

## DEMODULATION BY COMPLEX-VALUED WAVELETS FOR STOCHASTIC PATTERN RECOGNITION

JOHN DAUGMAN

*The Computer Laboratory, University of Cambridge, Cambridge, UK*  
*john.daugman@cl.cam.ac.uk*

Received 18 November 2002

Samples from stochastic signals having sufficient complexity need reveal only a little unexpected shared structure, in order to reject the hypothesis that they are independent. The mere failure of a test of statistical independence can thereby serve as a basis for recognizing stochastic patterns, provided they possess enough degrees-of-freedom, because all unrelated ones would pass such a test. This paper discusses exploitation of this statistical principle, combined with wavelet image coding methods to extract phase descriptions of incoherent patterns. Demodulation and coarse quantization of the phase information creates decision environments characterized by well-separated clusters, and this lends itself to rapid and reliable pattern recognition.

*Keywords:* Demodulation; phase; Gabor wavelets; iris; biometric; pattern recognition.

### 1. Introduction

The central issue in pattern recognition is the relationship between within-class variability and between-class variability. These are determined by the dimensions of variation (degrees-of-freedom) spanned by the pattern classes. Ideally the within-class variability should be small and the between-class variability large, since this creates optimal separation among the pattern classes. The reliability of pattern recognition decisions depends upon the separation, or amount of overlap, among the different pattern classes; we desire that the spacings between the clusters be larger than the diameters of the clusters. This statistical separation in turn depends partly upon the representation chosen for defining the classes. It is desirable to find image representations which lend themselves optimally to these statistical requirements of pattern recognition. This paper discusses the coupling of wavelet image coding with a test of statistical independence on extracted phase information, in order to obtain a demonstrably robust and reliable algorithm for recognizing stochastic patterns of high dimensionality.

## 2. Complex-Valued 2D Wavelets for Image Analysis

The use of multi-resolution two-dimensional wavelets in image analysis and computer vision has attracted much interest in recent years. The particular family of 2D wavelets that will be discussed here are closed under dilation, translation, and rotation, and they also form a complete image basis, or frame. They are also self-Fourier (equivalent in form to their own Fourier transforms), and they compose a set that is also closed under convolution (the convolution of any two members of the set is also a member of the set).

If we take  $\Psi(x, y)$  to be any chosen generic 2D wavelet, which may be called a mother wavelet, then we can generate from this function a complete self-similar family of parametrized daughter wavelets  $\Psi_{mpq\theta}(x, y)$  through the generating relation

$$\Psi_{mpq\theta}(x, y) = 2^{-2m}\Psi(x', y'), \quad (2.1)$$

where the substituted variables  $(x', y')$  incorporate dilations of the wavelet in size by  $2^m$ , translations in position  $(p, q)$ , and rotations through angle  $\theta$ :

$$x' = 2^{-m}[x \cos(\theta) + y \sin(\theta)] - p, \quad (2.2)$$

$$y' = 2^{-m}[-x \sin(\theta) + y \cos(\theta)] - q. \quad (2.3)$$

A particular choice for  $\Psi(x, y)$  which possesses several interesting and useful properties is the complex-valued 2D ‘‘Gabor’’ wavelet, so named because it is a generalization<sup>2</sup> of the 1D elementary functions originally discussed by Gabor.<sup>8</sup> The 2D wavelet form is defined as follows:

$$\Psi(x, y) = e^{-\pi[(x-x_0)^2/\alpha^2+(y-y_0)^2/\beta^2]} e^{-2\pi i[u_0(x-x_0)+v_0(y-y_0)]}, \quad (2.4)$$

where  $(x_0, y_0)$  specify wavelet position,  $(\alpha, \beta)$  specify effective width and length, and  $(u_0, v_0)$  specify a modulation wave-vector which can be interpreted in polar coordinates as spatial frequency  $\omega_0 = \sqrt{u_0^2 + v_0^2}$  and orientation (or direction)  $\theta_0 = \arctan(v_0/u_0)$ . Plots of the real and imaginary parts of such wavelets are shown in Fig. 1.

The 2D Fourier transform  $F(u, v)$  of the 2D Gabor wavelet  $\Psi(x, y)$  has exactly the same functional form (i.e. this family of wavelets is self-Fourier), with parameters just interchanged or inverted:

$$F(u, v) = e^{-\pi[(u-u_0)^2\alpha^2+(v-v_0)^2\beta^2]} e^{2\pi i[x_0(u-u_0)+y_0(v-v_0)]}. \quad (2.5)$$

Thus the 2D Fourier power spectrum of the  $\Psi(x, y)$  wavelet,  $F(u, v)F^*(u, v)$ , is a bivariate Gaussian centered on  $(u_0, v_0)$ . Its spectral energy peak is at orientation  $\theta_0$  and spatial frequency  $\omega_0$  as defined above, and so it serves to extract image structure in a particular band of the 2D Fourier spectrum. The effective support of this 2D spectral band-area is an ellipse centered at  $(u_0, v_0)$  and whose principle axes are  $(1/\alpha, 1/\beta)$ .




Fig. 1. Pattern encoding by phase demodulation. Image structure is extracted as phase sequences by projection onto multi-resolution complex-valued 2D wavelets.

It is noteworthy that as consequences of the similarity, shift, and modulation theorems of 2D Fourier analysis, together with the rotation isomorphism of the 2D Fourier transform, all the effects of the generating relation (2.1) applied to a 2D Gabor mother wavelet  $\Psi(x, y)$  to produce a new daughter wavelet  $\Psi_{mpq\theta}(x, y)$  will have corresponding or reciprocal effects on its 2D Fourier transform  $F(u, v)$  without any change in functional form.<sup>3</sup> This set of 2D wavelets and their 2D Fourier transforms is closed under the transformations of dilation, translation, rotation, and convolution with any member of the set. Two further interesting properties of these wavelets are the fact that they achieve the lower bound of the Weyl–Heisenberg “uncertainty relation” for conjoint resolution in the 2D space and 2D Fourier domains,<sup>3</sup> and the fact that they form an excellent model for the receptive field profiles of individual neurones (the so-called “simple cells”) in the visual cortex of mammalian brains.<sup>2,3</sup> Any image can be represented completely in terms of such wavelets used as expansion functions. A complication arises from the fact that these wavelets are non-orthogonal (the inner product of any two of them is in general nonzero), and so the expansion coefficients required do not correspond simply to the inner product projections of the image onto the wavelets. A relaxation network solution to this problem of obtaining the correct expansion coefficients was presented in Ref. 4.

Because these wavelets are complex valued, it is possible to use the real and imaginary parts of their convolution  $(*)$  with an image  $I(x, y)$  to extract a description of image structure in terms of local modulus and phase. These correspond, respectively, to an “amplitude modulation” function  $A(x, y)$  and a “phase

modulation” function  $\phi(x, y)$  of the spatial image coordinates:

$$A(x, y) = \sqrt{(\operatorname{Re}\{\Psi(x, y) * I(x, y)\})^2 + (\operatorname{Im}\{\Psi(x, y) * I(x, y)\})^2}, \quad (2.6)$$

$$\phi(x, y) = \tan^{-1} \frac{(\operatorname{Im}\{\Psi(x, y) * I(x, y)\})}{(\operatorname{Re}\{\Psi(x, y) * I(x, y)\})}. \quad (2.7)$$

These polar descriptors are shown in the phasor diagram of Fig. 1. The modulus of the phasor in the complex plane represents a given patch of an image in terms of its local contrast amplitude  $A(x, y)$ , and the angle of this phasor represents the patch in terms of its local phase  $\phi(x, y)$ . This phase angle can be quantized very coarsely, as shown in Fig. 1 using only two bits, for building descriptions of a pattern that lend themselves to efficient and reliable pattern recognition by the simple failure of a test of statistical independence.

We turn now to illustrating this principle in a particular application in image analysis and computer vision, namely the recognition of iris patterns in a person’s eye as a reliable method of automatic personal identification.

### 3. Biometric Recognition of Persons by Iris Patterns

The highest density of biometric degrees-of-freedom (forms of variability among individuals) which are both stable over time and easily measured, is to be found in the complex texture of the iris pattern of the eye. This protected internal organ, whose pattern can be encoded from distances of up to almost a meter, reveals about 250 independent degrees-of-freedom of textural variation across individuals. One way to calibrate the “information density” of the iris is by its human-population entropy per unit area. As we will see, this works out to about 3.2 bits per square millimeter on the iris, based upon 9.1 million paired IrisCode comparisons that have been performed using these algorithms.




Fig. 2. Example of a human iris pattern, imaged in near infrared light at a distance of 30 cm. Such patterns have high statistical dimensionality and can serve as unique, reliable identifiers.

### 3.1. Properties of the iris

The iris is composed of elastic connective tissue, the trabecular meshwork, whose prenatal morphogenesis is completed during the 8th month of gestation. It consists of pectinate ligaments adhering into a tangled mesh revealing striations, ciliary processes, crypts, rings, furrows, a corona, sometimes freckles, vasculature, and other features. During the first year of life a blanket of chromatophore cells often changes the color of the iris, but the available clinical evidence indicates that the trabecular pattern itself is stable throughout the lifespan. Because the iris is a protected internal organ of the eye, behind the cornea and the aqueous humor, it is immune to the environment except for its pupillary reflex to light. (The elastic deformations that occur with pupillary dilation and constriction are readily reversed mathematically by the algorithms for localizing the inner and outer boundaries of the iris.) Pupillary motion, even in the absence of illumination changes (termed *hippus*), and the associated elastic deformations in the iris texture, provide one test against photographic or other simulacra of a living iris in high security applications. There are few systematic variations in the amount of detectable iris detail as a function of ethnic identity or eye color; even dark-eyed persons reveal plenty of iris detail when imaged with infrared light. Further discussion of anatomy, physiology, and clinical aspects of the iris may be found in Ref. 1.

### 3.2. Localizing irises and analyzing their patterns

The two-dimensional modulations which create iris patterns are extracted by *demodulation*<sup>7</sup> with complex-valued 2D wavelets, as discussed above and illustrated in Fig. 1, albeit in polar rather than cartesian coordinates.

First it is necessary to localize precisely the inner and outer boundaries of the iris, and to detect and exclude eyelids if they intrude. These detection operations are accomplished by integro-differential operators of the form

$$\max_{(r, x_0, y_0)} \left| G_\sigma(r) * \frac{\partial}{\partial r} \oint_{r, x_0, y_0} \frac{I(x, y)}{2\pi r} ds \right|, \quad (3.1)$$

where contour integration parametrized for size and location coordinates  $r, x_0, y_0$  at a scale of analysis  $\sigma$  set by some blurring function  $G_\sigma(r)$  (e.g. a Gaussian of scale  $\sigma$ ) is performed over the image data array  $I(x, y)$ . The result of this optimization search is the determination of the circle parameters  $r, x_0, y_0$  which best fit the inner and outer boundaries of the iris.

Then a doubly-dimensionless coordinate system is defined which maps the tissue in a manner that is invariant to changes in pupillary constriction and overall iris image size, and hence also invariant to camera optical zoom factor and distance to the eye. This coordinate system is pseudo-polar, although it does not assume concentricity of the inner and outer boundaries of the iris since the pupil is normally somewhat nasal, and inferior, in the iris. The coordinate system compensates automatically for the stretching of the iris tissue as the pupil dilates. The inner




Fig. 3. Isolation of an iris for encoding, and its resulting “IrisCode.”

and outer boundaries of this dimensionless and non-concentric coordinate system are illustrated graphically in Fig. 3, together with a phase-demodulation IrisCode indicated in the top left as a bit stream.

The detailed pattern is encoded into a 256-byte “IrisCode” by demodulating it with 2D Gabor wavelets. Each resulting phasor angle in the complex plane is quantized to the quadrant in which it lies for each local element of the iris pattern, and this operation is repeated all across the iris, at many different scales of analysis. Such local phase quantization is described by the following conditional integral equations, in which each code bit  $h$  is represented as having both a “real part”  $h_{\text{Re}}$  and an “imaginary part”  $h_{\text{Im}}$ , with  $h = h_{\text{Re}} + ih_{\text{Im}}$ , and the raw iris image pixel data is given in the dimensionless pseudo-polar coordinate system  $I(\rho, \phi)$ :

$$h_{\text{Re}} = \begin{cases} 1 & \text{if } \text{Re} \int_{\rho} \int_{\phi} e^{-i\omega(\theta_0-\phi)} e^{-(r_0-\rho)^2/\alpha^2} e^{-(\theta_0-\phi)^2/\beta^2} I(\rho, \phi) \rho d\rho d\phi \geq 0, & (3.2) \\ 0 & \text{if } \text{Re} \int_{\rho} \int_{\phi} e^{-i\omega(\theta_0-\phi)} e^{-(r_0-\rho)^2/\alpha^2} e^{-(\theta_0-\phi)^2/\beta^2} I(\rho, \phi) \rho d\rho d\phi < 0. & (3.3) \end{cases}$$

$$h_{\text{Im}} = \begin{cases} 1 & \text{if } \text{Im} \int_{\rho} \int_{\phi} e^{-i\omega(\theta_0-\phi)} e^{-(r_0-\rho)^2/\alpha^2} e^{-(\theta_0-\phi)^2/\beta^2} I(\rho, \phi) \rho d\rho d\phi \geq 0, & (3.4) \\ 0 & \text{if } \text{Im} \int_{\rho} \int_{\phi} e^{-i\omega(\theta_0-\phi)} e^{-(r_0-\rho)^2/\alpha^2} e^{-(\theta_0-\phi)^2/\beta^2} I(\rho, \phi) \rho d\rho d\phi < 0. & (3.5) \end{cases}$$

### 3.3. Independence and the degrees-of-freedom in IrisCodes

It is important to establish that there exists independent variation in iris patterns, across populations and across positions in the IrisCode. This is confirmed by tracking the probability of a bit being set, as shown in Fig. 4. If there were any systematic

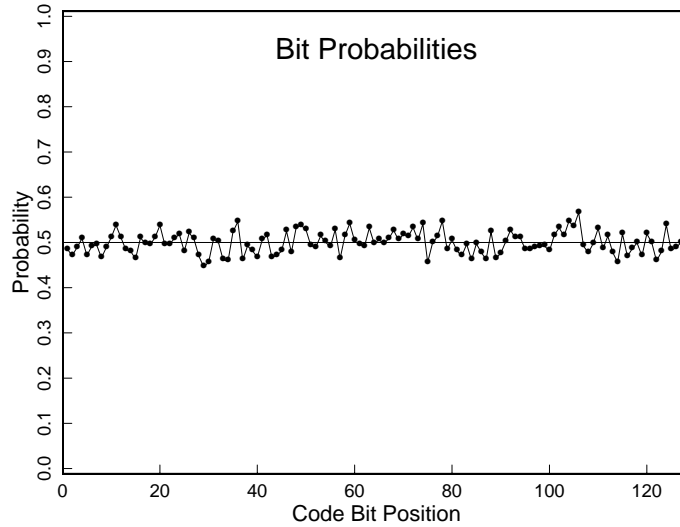


Fig. 4. Equiprobability of set bits across IrisCode bit positions.

correlations among irises, this plot would not be flat. The fact that it is flat at a value of 0.5 means that any given bit in an IrisCode is equally likely to be set or clear, and so IrisCodes are bit-wise maximum entropy codes.

The histogram in Fig. 5 compares different eyes' IrisCodes by Exclusive ORing them to detect the fraction of their bits that disagree. Since any given bit is equally likely to be set or clear, an average Hamming Distance fraction of 0.5 would be expected. The observed mean was 0.499 in comparisons between 9.1 million different pairings of IrisCodes captured in field tests of these algorithms in Britain, the USA, Japan and Korea. Altogether there were 4,258 different iris images in the combined database, including 10 each of one subset of 70 eyes. Excluding those duplicates of  $(700 \times 9)$  same-eye comparisons, and not double-counting pairs, and not comparing any image with itself, the total number of unique pairings between different eye images whose HDs could be computed was  $((4,258 \times 4,257 - 700 \times 9)/2) = 9,060,003$ . The standard deviation of this distribution, 0.0317, indicates that the underlying number of degrees-of-freedom in such comparisons is  $N = pq/\sigma^2 = 249$ . This indicates that within any given IrisCode, only a small subset of the 2,048 bits computed are independent of each other, due to the large correlations (mainly radial) that exist within any given iris pattern. (If every bit in an IrisCode were independent, then the distribution in Fig. 5 would be very much sharper, with an expected standard deviation of only  $\sqrt{pq/N} = 0.011$ ; thus the Hamming Distance interval between 0.49 and 0.51 would contain most of its area.) The solid curve fitted to the data is a binomial distribution with 249 degrees-of-freedom; this is the expected distribution from tossing a fair coin 249 times in a row, and tallying up the fraction of heads in each such run. The factorials which dominate the tails of




Fig. 5. Histogram of raw Hamming Distances between 9.1 million pairs of unrelated IrisCodes. The solid curve is a binomial distribution with 249 degrees-of-freedom.

such a distribution make it astronomically improbable that two different IrisCodes having these many degrees-of-freedom could accidentally disagree in much fewer than half their bits. For example, the chances of disagreeing in only 25% or fewer of their bits (achieving a Hamming Distance below 0.25, or equivalently the chances of getting fewer than 25% heads in 249 coin tosses) are less than one in  $10^{15}$ . Thus the observation of a match even of such poor quality (25% of the bits disagree) is extraordinarily compelling evidence of identity.

### 3.4. *Genetically identical irises*

Whereas the striking visual similarity of identical twins reveals the genetic penetrance of overall facial appearance, a comparison of genetically identical irises reveals that iris texture is a phenotypic feature, not a genotypic feature. A convenient source of genetically identical irises are the right and left pair from any given person. Such pairs have the same genetic relationship as the four irises of two identical twins, or indeed in the probable future, the  $2N$  irises of  $N$  human clones. Eye color of course has high genetic penetrance, as does the overall statistical quality of the iris texture, but the textural details are uncorrelated and independent even in genetically identical pairs. This is shown in Fig. 6, comparing 648 right/left iris pairs from 324 persons.



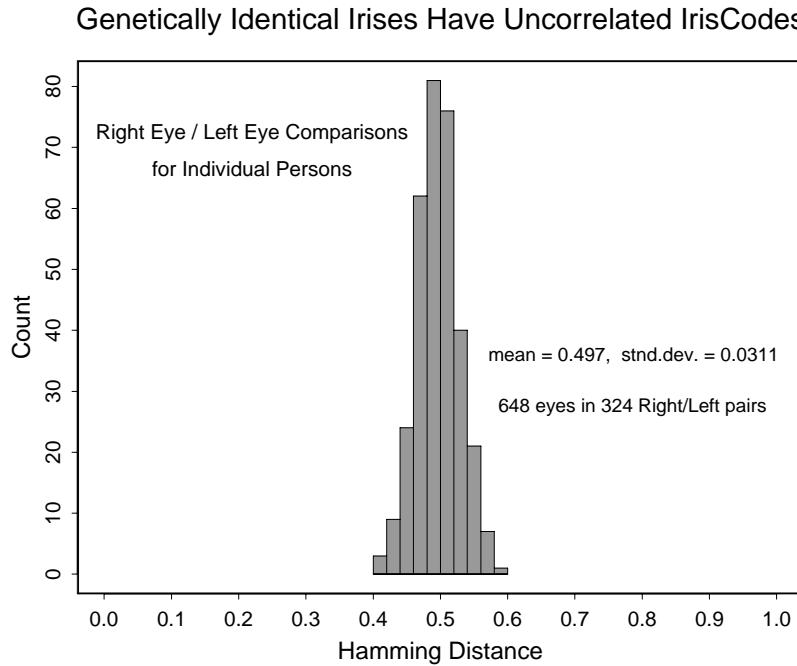


Fig. 6. Histogram of raw Hamming Distances between IrisCodes computed from 324 pairs of genetically identical irises (648 eyes in right/left pairs). This distribution is statistically indistinguishable from Fig. 5, which compared unrelated irises.

The mean Hamming Distance is 0.497 with standard deviation 0.031, indicating 259 degrees-of-freedom between genetically identical irises. These results are statistically indistinguishable from those shown in Fig. 5 for genetically unrelated irises. This shows that the detailed phase structure extracted from irises by phasor demodulation is entirely phenotypic, so performance is not limited (as it is for face recognition or DNA) by the birth rate of identical twins.

#### 4. Statistical Recognition Principle

The principle of operation underlying this approach to pattern recognition is the *failure* of a test of statistical independence. Samples from stochastic sequences with sufficient complexity need reveal only a little unexpected agreement, in order to reject the hypothesis that they are independent. For example, in two runs of 1,000 coin tosses, agreement rates between their paired outcomes higher than 56% or lower than 44% are extremely improbable: the odds against a higher or lower rate of agreement are roughly 10,000 to 1. The failure of a test of statistical independence can thereby serve as a basis for recognizing patterns very reliably, provided they possess enough degrees-of-freedom. Since the bitwise comparisons of unrelated IrisCodes are Bernoulli trials, their combinatorics generate Hamming

Distance distributions which are binomials, even though the underlying Bernoulli trials are correlated (see Ref. 9) due to internal correlations within each iris. With so many degrees-of-freedom, the binomial distributions have attenuation rates that are dominated by large factorials. For this key reason, iris patterns allow recognition decisions about personal identity to be made with astronomic confidence levels. The practical importance of such high odds against a chance False Match is that it permits huge databases (even of “planetary” size) to be searched exhaustively, without accumulating significant probability of a False Match despite the large number of opportunities.

#### 4.1. *Extreme-value distribution for rotated IrisCodes*

Because the computation of an IrisCode is preceded by localization of the inner and outer boundaries of the iris and the creation of a doubly-dimensionless coordinate system, an IrisCode is invariant under translations and dilations (changes in object position or size within the image), and also invariant under changes in the pupil diameter relative to the iris diameter. However, the phasor information scrolls in phase as the iris is rotated, due to tilt of the head or camera or due to torsional rotation of the eye in its socket. Therefore all iris comparisons need to be repeated over a range of relative rotations, keeping only the best match. This amounts to sampling the distribution of Fig. 5 many times and keeping only the smallest value, which leads to the extreme-value distribution given in Fig. 7.

The raw binomial distribution shown earlier in Fig. 5 had the form:

$$f(x) = \frac{N!}{m!(N-m)!} p^m q^{(N-m)}, \quad (4.1)$$

where  $N = 249$ ,  $p = q = 0.5$  and  $x = m/N$  is the Hamming Distance. Let  $F_0(x)$  be its cumulative from the left, up to  $x$ :  $F_0(x) = \int_0^x f(x)dx$ . When only the smallest of  $n$  samples from such a distribution is kept, the resulting extreme-value distribution (derived in Ref. 5) has density  $f_n(x)$ :

$$f_n(x) = n f(x) [1 - F_0(x)]^{n-1} \quad (4.2)$$

which is the solid curve superimposed upon the empirical histogram in Fig. 7. Figure 8 shows this extreme value probability density function and marks off the cumulatives under its left tail up to various points, illustrating that, for example, finding accidental agreement of two unrelated IrisCodes in even 65% or more of their bits (a Hamming Distance of 0.35 or lower) has very small probability (1 in 133,000). This illustrates that we can tolerate a huge amount of corruption in iris images due to poor resolution, poor focus, occluding eyelashes and eyelids, contact lenses, specular reflections from the cornea or from eyeglasses, camera noise, etc. We can accept matches of very poor quality, say up to 30% of the bits being wrong, and still make decisions about personal identity with very high confidence.




Fig. 7. Histogram of Hamming Distances between unrelated IrisCodes computed after comparisons in multiple ( $n = 7$ ) relative rotations, keeping only each best match.

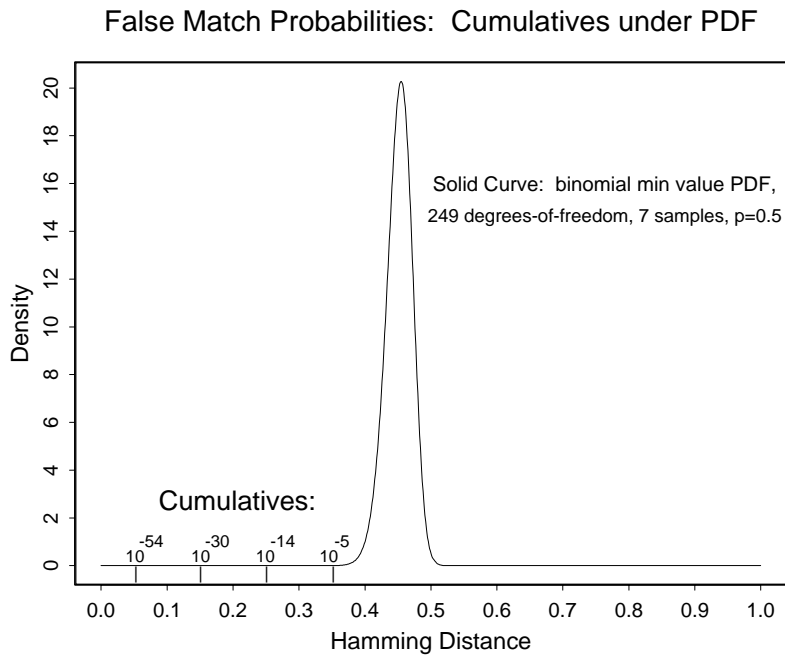


Fig. 8. Theoretical density function for the derived binomial distribution for best IrisCode matches after multiple ( $n = 7$ ) relative rotations.




Fig. 9. Decision environment for automatic personal identification based on iris patterns, using these algorithms in typical operational conditions.

#### 4.2. *Decidability of iris-based personal identification*

The overall *decidability* of the task of recognizing persons by their iris patterns is revealed by comparing the Hamming Distance distributions for “same” versus “different” irises. To the degree that one can confidently decide whether an observed sample belongs to the left (same) or the right (different) distribution in Fig. 9, this recognition task can be reliably performed.

For such a decision task, the Decidability Index  $d'$  measures how well separated the two distributions are, since recognition errors are caused by their overlap. If their two means are  $\mu_1$  and  $\mu_2$ , and their two standard deviations are  $\sigma_1$  and  $\sigma_2$ , then  $d'$  is defined as

$$d' = \frac{|\mu_1 - \mu_2|}{\sqrt{(\sigma_1^2 + \sigma_2^2)/2}}. \quad (4.3)$$

This measure of decidability (or detectability) is independent of how liberal or conservative is the acceptance threshold used. Instead it reflects the degree to which any improvement in (say) the False Match error rate must be paid for by a worsening of the Failure-to-Match error rate. The measured decidability is  $d' = 7.3$  for this particular Decision Environment, which was created by non-ideal image capture conditions. Different cameras were often used for enrollment and for recognition tests, and so the left-side distribution for “same” iris comparisons has a higher mean

Table 1. False Match probability as a function of decision criterion.

HD criterion	Odds of false match
0.26	1 in $10^{13}$
0.27	1 in $10^{12}$
0.28	1 in $10^{11}$
0.29	1 in 13 billion
0.30	1 in 1.5 billion
0.31	1 in 185 million
0.32	1 in 26 million
0.33	1 in 4 million
0.34	1 in 690,000
0.35	1 in 133,000
0.36	1 in 28,000
0.37	1 in 6,750
0.38	1 in 1,780
0.39	1 in 520
0.40	1 in 170

and a higher right tail than it does when image capture conditions are sufficiently ideal that a given iris always produces very similar IrisCodes.

The left-side distribution of Hamming Distances seen in Fig. 9 is not stable, since it depends on the imaging conditions; therefore the Failure-to-Match rate is not a fixed property of these algorithms. But the right-side distribution is extremely stable, almost independent of the imager, and is essentially a property of these algorithms. Therefore, only cumulatives of the right-side distribution (from 0) will be computed here. By calculating the area under the curve fitted to the observed distribution of Hamming Distances for different eyes (Fig. 7), we can compute the theoretical False Match rate as a function of the decision criterion employed. This is tabulated in Table 1, for various Hamming Distance acceptance thresholds.

#### 4.3. Identification versus verification

Because the probabilities of False Matches are so low even at rather high Hamming Distances, as shown in Table 1, the standard use of these algorithms is to perform exhaustive searches through very large databases for *identification* of a presenting iris pattern, rather than merely a one-to-one comparison for *verification*. Clearly, exhaustive search identifications are far more demanding than mere verifications, since the probabilities of a False Match in any single comparison are increased monotonically with the size of the exhaustive search database. More precisely, if  $P_1$  is the probability of a False Accept in a single (one-to-one) verification trial with an imposter, then  $P_N$ , the probability of getting any False Matches in identification trials after searching exhaustively through a database of  $N$  different IrisCodes, is:

$$P_N = 1 - (1 - P_1)^N. \quad (4.4)$$

This is an extremely demanding relationship. For example, even if  $P_1$  were 0.001 (which is much better than published test results for face recognition or most other biometrics), then even after searching through a database of merely  $N = 200$  persons, the probability of getting one or more False Matches among those persons would be  $P_N = 0.181$ . But for iris recognition, the cumulatives (see Table 1) under the distribution in Fig. 7 indicates that if we use an acceptance Hamming Distance criterion of 0.27 (i.e. allowing up to 27% of the bits in two IrisCodes to disagree while still accepting them as a match), the False Accept probability in single trials is  $10^{-12}$ . Even after diluting down these odds by performing an exhaustive search over the total number of human irises on the planet, roughly  $10^{10}$ , the chances of any False Match among them all would still be only about 1%. This is an extraordinary statistical situation for a recognition system, and it reveals the power of reducing a pattern recognition problem to a test of statistical independence.

#### 4.4. Countermeasures against subterfuge

There are several ways to confirm that a living iris is being imaged, and not (for example) a photograph, a videotape, or a fake iris printed onto a contact lens, glass eye, or other artifice. One obvious method is to track the ratio of pupil diameter to iris diameter, either when light levels are changing, or even under steady illumination. The pupil can be driven larger or smaller by programmed random changes in light level, with a response time constant of about 250 msec for constriction and about 400 msec for dilation. But even without programmed illumination changes, the disequilibrium between excitatory and inhibitory signals from the brainstem to the enervation of the pupillary sphincter muscle<sup>1</sup> produces a steady-state small oscillation at about 0.5 Hz termed *hippus*. Since the algorithms described here must constantly track both the pupil boundary and the iris boundary anyway,<sup>5</sup> it is routine to monitor the amount of hippus. Its coefficient of variation is normally at least 3%.

Other tests to exclude a photograph of somebody else's iris involve tracking eyelid movements, or examining corneal reflections of infrared LEDs illuminated in random sequences. Still further measures could test for the characteristic spectral signature of living tissue in infrared illumination. Hemoglobin in oxygenated blood has an absorption band in near infrared wavelengths, whereas printers' dyes and emulsions and the reflectance properties of photographic papers are often completely ineffective for infrared light.

Finally, certain vanity contact lenses are available in the USA with fake iris patterns printed onto them (for the purpose of changing one's apparent eye color). The fact that such a fake "iris" is floating on the spherical, external surface of the cornea, rather than lying in an internal plane within the eye, lends itself to optical detection; likewise the fact that the printed iris pattern does not undergo any distortions when the pupil changes in size, as does a living iris pattern. Moreover, the printing process itself creates a characteristic signature that can be detected,




Fig. 10. Illustration of one countermeasure against subterfuge: detecting a printed iris pattern on a contact lens by the 2D Fourier domain artifacts of printing.

as illustrated in Fig. 10. The panels show a natural iris, and a fake one printed onto a contact lens, together with their 2D Fourier power spectra. (The central square of each Fourier spectrum has been blanked out to prevent its domination.) The dot matrix printing process generates four points of spurious energy in the Fourier plane, corresponding to the directions and periodicities of coherence in the printing dot matrix, whereas a natural iris does not have these spurious coherences.

## 5. Execution Speeds

On a 300 MHz RISC processor, the execution times for the critical steps in iris recognition are as shown in Table 2, with optimized integer code.

Once an IrisCode has been computed, it is compared exhaustively against all enrolled IrisCodes in the database, in search of a match. The search process is facilitated and accelerated by vectorizing the Exclusive-OR comparisons to the word-length of the machine, since two integers of such length (say 32 bits) can have all of their bits XOR'd at once in a single machine instruction. Thus the elementary integer XOR instruction is an extremely efficient way to detect and tally up the total number of bits that disagree (i.e. the Hamming Distance) between

Table 2. Execution speeds for the critical steps in iris recognition.

Operation	Execution Time
Assessing image focus	15 msec
Scrubbing specular reflections	56 msec
Localizing the eye and iris	90 msec
Fitting the pupillary boundary	12 msec
Detecting and fitting both the eyelids	93 msec
Removing eyelashes and contact lens artifacts	78 msec
Demodulation and IrisCode creation	102 msec
XOR comparison of any two IrisCodes	10 $\mu$ s

two IrisCodes. Ergodicity (representativeness of subsamples) and commensurability (universal format of IrisCodes) facilitate extremely rapid comparisons in searches through large databases. On a 300 MHz processor the rate of raw comparisons is about 100,000 IrisCodes per second, and this rate could be further accelerated using dedicated PLA hardware to many millions of persons per second if such large databases of IrisCodes were ever enrolled.

## 6. Current Usage of this Technique

All current publicly deployed systems for iris recognition use the algorithms described here, by software license of the executable binary code. These include products or test systems deployed by: Panasonic, LG, Oki, Iridian, IriScan, Sensor-Sarnoff, EyeTicket, British Telecom, US Sandia Labs, UK National Physical Laboratory, Unisys, Siemens, IBM, and Schiphol-Group. Current major applications are airport security (including installations at London-Heathrow, Amsterdam-Schiphol, New York JFK, Charlotte, and Toronto); automatic passenger processing and immigration control in lieu of passports; bank cash machine access; database and network access; building entry control; nuclear power station security; computer login validation; prison controls; electronic commerce; and various government applications.

## References

1. F. H. Adler, *Physiology of the Eye: Clinical Application* (The C.V. Mosby Company, 1965), 4th edition.
2. J. G. Daugman, *Two-dimensional spectral analysis of cortical receptive field profiles*, *Vision Res.* **20** (1980) 847–856.
3. J. G. Daugman, *Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters*, *J. Opt. Soc. Am.* **A2** (1985) 1160–1169.
4. J. G. Daugman, *Complete discrete 2D Gabor transforms by neural networks for image analysis and compression*, *IEEE Trans. Acoustics, Speech Signal Processing* **36** (1988) 1169–1179.
5. J. G. Daugman, *High confidence visual recognition of persons by a test of statistical independence*, *IEEE Trans. Pattern Anal. Machine Intel.* **15** (1993) 1148–1161.



6. J. G. Daugman, United States Patent No. 5, 291, 560 (issued March 1994); *Biometric Personal Identification System Based on Iris Analysis*, U.S. Government Printing Office, Washington DC, 1994.
7. J. G. Daugman and C. J. Downing, *Demodulation, predictive coding, and spatial vision*, *J. Opt. Soc. Am.* **A12** (1995) 641–660.
8. D. Gabor, *Theory of communication*, *J. Inst. Electr. Eng.* **93** (1946) 429–457.
9. R. Viveros, K. Balasubramanian and N. Balakrishnan, *Binomial and negative binomial analogues under correlated Bernoulli trials*, *The Am. Statist.* **48** (1984) 243–247.