HOW IRIS RECOGNITION WORKS

John Daugman, OBE

University of Cambridge, The Computer Laboratory, Cambridge CB3 0FD, U.K.

ABSTRACT

The principle that underlies the recognition of persons by their iris patterns is the failure of a test of statistical independence on texture phase structure as encoded by multi-scale quadrature wavelets. The combinatorial complexity of this phase information across different persons spans about 249 degrees of freedom and generates a discrimination entropy of about 3.2 bits/mm^2 over the iris, enabling real-time decisions about personal identity with extremely high confidence. Algorithms first described by the author in 1993 have now been tested in several independent field trials and are becoming widely licensed. This presentation will review how the algorithms work and will present the results of 9.1 million comparisons among different eye images acquired in trials in Britain, the USA, Korea, and Japan.

1. INTRODUCTION

Biometric identification systems all rely upon forms of random variation among persons. The more complex the randomness the better, because more dimensions of independent variation produce signatures having greater uniqueness. But while desiring maximal between-person variability, biometric templates also need minimal within-person variability across time and conditions. In the case of face recognition for example, difficulties arise from the fact that faces are changeable social organs displaying a variety of expressions, as well as being active 3D objects whose projected images vary with pose and viewing angle, illumination, accoutrements, and age. Against this large within-person variability, between-person face variability is barely sufficient: different frontal faces possess the same canonical set of features, always in basically the same canonical geometry.

When imaging can be done at distances of about a meter or less, iris patterns become interesting as an alternative approach to reliable visual recognition of persons, especially when there is a need to search very large databases without making false matches despite so many opportunities. Although small (11 mm) and sometimes problematic to image, the iris has the great mathematical advantage that its pattern variability among different persons is vast. Methods for encoding and recognizing iris patterns were first described in Daugman (1993). These algorithms, released as executables, have been the basis of all iris recognition systems so far deployed for public use or in documented trials, including those by Panasonic, Oki, IBM, British Telecom, LG, Unisys, US Sandia Labs, UK National Physical Laboratory, the NBTC of SJSU, EyeTicket, Schiphol Group, NCR, Siemens, IriScan, Iridian, and Sensar. Airport deployments have recently begun at Heathrow, Schiphol, JFK, Dulles, Charlotte, and are mandated for all 11 Canadian international airports. This paper focuses on the encoding and matching process, and presents exhaustive cross-comparisons across pooled databases. Other aspects, such as real-time eye localization, the achievement of object invariances, and iris morphogenesis, are described in other papers by the author.

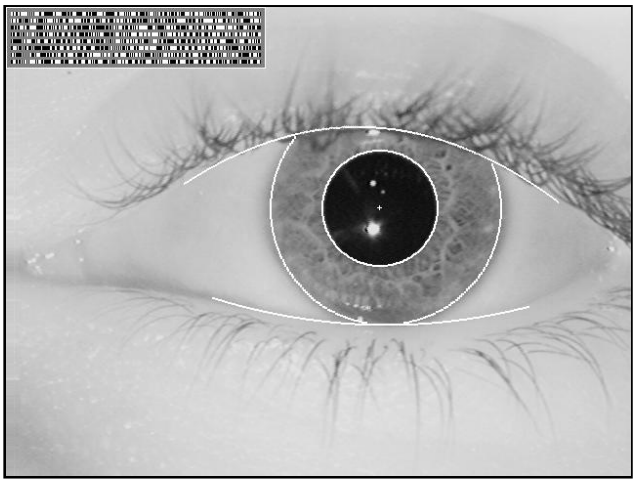

Fig. 1. Example of an iris pattern, imaged in near infrared at a distance of about 35 cm. The outline overlay shows the results of iris, pupil, and eyelid localization steps, and the bit stream in the corner depicts the computed IrisCode.

2. IRIS PHASE CODING BY 2D WAVELET DEMODULATION

After the pupillary and iris/sclera boundaries have been located, any occluding eyelids detected, and reflections or eyelashes excluded, the isolated iris is mapped to size-invariant coordinates and demodulated to extract its phase information using quadrature 2D Gabor wavelets (Daugman 1985, 1988, 1993). The demodulation process is illustrated in Fig

2. It amounts to a patch-wise phase quantization of the iris texture, by identifying in which quadrant of the complex plane each resultant phasor lies when a given area of the iris is projected onto complex-valued 2D Gabor wavelets:

$$h_{\{Re,Im\}} = \mathrm{sgn}_{\{Re,Im\}} \int_{\rho}\int_{\phi} I(\rho,\phi) e^{-i\omega(\theta_0-\phi)} \cdot e^{-(r_0-\rho)^2/\alpha^2} e^{-(\theta_0-\phi)^2/\beta^2} \, \rho \, d\rho \, d\phi \quad (1)$$

where $h_{\{Re,Im\}}$ can be regarded as a complex-valued bit whose real and imaginary parts are either 1 or 0 (sgn) depending on the sign of the 2D integral; $I(\rho,\phi)$ is the raw iris image in a dimensionless polar coordinate system that is size- and translation-invariant, and which also corrects for pupil dilation; α and β are the multi-scale 2D wavelet size parameters, spanning an 8-fold range from 0.15mm to 1.2mm on the iris; ω is wavelet frequency, spanning 3 octaves in inverse proportion to β; and (r_0, θ_0) represent the polar coordinates of each region of iris for which the phasor coordinates $h_{\{Re,Im\}}$ are computed. Such a phase quadrant coding sequence is illustrated for one iris by the bit stream depicted in Fig 1. A desirable feature of the phase code, explained in Fig 2, is that it is a cyclic or Gray code: in rotating between any adjacent phase quadrants, only a single bit changes, unlike a binary code in which two bits may change, making some errors arbitrarily more costly than others. Altogether 2,048 such phase bits (256 bytes) are computed for each iris, and an equal number of masking bits are also computed to signify whether any iris region is obscured by eyelids, contains any eyelash occlusions, specular reflections, boundary artifacts of hard contact lenses, or poor signal-to-noise ratio and thus should be ignored in the demodulation code as artifact.

Only phase information is used in the IrisCode because amplitude information is not very discriminating, and it depends upon extraneous factors such as imaging contrast, illumination, and camera gain. The phase bit settings which encode the sequence of projection quadrants as shown in Fig 2 capture the information of wavelet zero-crossings, as is clear from the sign operator in (1). The extraction of phase has the further advantage that phase angles are assigned regardless of how poor the image contrast is, as with very poorly focused images. In such cases the phase bits are set largely on the basis of random CCD noise and have binomial run-length statistics. This advantage of phase encoding prevents different poorly focused irises from being confused with each other, as happens, for example, with poorly focused face images.

3. THE TEST OF STATISTICAL INDEPENDENCE: COMBINATORICS OF PHASE SEQUENCES

The key to iris recognition is the failure of a test of statistical independence which involves so many degrees-of-freedom

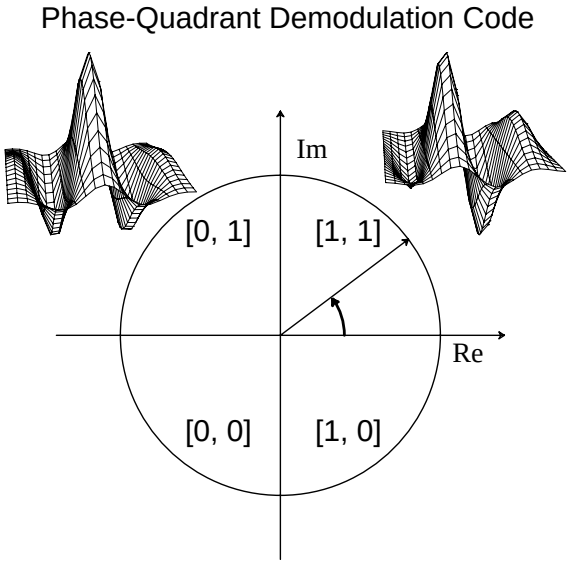

Fig. 2. The phase demodulation process used to encode iris patterns. The angle of each projection phasor is quantized to its quadrant, setting two bits of phase information. This process is repeated all across the iris with many wavelet sizes, frequencies, and orientations, to extract 2,048 bits.

that the test is virtually guaranteed to be passed whenever the IrisCodes for two different eyes are compared, but to be uniquely failed when any eye's IrisCode is compared with another version of itself. The test of statistical independence is implemented by the simple Boolean Exclusive-OR operator (XOR) applied to the 2,048 bit phase vectors that encode any two iris patterns, masked (AND'ed) by both of their corresponding mask bit vectors to prevent non-iris artifacts from influencing iris comparisons. The XOR operator $\otimes$ detects disagreement between any corresponding pair of bits, while the AND operator $\cap$ ensures that the compared bits are both deemed to have been uncorrupted by eyelashes, eyelids, specular reflections, or other noise. The norms ($||\ ||$) of the resultant bit vector and of the AND'ed mask vectors are then measured in order to compute a fractional Hamming Distance (HD) as the measure of the dissimilarity between any two irises, whose two phase code bit vectors are denoted {*codeA, codeB*} and whose mask bit vectors are denoted {*maskA, maskB*}:

$$\mathrm{HD} = \frac{||(codeA \otimes codeB) \cap maskA \cap maskB||}{||maskA \cap maskB||} \quad (2)$$

The denominator tallies the total number of phase bits that mattered in iris comparisons after artifacts such as eyelashes and specular reflections were discounted, so the resulting HD is a fractional measure of dissimilarity; 0 would represent a perfect match. The Boolean operators $\otimes$ and $\cap$ are applied in vector form to binary strings of length up to the word length of the CPU, as a single machine instruc-

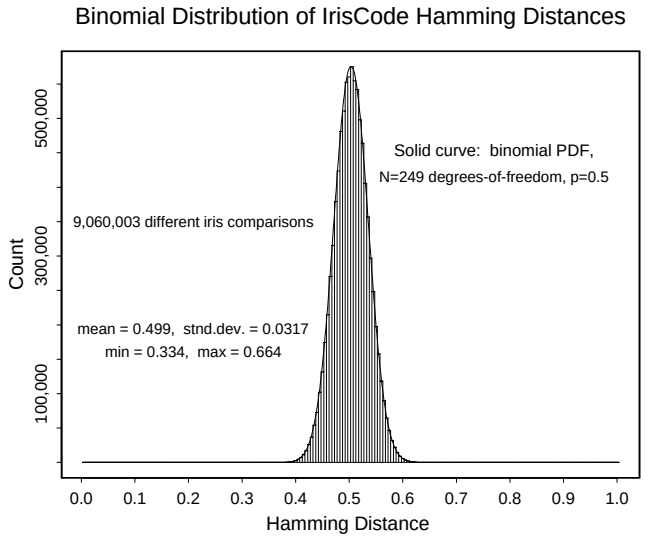

Fig. 3. Distribution of Hamming Distances obtained from all 9.1 million possible comparisons between different pairs of irises in the database. The histogram forms a perfect binomial distribution with $p = 0.5$ and $N = 249$ degrees-of-freedom, as shown by the solid curve (3). The data implies that it is extremely improbable for two different irises to disagree in less than about a third of their phase information.

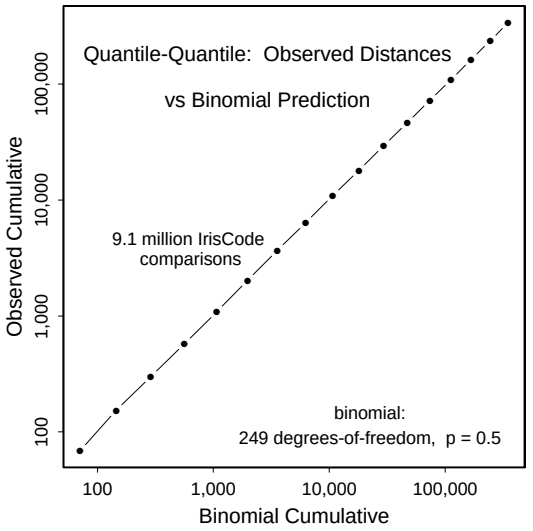

Fig. 4. Quantile-quantile plot of the observed cumulatives under the left tail of the histogram in Fig 3, versus predicted binomial cumulatives. Their close agreement over several orders of magnitude strongly confirms the binomial model for phase bit comparisons between different irises.

tion. Thus for example on an ordinary 32-bit machine, any two integers between 0 and 4 billion can be XOR'ed in a single machine instruction to generate a third such integer, each of whose bits in a binary expansion is the XOR of the corresponding pair of bits of the original two integers. This implementation of (2) in parallel 32-bit chunks enables extremely rapid comparisons of IrisCodes when searching through a large database to find a match. On a 300 MHz CPU, such exhaustive searches are performed at a rate of about 100,000 irises per second.

Because any given bit in the phase code for an iris is equally likely to be 1 or 0, and because different irises are uncorrelated, the expected proportion of agreeing bits between the codes for two different irises is HD = 0.500. The histogram in Fig 3 shows the distribution of HDs obtained from 9.1 million comparisons between different pairings of iris images acquired by licensees of these algorithms in the UK, the USA, Japan, and Korea. Their observed mean HD was $p = 0.499$ with standard deviation $\sigma = 0.0317$; their full distribution in Fig 3 corresponds to a fractional binomial having $N = p(1-p)/\sigma^2 = 249$ degrees-of-freedom, as shown by the solid curve. The extremely close fit of the theoretical fractional binomial to the observed distribution is a consequence of the fact that each comparison between two phase code bits from two different irises is essentially a Bernoulli trial, albeit with correlations among the "coin tosses."

In any given IrisCode, only small subsets of phase bits are mutually independent, due to the internal correlations, especially radial, within an iris. (If all $N = 2,048$ phase bits were independent, then the distribution in Fig 3 would be very much sharper, with an expected standard deviation of only $\sqrt{p(1-p)/N} = 0.011$ and so the HD interval between 0.49 and 0.51 would contain most of the distribution.) Bernoulli trials that are correlated (see Viveros et al., 1984) remain binomially distributed but with a reduction in N, the effective number of tosses, and hence an increase in the σ of the normalized HD distribution. The form and width of the HD distribution in Fig 3 tell us that the amount of difference between the phase codes for different irises is distributed equivalently to runs of 249 tosses of a fair coin (Bernoulli trials with $p = 0.5, N = 249$). Expressing this variation as a discrimination entropy (see Cover and Thomas 1991) and using typical iris and pupil diameters of 11mm and 5mm respectively, the observed amount of statistical variability among different iris patterns corresponds to an information density of about 3.2 bits/mm^2 on the iris.

The theoretical binomial distribution plotted as the solid curve in Fig 3 has the fractional functional form

$$f(x) = \frac{N!}{m!(N-m)!}\, p^m (1-p)^{(N-m)} \qquad (3)$$

where $N = 249$, $p = 0.5$, and $x = m/N$ is the outcome fraction of N Bernoulli trials (e.g. coin tosses that are "heads" in each run). In our case, x is the HD, the fraction of phase bits that happen to disagree when two different irises are compared. To validate such a statistical model we must also study the behaviour of the tails, by examining quantile-quantile plots of the observed cumulatives versus the theoretically predicted cumulatives from 0 up to sequential points in the tail. Such a "Q-Q" plot is given in Fig 4. The perfect linear relationship indicates precise agreement between model and data, over a range of almost four orders of magnitude. It is clear from both Figs 3 and 4 that the

extremely rapid attenuation of the tails of binomial distributions when N is so large makes it extremely improbable that two different irises might disagree by chance in fewer than at least a third of their bits. Computing the cumulative of $f(x)$ from 0 to 0.333 indicates that the probability of such an event is about 1 in 16 million. (Of the 9.1 million iris comparisons plotted in the histogram of Fig 3, the smallest Hamming Distance observed was 0.334.) The binomial cumulative from 0 to just 0.300 is 1 in 10 billion, roughly the number of human eyes on the planet. Thus, even the observation of a relatively poor degree of match between the phase codes for two different iris images (say, 70% agreement or HD = 0.300) would still provide extraordinarily compelling evidence of identity, because this test of statistical independence is still failed so convincingly.

4. UNIQUENESS OF FAILING THE TEST OF STATISTICAL INDEPENDENCE

The statistical data and theory presented above show that we can perform iris recognition successfully just by a test of statistical independence. Any two different irises are statistically "guaranteed" to pass this test of independence, and any two images that fail this test (i.e. produce HD $\leq$ 0.32) must be images of the same iris. This means that we need only demand a very forgiving degree of match – allowing up to about 32% of the bits to disagree – yet can still make identifications with extremely high confidence.

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

Table 1: Cumulatives giving false match probabilities.

It is informative to calculate the significance of any observed HD matching score, in terms of the likelihood that it could have arisen by chance from two different irises. The confidence level or error probability corresponds to the cumulative under the left tail of the histogram in Fig 3 up to the criterion employed, also factoring in comparisons in many tilt orientations. Table 1 enumerates the cumulatives (false match probabilities) as a function of HD decision criterion, in the range between 0.26 and 0.40. Calculation of the large factorial terms in (3) was done with Stirling's approximation which errs by less than 1% for $n \geq 9$:

$$n! \approx \exp(n\ln(n) - n + \frac{1}{2}\ln(2\pi n)) \qquad (4)$$

The practical importance of the astronomic odds against a false match when the match quality is better than about HD $\leq$ 0.32 is that very large databases can be searched exhaustively without succumbing to any of the many opportunities for suffering a false match. The requirements of operating in one-to-many "identification" mode are vastly more demanding than operating merely in one-to-one "verification" mode (in which an identity must first be explicitly asserted, which is then verified in a yes/no decision by comparison against just the single nominated template). This identification power, which means users need not even bother to assert an identity, is one of the main advantages of iris recognition as a biometric.

5. REFERENCES

Cover, T., and Thomas, J. (1991) *Elements of Information Theory.* New York: Wiley.

Daugman, J. (1985) Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Journal of the Optical Society of America A* 2(7): 1160-1169.

Daugman, J. (1988) Complete discrete 2D Gabor transforms by neural networks for image analysis and compression. *IEEE Trans. Acoustics, Speech, and Sig. Proc.* 36(7): 1169-1179.

Daugman, J. (1993) High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. Pattern Analysis and Machine Intell.* 15(11): 1148-1161.

Daugman, J. (1994) U.S. Patent No. 5,291,560: *Biometric Personal Identification System Based on Iris Analysis.*

Daugman, J. and Downing, C. (1995) Demodulation, predictive coding, and spatial vision. *Journal of the Optical Society of America A* 12(4): 641-660.

Viveros, R., Balasubramanian, K., and Balakrishnan, N. (1984) Binomial and negative binomial analogues under correlated Bernoulli trials. *The American Statistician* 48(3): 243-247.

website: http://www.CL.cam.ac.uk/users/jgd1000/