and ϖ . The result of one such convolution integral is a complex number N . Two bits are obtained from N by extracting its real and imaginary parts, then retaining only the sign of each part. Equations (4)-(7) in the patent [1] describe this analytic approach.
- The resulting code for an iris contains 2048 bits; this is clearly some combination of 1024 separate convolutions on the pixel values found in the iris bands.

Pixel Analysis of a Typical Iris Partitioning

See figures 1 and 2. Figure 1 is taken from the Daugman patent [1]. It shows an iris image and a set of eight concentric rings. The four inner rings (which we'll index 1..4) are nearly continuous, covering approximately 330° in angle. The next two rings (5 and 6) cover approximately 226° in all, in two segments. The last two rings (7 and 8) cover approximately 180°, also in two segments. The box in the upper left contains a graphic depiction of 2048 code bits reduced from this image.

The width of each (measured radially) is approximately 13 pixels. Their lengths (measured along the centerline periphery) are given in Figure 2, in pixels. The area, in square pixels, is also given in Figure 3. The total lengths and areas are roughly equal among the rings, with an average length of 856 pixels and an average area of 10,000 sq. pixels. Although the inner rings are more nearly complete, their radii are smaller.

For other iris images, these figures will clearly vary, depending on the camera distance and the visual accommodation of the candidate's pupil to incident lighting. This image appears to have been taken in relatively strong incident lighting, since the pupil is relatively small. An image taken under low lighting conditions will show a larger pupil diameter and a correspondingly smaller iris region. One might expect the average length of the inner bands to increase somewhat, perhaps to 1000 pixels, as the pupil expands. The area might decrease by a factor of two or more, since it's more sensitive to the ring width.

An examination of the image suggests that the iris patterns have a strong radial component. By this we mean that the pixels along a radius are more strongly correlated than those along a circle. This observation is supported by the known structure and growth patterns of the iris filaments. The pattern is formed from a striated trabecular meshwork of

chromatophore and fibroblast cells. This region of the eye is made up of continuous cell matter at birth, and then tears into its characteristic filamentary pattern shortly after birth.




Figure 1. Typical iris image (from the Daugman patent [1])

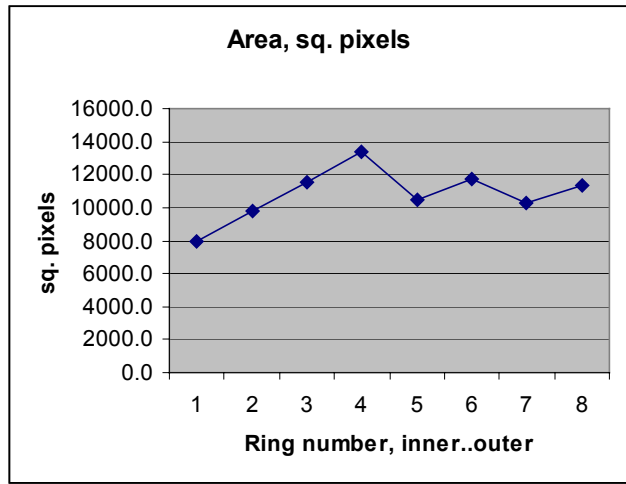
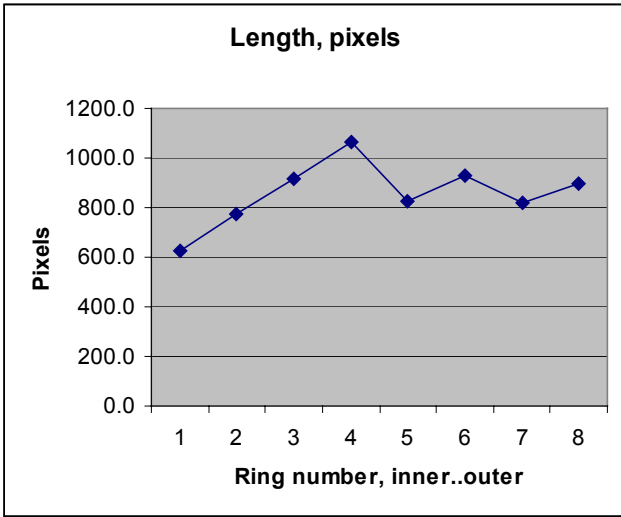


Figure 2. Length of rings and area of rings, from the image in Figure 1. Assumes image is 640 x 480 pixels.

Focus Quality

Daugman is aware that a good focus is vital to the overall process, but he does not describe any tools for judging the quality of the focus. The IriScan 2100 unit software [6] is nevertheless designed to judge the focus quality somehow by examining successive video images. When it decides on an image with an appropriate focus, it grabs that image and submits it for subsequent processing. There is no turning back once the image is grabbed. The subsequent algorithmic steps are carried out with few or no quality checks along the way. A proof of that can be found in several images in which the focus is clearly very poor, yet DA was run through to completion with no complaints, yielding a spurious iris code [4].

The IriScan company [6] appears to be aware of this problem. They urge a careful visual inspection of each enrollment image to ensure that it has a good focus and that DA shows a normal result. However, this requires a trained person in addition to the candidate. For recognition, only the presumably untrained candidate is present, and acquiring a quality image is therefore primarily the responsibility of the unit and its software. If the focus is poor during recognition, the system will likely reject the candidate and demand another image. We feel that this phenomenon is responsible for the relatively high level of false rejections (about 5-10%) we've seen with the 2100 system [4, 7].

The unit sometimes grabs a frame in which the best focus is on the eyelashes or the rim of eyeglasses. Since the depth of focus is small, this yields an out-of-focus iris. Other sources of false rejection include stray light falling into the iris or pupillary area, and reflections on the lens or frame of eyeglasses. Yet another source of false rejection occurs when the candidate's eye shifts suddenly during capture. This results in a Venetian-blind effect with the 2100, which uses the NTSC interlaced video standard [4].

These aspects of the IriScan system clearly need engineering improvement. The frame-grabbing system should somehow determine a better focus on the *iris*, not on the image in general. The subsequent DA should also examine the general quality of the focus through the encoding found in the iris bands, and reject an image that fails to meet a reasonable quality standard. For example, it's reasonable to expect a certain level of variability of intensities above a background noise level throughout the iris bands. Failing to find that is an indication of a poor focus condition or a failure to correctly locate the pupil.

Gabor Filtering

Assuming that the image is in good focus (or not, as the case may be), and that we have a reasonable estimate of the pupillary and sclera diameters and centers, the iris area is next subdivided (in DA) into 8 concentric rings. Each ring has the same radial size, but may cover different angular regions. (See Figure 1). The pixel width of a ring and its angular span will vary from image to image depending on the position of the upper eyelid. DA as implemented on the 2100 has an extension that looks for the upper eyelid, reducing the ring spans as needed so that eyelid pixels are not included in the coding.

A linear transform then maps the pixels in each ring from a polar coordinate to a Cartesian coordinate system. DA then applies a set of Gabor filters to the now-rectangular space of transformed pixels, looks at the signs of the real and imaginary parts, and generates a binary code representing the core patterns of the iris.

We expect a strong correlation of pixels along the radial direction, but a weak correlation in the angular direction. This is because the iris is formed of radial fibres that appear as radial streaks in any image. Although DA claims that a 2D image filtering is used, my speculation is that in fact a 1D filter is applied, one within each of the 8 rings. The subdivision of the iris into eight radial regions is probably sufficient partitioning of the radial information, given the strong radial correlation.

DA yields 2048 bits of code. Since each convolution yields 2 bits, this implies 1024 convolutions. With 8 bands, this is 128 convolutions per band. The question is how are these 128 convolutions partitioned? We can have one radial partition with 128 angular, or two radial with 64 angular, four radial with 32 angular, etc. (There's no such thing as a fractional partitioning, so only these combinations of the factors of 128 are feasible).

Also look at the number of pixels involved in each ring. There are about 13 pixels radially and 800 angularly (see Figure 2), for a ratio of 1:62. Suppose that the relative correlation in the angular direction to the radial direction is at least 4:1 (this can be measured experimentally, and seems reasonable from typical iris images). Then we should use $4 \times 62 = 248$ times as many convolutions in the angular direction to each one in the radial direction. This strongly suggests that DA uses a *single* partition in each band radially, and all 128 *angularly*, and that is essentially a 1D convolution. There's nothing to be gained by folding a Gabor filter along the radius; we might as well compute an average along each radius instead.

This is computationally more efficient. By choosing a *fixed* number of pixels in each band, we can also greatly simplify each convolution. (Note that the variation in band length is not very large, something like 800 ± 200). The filtering can be done in integer mode, as Daugman suggests, using normalized tables of Gabor coefficients as a relatively modest exercise in computer arithmetic.

Boundary Conditions

An issue not discussed in DA is how to deal with the endpoint boundaries of the rings. In the real domain, a Gabor filter extends infinitely far in each direction, and the data is also assumed to extend infinitely far. (The convolution is also a real-domain integration, of course). In the finite domain, the filter, the source data, and the integration are finite and discrete. The discreteness property makes it possible to compute the convolutions with a finite number of algorithm steps, but also introduces a certain level of quantization noise in the result. Also, a discrete Gabor filter may not average to zero, so some adjustment to the filter values is important to yield a zero average (Daugman alludes to this point).

But what should be done about the end regions of the rings? Given that each Gabor filter has a finite domain, i.e. $\pm d$ pixels, we can simply define the convolution center range as $(d, D-d)$, where the ring domain is $(0, D)$. This ensures that each convolution sum operates over a valid set of pixels. We are arguably losing some information by not carrying the filtering

out to each end. On the other hand, trying to extend the range beyond the endpoints would require some extrapolation of the range, which can't be done safely.

One-dimensional Range Partitioning

Algorithm DA now requires an additional partitioning of the 128 degrees of freedom. We're assuming that we have a one-dimensional array of (averaged) pixels, of approximate dimension 800. Note that since a mapping from Cartesian (the raw image) to polar form (an iris ring) is required anyway, we could simply fix this dimension for the sake of all convolutions and use linear interpolation to fill in the eight linear arrays.

The degrees of freedom we have are in the frequency ω , position θ and extent β of the Gabor filters (see Eqs. 4-7 in the patent [1]). Frequency and extent are continuous variables. The number of positions (implied in θ) is a discrete variable, and the product of this with the *total number* of other variable choices must be 128.

We introduce a Gabor angular wavelength λ , which is related to the Gabor frequency by the wave equation $\lambda = 2\pi/\omega$. We can also measure the wavelength in pixels, which makes it easy to estimate a range of wavelengths to use.

One Frequency?

Our guess is that only *one* frequency (or wavelength) and extent is used for all the Gabor wavelets in DA. That makes the computational task especially easy, but this may not be optimal. (Daugman suggests that several frequencies are used, and recommends a set of frequencies with a constant logarithmic interval).

Given that 128 positions are required in a span of 800 pixels, we might expect that the wavelength should be in the neighborhood of $800/128 = 6.25$ pixels. Note that this is about half the *width* of each ring. If we use a shorter wavelength, the wavelets will not overlap very much, leaving some pixel area uncovered, but will yield less correlation of neighboring wavelets. A shorter wavelength will also make the filtering more sensitive to noise. A longer wavelength will increase the wavelet overlap, causing an increase in the correlation of neighboring wavelets, and also provide better noise filtering.

The Gabor extent should be considered an independent variable, although Daugman suggests that the extent be approximately equal to the wavelength. If the extent is much larger than the wavelength, then many cycles are considered at any one position, and in the limit approaches a Fourier filter. Conversely, a small extent is little better than a unit impulse, averaging a few values, but not contributing much in the way of a frequency-based filtering.

There is clearly some optimal wavelength and extent which can only be judged through an empirical study using the algorithm with varying wavelength and a representative sample of iris images.

Or Multiple Frequencies?

The use of multiple frequencies obviously introduces more dimensional variations in the overall problem of optimization. Yet the space of variations is not necessarily that large. Given 128 degrees of freedom, we can only choose a power of two for the number of frequencies. For 2 frequencies, 64 positions are possible (12.5 pixels each). With 4 frequencies, 32 positions (25 pixels each) are possible. We can probably assume that the extent of a wavelet should vary directly with its wavelength, although this, too, is an issue that can only be resolved through a careful optimization study.

There's no *a priori* reason to choose one *vs.* several wave lengths. A useful experiment is to explore the space of possibilities for optimality, *i.e.* find a combination that yields the least cross-over probability of the authentic and imposter distributions.

Pupil Finding

Given some arbitrary image, DA apparently runs a fairly large number of experiments computing expression (2) in the patent [1]. This computes the gradient of the average pixel intensity around a circular band of radius r centered at some point (x, y) . Three variables are involved here, and the problem is to explore this three-dimensional space for a gradient maximum. One can also impose some constraints, for example, that the pupil center be within a certain sub-rectangle of the image, and that its radius also be within certain upper and lower bounds. One might also require that the field *within* the pupil be uniformly dark, at least over (say) 80% of the pupil. That still leaves a large space of parameters.

A group of senior CmpE students, led by David Yang, investigated this problem in 1999 and were able to confirm Daugman's approach, but not improve upon it [5]. They used a grid of (x, y) pairs (220 x 145 pixels, with 10 pixel steps) with a set of radii r (minimum 40, maximum 80, with 3 pixel steps). Thus approximately 4400 candidate circle-pairs were examined in each image. Every point on each circle is used, through a Bresenham algorithm. The circle-pair that yields the largest difference is found. They didn't report on how well this worked with a large class of images, but they apparently looked at a few different images.

They also investigated the problem of eliminating bright reflections. A bright reflection within the pupillary area can sometimes cause a false maximum gradient to appear, so something needs to be done about it. A "bright reflection" (from their point of view) is a circular area of at least 3 pixels radius in which all the pixels have the value 255 (maximum white). These are just changed to 0. Note that if the bright area has some pixels at (say) 250-254 value, which is still "bright"

compared to the pupillary intensities, these are not affected. My feeling about this approach is that it is crude and very likely to fail on eye images in which the contrast between the pupil and the iris is low. This approach is evidently used by IriScan in their model 2100 software, and we've seen cases in which it fails to properly locate the pupil.

Suggested Research Topics

1. Lam, Pham and Yang [5] wrote some software that locates the pupil. It needs to be worked over to clean it up and make it available as a reliable documented software module.
2. Complete our own software to the point of delivering a binary code for an iris pattern. This code will almost certainly not be compatible with those delivered by the 2100 unit (or other commercial systems), but can be used to find an optimal Gabor filtering, and to independently evaluate the quality of this system for identification and recognition.
3. Investigate whether a better approach to eliminating bright spots in the pupil can be found. These spots don't interfere with DA as a rule, but occasionally cause a failure in pupil finding and a subsequent failure of DA.
4. Investigate the problem of establishing that the focus conditions in capture are suitable. We feel that this requires pushing the algorithm through pupil and sclera location, then applying some selection criterion to the pupillary area and the iris area. This will not necessarily reduce the number of false recognitions, but should at least assign responsibility for a sure failure to the capture conditions.
5. Investigate what Gabor filtering conditions seem optimal for the system, given a representative sample of iris images (which we have).
6. Measure the relative correlation of the radial patterns vs. the angular patterns. My surmise is that the angular pattern correlation is at least four times that of the radial patterns.
7. Independently evaluate the iris scanning technology through our own tools, and be in a position to assess the overall quality of the technology, along with caveats on the conditions that must be met for best results.
8. Investigate other issues for which the iris scanning technology may be sensitive, i.e. head rotation, stray lighting, influence of focus, influence of eyelashes, etc.
9. Investigate the maximum practical distance at which an iris image can be reliably acquired. This comes down to such fundamentals as the optical power, resolution and angular stability of a suitable telescope, coupled with problems with air turbulence and lighting conditions. The commercial IriScan equipment [5] operates with a built-in infrared illuminator positioned at a critical angle with respect to the candidate's eye. There are no studies of iris technology using various forms of ambient illumination, nor at distances greater than a few feet.

References

- [1] United States Patent 5,291,560, issued Mar. 1, 1994, to John G. Daugman.
- [2] John G. Daugman, *High Confidence Visual Recognition of Persons by a Test of Statistical Independence*, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. 15, No. 11, November 1993.
- [3] John G. Daugman and Cathryn J. Downing, *Demodulation, predictive coding, and spatial vision*, Journal of the Optical Society of America A, Vol 12, No. 4, April 1995.
- [4] William A. Barrett, *IriScan Evaluation*, Biometrics Test Center, San Jose State University, March 11, 1999. Available on request from the author.
- [5] Lucian Lam, Giao Pham, and David Yang, *Human Iris Pattern Classification*, written to satisfy the requirements of a senior computer engineering project requirement. Available on request from the author at SJSU.
- [6] IriScan, Inc., Marlton, New Jersey. See their web site at <http://www.iriscan.com> for more information. Their model 2100 is an early product of this company, and is no longer manufactured. IriScanTM is a registered trademark of IriScan, Inc.
- [7] F. Bouchier, J. Ahrens, and G. Wells, *Laboratory Evaluation of the IriScan Prototype Biometric Identifier*, Sandia National Laboratories, SAND96-1033, April 1996.