

Iris Recognition

The colored part of the eye contains delicate patterns that vary randomly from person to person, offering a powerful means of identification

John Daugman

Gwendolen. *What is your name, now that you have become someone else?*

Ernest. *Lady Bracknell, I hate to seem inquisitive, but would you kindly inform me who I am?*

These snippets come from Oscar Wilde's *The Importance of Being Earnest*, a Victorian-era comedy about mistaken identity and the assumption of different identities when convenient. In real life, of course, establishing or proving a person's identity can be very difficult if it must be done instantly, automatically and with a high level of certainty. Conventional methods of identification use documents, tokens or secret knowledge such as login passwords or personal identification numbers. These methods first appeared thousands of years ago—Bronze Age amulets and the military passwords of Roman centurions for example. But the trouble with such methods is that they really only test

whether the secret knowledge or special possession is present, not whether its rightful holder is. So an alternative approach, called biometric identification, is receiving increasing attention.

Biometrics, the measurement of anatomical or behavioral characteristics, seeks to link identity with some complex and unique biological trait. Familiar examples include fingerprints, facial appearance, voice spectrogram, hand geometry, retinal vasculature and DNA sequence. Although each of these attributes has certain universal features, each also shows a significant amount of random variation.

Following a fundamental principle of pattern recognition, biometric identification is reliable only if the disparity between people in the trait measured is significantly greater than the variation presented by a given person over time or across conditions. And therein lies the rub. An acknowledged problem with face recognition, for example, is that even frontal views of a given face can show greater variation (say, from changes in expression or angle of illumination) than do images of different faces. Sandy Pentland of the Massachusetts Institute of Technology and Jonathon Phillips of the National Institute of Standards and Technology recently reported that for images taken at least one year apart, the best face-recognition algorithms have error rates in the range of 43 to 50 percent, even when discounting the effects of changing pose, viewing angle and accoutrements.

A further problem in face recognition is that different faces have much in common, and indeed they can be almost indistinguishable for persons who are genetically the same (identical twins, or in the likely future, clones). Given the frequency of identical twins

in the general population, any identification scheme based on a genetically determined trait or indeed on DNA itself will fail to make the correct identification at least one percent of the time.

So it is desirable to base biometric identification on traits that show great variability among different persons, are stable over time, and are epigenetic (not genetically determined). Several years ago I was prompted by ophthalmologists to try to develop mathematical algorithms that could achieve this goal using the random patterns in the iris of the eye. I reluctantly agreed to the challenge. The algorithms I developed are the basis of all iris recognition systems now in use or undergoing public trials. Some of these field tests have involved millions of iris comparisons, yet there has never been a single false match recorded. I will describe here how these algorithms work and why this fascinating organ of the eye is able to act as such a discriminating identifier.

Windows to the Soul

There is a long tradition of regarding eyes as windows to the soul. For example, in Shakespeare's *The Merchant of Venice*, Portia's suitor Bassanio recalls how "Sometimes from her eyes I did receive fair speechless messages." And perhaps because we all instinctively understand that pupil size is related to

Figure 1. Irises show richly detailed geometric patterns, as this close-up image of the pupil and surrounding iris tissue reveals. As early as 1936, people toyed with the notion of using the iris for identification, because the arrangement of features is unique to the individual. But only recently has a practical scheme for iris recognition emerged, one that uses a computer algorithm the author developed for this purpose.

John Daugman is a tenured member of the faculty at the University of Cambridge, where he teaches courses in neural computing, information theory, and computer vision. Before coming to Cambridge he held the Toshiba Endowed Chair at the Tokyo Institute of Technology, and before that he taught at Harvard University, where he also received his B.A. and Ph.D. degrees. He is a co-founder and board member of Iridian Technologies, a company created to exploit the iris recognition algorithms described in this article. The awards he has received include the Presidential Young Investigator Award of the National Science Foundation; the Information Technology Award and Medal of the British Computer Society; and the Order of the British Empire from Her Majesty Queen Elizabeth II. Address: The Computer Laboratory, University of Cambridge, Cambridge CB2 3QG United Kingdom. Internet: John.Daugman@cl.cam.ac.uk

© Sigma Xi, The Scientific Research Society.
Reproduction with permission only.



emotional state and arousal, Italian women of the Renaissance cosmetically applied atropine to their eyes to enlarge their pupils. For this reason, charmingly, the plant that is the natural herbal source of atropine is named *belladonna*. But beauty aside, the eyes can truly serve as windows to a person's identity. Indeed the first automated ocular biometric was a retinal method patented by Robert Hill in 1978, using close-up ophthalmoscopic imaging of the pattern of internal blood vessels at the very back of the eye.

Although eye color was studied as an identifier by the 19th-century French physician Alphonse Bertillon, the notion that the complex patterns displayed in the iris could serve as a kind of optical fingerprint was first proposed in 1936. Frank Burch, a prominent eye surgeon from St. Paul, made this suggestion in an address to his medical colleagues at the annual meeting of the American Academy of Ophthalmology. His concept became reproduced in standard ophthalmology textbooks, yet even a half-century later, no practical means for iris recognition had emerged. Then in 1989 two ophthalmologists, Leonard Flom and Aran Safir, brought the idea to my attention. Two years earlier Flom and

Safir had patented Burch's concept, but they had no algorithms or mathematical theory to make it a reality, so they asked me to give it a try. At that time I was teaching at Harvard University and was soon to be moving to the University of Cambridge, so I was not eager to engage in any new projects—particularly one only superficially related to the core of my research. But Flom and Safir intrigued me with photographs they had obtained in their clinical practice.

It was apparent from those pictures that irises show complex random patterns, created by arching ligaments, furrows, ridges, crypts, rings, sometimes freckles, a corona and a zigzag collarette. Muscles constrict and dilate the pupil to control the amount of light entering the eye, since the back of the iris is covered with dense black pigment. The front layer of the iris also contains scattered granules of melanin pigment, and it is the density of melanin in this layer that determines eye color. Blue irises result from an absence of such pigment, since long-wavelength light (the red part of the spectrum) penetrates into the tissue and is absorbed, whereas shorter wavelengths are scattered over a wide angle. (The same phenomenon acting

in the atmosphere accounts for the color of the sky.) A thick layer of melanin makes brown irises, but they too reveal a lot of texture and detail when illuminated with light in the near infrared band (700 to 900 nanometers is best, just outside the visible range so that illumination is not unpleasant).

The iris begins to form in the third month of gestation, and the structures creating its pattern are largely complete by the eighth month, although the accumulation of pigment can continue into the first postnatal years. Whereas the overall anatomical form, physiology and color of the iris are genetically determined, the finer details develop in a haphazard fashion. Identical twins, for example, have uncorrelated iris patterns. It is easy to demonstrate that genes do not specify this distinctive texture in detail, simply by looking into a mirror to compare one's own left and right eyes—which are, needless to say, genetically identical.

Situated behind the cornea and the aqueous humor, the iris is an internal organ, yet one that is externally visible and that can be imaged well even from distances up to about a meter. So it is well protected from the environment, and stable over time except for the distortions caused by constriction and dilation of the pupil. Clinical photographs of eyes spanning decades attest to the stability of the iris pattern, even in minute detail.

There is a popular occult belief that iris features can reveal one's changing health (certain points in the iris being associated with particular organs in the body) and that their appearance reflects personality and can be used to read the future. Such beliefs and practices, called "iridology," are occasionally taken seriously enough to be submitted to blind scientific tests. The conclusions have always been the same. In one such study, Lennart Berggren, an ophthalmologist at Uppsala University Hospital in Sweden, concluded that iridology should be dismissed as "a deceptive practice and a medical fraud."

The available evidence to date indicates that iris patterns are indeed as fixed as one's fingerprints, provided that one inverts mathematically the deformations caused by changing pupil size. Although some eyedrops (for example, in the treatment of glaucoma) can cause pigment blotches and coloration changes, these effects on

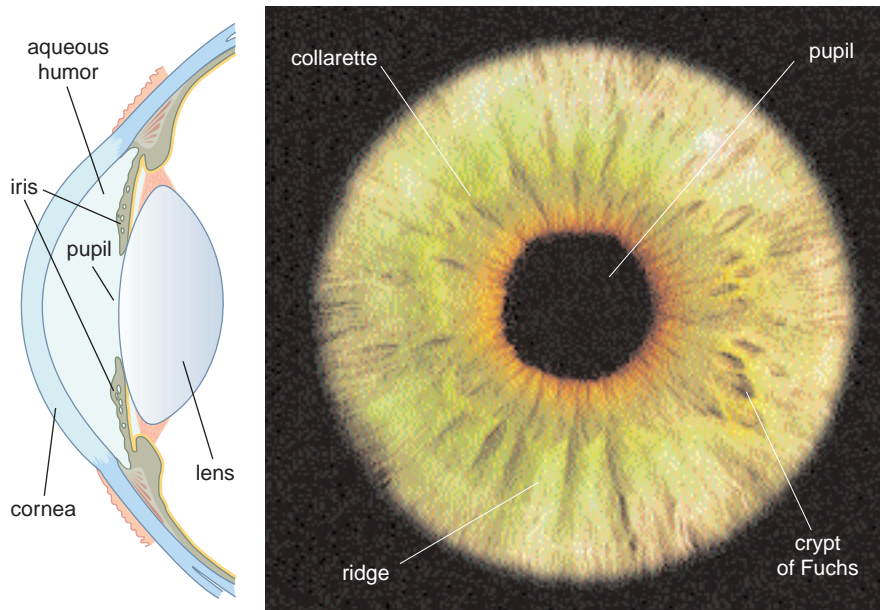


Figure 2. Anatomy of the eye mimics the workings of a camera, with the heavily pigmented iris providing a variable aperture, allowing more or less light to pass through the lens. Situated behind the cornea and aqueous humor, the iris is an internal yet easily visible organ. Although the detailed pattern of iris tissue varies from person to person, certain anatomical features are common. For example, in the middle of many irises lies a zigzag, largely circular structure called the collarette. Running radially and blending into the collarette one often finds slightly raised linear "ridges." Dark oval areas of relatively deep relief seen in some subjects are known as "crypts of Fuchs."

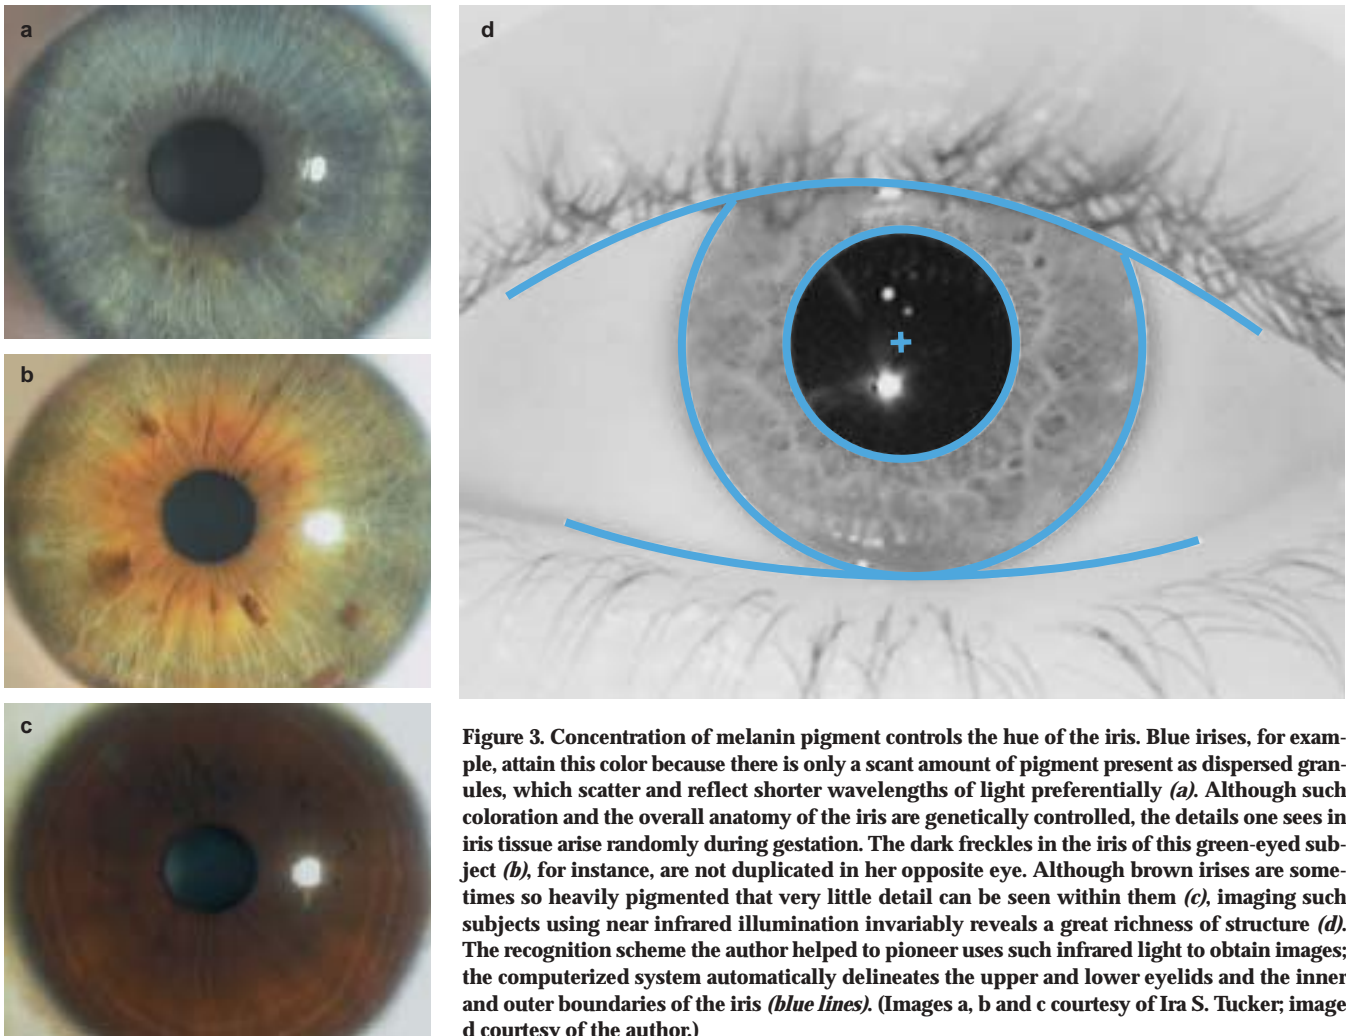


Figure 3. Concentration of melanin pigment controls the hue of the iris. Blue irises, for example, attain this color because there is only a scant amount of pigment present as dispersed granules, which scatter and reflect shorter wavelengths of light preferentially (a). Although such coloration and the overall anatomy of the iris are genetically controlled, the details one sees in iris tissue arise randomly during gestation. The dark freckles in the iris of this green-eyed subject (b), for instance, are not duplicated in her opposite eye. Although brown irises are sometimes so heavily pigmented that very little detail can be seen within them (c), imaging such subjects using near infrared illumination invariably reveals a great richness of structure (d). The recognition scheme the author helped to pioneer uses such infrared light to obtain images; the computerized system automatically delineates the upper and lower eyelids and the inner and outer boundaries of the iris (blue lines). (Images a, b and c courtesy of Ira S. Tucker; image d courtesy of the author.)

melanin are invisible in the infrared wavelengths used for unobtrusive imaging. Infrared illumination also allows selective filtering to eliminate corneal reflections of the ambient environment, which may obscure the iris. The predominant texture seen in visible light is created by the elastic meshwork of ligaments at the front of the iris, whereas infrared light reveals deeper and somewhat coarser features. Some of the fine details of blue irises are lost when illuminating in this wavelength band, but it is a good choice for iris recognition because it ensures that irises of even the most dark-eyed subjects will reveal sufficiently rich features for making a unique identification.

The IrisCode

Although it is small (11 millimeters in diameter on average) and sometimes problematic to image, the iris proves quite amenable to automatic biometric analysis. Its prime mathematical advantage is that its pattern variability among different persons is enormous.

But another advantage is that it is fairly easy for computer vision algorithms to find eyes in faces and track them. The distinctive annular shape of this organ makes it easy to isolate its inner and outer boundaries, and also to detect any occlusion by the upper and lower eyelids. Once detected, all of these boundaries are given a simple mathematical description, and so the iris region is isolated.

Because the apparent size of the iris depends on distance to the camera and the magnifying power of the optics employed, and also because the diameter of the pupil within the iris can dilate over a wide range (from about 2 to 8 millimeters), any biometric encoding must begin with construction of a size-invariant representation. The method for iris recognition that I devised in 1990 and patented in 1994 starts by creating a doubly dimensionless polar-coordinate system. The angular coordinate is intrinsically dimensionless—that is, size-invariant. And in this case the radial coordinate is also, because the

span from the inner to the outer boundary of the iris is made to correspond to the unit interval from 0 to 1. The grid circles of this coordinate system are not strictly polar and concentric, however, because the pupil is usually not exactly centered within the iris (it tends to lie slightly below center and toward the nose). This dimensionless coordinate system has the added advantage that it automatically compensates for the elastic deformations that affect the iris pattern as the pupil changes size.

These first computational steps, isolating the relevant portion of the image and creating a suitable coordinate system for it, were rather obvious requirements. But extracting and encoding the complex patterns in the iris was a more challenging question. At the time I took on this problem, I had been working on an emerging mathematical scheme for processing visual information and for pattern recognition, using two-dimensional *wavelets*. The complex natural patterns in irises seemed like an in-




Figure 4. Two-dimensional Gabor wavelets used to encode iris patterns are oscillating functions restricted to compact, bell-shaped envelopes. The wavelets employed in the author's algorithm either have an even symmetry (left) or an odd symmetry (right), just as do the ordinary cosine and sine functions. Gabor wavelets such as these provide an efficient way to encode both the position and the spatial frequency content of the features in an image.

triguing opportunity to test the theory.

The wavelets are, in essence, sinusoids restricted to compact, bell-shaped envelopes in two spatial dimensions. The sinusoids have both a spatial frequency and a direction of oscillation, but the envelopes give them something else that pure sinusoids lack: locality. The oscillations rapidly die away outside of their localizing envelopes. The analysis of signals or images in terms of wavelets has something in common with Fourier analysis. Just as the French mathematician Joseph Fourier conjectured 200 years ago that any function could be made from a superposition of sinusoids with suitably chosen amplitudes and phases, so also can any function be represented as a superposition of localized wavelets. In fact, all of Fourier analysis can be regarded as just a special case of wavelet theory—namely when the wavelets' envelopes are made so broad that they no longer restrict the sinusoids. But imposing the localizing envelopes adds a profound enrichment over classical frequency analysis: It allows us to assign a phase and an amplitude to any point in a signal, rather than only being able to assign phases and amplitudes to global Fourier frequency constituents.

There are many different kinds of wavelets studied today, and their powerful mathematical properties have revolutionized signal analysis and image processing in little more than a decade. The ones that I had been developing for two-dimensional patterns (and chose to apply to iris recognition) were first proposed in one-dimensional form in 1946 by the Hungarian physicist Dennis Gabor. These have the unique mathemat-

ical property that their precision in localizing the position of some feature in a signal, and their simultaneous specificity about the frequency composition of that feature, achieve the greatest levels that are simultaneously possible under the famous Heisenberg uncertainty principle. When used for image analysis and coding, such wavelets therefore extract information about both *what* and *where* with the greatest possible joint resolution.

Applying the two-dimensional form of these wavelets to iris patterns allows us to encode them into a bit stream representing the iris phase sequence—almost like a DNA sequence, but specifying local phase angles instead of amino acids. This encoding process may be viewed as a kind of *demodulation*, because it regards the iris texture as being made up of modulated wavelets and then determines what their phases would be. When a patch of an iris is projected onto (multiplied by) each of these complex-valued wavelets, the resulting coefficients specify a point in the complex plane having real and imaginary axes. The angle of this “phasor” is an estimate of the local phase for that patch of the iris. I quantize this phasor angle to one of the four quadrants of the complex plane and assign a pair of bits to identify which quadrant it is. These phase-sequence bits constitute a complete IrisCode when the demodulation process has been repeated across all parts of the iris, using wavelets having a number of different sizes, spatial frequencies and orientations. A total of 2,048 bits (256 bytes) of phase information is extracted.

IrisCodes are rapidly compared by computing their logical “exclusive-or” (XOR) in parallel chunks as a single computer instruction. The elementary Boolean operator XOR detects disagreement between pairs of bits: The four possible bit pairings—00, 01, 10 and 11—respectively have XOR values of 0, 1, 1 and 0. Standard programming languages allow a computer with (say) a 32-bit central processing unit to generate from two binary strings of length 32 bits a third string of the same length, each of whose bits is the XOR of the two corresponding bits in the two input strings—and to do this on all of the bits in parallel as a single operation. This capability allows IrisCodes to be fully compared with each other at a rate of about 100,000 per second on a desktop PC when searching for a match. The total number of “1” bits in the XOR resultant string is a measure of their dissimilarity (Hamming distance), which is compared against a predetermined threshold in order to make a decision about whether or not two IrisCodes are statistically independent.

Test of Statistical Independence

When different irises are compared in this manner, the fact that their patterns possess so much random variation makes it nearly impossible for them to agree by chance in more than about two-thirds of their phase-sequence bits. Any given bit in an IrisCode is equally likely to be a 1 or a 0, and so the chance of agreement between two bits for different irises is 50–50. Getting a large number of agreeing bits by chance has the same probability distribution as tossing a coin a large number of times and getting significantly more than expected of a certain outcome. To be precise: If one tosses a coin whose probability of “heads” is p (which need not necessarily be 0.5) a total of N times, then the probability of getting m heads out of the N tosses, that is to say a fraction of $x=m/N$ heads, is given by the binomial distribution $f(x)$:

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

The binomial distribution has extremely rapidly decaying tails when N is large. For example, whereas the odds against getting all “heads” in 10 tosses of a fair coin are 1,024 to 1, the odds against getting all “heads” in 100 tosses are an astronomical 2^{100} , or approxi-

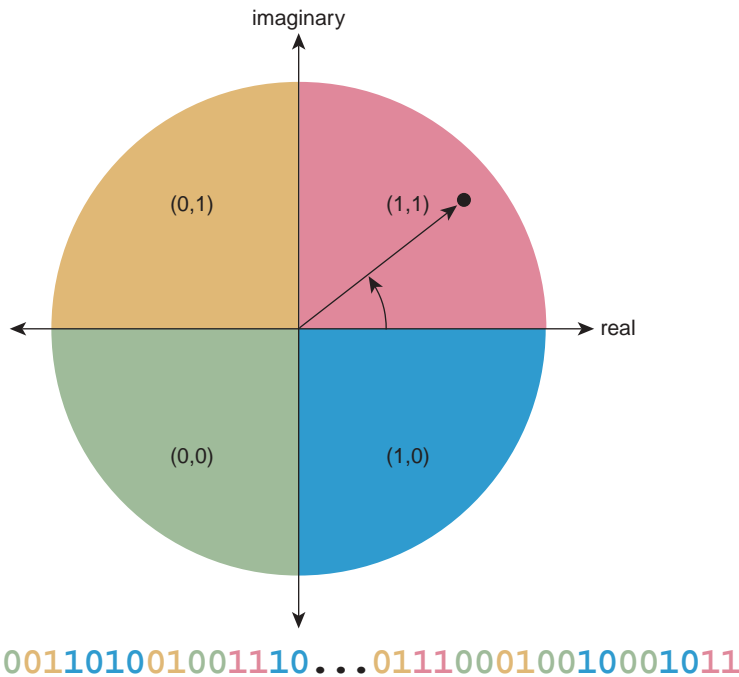


Figure 5. Multiplication of the signal (a two-dimensional patch of the iris image under study) with a pair of Gabor wavelets such as shown in Figure 4 specifies a point or position vector in the complex plane. The real component of this vector represents the correspondence between the wavelet of even symmetry and the signal; the imaginary component represents the match between the wavelet of odd symmetry and the signal. The author’s method discards the amplitude of these components and retains only a coarse measure of the phase angle, recorded with two bits that specify the quadrant in which the point falls (*red, yellow, green or blue*). A string of 1,024 such pairs (*bottom*) constitutes what the author has dubbed the “IrisCode.”

mately 10^{30} (a thousand billion billion billion), to 1.

The histogram in Figure 7 shows the results of comparing 2.3 million different pairs of irises, using images obtained during public trials of these algorithms conducted by licensees in several government testing bodies and companies in the U.K., the U.S. and Japan. The abscissa is Hamming distance (HD), the fraction of bits that disagree between any two different IrisCodes. As expected, this distribution is centered near 0.5, having an actual mean of 0.499. The solid curve that is also plotted, fitting the count distribution almost perfectly, is a plot of the binomial distribution with parameters $p=0.5$ and $N=244$. This is the distribution one would expect to get if one tossed a fair coin 244 times in a row and counted the fraction $x=m/N$ of “heads” in each of many such repeated runs. The reason for the near-perfect agreement between the raw data and the theoretical binomial distribution is that any comparison between two bits from different IrisCodes is essentially what statisticians call a Bernoulli trial (and everyone else calls a coin toss).

But these are runs of correlated Bernoulli trials rather than independent ones, because in the IrisCode for any given iris there are some internal correlations among the bits. These arise because iris patterns themselves are auto-correlated, especially radially. For example, an existence of radial furrows in an iris means that the structure near the pupil predicts to a certain extent the structure away from the pupil. Bernoulli trials that are correlated still generate binomial distributions, but with a reduction in N , their effective number of “tosses.”

It is important to validate further such a statistical model by examining the cumulatives of the observed data up to successive points on the tail of the distribution and comparing these with the cumulatives predicted by binomial theory. The straight line relationship in such a “quantile-quantile” plot (Figure 8) reveals very precise agreement between model and data, over a range of more than three orders of magnitude. It is clear that it is extremely improbable that two different irises might disagree by chance in fewer than at least a third of their bits. (Of

the 2.3 million iris comparisons plotted in Figure 7, the smallest HD observed was 0.353.) Computing the cumulative of the theoretical binomial distribution $f(x)$ from 0 to 0.333 indicates that the probability of such an event is about 1 in 10 million. The cumulative from 0 to just 0.300 is 1 in 7 billion. Observing such small HDs would allow one to reject the hypothesis of statistical independence, a hypothesis that requires HDs to be centered around 0.5 and to be binomially distributed. It is the rejection of this hypothesis of statistical independence by observing significantly smaller HDs that implies that two IrisCodes originate from the same iris, thereby permitting a positive recognition decision. Even the observation of a rather poor degree of match between the IrisCodes for two different iris images (say, only 70 percent agreement or HD = 0.300) would still provide extraordinarily compelling evidence of identity, because the test of statistical independence would still be failed so convincingly, indeed at the 10^{-9} level.

As a side question, I also compared genetically identical eyes in the same manner, in order to discover the degree to which their iris patterns were correlated and hence genetically determined. A convenient source of genetically identical irises is the right and left pair from each person; such pairs have the same genetic relationship as the four irises of monozygotic twins, or




Figure 6. Bits in the IrisCode have an equal probability of being set to 1, no matter what their position. (For clarity, only 128 representative bits of the complete 2,048-bit code are shown here.) Because the probability of being set to 1 is roughly 0.5 in all cases, each bit carries the most information theoretically possible, allowing the code to be extremely compact. Such short codes permit a computer to compare a large number of them rapidly.

indeed the prospective $2N$ irises of N clones. Although eye color is of course strongly determined genetically, as is overall iris appearance, the detailed patterns of genetically identical irises are as uncorrelated as the iris patterns of unrelated eyes. Using the same methods as described above, right-left iris pairs from 324 persons were examined. Their mean HD was 0.497 with standard deviation 0.031, and their distribution was statistically indistinguishable from the distribution for unrelated eyes. A set of six pairwise comparisons among the eyes of actual monozygotic twins also yielded a result (mean HD = 0.507) expected for unrelated eyes. It appears that the phenotypic random patterns visible in the human iris are not genetically determined but are almost entirely epigenetic.

I described earlier how invariance to iris size (hence distance, and optical magnification) is achieved; also invariance to pupil dilation (by the deforming polar-coordinate system); and of course, invariance to iris position (by localizing the iris within the two-dimensional image). But a further invariance is still required. The orientations at which persons' irises are visible to cameras in practical implementations of this technology are unpredictable. Persons' heads may be tilted; video-camera imaging through pan-tilt mirror systems can produce tilt; and the eye itself can rotate about 15 degrees around its optical axis in its socket. For all of these reasons, it is important to conduct iris comparisons in a way that allows for numerous possible relative rotations between the presenting and the enrolled iris images.

The most computationally efficient way to deal with this factor, given the enormous speed of the search engine, is simply to perform iris comparisons at multiple relative orientations. IrisCodes can be scrolled in cylindrical coordinates to represent the effects of tilted imaging. The inherent odds against accidental matches occurring by chance are so astronomical that even this large increase in the number of candidate matches can be easily accommodated without observing a degradation in performance.

The resulting new distribution of HDs, among the same set of 2.3 million different iris pairings considered earlier but after performing seven relative rotations between each pair and recording their best match, is shown in Figure 9.




Figure 7. Distribution of Hamming distances (a quantitative measure of the similarity between the IrisCode bit strings) was obtained for 2.3 million pairings of different people's eyes, which were examined during public trials of the author's system for iris recognition (orange bars). The statistical results obtained match the theoretical binomial distribution for a probability of 0.5 and 244 degrees of freedom (purple line). Although the numbers are smaller and the resulting distribution coarser, a comparison of right and left eyes for 324 people follows exactly the same statistical distribution, demonstrating that the features captured by the IrisCode are not genetically determined.

Because only the best match (smallest HD) is plotted after the comparisons in several orientations, this distribution is biased towards lower HDs, with a mean of 0.458, and it is sharper. The solid curve in this histogram is the theoretical prediction for such "extreme-value" repeated sampling from a binomial, and once again it very precisely agrees with the data histogram. Marked off under the left tail of the distribution are its theoretical cumulatives from zero up to various points. These infinitesimally small numbers represent the probability that a sample from this distribution could occur anywhere to the left of each such point. They therefore represent the probability that any two different irises might generate an HD smaller than such a score. For example, the probability that two different irises might by chance produce an HD smaller than 0.25 (that is, agree in more than 75 percent of their phase sequence) is a minuscule 10^{-14} .

The fact that the odds against a false match remain so astronomical has great practical significance, because in searching very large databases, the

confidence of still avoiding any false match despite so many opportunities must be diluted by the number of such independent comparisons being performed. For example, using a decision criterion of HD = 0.31, the single-trial probability of a false match is 1 in 127




Figure 8. Quantile-quantile plot for the left tail of the orange distribution shown in Figure 7 demonstrates excellent agreement over several orders of magnitude between observed and theoretical cumulative values. The correspondence indicates that the empirical results are indeed binomially distributed.




Figure 9. Probability distributions for the comparison of two images of the same iris taken under different conditions (*blue*) is distinct from the distribution obtained from pairs of different irises (*pink*). The latter distribution is shifted slightly to the left compared with the one shown in Figure 7, because the test displayed here examined multiple relative rotations and kept only the best match. These results follow theoretical expectations for such “extreme-value” sampling (*red line*). The theoretical curve is useful for predicting the astronomically small probabilities that two unrelated irises will by chance produce a Hamming distance smaller than the indicated value.

million; but this implies that if the size of the database being searched were 1 million people, then the probability of getting at least one false match among them is nearly 1 percent (barely acceptable). For this reason, decision thresholds are adaptive to take into account the size of the search database. Although the total image processing time (from raw image to extracted IrisCode) is less than one second, the search speed is about 100,000 IrisCodes per second and so larger database searches should be farmed out into parallel chunks of about this size.

Iris recognition is based on the *failure* of a test of statistical independence between a presenting iris image and a previously enrolled one. This test is deemed to have been failed if the amount of disagreement between phase sequences is less than about a third, because the chance of this happening for two different irises is well below one in a million. But how closely will two different images of the same eye, taken at different times using different cameras, in different environments, and under different illumina-

tion conditions, agree with each other? Data from 7,070 such comparisons are shown in the left distribution of Figure 9. Although it is rare to get a perfect match (HD = 0.00) between different images of the same eye across such conditions, still the amount of disagreement introduced by all these confounding factors is almost always less than 34 percent. Therefore these two distributions for “same” and for “different” iris comparisons can very reliably be separated. The amount of overlap between them on either side of any chosen decision criterion (for example, HD = 0.33) gives the error rates for false match (“false accept”) and for failure to match (“false reject”).

As has been confirmed by all of the independent field trials of these algorithms to date, the dual distributions show that the false accept probability is well below one in a million (indeed, no false match has ever actually been recorded), while the false reject rate remains below about 1 percent. Another quantitative way to calibrate the power of decision making for this type of two-choice (same versus different) task is

using the metric called d' , for decidability, which is defined as the difference between the means of the two distributions divided by the square root of their average variance. The advantage of this metric is that it does not depend on any particular choice about where to put the decision threshold, as do statistics about absolute error rates. Instead, it reflects the degree to which any improvement in one error rate must be paid for by a worsening of the other error rate—in other words, the amount of overlap between the two distributions. This measure of separability gives a rating of 7.3 for the dual distributions shown in Figure 9, a value higher than for any other known biometric. Such powerful decision making in personal identification, based on iris pattern combinatorics and a simple test of statistical independence, reveals (as Oscar Wilde might have said) the importance, and the benefits, of being random.

Bibliography

- Berggren, L. 1985. Iridology: A critical review. *Acta Ophthalmologica* 63(1):1-8.
- 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. 1993. High confidence visual recognition of persons by a test of statistical independence. *IEEE Transactions Pattern Analysis and Machine Intelligence* 15(11):1148-1161.
- Daugman, J. 1994. Biometric Personal Identification System Based on Iris Analysis. U.S. Patent Number 5,291,560.
- Pentland, A., and T. Choudhury. 2000. Face recognition for smart environments. *Computer* 33(2):50-55.
- Phillips, P. J., A. Martin, C. L. Wilson and M. Przybocki, 2000. An introduction to evaluating biometric systems. *Computer* 33(2):56-63.
- Wilde, O. 1895. *The Importance of Being Earnest*. London: Leonard Smithers.

Links to Internet resources for “Iris Recognition” are available on the *American Scientist* Web site:

<http://www.americanscientist.org/articles/01articles/daugman.html>