Probing the Uniqueness and Randomness of IrisCodes: Results From 200 Billion Iris Pair Comparisons

Chances are assessed of making false matches using iris recognition when huge numbers of individuals are enrolled and massive database searches are performed.

By John Daugman

ABSTRACT | Recent large-scale deployments of iris recognition for border-crossing controls enable critical assessment of the robustness of this technology against making false matches, since vast numbers of cross comparisons become possible within large databases. This paper presents results from the 200 billion iris cross comparisons that could be performed within a database of 632 500 different iris images, spanning 152 nationalities. Each iris pattern was encoded into a phase sequence of 2048 bits using the Daugman algorithms. Empirically analyzing the tail of the resulting distribution of similarity scores enables specification of decision thresholds, and prediction of performance, of the iris recognition algorithms if deployed in identification mode on national scales.

KEYWORDS | Biometrics; demodulation; face recognition; Gabor wavelets; iris recognition; large-scale search; personal identification; texture analysis

I. INTRODUCTION

Governments in a number of countries are contemplating schemes for biometrically-enabled national identity cards. One purpose would be to detect fraudulent multiple identities. Other purposes include expedited immigration controls using biometric passports, allowing automated border crossing; and security screening, searching against watch-lists at ports of entry. Some of these applications require "all-against-all" exhaustive comparisons, such as the search for multiple identities across a national population, which entails a total number of cross comparisons that scales as the square of the national population. In a country like the U.K., with about 60 million persons, such a search for multiple enrolled identities would entail 1800 trillion (1.8×10^{15}) different pairings, or almost 2 peta-comparisons. Can any biometric identification system possibly survive so many comparisons between different pairings of persons, without making false matches?

Other biometric deployment applications are somewhat less ambitious numerically, yet still daunting. The U.K. has launched Project IRIS (Iris Recognition Immigration System), in which at least 1 million frequent travelers to the U.K. from abroad will be routinely able to enter the country without presenting a passport or explicitly asserting their identity. Instead, their iris patterns are captured by video cameras for comparison exhaustively against the enrolled database of authorized persons. That application will thus require a trillion iris comparisons by the time each enrolled traveler has used the system just once. Another application that has been introduced in the Middle East involves "watch-list screening," in which foreigners entering a country through all ports of entry are compared exhaustively against all persons registered in a watch-list. The daily volume of such cross comparisons during periods of high travel

approaches 10 billion per day. How can any biometric identification system survive such large numbers of opportunities for making erroneous matches?

Clearly this question hinges on the precise nature of the tail of the distribution of similarity scores obtained when different persons are compared biometrically. In effect one is drawing samples from that distribution billions of times per day, in the application last mentioned, and the key issue is the likelihood of drawing samples sufficiently far out along the tail to constitute false matches due to accidental similarity between different biometric templates. Extreme value theory (and common statistical intuition) teaches that the likelihood of drawing extreme samples from the tail of a probability distribution accumulates with the number of samples drawn. The cumulatives under the tail of that distribution, up to the imposed similarity threshold for match decisions, constitute the probability of making false matches. The purpose of this paper is to probe those questions in the case of a particular biometric method that was developed by the author, iris recognition, using a large database from the Middle East which allowed 200 billion different pair comparisons to be made.

II. THE DATABASE

In 2001 the United Arab Emirates (UAE) launched a national border-crossing security system that is today deployed at all 27 of the UAE's air, land, and sea ports. All foreign nationals who possess a visa to enter the UAE must look at an iris camera, installed at immigration desks. Algorithms developed by the author locate the eyes and compute IrisCodes from the random texture that is visible in the iris, when illuminated by infrared light in the near-IR band (700–900-nm). The database against which the visitors are checked is a "negative watch-list" of persons deemed untrustworthy or who have been denied entry for a variety of reasons, including security concerns, past violations, previous imprisonment, traveling under false documents, or work permit violations. Most persons who reside and work in the UAE are not UAE nationals but foreign nationals. Many who had overstayed their work permits or committed other violations were expelled, under an amnesty program in lieu of other sanctions. The total number of persons in the watch-list as of June 2005 was about 316 250, spanning 152 nationalities. Both irises of each person were enrolled in the database, which thus consisted of about 632 500 different iris patterns.

On a typical day some 6000 persons enter the UAE who must submit to the iris recognition check with both eyes. This generates about 7.2 billion iris comparisons daily. Since the start of the deployment, nearly 5 trillion (5 million million) iris comparisons have been performed, over a networked architecture ("IrisFarm," developed by IrisGuard UK) that computes IrisCodes at the local ports

of entry and sends them over a variety of communication channels to the central cold storage database at the Abu Dhabi Police General Directorate for exhaustive comparison. All iris images enrolled in this database were captured using LG-2200 and LG-3000 iris cameras from LG (Korea). In live operation, each presenting iris is compared exhaustively against all in the watch-list database in less than 2 s. To date, some 47 000 persons have been caught trying to enter the UAE under false travel documents, by this iris recognition system. The Abu Dhabi Directorate of Police report that so far there have been no matches made that were not eventually confirmed by other data.

It was desired to exploit this large enrollment database of IrisCodes in order to understand better the statistical powers of iris recognition. For example, by computing the similarities between all possible pairings of different irises in the database, much could be learned about the robustness of these algorithms against making any false matches, when there are such vast numbers of opportunities for error. The total number of different pairings of irises that can be made in a database of 632 500 is more than 200 billion: 200 027 808 750. The UAE Minister of Interior, H.R.H. Sheikh Saif Bin Zayyed, therefore made the enrollment database of IrisCodes available to the University of Cambridge for detailed analysis and dissemination of results.

III. THE ALGORITHMS

The iris recognition algorithms that are used in all public deployments of this technology to date, such as the UAE deployment, have been described previously by Daugman [2], [3], [5], [6] and they will be only briefly summarized here. The basic principle behind these algorithms is the failure of a test of statistical independence. Persons are recognized by their iris patterns because they uniquely fail a test of independence against earlier enrolled descriptions ("IrisCodes") derived from themselves, but they are statistically guaranteed to pass that same test of independence against everybody else.

A. Iris Texture

The key to biometric identification is random variation among different persons, since this is the origin of uniqueness, and the basis for discrimination. Facial appearance does vary among people, but the basic geometry of facial features is always the same; the residual distinguishing dimensions are secondary modulations that are not significantly larger than the within-person variations due to illumination geometry, emotional expression, pose angle, or aging. Fingerprints show more between-person variation; but rather than having a multiscale texture like iris patterns, their scale of structuring is fixed by the (roughly 0.5 mm) spacing of the ridge flow.
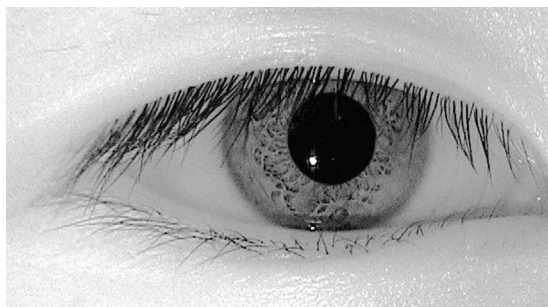
Fig. 1. *Even dark brown eyes reveal rich iris texture when illuminated in near-infrared light (700–900-nm band). The randomness of this texture and its complexity, spanning at least 3 octaves in usable scales of analysis, enables the discriminating power of the IrisCode.*

In contrast, the iris displays a rich two-dimensional random structure [4] spanning several octaves in scales of analysis. If imaged with visible wavelengths of light, the most striking and rich patterns are seen in persons (mainly of European origin) whose apparent eye color is blue, gray, or green, while little structure seems evident in the 90% or more of the world's peoples who have dark brown eyes. But when illuminated with near-infrared light in the 700–900-nm band (used in all current iris recognition cameras), even dark brown irises reveal rich texture. Fig. 1 illustrates the almost lunar surface appearance of such an iris, whose anterior layer is richly cratered with relief features as well as the arching ligaments of connective tissue more typically visible in light-colored eyes. The eyelashes which occlude parts of the iris are readily detected by statistical inference of bimodality in the iris pixel histogram, allowing those pixels that are deemed to form a separate darker population of pixels to be excluded from influencing the computed code for iris texture. The resolution of current iris cameras is typically in the range of 5 lp/mm, with the imaged iris diameter usually in the range of 150–220 pixels. Acquisition distances range from several centimeters to, recently, 3 m or more (see accompanying paper: "Iris on the Move").

B. Segmentation and Normalization

Critical steps in iris recognition include segmentation or extraction of the visible iris portion of the image, with exclusion of any obscuring elements such as eyelids, eyelashes (notable in Fig. 1 above), and reflections from the cornea or possibly from eyeglasses. These steps are performed using a variety of boundary and region detection and active contour techniques. The eyelid boundaries may be described as quadratic or cubic splines, whose parameters can be estimated by statistical model-fitting techniques. Fig. 2 illustrates detection and demarcation of the four boundaries of the iris (pupil, limbus, upper and lower eyelids).

Both the inner and outer boundaries of the iris are often significantly nonround. Both boundaries may be oval or pear-shaped, and pupils can have particularly irregular shapes. It is therefore important not to force them to be described as circles, despite the simplifications which that would allow in the coordinate system. Fortunately, in infrared light, the pupillary boundary is always a very strong, high-contrast signal (unlike the case for dark-eyed persons in visible light). On the other hand, the outer boundary of the iris is generally a very weak, low-contrast signal in infrared light (again, unlike the case for visible wavelengths), because the sclera contains much blood, and hemoglobin in blood absorbs strongly in the near infrared spectrum, making the sclera often as dark as the iris.

A fundamental principle of shape description in computer vision teaches that one should use weak constraints when the data is strong, and strong constraints when the data is weak. Therefore, the pupillary boundary is fitted with an active contour or "snake" [1], [9] having many degrees of freedom, whereas the iris outer boundary is fitted by a snake having comparatively few degrees of
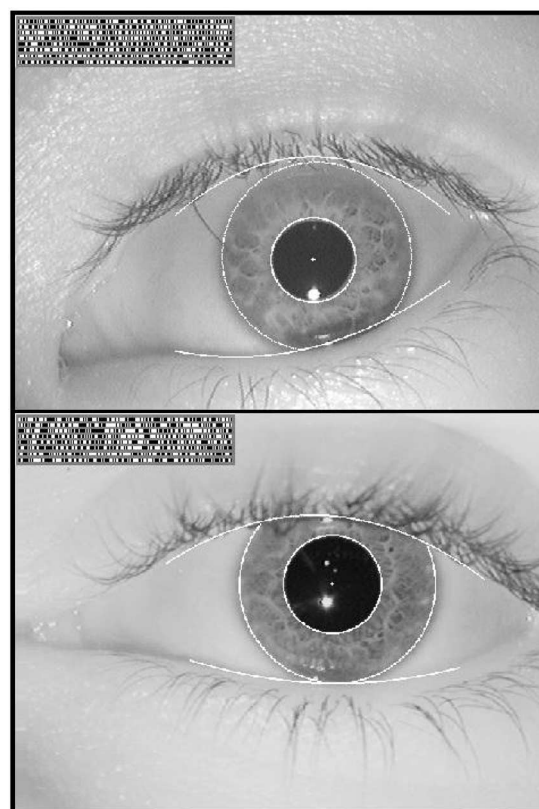

Fig. 2. *Isolation of the iris from the rest of the image. The white graphical overlays signify detected iris boundaries resulting from the segmentation process. For these images, circular inner and outer iris boundaries were assumed, but this is an overly restrictive assumption.*

freedom. The points between these inner and outer boundary contours are interpolated linearly by a homogeneous rubber sheet model, which automatically reverses the iris pattern deformations caused by pupillary dilation or constriction. Under assumptions of uniform iris elasticity (which may be questionable), this normalization maps the iris tissue into a doubly-dimensionless coordinate system.

The homogeneous rubber sheet model assigns to each point in the iris, regardless of iris size in the image and of pupillary dilation, a pair of dimensionless real coordinates (r, θ) where r lies in the unit interval $[0,1]$ and θ is the angular variable, cyclic over $[0, 2\pi]$. The remapping or normalization of the iris image $I(x, y)$ from raw coordinates (x, y) to a doubly dimensionless and nonconcentric coordinate system (r, θ) can be represented as

$$I(x(r, \theta), y(r, \theta)) \rightarrow I(r, \theta) \qquad (1)$$

where $x(r, \theta)$ and $y(r, \theta)$ are defined as linear combinations between the set of pupillary boundary points $(x_p(\theta), y_p(\theta))$ determined by the internal active contour, and the set of outer boundary points along the limbus $(x_s(\theta), y_s(\theta))$ determined by the external active contour describing the iris/sclera boundary

$$x(r, \theta) = (1 - r)x_p(\theta) + rx_s(\theta) \qquad (2)$$
$$y(r, \theta) = (1 - r)y_p(\theta) + ry_s(\theta). \qquad (3)$$

This homogeneous rubber sheet model maps the iris into a dimensionless, normalized coordinate system that is size-invariant, and therefore invariant to changes in the target distance and the optical magnification factor, as well as invariant to the position of the eye in the image, and invariant to the pupil dilation (assuming uniform iris elasticity). It is not strictly polar, because it makes no assumption that the pupil and iris are concentric (indeed the pupil's actual center is usually nasal, and inferior, to the center of the iris), nor even that their boundaries are circular. These flexible contours are illustrated in Fig. 3. An important aspect of this method of segmentation and normalization is that it does not introduce unnecessary cuts in the intrinsically continuous and cyclic angular variable, which would interrupt subsequent convolution, as occurs in other methods that explicitly unwrap the iris into a rectangular domain.

C. The IrisCode

The normalized and dimensionless iris mapping is encoded into an IrisCode through a process of demodulation that extracts phase sequences. This encoding
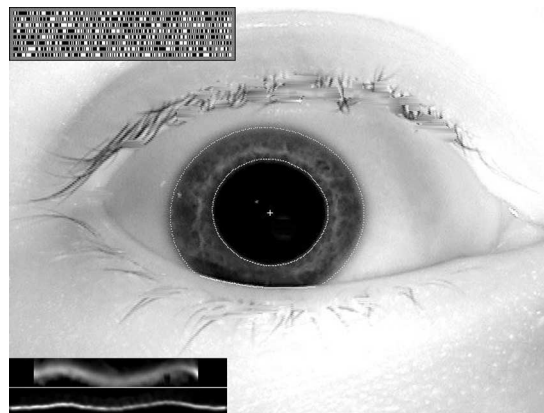

Fig. 3. *Illustration of "snakes," or active contours, to fit the inner and outer boundaries of the iris. These are generally noncircular, as illustrated by the curvature maps in the lower left; those would be flat and straight if the iris inner and outer boundaries were circles.*

process is illustrated in Fig. 4. It amounts to a patchwise phase quantization of the iris pattern, 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 two-dimensional (2-D) Gabor [7] wavelets

$$h_{\{Re, Im\}} = \text{sgn}_{\{Re, Im\}} \int_\rho \int_\phi I(\rho, \phi) e^{-i\omega(\theta_0 - \phi)}$$
$$\times \, e^{-(r_0 - \rho)^2 / \alpha^2} e^{-(\theta_0 - \phi)^2 / \beta^2} \rho d\rho d\phi. \qquad (4)$$
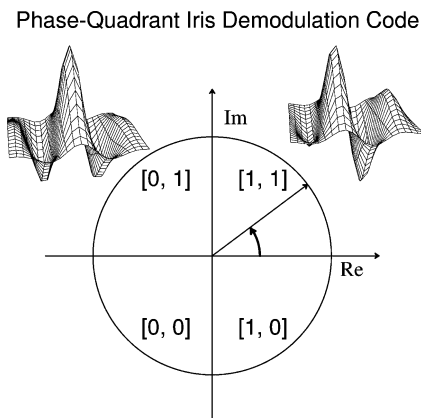

Fig. 4. *Extracting phase sequences to encode iris patterns. Local regions of an iris are projected (4) onto quadrature 2-D Gabor wavelets, generating complex-valued coefficients whose real and imaginary parts specify the coordinates of a phasor in the complex plane.*

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 2-D integral; $I(\rho, \phi)$ is the raw iris image in the dimensionless coordinate system; α and β are the multiscale 2-D wavelet size parameters, spanning an eightfold range from 0.15 to 1.2 mm on the iris; ω is wavelet frequency, spanning three octaves in inverse proportion to β; and (r_0, θ_0) represent the dimensionless coordinates of each region of iris for which the phasor bits $h_{\{Re,Im\}}$ are computed. Such phase quadrant coding sequences are portrayed by the barcode-like bit streams in Figs. 2 and 3. A desirable feature of the encoding as indicated in Fig. 4 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 2048 such phase bits (256 bytes) are computed for each iris; but in addition 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 as artifact.

D. Matching Engine

In comparing any two IrisCodes, the phase data bits are Exclusive or'ed ($\otimes$) to detect disagreement and thereby to assess the similarity between the two iris patterns. But the masking bits are also and'ed ($\cap$) with this xor combination in order to restrict the comparison to those bits that for both irises are deemed not to be corrupted by eyelashes, eyelids, or reflections. The norms ($\| \|$) of the resultant xor'ed data vectors and of the and'ed mask vectors are then computed in order to derive the raw Hamming distance HD_{raw}, as the fraction of bits (deemed significant) that disagree between two irises. For any two different irises, statistical independence creates the expectation $HD_{raw} = 0.5$ for this score. If we denote the two iris phase data vectors as $\{codeA, codeB\}$ and their associated mask vectors as $\{maskA, maskB\}$, then their raw dissimilarity score is

$$HD_{raw} = \frac{\|(codeA \otimes codeB) \cap maskA \cap maskB\|}{\|maskA \cap maskB\|}. \quad (5)$$

But different people expose different amounts of iris between their eyelids, and the amount visible depends also on occluding eyelashes, reflections, and other circumstances. Therefore, the number of bits available for comparison between two different IrisCodes is quite variable. A close match (say a Hamming distance of $HD_{raw} = 0.10$) based on only few compared bits is much less indicative of identity than an apparently poorer match (say $HD_{raw} = 0.20$) based on a large number of compared

bits. This requires a renormalization of any observed raw Hamming distance score HD_{raw} into one HD_{norm} whose deviation from statistical independence ($HD_{raw} = 0.50$) has been rescaled for statistical significance, based on the number of bits n that were actually compared between the two IrisCodes

$$HD_{norm} = 0.5 - (0.5 - HD_{raw})\sqrt{\frac{n}{911}}. \quad (6)$$

The parameters in the above equation influence the standard deviation of the distribution of normalized Hamming distance scores, and they give the distribution a stable form which permits a stable decision rule to be employed.

E. Execution Speed of the Algorithms

Execution of all the image processing steps, from localizing the iris and all of its boundaries (including the eyelids), to eyelash exclusion and computing the IrisCode, occurs faster than the video frame rate. On a 3-GHz PC, slightly more than 30 frames/s are fully processed. The match search speed with a memory-resident database of IrisCodes exceeds 1 million IrisCodes/sec in an unguided exhaustive search.

IV. RESULTS FROM UNROTATED IRIS PAIR COMPARISONS

Fig. 5 shows all cross comparison similarity scores obtained from making all possible pair comparisons among the 632 500 different irises. N different objects can generate a total of $N \cdot (N - 1)/2$ different pairings, which for the UAE database means 200 027 808 750, or about 200 billion different pair comparisons. The vast majority of IrisCodes from different eyes disagreed in roughly 50% of their bits, as expected since the bits are equiprobable and uncorrelated between different eyes. Very few pairings of IrisCodes could disagree in fewer than 35% or more than 65% of their bits, as is evident from the distribution in Fig. 5. The smallest and largest Hamming distances found in this set of 200 billion simple comparisons of different IrisCodes were around 0.26 and 0.75 respectively.

The solid curve that fits the data very closely in Fig. 5 is a binomial probability density function. This theoretical form was chosen because comparisons between bits from different IrisCodes are Bernoulli trials, or conceptually "coin tosses," and Bernoulli trials generate binomial distributions. If one tossed a coin whose probability of "heads" is p in a series of N independent tosses and counted the number m of "heads" outcomes, and if one tallied this fraction $x = m/N$ in a large number of such repeated runs
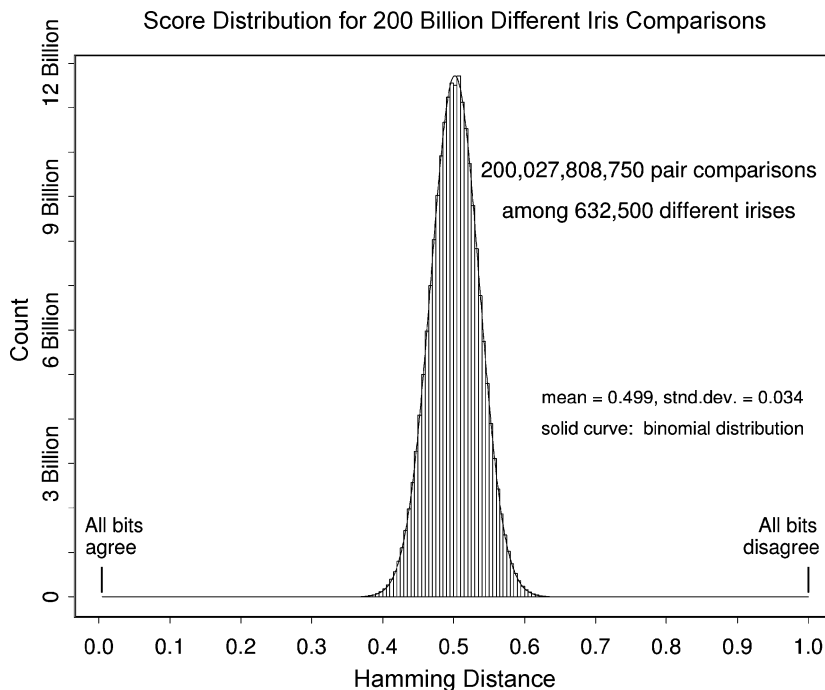
Fig. 5. *Distribution of* HD$_{norm}$ *Hamming distance scores for 200 billion different pairings of iris patterns, without relative rotations. The solid curve fitting the histogram is a binomial distribution (7).*

of N tosses, then the expected distribution of x would be as per the solid curve in Fig. 5

$$f(x) = \frac{N!}{m!(N-m)!} p^m (1-p)^{(N-m)}. \qquad (7)$$

The analogy between tossing coins and comparing bits between different IrisCodes is deep but imperfect, because any given IrisCode has internal correlations arising from iris features, especially in the radial direction. Further correlations are introduced by the 2-D Gabor wavelet filters: their lowpass aspect introduces correlations in amplitude, and their bandpass aspect introduces correlations in phase, both of which linger to an extent that is inversely proportional to the filter bandwidth. The effect of these correlations is to reduce the value of the distribution parameter N to a number significantly smaller than the number of bits that are actually compared between two IrisCodes; N becomes the number of effectively independent bit comparisons. The value of p is very close to 0.5 (empirically 0.499 for the UAE database), because the states of each bit are equiprobable *a priori*, and so any pair of bits from different IrisCodes are equally likely to agree or disagree.

The apparently binomial functional form that describes the distribution of dissimilarity scores [Fig. 5 and (7)] for comparisons between different iris patterns is

key to the robustness of this technology in large-scale search applications. The tails of the binomial attenuate extremely rapidly, because of the dominating central tendency caused by the factorial terms in (7). Rapidly attenuating tails are critical for a biometric to survive the vast numbers of opportunities to make false matches without actually making any, when applied in an "all-against-all" mode of searching for matching or multiple identities, as is contemplated in some national ID card projects in the U.K. and elsewhere in Europe.

V. CLOSEST MATCH SCORES AFTER MULTIPLE ROTATIONS

When IrisCodes are compared in a search for a match, it cannot be known precisely what was the amount of head tilt, camera tilt, or eye rotation when the IrisCodes were obtained. Therefore, it is necessary to make comparisons over a reasonable range of relative tilts (rotations) between every pair of IrisCodes, keeping the best match as their similarity score. This generates an *extreme value distribution* that is skewed towards lower Hamming distances, even between unrelated irises, because of the increased opportunities to get a closer match just by chance.

The new distribution after k rotations of IrisCodes in the search process still has a simple analytic form that can be derived theoretically. Let $f_0(x)$ be the raw density distribution obtained for the HD$_{norm}$ scores between different

irises after comparing them in only a single relative orientation; for example, $f_0(x)$ might be the binomial defined in (7). Then $F_0(x)$, the cumulative of $f_0(x)$ from 0 to x, becomes the probability of getting a false match in such a test when using an HD_{norm} acceptance criterion at x

$$F_0(x) = \int_0^x f_0(x)dx \qquad (8)$$

or equivalently

$$f_0(x) = \frac{d}{dx}F_0(x) \qquad (9)$$

Clearly, then, the probability of *not* making a false match when using decision criterion x is $1 - F_0(x)$ after a single test, and it is $[1 - F_0(x)]^k$ after carrying out k such tests independently at k different relative orientations. It follows that the probability of a false match after a "best of k" test of agreement, when using HD_{norm} criterion x, regardless of the actual form of the raw unrotated distribution $f_0(x)$, is

$$F_k(x) = 1 - [1 - F_0(x)]^k \qquad (10)$$

and the expected density $f_k(x)$ associated with this cumulative is

$$\begin{aligned} f_k(x) &= \frac{d}{dx}F_k(x) \\ &= kf_0(x)[1 - F_0(x)]^{k-1} \end{aligned} \qquad (11)$$

(11) for the extreme value distribution is the solid curve in Fig. 6, fitting the distribution of the set of 200 billion IrisCode comparisons after $k = 7$ relative rotations of each pair.

VI. OBSERVED FALSE MATCH RATES

The cumulative scores under the left tail of the distribution shown in Fig. 6, up to various Hamming distance thresholds, reveal the false match rates among the 200 billion iris comparisons if the identification decision policy used those thresholds. These rates are provided in the following Table 1. No such matches were found with Hamming distances below about 0.26; but the table has been extended down to 0.22 using (7)–(10) for the theoretical cumulative of the extreme value distribution of multiple samples from the binomial (plotted as the solid curve in Fig. 6), in order to extrapolate the theoretically expected false match rates for such decision policies. These false match rates, whether observed or theoretical, also serve as confidence levels that can be associated with a given quality of match.
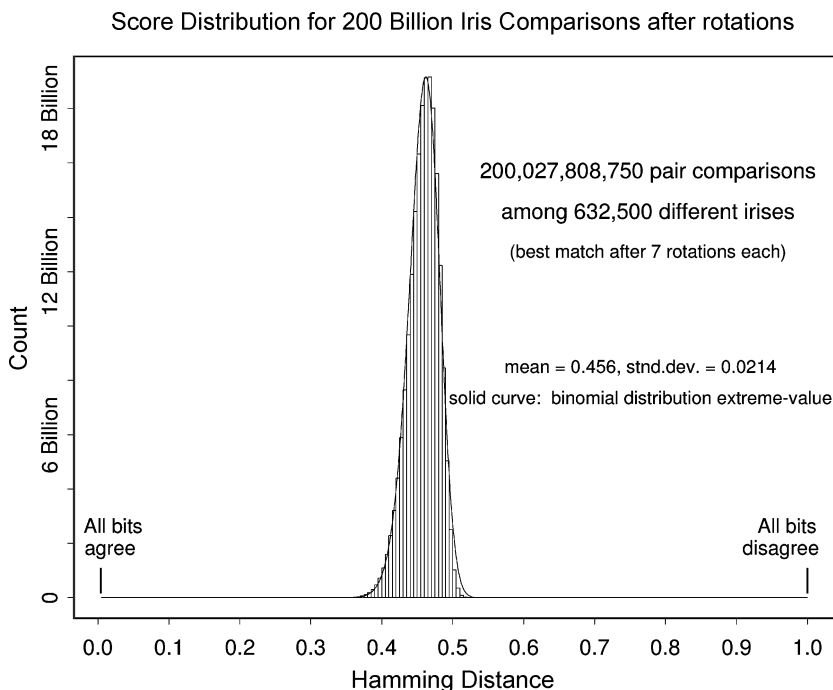

Score Distribution for 200 Billion Iris Comparisons after rotations

200,027,808,750 pair comparisons
among 632,500 different irises
(best match after 7 rotations each)

mean = 0.456, stnd.dev. = 0.0214
solid curve: binomial distribution extreme-value

All bits agree

All bits disagree

Fig. 6. *Distribution of dissimilarity scores between the same 200 billion iris pair comparisons as given in Fig. 5, but showing only the best match after seven relative rotations of each pair because of uncertainty about actual iris orientation. The solid curve is (11).*

Table 1 The False Match Rates, Either Observed in the Distribution of Scores or Predicted Theoretically From Eqns (7)–(10), Are Tabulated as a Function of Possible Decision Policy Match Criteria Imposed on the Normalized Hamming Distance Scores HD_{norm}

HD Criterion	Observed False Match Rate
0.220	0 (theor: 1 in 5×10^{15})
0.225	0 (theor: 1 in 1×10^{15})
0.230	0 (theor: 1 in 3×10^{14})
0.235	0 (theor: 1 in 9×10^{13})
0.240	0 (theor: 1 in 3×10^{13})
0.245	0 (theor: 1 in 8×10^{12})
0.250	0 (theor: 1 in 2×10^{12})
0.255	0 (theor: 1 in 7×10^{11})
0.262	1 in 200 billion
0.267	1 in 50 billion
0.272	1 in 13 billion
0.277	1 in 2.7 billion
0.282	1 in 284 million
0.287	1 in 96 million
0.292	1 in 40 million
0.297	1 in 18 million
0.302	1 in 8 million
0.307	1 in 4 million
0.312	1 in 2 million
0.317	1 in 1 million

The U.S. Department of Homeland Security (DHS) recently sponsored independent testing in the USA of the same Daugman algorithms as reported here for the UAE database. In a total of 1 707 061 393 (1.7 billion) cross comparisons between different irises, the smallest Hamming distance observed in the DHS test [8] was in the range of 0.28, consistent with the above table for that number of comparisons.

VII. CONCLUSION AND DECISION POLICY RECOMMENDATIONS

The requirements of biometric operation in "identification" mode by exhaustively searching a large database 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).

If P_1 is the false match probability for single one-to-one verification trials, then $(1 - P_1)$ is the probability of not making a false match in single comparisons. The likelihood of successfully avoiding this in each of N independent attempts is therefore $(1 - P_1)^N$, and so P_N, the probability of making at least one false match when searching a database containing N different patterns, is

$$P_N = 1 - (1 - P_1)^N. \qquad (12)$$

Observing the approximation that $P_N \approx NP_1$ for small $P_1 \ll (1/N) \ll 1$, when searching a database of size N an identifier needs to be roughly N times better than a verifier to achieve comparable odds against making false matches. In effect, as the database grows larger and larger, the chance probability of making a false match also grows almost in proportion. These chances also grow in proportion to the number of independent searches that are conducted against the database. To survive successfully so many opportunities to make false matches, the decision threshold policy must be adaptive to both of these factors. Fortunately, the algorithms for iris recognition generate extremely rapidly attenuating tails for the HD_{norm} distribution $f_k(x)$ because of the underlying binomial combinatorics. This felicitous property enables very large databases to be accommodated, and large numbers of searches to be conducted against them. The rule to be followed for decision policy threshold selection is to multiply the size of the enrolled database times the number of searches to be conducted against it in a given interval of time, and then to determine from Table 1 what Hamming distance threshold will correspond to the risk level that is deemed to be acceptable.

For example, in the U.K. with a national population of about 60 million, an "all-against-all" comparison of IrisCodes (totaling about 10^{15} pairings) as envisioned to detect any multiple identities when issuing the proposed biometric identity cards, could be performed using a decision threshold as high as 0.22 without expecting to make any accidental false matches. At this threshold, the false nonmatch rate, using today's better iris cameras (assuming good acquisition and cooperative subjects) would be below 1%. In everyday biometric transactions in which an identity is first asserted and then verified without exhaustive database search, matches with a very forgiving Hamming distance as high as about 0.32 could be accepted safely. ∎

REFERENCES

[1] A. Blake and M. Isard, *Active Contours*. Heidelberg, Germany: Springer-Verlag, 1998.

[2] J. G. Daugman, "High confidence visual recognition of persons by a test of statistical independence," *IEEE Trans. Pattern Anal. Mach. Intell.,* vol. 15, no. 11, pp. 1148–1161, Nov. 1993.

[3] ——, "Statistical richness of visual phase information: Update on recognizing

persons by their iris patterns," *Int. J. Comput. Vis.,* vol. 45, no. 1, pp. 25–38, 2001.

[4] J. G. Daugman and C. J. Downing, "Epigenetic randomness, complexity, and singularity of human iris patterns," *Proc. Royal Soc. Lond. B, Biol. Sci.,* vol. 268, pp. 1737–1740, 2001.

[5] J. G. Daugman, "The importance of being random: Statistical principles of iris recognition," *Pattern Recognition,* vol. 36, pp. 279–291, 2003.

[6] ——, "How iris recognition works," *IEEE Trans. Circuits Syst. Video Technol.,* vol. 14, no. 1, pp. 21–30, Jan. 2004.

[7] D. Gabor, "Theory of communication," *J. Inst. Electr. Eng.,* vol. 93, pp. 429–457, 1946.

[8] International Biometric Group, *Independent Testing of Iris Recognition Technology,* 2005.

[9] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. Comput. Vis.,* vol. 1, pp. 321–331, 1988.

ABOUT THE AUTHOR

John Daugman received his degrees from Harvard University, Cambridge, MA.

He taught on the Harvard faculty before coming to Cambridge University, Cambridge, U.K., in 1991, where he is tenured in the Computer Laboratory. He has held visiting Professorships at the University of Groningen (the Johann Bernoulli Chair of Mathematics and Informatics) and at the Tokyo Institute of Technology (the Toshiba Endowed Chair). He serves or has served as Associate Editor of journals such as IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, *Network: Computation in Neural Systems*, *International Journal of Wavelets Multiresolution and Information Processing*, *Cognitive Brain Research*, and *Journal of Computation and Mathematics*. His areas of research and teaching are computer vision and biological vision, pattern recognition, information theory, and neural computing. He is the inventor of iris recognition, and his algorithms currently underlie all public deployments of this technology worldwide.

Awards for his scientific and technical work include the U.S. Presidential Young Investigator Award (National Science Foundation), the Information Technology Award and Medal of the British Computer Society, the "Millennium Product" Award of the U.K. Design Council, the Smithsonian and "*Time* 100" Innovators Award, and the OBE, the Order of the British Empire.