

Understanding Biometric Entropy and Iris Capacity: Avoiding Identity Collisions on National Scales

John Daugman

John.Daugman@CL.cam.ac.uk

*Department of Computer Science and Technology
Cambridge University
Cambridge CB3 0FD, United Kingdom*

Corresponding Author: John Daugman

Copyright © 2024 John Daugman. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract

The numbers of persons who can be enrolled by their iris patterns with no identity collisions is studied in relation to the biometric entropy extracted, and the decision operating threshold. The population size at which identity collision becomes likelier than not, given those variables, defines iris “capacity.” The general solution to this combinatorial problem is derived, in analogy with the well-known “birthday problem.” Its application to unique biometric identification on national population scales is shown, referencing empirical data from US NIST (National Institute of Standards and Technology) trials involving 1.2 trillion (1.2×10^{12}) iris comparisons. The entropy of a given person’s two iris patterns suffices for global identity uniqueness.

Keywords: Biometric, Entropy, Identity, Iris, Pattern recognition

1. INTRODUCTION

Applicants for Cambridge University undergraduate studies in mathematics or computer science are asked sometimes in their College interviews to reason about the “birthday problem”: how many people, chosen at random, must be assembled until it becomes more likely than not that at least one pair of them have the same birthday? Some students are surprised that the answer is only 23 people. Although arriving at the exact number requires a calculator, the reasoning is that N people make $N(N-1)/2$ possible pairings. Given that each pairing has probability $1/365$ of sharing their birthday and $364/365$ of not, the probability that *none* of the pairings share a birthday is approximately $(364/365)^{N(N-1)/2}$, which is < 0.5 once $N \geq 23$.

There is a clear analogy with biometric collision avoidance, which we can formulate as the:

Biometric birthday problem: if some biometric technology is operating with a verification FMR (“one-to-one” False Match Rate), how many people, chosen at random, must be assembled until it becomes more likely than not that at least one pair of them have a biometric collision (are falsely matched to each other)?

A good example is face recognition, tested across a broad variety of scenarios and using a wide range of image quality, for which a good performance benchmark corresponds to making just one verification False Match in 1,000 non-mated comparisons [1–3]. That accuracy standard is better than human (even “super-recogniser”) performance in some circumstances [3]. Face recognition algorithms have improved greatly in recent years, in terms of Rank-1 identification rates [1, 2], in test protocols in which a correct match does always exist within a search gallery that is populated also with other “distractors”. But even in the recent tests, the best algorithms do still make some False Matches to distractor images even when there are only 100 distractors [1, 2], despite the presence of a correct match within the gallery, that should instead actually be returned at Rank-1.

Let us now consider the “biometric birthday problem” for a face recognition algorithm performing at $FMR = 0.001$ when examining a gallery of non-mated faces. How large must this gallery get before False Matches become likelier than not, in all-versus-all comparisons? The answer: just 38. That number creates $38 \cdot 37/2 = 703$ possible pairings to consider, and $(1 - 0.001)^{703} = 0.495$ so False Matches are then already likelier than not. When waiting at Passport Control (or some other such queue), it is entertaining to turn around, look at the first 38 persons standing behind oneself, and try to spot the pair of facial doppelgängers [4], among them.

Biometric deployments at a national or even prospectively at the planetary scale face a massively challenging biometric “birthday problem” if they need to search for any duplicate identities, as was necessary in India when all 1.4 billion citizens were recently enrolled in a national ID programme for welfare distribution, government services, and subsidies (UIDAI: Unique IDentification Authority of India) [5]. Because enrollees had an incentive to acquire multiple identities and thereby issuance of multiple subsidies, every new enrollment had to be compared against all existing enrollments before an Aadhaar would be issued. This amounts to a search for identity collisions, all-versus-all, among an astronomical $N(N - 1)/2$ pairings of persons. Obviously any attempt to do this by face recognition would drown in False Matches from the very beginning. There simply is not enough entropy, or randomness, in human face structure; the necessary functional purposes of major facial features (mouth, nose, ocular areas) constrain their possible randomness. The bilateral symmetry normally present in a face further reduces its entropy by half. The key idea, the fundamental factor underlying the power of biometric identification, is entropy [6, 7].

Weak biometrics may be sufficient to enable “one-to-one” verification; stronger biometrics may enable identification in a search database of size N , “one-to-few” or “one-to-many” depending on N ; but de-duplication applications exemplify the birthday problem in that they are essentially “all-versus-all”, and the number of False Match opportunities they must survive grows massively with N . In such deployments on a national scale, falsely detected or undetected identity collisions (even if few in percentage) would lead to reduced public confidence in and acceptance of the system, its impaired functionality, and legal problems caused both by undetected duplicates and falsely detected ones. TABLE 1 presents, for a broad range of FMR levels spanning 15 orders-of-magnitude, how large N can get before collisions become likelier than not. TABLE I clearly shows that the demands for a minuscule FMR become extremely daunting once the population size N is even that of a small town, let alone a population of national, continental, or of planetary scale.

Table 1: Accuracy Requirements for Biometric Collision Avoidance

Verification FMR	Critical Population Size N
0.001	38 persons
0.0001	119 persons
10^{-5}	373 persons
10^{-6}	1,177 persons
10^{-9}	37,229 persons
10^{-12}	1.2 million persons
10^{-15}	37 million persons
10^{-18}	1.2 billion persons

2. GENERAL SOLUTION FOR POPULATION BOUNDS

The number of pairings possible among N persons is $N(N - 1)/2$ because each person can be paired with $N - 1$ others, but half of these are redundant (e.g. Alice and Bob, then also Bob and Alice); hence the halving. If a biometric technology is operating at some verification False Match Rate FMR, then the probability of a given pairing *not* resulting in a False Match is $(1 - \text{FMR})$, and the probability that *none* of the possible pairings do so is $(1 - \text{FMR})^{N(N-1)/2}$. For what value of N does this expression become < 0.5 , and therefore a biometric collision becomes likelier than not?

We will invoke a property of the base e “natural logarithm” function $\log_{e=2.718...}(\)$, commonly denoted $\ln(\)$. We seek:

$$(1 - \text{FMR})^{N(N-1)/2} < 0.5 \tag{1}$$

$$\ln\left((1 - \text{FMR})^{N(N-1)/2}\right) < \ln(0.5) \tag{2}$$

$$\frac{N(N - 1)}{2} \ln(1 - \text{FMR}) < -0.693 \tag{3}$$

Now using the power series expansion

$$\ln(1 + x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \dots, \tag{4}$$

we have $\ln(1 + x) \approx x$ for small $|x|$, whether $x \geq 0$ or $x < 0$. Basically this reflects the fact that the logarithm function is linear near where it crosses 0 at $\log(1)$, and the slope of this line is 1 if the base of the logarithm is e . Thus for any small FMR (say < 0.01), which also entails that $N^2 \gg N$, we have

$$-\frac{N(N - 1)}{2} \text{FMR} \lesssim -0.693 \tag{5}$$

$$N^2 \gtrsim 1.386/\text{FMR} \tag{6}$$

$$N \gtrsim \sqrt{1.386/\text{FMR}} \tag{7}$$

This general (but approximated) solution can be confirmed by evaluating (1) exactly, using for N each of the corresponding FMR cases tabulated in TABLE 1, insofar as the available tools of calculation can handle the combinatorial exponents required in (1) when N is large.

3. BIOMETRIC ENTROPY TO THE RESCUE

Entropy measures the complexity and randomness [6], that is present in (and between) random variables. Facial structure has limited capacity for randomness. The major facial features have a canonical standard configuration, usually with bilateral symmetry; the eyes are normally on opposite sides of the nose. Much greater randomness is found in iris patterns, and this is the origin of their legendary resistance to False Matches. Although often there do exist strong radial correlations within an iris, with mutual information as large as 0.3 bits per bit across radius [8], and also IrisCode bits at adjacent or nearby angles but a shared radial coordinate have “sticky oscillator” correlations that reduce their entropy as much as 0.5 bits per bit [7], nevertheless the remaining entropy is vast. FIGURE 1 illustrates this graphically in the bit streams that constitute the IrisCodes of four different eyes. How IrisCodes are computed has been revealed previously [9]. The two bit values are equiprobable, so when bits in IrisCodes from two different eyes are compared by XOR (Exclusive-OR) to detect whether they agree or disagree, these outcomes again are equiprobable, amounting to the toss of a fair coin.

The non-independence among the bits in a given IrisCode reduces their collective entropy from what would have been a maximum of 2,048 bits (if each bit corresponded to an independent “fair coin toss” Bernoulli trial) to only about 245 bits. Modelled as a “sticky oscillator” Markov process [7], IrisCode bits exhibit a phase coherence that can persist across several bits. Despite such losses in entropy, enough entropy remains that the collision probability between two IrisCodes from different eyes attenuates by astronomical factors, for small reductions in the tolerated fraction of disagreeing bits.

4. DISCUSSION

A good way to understand this effect intuitively is to consider tossing a fair coin in runs of 245 tosses, tallying each run’s fraction of heads. The total number of possible outcome sequences is 2^{245} and each of these has the same probability, namely $p_i = 2^{-245}$ (including, say, the “all heads” sequence). The entropy [6], contained in these possible sequences is:

$$H = - \sum_i p_i \log_2(p_i) \tag{8}$$

$$= - \sum_{i=1}^{2^{245}} 2^{-245} \log_2(2^{-245}) = 245 \text{ bits.} \tag{9}$$

The vast majority of these sequences will have a nearly equal mix of heads and tails. The fraction of possible sequences that have (say) fewer than 30% heads is less than one-billionth of the total. This combinatorial property when large entropy (245 bits) exists in a random variable is ultimately the reason why, for iris recognition, a match between two IrisCodes can be accepted even when (say)



Figure 1: Representation of the IrisCodes [9], produced by four different eyes. The eight rows within each can be regarded as eight concentric rings, each encoding a $[0, 2\pi]$ traversal around the iris. (Eyelid masking is not shown.)

30% of their bits disagree due to problematic image acquisition. Despite such a lenient criterion being so tolerant of noisy bits, the probability that such an accepted match would actually be a False Match is, indeed, less than 1 in a billion.

The huge exponents appearing in (9) (note that $2^{245} \approx 10^{74}$) are key to understanding why sufficient entropy is the basis for biometric collision avoidance even at a planetary scale. A detailed tabulation of the relevant probability distributions, both densities and their cumulatives [9], with and without selecting for best matches after multiple image rotations to compensate for unknown head and camera tilt, is provided at [10], as a function of Hamming distance HD (fraction of bits that disagree in IrisCodes from two different eyes). This probability table enables us to predict how tolerant we can be of poor image acquisition (how large a fraction HD of disagreeing bits we can tolerate and still declare a match), without resulting in False Matches. The TABLE [10], shows for acceptance criteria HD the resulting False Match probability, and its \log_{10} (last two columns).

TABLE 2 extracts coarser HD increments of 0.01 from [10] (first column), showing the corresponding FMR predictions (second column). By 2003 image databases were only large enough to perform about 10 million iris cross-comparisons [9] but distribution parameters could be estimated, implying 249 bits of entropy (slightly more than 245), predicting FMR performance very similar to what is shown in TABLE 2. No False Matches were observed below roughly the HD = 0.33 criterion, for the small databases available. The predicted FMR values were generally dismissed with incredulity [11], because such FMR performance was unknown in other biometrics. But subsequently, other NIST researchers did actually perform billions [12], and then more than a trillion iris comparisons [13], obtaining FMR values in good agreement with those predictions, as reported in column 3.

Table 2: False Match Rates Predicted in [10], and as Measured by NIST [12], with 1.16 Billion Iris Comparisons, and [13], with 1.2 Trillion Iris Comparisons

<i>HD criterion</i>	<i>FMR predicted in [10]</i>	<i>NIST [12, 13] measured FMR</i>
0.36	1 in 24,000	1 in 25,000
0.35	1 in 110,000	1 in 71,000
0.34	1 in 556,000	1 in 476,000
0.33	1 in 3.1 million	1 in 3.4 million
0.32	1 in 20 million	1 in 24 million
0.31	1 in 137 million	1 in 165 million
0.30	1 in 1.1 billion	1 in 2 billion
0.29	1 in 9 billion	(not measured)
0.28	1 in 92 billion	1 in 40 billion

An important cause of skepticism about the FMR performance levels shown in TABLE 2, before they were eventually confirmed by NIST, was the existence of ‘ground-truth’ errors in early biometric databases that had created illusory identity collisions. Apart from sloppy and naïve data collection, (e.g. incentivising paid student volunteers to change names and thereby enroll multiple times), there is an inherent risk in estimating FMR by *intra*-dataset cross-comparisons. If even just one of N subjects is enrolled under two different identities, whether deviously or just through an

innocent clerical error, the estimated FMR then cannot be better than $2/N^2$. The measured threshold calibration of FMR such as tabulated in TABLE 2, must then approach a floor, corresponding to this illusory FMR, which cannot be reduced by any reasonable change in threshold, and indeed NIST [12] demonstrated this problem for (university-sourced) *intra*-dataset comparisons.

NIST overcame this problem by performing *inter*-dataset comparisons: if two disjoint populations, of sizes (say) N and M in geographically remote places can be biometrically enrolled, then $N \times M$ inter-comparisons become possible without the contaminating effect of ground-truth errors. NIST [13] acquired enrollment datasets for two populations “very well separated geographically and occupationally,” one having 3.9 million iris images used as the gallery, and the other having 315,000 iris images used as probes to search against this entire gallery, asserting there was zero likelihood of co-membership. Thereby NIST performed $N \times M = 1.2$ trillion IrisCode comparisons, leading to the FMR results shown in column 3 of TABLE 2 (from [13] p. 61), for various HD threshold criteria. This close confirmation of theory (column 2), manipulating FMR over a larger than million-fold range, is striking.

5. DEMOGRAPHIC SPECIFIC APPLICATION

Iris pattern entropy differs somewhat across ethnic groups [14]. For example, the anterior layer of the iris in persons of Sub-Saharan African descent contains a thick blanket of melanocytes [15] creating a coarser texture of crypts and craters, than the finer fibrous details typically visible in an iris of persons descended from more northern regions. FIGURE 2 illustrates these entropy differences in samples from three demographics: West African; Irish-American; and Nordic.

Using image databases having particular ethnic demographics, it is possible to estimate quantitatively their characteristic entropies. Such calculations are needed in order to understand how many persons can be enrolled before identity clashes in “all-versus-all” cross-comparisons (at a given acceptance operating criterion), start to become likely. FIGURE 3 illustrates this process for a new West African database of iris images [14] “AFHIRIS”, plotting the distribution of Hamming distances (HD, fraction of bits that disagree) between all possible pairings of IrisCodes for different eyes. The red curve is a plot of the following probability distribution $\text{prob}(\text{HD})$ for the fraction of Heads (HD) in a run of N tosses of a coin whose probability of Heads is p :

$$\text{prob}(\text{HD}) = \frac{N!}{m!(N-m)!} p^m (1-p)^{(N-m)} \quad (10)$$

where in this case $N = 228$, $p = 0.5$, and $\text{HD} = m/N$ is the outcome fraction of N Bernoulli trials (e.g. observing m Heads within a run of N coin tosses). Measuring the std dev σ for an empirical distribution of HD scores from independent pairings tells us the equivalent number of tosses of a coin (having probability p of Heads), namely $N = p(1-p)/\sigma^2$. The empirical distribution has $\sigma = 0.0331$, with $p \approx 0.5$ (mean HD) so each toss adds 1 bit of entropy. Therefore we estimate AFHIRIS biometric entropy as $N = 228$ bits. The fit in FIGURE 3, between the empirical distribution data and the theoretical probability density curve (10) seems excellent.

As was visible in FIGURE 2, and investigated in [14], biometric entropy in iris patterns varies among ethnic groups. The range observed spans from about 225 bits to 265 bits. Those values impact the False Match Rates for any given operating point (with higher entropy reducing the




Figure 2: Entropy differences in iris patterns from three different demographic groups: West African (top); Irish-American (middle); and Nordic (bottom).

Comparison of Empirical AFHIRIS and Theoretical Density Distributions

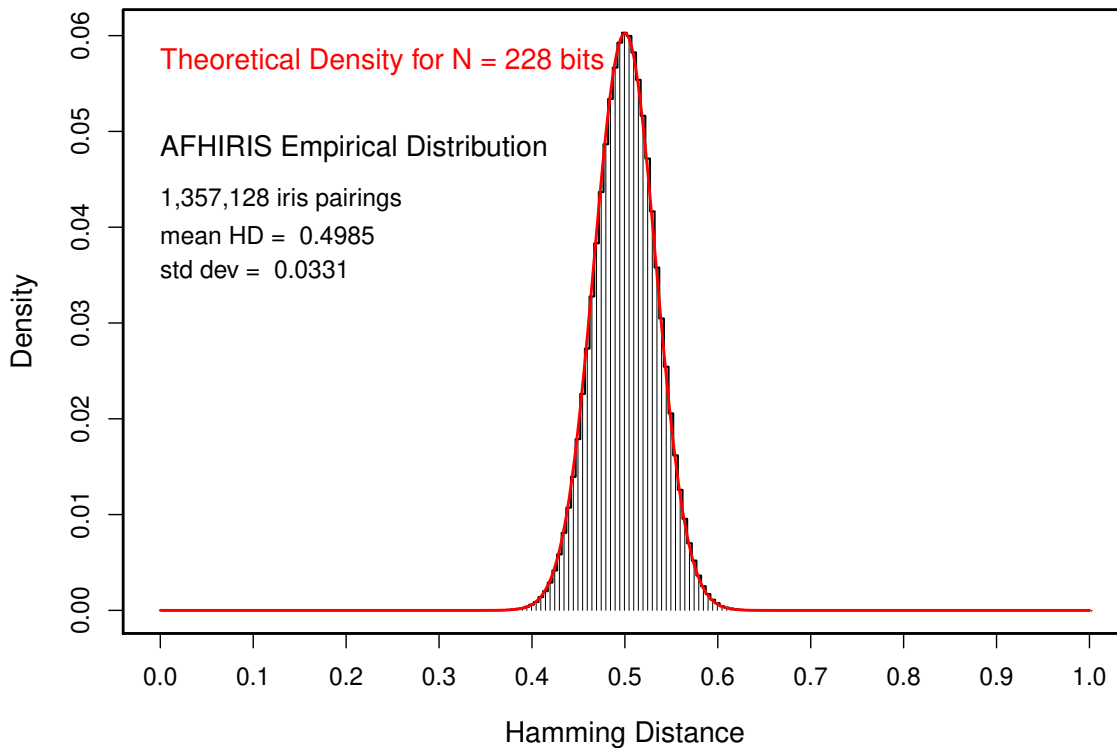


Figure 3: Empirical histogram of “all-versus-all” cross-comparison Hamming distance scores observed in the West African iris image database AFHIRIS, superimposed with the theoretical binomial probability density distribution (red curve) which plots (10) using parameters $p = 0.5$ and $N = 228$.

Table 3: Numbers of Persons Enrollable, With All-Versus-All Iris Cross-Comparisons Unlikely To Have Any Identity Collisions, For Two Operating Points. Single Eye Enrollment Presumed.

<i>Encoded Iris Entropy</i>	$HD_{\text{threshold}} = 0.28$	$HD_{\text{threshold}} = 0.24$
225 bits	134,000 persons	16 million persons
235 bits	222,000 persons	32 million persons
245 bits	370,000 persons	66 million persons
255 bits	615,000 persons	136 million persons
265 bits	1.02 million persons	278 million persons

FMR), and therefore they also affect how large a population of persons can be enrolled without identity collisions in all-versus-all cross-comparisons. Such a concept is sometimes called biometric “capacity” [16] for a given modality and operating point. We can now apply the framework that was introduced at the beginning of this paper, the “biometric birthday problem,” to calculate iris capacity across this observed range of entropies. For any given estimate of biometric entropy, the FMR at a given operating point can be calculated as described in [9] and tabulated in [10] (for the case of $N = 245$ bits of entropy). Using (7) we arrive at the numbers of persons who can be enrolled while identity collision still remains unlikely. These numbers are presented in TABLE 3, for two different HD operating thresholds and five estimates of entropy, always assuming single eye enrollment, to illustrate the combined effects of these variables.

A way to estimate the scalability of face recognition systems was proposed by [16]. They defined “face capacity” in terms of packing bounds: the ratio of the total volume in a representation space, to the volume that is required to represent individual faces in it (as separate spheres or ellipsoids). This yields an extreme upper bound estimate of capacity, because there is no way to ensure that the spheres or ellipsoids for different faces do not overlap. Such collisions or overlaps certainly occur for identical twins, and even for unrelated persons who are facial *doppelgängers* (as illustrated in this collage [4] of examples.) Recent tests by NIST [2] show that current face recognition algorithms fail completely to distinguish between identical twins. About 1% of persons have an identical twin, so in any sufficiently broad sample, face representations must suffer identity clashes for at least those 1%. By contrast, it is well-known that the IrisCode produces as much distance between the encoded iris patterns of identical twins (or indeed between the two eyes of any given person) as between unrelated eyes [9].

6. CONCLUSION

Iris recognition is perhaps unique among biometrics in having clear mathematical foundations, enabling strong predictions about IrisCode collision likelihood as a function of the decision threshold. As shown in TABLE 2, for decision criteria in which no more than about 31% of the IrisCode bits are allowed to disagree when declaring a match (which is a very noise-tolerant criterion), the predicted FMR attenuates by almost a factor of 10 for each additional 1% reduction in the tolerated amount of bit disagreement. This extraordinary fact seems not to be generally understood or appreciated;

but it is a direct result of using high-entropy random variables in biometric codes. A critical lesson emerging here is the same as a lesson from cryptography: the great power of randomness, if you can get enough of it.

As confirmed independently by NIST in [13], the slope of the IrisCode Decision Error Trade-off curves is so flat that the FMR can be lowered by a factor of 10,000 to 100,000 while not even doubling the False non-Match rate (FnMR). A consequence of this relationship is that only small costs in increased FnMR need be paid, by lowering HD threshold, in order to increase greatly the size of a biometrically enrolled population without suffering collisions. Thus for IrisCodes from any two different eyes, the probability of $HD \leq 0.29$ is about 10^{-10} . If we also exploit the fact that a person's two eyes generate IrisCodes that are almost completely independent, specifying 0.29 as a match criterion *binocularly* would yield a fusion FMR of about 10^{-20} . Equation (7) shows us that this is how the planetary human population can survive the "biometric birthday problem": it is unlikely that even a single pairing among 12 billion persons (despite the vast numbers of possible pairings) would disagree in $\leq 29\%$ of their IrisCode bits for both pairs of eyes. Thus speaks biometric entropy.

References

- [1] Kemelmacher-Shlizerman I, Seitz SM, Miller D, Brossard E. The Megaface Benchmark: 1 Million Faces for Recognition at Scale. *Recog Int Conf. Comp Vision & Patt.* 2016:4873-4882.
- [2] <https://nvlpubs.nist.gov/nistpubs/ir/2018/NIST.IR.8238.pdf>
- [3] Phillips PJ, Yates AN, Hu Y, Hahn CA, Noyes E, et al. Face Recognition Accuracy of Forensic Examiners, Superrecognizers, and Face Recognition Algorithms. *Proc Natl Acad Sci U S A. Proceedings of the Nat' l.* 2018;115:6171-6176.
- [4] <http://www.CL.cam.ac.uk/users/jgd1000/Doppelganger-photos.pdf>
- [5] Aiyar S. Aadhaar: A Biometric History of India's 12-Digit Revolution. New Delhi: Westland Publications. 2017.
- [6] Cover T, Thomas J. *Elements of Information Theory*. 2nd Ed. New York: Wiley Interscience; 2006.
- [7] Daugman J. Information Theory and the Iriscode. *IEEE Trans. Inf. Forensics Secur.* 2015;11:400-409.
- [8] Daugman J, Downing C. Radial Correlations in Iris Patterns, and Mutual Information Within IrisCodes. *IET Biom.* 2019;8:185-189.
- [9] Daugman J. The Importance of Being Random: Statistical Principles of Iris Recognition. *Pattern Recognit.* 2003;36:279-291.
- [10] <http://www.CL.cam.ac.uk/users/jgd1000/IrisCumulatives.pdf>
- [11] <https://www.proceedings.com/04260.html>
- [12] Grother PJ, Tabassi E, Quinn GW, Salamon WJ. "IREX-I: Performance of Iris Recognition Algorithms on Standard Images." NIST Interagency Report 7629, 2009. URL: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=903606

- [13] Grother P, Quinn GW, Matey JR, Ngan M, Salamon W, et al. 'Irex-III: Performance of Iris Identification Algorithms.' NIST interagency report. 2012:7836.
- [14] Daugman J, Downing C, Akande ON, Abikoye OC. Ethnicity and Biometric Uniqueness: Iris Pattern Individuality in a West African Database. *IEEE Trans Biom Behav Identity Sci IEEE*, translator. 2023;6:79-86.
- [15] Snell R, Lemp M. *Clinical Anatomy of the Eye*. 2nd Ed. London: Blackwell Publishing Science. 1998.
- [16] Gong S, Boddeti VN, Jain AK. On the Capacity of Face Representation. 2017. ArXiv Preprint: <https://arxiv.org/pdf/1709.10433.pdf>