Black-box Collision Attacks on Widely Deployed Perceptual Hash Functions

Diane Leblanc-Albarel[1][0000−0001−5979−8457] and Bart Preneel[2][0000−0003−2005−9651]

[1] KU Leuven, Belgium, diane.leblanc-albarel@kuleuven.be
[2] KU Leuven, Belgium, bart.preneel@esat.kuleuven.be

Abstract. Perceptual hash functions identify multimedia content by mapping similar inputs to similar outputs. They are widely used for detecting copyright violations and illegal content but lack transparency, as their design details are typically kept secret. Governments are considering extending the application of these functions to Client-Side Scanning (CSS) for end-to-end encrypted services: multimedia content would be verified against known illegal content before applying encryption. In 2021, Apple presented a detailed proposal for CSS based on the NeuralHash perceptual hash function. After strong criticism pointing out privacy and security concerns, Apple has withdrawn the proposal, but the NeuralHash software is still present on Apple devices. Brute force collisions for NeuralHash (with a 96-bit result) require 2^{48} evaluations. Shortly after the publication of NeuralHash, it was demonstrated that it is easy to craft two colliding inputs for NeuralHash that are perceptually dissimilar. In the context of CSS, this means that it is easy to falsely incriminate someone by sending an innocent picture with the same hash value as illegal content. This work shows a more serious weakness: when inputs are restricted to a set of human faces, random collisions are highly likely to occur in input sets of size 2^{16}. Unlike the targeted attack, our attacks are black-box attacks: they do not require knowledge of the design of the perceptual hash functions. In addition, we show that the false negative rate is high as well. We demonstrate the generality of our approach by applying a similar attack to PhotoDNA, a widely deployed perceptual hash function proposed by Microsoft with a hash result of 1152 bits. Here we show that specific small input sets result in near-collisions, with similar impact. These results imply that the current designs of perceptual hash function are completely unsuitable for large-scale client scanning, as they would result in an unacceptably high false positive rate. This work underscores the need to reassess the security and feasibility of perceptual hash functions, particularly for large-scale applications where privacy risks and false positives have serious consequences.

Keywords: Perceptual Hashing, Collisions, Client-Side Scanning, NeuralHash, PhotoDNA, CSAM detection

1 Introduction

The goal of perceptual hash functions is to identify perceptually similar multimedia content such as images, videos, or sounds. Unlike cryptographic hash functions, that produce completely different outputs for even slight changes in input, perceptual hash functions should ensure that *visually* or *audibly* similar inputs produce close or identical hash values.

The primary applications of perceptual hash functions include the identification of copyright violations [28, 52], the detection of problematic content [33], such as Child Sexual Abuse Material (CSAM) [2] or terrorist content [15], and biometric authentication [32]. Comparing hash values is more efficient and allows for comparison without access to the illegal content itself.

The US National Center for Missing and Exploited Children (NCMEC) has reported a significant increase in CSAM detections in recent years [45]. Therefore NCMEC, along with researchers such as Bursztein et al. [11] and Farid [25] have urged the development of automated detection techniques, such as perceptual hash functions, to combat the proliferation of CSAM. In addition, centralized detection methods that were used so far have been hampered by the increased deployment of end-to-end encryption.

In response, the UK [58] and the European Union [23] have launched regulatory proposals to mandate Client-Side Scanning (CSS): CSS performs local detection of illegal content. By scanning content directly on user devices before it is encrypted, users can be reported and authorities can, in principle, take actions before the content is shared [2, 38, 40].

However, CSS has faced heavy criticism from a broad range of groups, including academics, industry, and NGOs (see Abelson et al. [1] for an overview). Experts have pointed out that the mechanism is a highly intrusive tool for mass surveillance that undermines the fundamental right to privacy and creates a chilling effect. Moreover, it opens the door to abuse, as it may be impossible to verify which content is being detected: once such a system is deployed for CSAM detection, it could be easily extended to other criminal content, to identify whistleblowers or journalistic sources, and to detect critics of oppressive regimes. Although some researchers have proposed messaging solutions that would allow the detection of CSAM while preserving user privacy [8], the technologies currently deployed do not offer strong privacy guarantees. Furthermore, it is unclear whether these systems are effective due to the risk of false positives and the ease with which the mechanism can be circumvented. As governments continued to push for legislation, members of the EU Parliament [22] and experts have issued open letters and public statements warning the public about these ethical and privacy concerns [47, 49].

This paper addresses the need for an evaluation of perceptual hash functions for the detection of illegal content. One of the key criticisms is that perceptual hash functions are ineffective as it is easy to create false positives and false negatives; while designing a secure cryptographic hash function is known to be difficult, designing a secure perceptual hash function is much harder. The literature contains many designs, but most of them have serious weaknesses. Most

deployed designs are not published and rely instead on security by obscurity: by violating Kerckhoffs' principle, this approach impedes public evaluation and makes it impossible to create trust.

This work primarily focuses on a public construction for image hashing developed by Apple to detect CSAM, called NeuralHash [2]. In particular, it focuses on false positives, that is two images that are perceptually different but with identical hash values. While earlier work [53] has pointed out that it is easy to deliberately create false positives by manipulating images, this work considers a setting in which NeuralHash would be deployed on a large scale and users who share legitimate images would be falsely accused of sharing CSAM. In its Q&A report [3] on NeuralHash, Apple addresses the concern related to false positives, stating:

> *"Will CSAM detection in iCloud Photos falsely report innocent people to law enforcement? No. The system is designed to be very accurate, and the likelihood that the system would incorrectly identify any given account is less than one in one trillion per year."*

Yet, to our knowledge, no rigorous evaluation of NeuralHash has been conducted to confirm this statement.

In this paper, we analyze the efficiency and accuracy of NeuralHash when used on a large scale. We aim to provide insights into its effectiveness in real-world applications and to highlight potential risks associated with its mass deployment. Our main conclusion is that NeuralHash would lead to a huge number of false positives even if users do not use manipulated images. As part of our research, we have conducted responsible disclosure with Apple. They reproduced our findings, acknowledged the issue, and are investigating solutions.

Throughout this paper, we use NeuralHash as a representative example in our theoretical discussions to ensure clarity and consistency in analyzing hash properties. However, all theoretical results, including our propositions and proofs, apply to *any* perceptual hash function, not just NeuralHash.

To illustrate that our findings on NeuralHash extend to other perceptual hash functions, we extend our empirical analysis to Microsoft's PhotoDNA, one of the most widely used perceptual hash functions in CSAM detection. We replicate our experiments on PhotoDNA to confirm whether the same types of vulnerabilities—high false positive rates, false negatives, and hash collisions—are also present. Our results show that these security flaws are not unique to NeuralHash but are likely inherent to the design of perceptual hashing itself.

Section 2 provides background information on perceptual hashing, the NeuralHash and PhotoDNA algorithms, and defines what constitutes perceptually identical content. Section 3 analyzes the properties of NeuralHash and in particular the collision properties for different types of images. Quantitative results and examples of collisions for each type of image are shown in Section 4. Section 5 presents collision results for PhotoDNA, demonstrating the generality of our approach. Section 6 discusses the impact of NeuralHash and PhotoDNA collisions on the false positive rate when these functions are used at a large scale.

Section 7 examines the impact on large-scale CSAM detection using current databases, and Section 8 presents the conclusions.

Disclaimer. Child sexual abuse and exploitation are serious crimes that must be addressed in our digital society. We firmly condemn the creation and distribution of CSAM. This paper aims to inform the discussion regarding the efficacy and implications of using perceptual hashing and Client-Side Scanning on a large scale. Our analysis of NeuralHash and PhotoDNA is not intended as a critique of Apple Inc., Microsoft, or any initiatives aimed at mitigating the spread of CSAM; rather, it highlights the risks associated with their large-scale deployment. NeuralHash and PhotoDNA serve as a case study that underscores the critical need for accuracy and robustness in these systems to prevent false positives. We hope this work will encourage further investigation and dialogue aimed at combating CSAM while ensuring a fair balance between efficiency and user privacy.

2 Background

This section defines the notion of *perceptually identical* content and discusses the properties of perceptual hash functions, and their typical building blocks. Next it reviews the most important perceptual hash function designs and presents NeuralHash.

2.1 Perceptually Identical Content

While content can consist of images, video, audio, 3D objects or 3D-immersive environments, this paper limits itself to the first category which is the target of NeuralHash. Perceptually identical images are those that appear identical or nearly identical for a *human observer*. From this human perception, similarity between images can be characterized by the following factors:

- **Color Consistency:** Images with minor differences in brightness, contrast, or color balance but identical in overall color composition and layout are considered perceptually the same. For example, an image with slightly adjusted brightness levels remains perceptually identical to its original.
- **Structural Similarity:** Images that share the same shapes, edges, and textures, even if subjected to minor geometric transformations such as rotations, translations, or small distortions, are considered perceptually similar. For instance, an image and its slightly rotated version would be perceptually identical.
- **Content Preservation:** Images that maintain the same core content but differ in resolution or compression are also perceptually similar. For example, a high-resolution image and a compressed version with some loss of detail are perceived as the same image.

– **Noise Robustness:** Images with minor noise additions, such as random pixel variations or blurring, which do not significantly alter the perceived content, are considered perceptually the same. For instance, an original image and one with a limited amount of Gaussian noise are perceptually identical.

Mathematically, some of the perceptual similarity elements mentioned above are typically formalized using metrics that attempt to quantify human visual perception. Two widely used metrics are the *Structural Similarity Index Measure*, that measure the similarity between two images based on luminance, contrast, and structure [10, 62] and the *Peak Signal-to-Noise Ratio*, that measures the ratio between the maximum possible pixel value of the image and the number of corrupting pixel [13].

Even if widely used, these metrics cannot detect every case of perceptually similar images. Hence in this paper the perceptually similarity between images is evaluated taking into account the four properties identified above. Section 3 provides the exact definition used to classify images as similar.

2.2 Properties

While cryptographic hash functions and perceptual hash functions share some similarities, they have different purposes and thus different properties. Both hash functions efficiently process large inputs and reduce these to short outputs in a deterministic way.

The properties [21, 25] of perceptual hash functions in the context of image hashing are:

– **Pre-image resistance:** Similar to cryptographic hash functions, perceptual hash functions should be one-way, meaning it should be computationally infeasible to reconstruct an original input x from its hash value $H(x)$.
– **Second pre-image resistance:** Given x_1 and its hash value $H(x_1)$, it should be computationally infeasible to find an input x_2, perceptually different from x_1, with $H(x_1) = H(x_2)$.
– **Illegitimate-Collision Resistance:** Perceptual hash functions should ensure that perceptually different inputs produce different or sufficiently different hash values. More formally, it should be computationally infeasible to find two perceptually different inputs x_1 and x_2 for which $H(x_1) = H(x_2)$.
– **Accuracy:** Perceptual hash functions should produce the same (or similar) hash values for perceptually identical or very similar inputs. This property ensures that minor variations in the input (such as slight changes in brightness or minor cropping in an image) do not result in significantly different hash values.

2.3 Perceptual Hashing Process

A first generation perceptual hash functions does not involve deep learning techniques. Examples include pHash [66] and Microsoft's PhotoDNA [52]. These

designs typically rely on hand-crafted features and transformations to generate hash values that are robust to minor changes in input.

Deep perceptual hash functions [41, 42], on the other hand, are based on a Machine Learning (ML) model. They involve training deep neural networks to learn feature representations that capture the perceptual similarity of inputs. Deep perceptual hash functions are less used than non-deep ones as they are more recent. Research on deep perceptual hashing includes hashing for image retrieval [67], for label prediction [63], and for CSAM detection [2].

Techniques such as Locality-Sensitive Hashing (LSH) and Binary Reconstructive Embeddings (BRE) are also used in various perceptual hashing functions. LSH maps similar input items to the same hash bucket with high probability [26, 31], while BRE creates compact binary values for similar inputs [37].

Perceptual hash functions typically follow similar sequences of steps to generate hash values [21].

The first step is preprocessing, which prepares the input image by normalizing it into a standardized format. This often involves resizing the image to fixed dimensions, such as (360×360) pixels, and normalizing pixel values to a specific range, for example, $[-1, 1]$. Additional preprocessing techniques may include computing image gradients or converting the image to grayscale.

Deep perceptual hash functions extract features from the preprocessed image using an ML model, This model is typically trained using techniques such as contrastive learning [41], that helps the model differentiating between perceptually similar and dissimilar images.

The next step is hash value extraction, where specific features of the image are extracted to form the hash result. The features selected depend on the hash function application. Common methods for hash extraction include computing the average color of the image, its gradient, or using BRE or LSH.

Finally, the generated hash values are converted to binary strings and possibly compared to another hash value using a hash comparison technique. The Euclidean distance [21, 25, 52] is the most common distance metric for comparisons.

2.4 Perceptual Hash Functions Designs

Perceptual hash functions have been developed and deployed in various contexts since 2000 [21]. We mention below some of the most notable designs and their applications.

One of the earliest implementations of perceptual hashing for content identification was YouTube's Content ID [28]. Content ID identifies copyrighted material to assist copyright holders in managing their rights.

The 2010 thesis of Zauner [36, 66] provides a detailed introduction to perceptual hash function and introduces the open-source construction pHash, that inspired several other designs.

PhotoDNA [14] was developed by Microsoft and Farid in 2009. It is extensively used for content moderation [52] and on platforms such as Gmail, Twitter, Facebook, Reddit, and Discord for detecting illegal content, particularly CSAM.

Despite its widespread use, some successful attacks have been reported (see Section 2.7).

eGlyph [15], based on PhotoDNA, was implemented by the Counter Extremism Project (CEP), a nonprofit international policy organization combating extremist ideologies, with the help of Farid. In particular in 2018, eGlyph was used to detect extremist videos on YouTube.

In 2019, Facebook released PDQ and TMK+PDQF [24] as open-source perceptual hash functions based on pHash [17]. These functions are designed to enhance the identification and moderation of prohibited content, in particular CSAM content, across Facebook's platform and beyond. PDQ function is designed for images while TMK+PDQF targets videos.

Apple introduced in 2021 NeuralHash [2], a perceptual hash function specifically designed for CSAM detection. NeuralHash uses a convolutional neural network (CNN) to generate hash values from images, aiming to detect illegal content.

2.5 NeuralHash

This section presents NeuralHash: it explains why it was created, in which context it is supposed to work and describes briefly the algorithm.

Context and Deployment. In August 2021, Apple announced NeuralHash as a key component of its new CSAM detection system [2, 9]. This system was designed to identify known CSAM images stored in iCloud Photos by comparing on-device image hashes to a database of known CSAM hashes provided by child safety organizations.

The deployment of NeuralHash involves integrating the hashing algorithm directly into the iOS operating system. When an image is uploaded to iCloud Photos, NeuralHash generates a hash value representing the image. This hash value is then compared to a database of known CSAM hash values. If a match is found, the image is flagged and, if confirmed, the user is reported to the authorities.

Apple's approach to implementing NeuralHash raised significant public debate and controversy [48]. As a consequence, Apple has officially withdrawn the proposal and postponed the large-scale deployment of NeuralHash; in spite of this, NeuralHash has been added to all Apple devices.

Client-Side Scanning. NeuralHash was intended to work conjointly with CSS. We briefly describe the process presented in Figure 1.

First, the user device generates a long term secret key S, that will be used for threshold secret sharing as explained below. For an image I, the device computes the NeuralHash value h_I and derives K_I and k_I from h_I. Using the key S and the hash result h_I, the device computes a secret sharing key S_I.

The client device then sends to the server the value K_I, the encrypted secret sharing key $\{S_I\}_{k_I}$, and the doubly encrypted image $\{\{I\}_S\}_{k_I}$.

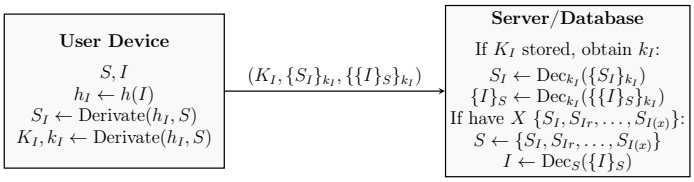

Fig. 1. Process designed by Apple to decrypt CSAM images with CSS

The server stores a database of pairs (K_I, k_I) for every problematic image. When receiving the values $(K_I, \{S_I\}_{k_I}, \{\{I\}_S\}_{k_I})$, the server checks whether K_I is in its database; if so, it gets the corresponding k_I. With k_I, the server decrypts $\{S_I\}_{k_I}$ to obtain S_I; it also decrypts the first layer of $\{\{I\}_S\}_{k_I}$ and thus retrieves $\{I\}_S$.

If the server collects at least X shares S_I, it can reconstruct the key S. Once S is known, the server can decrypt any image having a hash value present in its database by using S to decrypt $\{I\}_S$.

The threshold X is defined by Apple, but its numerical value is not disclosed. Apple claims [2] that X is chosen to ensure *"an extremely low probability (1 in 1 trillion) of incorrectly flagging a given account."*

In its technical report, Apple refers to the value we have called K_I as a *"cryptographic header"* derived from h_I. Apple does not provide any information on the size of K_I nor how it is obtained. It also does not specify in its report that the property $K_I \neq K_J$, $\forall\ I \neq J$ is guaranteed. If this property is not guaranteed, then a legitimate image that does not have a NeuralHash value flagged as problematic content can still generate a K_I known by the server and thus be flagged, even if neither the image nor its hash value is flagged. Similarly, a legitimate image with a hash value h_I equal or close to one in the database will lead to the use of a flagged pair (K_I, S_I) and thus result in the flagging of a legitimate image.

NeuralHash Computation. The NeuralHash algorithm, as deployed on user devices, consists of two primary components: a CNN and a LSH step. The main steps of the hash computation are described bellow.

The algorithm begins by the preprocessing stage that resizes the image to dimensions $(360 \times 360 \times 3)$ within a normalized pixel range of $[-1, 1]$. The resized image is then sent into the embedding CNN.

The CNN produces a vector z of 128 bits. The goal is that for perceptually similar images the vectors z are close and that for perceptually different images they lie far apart.

The LSH step multiplies the vector z of length 128 by a (128×96) matrix B to obtain a real vector y of length 96: $y = B \cdot z$. This step checks the position of the vector z relative to each hyperplane of the matrix B. Each vector is mapped

to a specific bucket, with similar vectors (and thus similar images) placed in the same or adjacent buckets.

The final output is a bitstring of length 96. If $y_i > 0$, the corresponding bit is set to 1, otherwise, it is set to 0.

For our experiments, we extracted the neural network as well as the hashing matrix used by NeuralHash from an iPhone with firmware version 16.2. We then used [65] to convert the NeuralHash model into ONNX format, allowing us to run NeuralHash in any script or device.

2.6 Microsoft PhotoDNA

Unlike NeuralHash, PhotoDNA relies solely on traditional image processing operations. While Microsoft has never publicly disclosed the exact algorithm of PhotoDNA, several leaks have allowed researchers to analyze and reconstruct equivalent implementations. Publicly available code and research enable the retrieval of corresponding values for known PhotoDNA input-output pairs. In this paper, we utilized the implementation available on GitHub [34]. This code has also been referenced in [5, 51, 52].

The computation of PhotoDNA follows these steps:

1. **Normalization:** The input image is first converted to grayscale and resized to a fixed resolution of 26x26 pixels, where each pixel is represented by a byte indicating the gray value.
2. **Overlapping Grid Segmentation:** To mitigate sensitivity to minor cropping, the 26x26 pixel image is divided into overlapping 6x6 grids with a stride of 4 pixels, resulting in 36 overlapping regions. This overlapping strategy is designed to improve robustness by ensuring that slight shifts in image alignment do not significantly alter the extracted features.
3. **Gradient Computation:** Within each 6x6 grid, gradients are computed based on pixel intensity differences. The gradient values capture directional intensity changes and are computed as follows:
   - The sum of all increasing horizontal pixel values.
   - The sum of all decreasing horizontal pixel values.
   - The sum of all increasing vertical pixel values.
   - The sum of all decreasing vertical pixel values.

   The result is four gradient values per grid, leading to a total of $36 \times 4 = 144$ values.
4. **Scaling and Equalization:** To ensure uniformity in hash representation, the computed gradients undergo a normalization process where extreme values are adjusted to fit within the range of 0 to 255. This step accounts for variations in scaling, blurring, and lighting conditions.

Note that the hash result composed of 144 values in the range 0 to 255 can also be represented as a classical hash value of 1152 bits. Moreover, with such a large hash value, the focus is not on exact collisions but rather on near-collisions [39] i.e., hashes that are close but not identical (see Section 5).

2.7 Related Work

Perceptual hash functions have been studied in academic papers for over two decades. A comprehensive introduction to the topic was provided by Farid in 2021 [25], detailing the fundamental techniques and challenges. Additionally, a survey of perceptual hashing techniques and their applications can be found in [21].

Recent research in perceptual hashing has primarily focused on the creation of collisions and on information leakage. The creation of collisions involves modifying one of two perceptually different images to produce the same hash value, effectively creating false positives. Information leakage, on the other hand, refers to the potential of inferring information about the input from its hash value.

One prevalent method of attacking perceptual hash functions is through gradient-based hash attacks. These attacks involve imperceptibly altering pixels to achieve specific outputs [12, 16, 19, 54, 55]. Various optimization techniques have been proposed to enhance the effectiveness of these attacks [27, 44, 50]. These attacks generally follow the same principles as classic image processing attacks, manipulating inputs to achieve specific outputs [7, 59, 60, 64, 68].

Some papers have analyzed the resistance of perceptual hash functions to image modifications and the potential to recover original images from their hash values. This includes research on the robustness of perceptual hash functions derived from pHash [20, 33] and attempts to reconstruct original images using Binary Reconstructive Embeddings (BRE) instead of Locality-Sensitive Hashing (LSH) [61].

Recent work has revealed a series of attacks that exploit vulnerabilities in the internal structure of NeuralHash and other perceptual hash functions. Notably, Struppek et al. [53] showed how to create second pre-images for NeuralHash, where input images were imperceptibly modified and published a proof-of-concept implementation (other tools for second pre-images can be found in [6, 35]). Struppek et al. also described classification attacks, in which hash values are used to categorize inputs, achieving a maximum success rate of 52% in some categories. Similarly, attacks on pHash have shown how images can be manipulated to produce specific hash values [29]. Additionally, partial inversion of PhotoDNA has been achieved using neural networks [5].

Despite this research on attacking perceptual hash functions and NeuralHash in particular, there remains a significant gap in the evaluation of these functions concerning real false positive rates and their performance when approached as black boxes. This gap underscores the necessity of our analysis, which aims to evaluate the performance, accuracy, and false positive rate of NeuralHash in the context of large-scale CSAM detection.

3 Analysis of Perceptual Hash Functions

This section presents the dataset of images used in the paper and analyses the NeuralHash output for several types of images.

3.1 Image Sets

We use two different image datasets. The first dataset is made from non-human images extracted from the PASS dataset [4], while the second dataset consists of 202 599 real celebrity face images from the CelebA dataset [43].

Our primary goal is to evaluate whether hashing face images with perceptual functions yields different results compared to non-human images. Specifically, we aim to analyze the number of collisions and the statistical properties of the individual bits of the hash values.

We selected face images for two main reasons: first, to approximate a use case related to CSAM detection, as such images often contain faces; second, because face images are common on mobile devices, making this a realistic use case scenario. In addition we modify these images with various levels of blurring, as it is plausible that in the context of CSAM production it is likely that blurring is used.

Image Types. The following six image types are used: non-human images, human face images, and human face images with varying levels of blurring. An example of each type of image is provided in Figure 2. Each set contains 202 599 images. Throughout the paper, the abbreviation *BF* stands for *blurred faces*. The six types of images are referred to as follows:

- *Non-human*: randomly selected images from the PASS [4] dataset.
- *Non BF*: images from the CelebA dataset [43] without any transformations.
- *Light BF*: CelebA images with light blur applied to the entire image.
- *Medium BF*: CelebA images with medium blur applied to the entire image.
- *High BF*: CelebA images with high blur applied to the entire image.
- *BF only*: CelebA images with blur applied only to the face region.

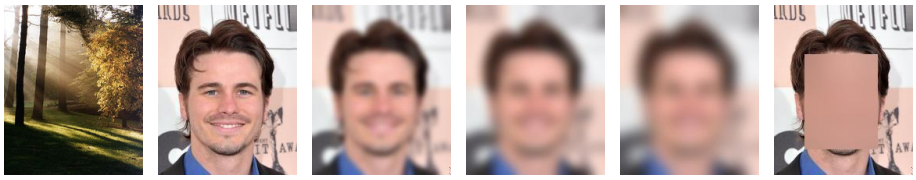

Fig. 2. Example of the 6 types of images in the experiments

Perceptually Identical Images. Out of 202 599 face images in the CelebA dataset, we found at least 4 425 images with duplicates or perceptually similar counterparts—images that should share the same hash. Figure 3 shows examples of identical and perceptually similar images.



(a) Identical images (b) Perceptually identical images (different filter) (c) Perceptually identical images (same person, very close pose)

Fig. 3. Perceptually identical images

3.2 Statistical Properties of the Hash Values

This section presents both a theoretical analysis and experimental results to assess the properties of perceptual hash functions. We analyze the statistical distribution of hash values and test the independence of hash bits using probabilistic models. In parallel, we conduct experiments on large-scale datasets to validate these theoretical findings, using NeuralHash as a primary example before extending our evaluation to PhotoDNA.

Distribution. We first experimentally analyse the distribution and in particular the uniformity and variance of NeuralHash bits values for different types of images. We then introduce three propositions that help to determine the theoretical number of hash function evaluations before obtaining an illegitimate collision.

Experiments. As stated in Section 2, the bit distribution of the output of a hash function must be uniform with a small variance to prevent information leakage about the input and to resist (2nd) pre-image and collision attacks. A perceptual hash function is also expected to have a uniform distribution to avoid leaking information about its input.

To compute the average distribution of Neuralhash bits for each image type, we hashed 30 000 randomly selected images from each set. For each image type, we count for each of the 96 bits the number of times the bit is equal to 1 and 0 respectively. The results for the Non-human type and the Non-blur face type are presented in Figure 4a, where the percentage of times each bit is equal to 1 is plotted. The expected result is that each bit is equal to 1 approximately 50% of the time, with a small variance.

For non-human images, the results meet expectations, with the percentage of bits equal to 1 uniformly distributed around 50%. In contrast, for non-blurred human face images, a significant number of bits deviate from being equal to 1, 50% of the time. For non-human images, none of the bits of the hash is less than 40% or more than 60% of the time equal to 1. In contrast, for human face images, 44 bits fall outside this range, with even 2 bits being less than 20% or more than 80% of the time equal to 1.

The results for the four remaining image types are presented in Figure 4b. The resulting hash values are even less uniformly distributed compared to the non-blurred face image type. For lightly blurred images (orange dots), 56 bits fall outside the $40\% - 60\%$ range. For images with only the face blurred (green dots), 59 bits fall outside the range. For medium blurred images (black dots), 67 bits fall outside the range, and for highly blurred images, 70 bits out of the 96 final bits fall outside the $40\% - 60\%$ range.

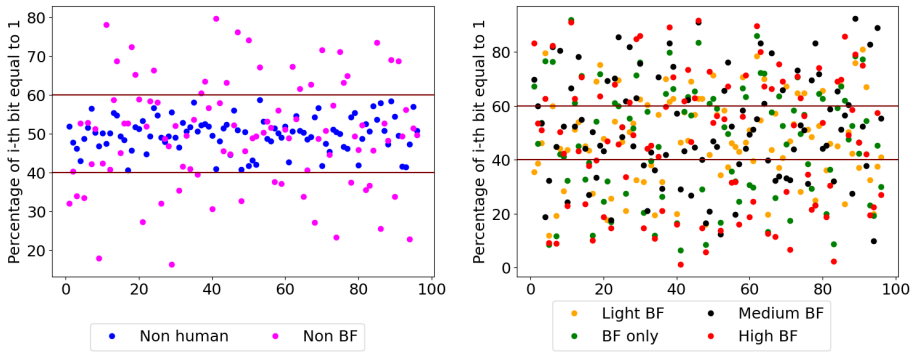

(a) Non-human vs., Non BF (b) Blurred human face

Fig. 4. Bit distribution for each image type

Additionally, Figure 10 in Appendix A shows a boxplot of the distribution of bit values for each image type: except for the non-human images, the distribution deviates strongly from the expected value.

Theoretical Collision Probability with Independence Hypothesis. An illegitimate collision corresponds to the event where the hash value of two perceptually different images are the same. For regular hash functions, this event should be exceedingly rare. For an uniformly distributed 96-bit hash value (NeuralHash case), where each bit is independent of the others, the birthday paradox states that the average number of hash values to compute before encountering an illegitimate collision is equal to 2^{48}.

Assuming that all bits of the hash value are mutually independent, and considering the observed distribution for each dataset, Proposition 1 provides the probability that two hash values are identical. Proposition 2 gives the expected number of hash values required before encountering the first illegitimate collision. To make this result more practical, Proposition 3 provides a tight bound on the expected number of hash values necessary before observing the first illegitimate collision.

Proposition 1. *Consider a hash value of n bits and a vector p for which the i-th element p_i denotes the probability that the i-th bit of the hash value is equal*

to 1 *(*$0 \leqslant p_i \leqslant 1$*). If* E_p^n *denotes the event that two hash values are equal, then*

$$\mathbb{P}(E_p^n) = \prod_{i=1}^{n} (p_i^2 + (1-p_i)^2)\,.$$

Proof. Let $H_{i,j}^n$ denote the first j bits of the i-th hash value of n bits, and let a_i^n, b_i^n be the first i bits of two given hash values of size n. E_p^n is thus the event that the two strings a_n^n and b_n^n are equal.

We have:

$$\mathbb{P}(E_p^n) = \sum_{h=H_{i,n}^n} \mathbb{P}(a_n^n = h)\mathbb{P}(b_n^n = h) = \sum_{h=H_{i,n}^n} \mathbb{P}(a_n^n = h)^2\,.$$

Considering the sum in detail:

$$\begin{aligned}\mathbb{P}(E_p^n) &= \sum_{h=H_{i,n}^n} \mathbb{P}(a_n^n = h \mid h(n) = 1)^2 + \sum_{h=H_{i,n}^n} \mathbb{P}(a_n^n = h \mid h(n) = 0)^2 \\ &= p_n^2 \sum_{h=H_{i,n-1}^n} \mathbb{P}(a_{n-1}^n = h)^2 + (1-p_n)^2 \sum_{h=H_{i,n-1}^n} \mathbb{P}(a_{n-1}^n = h)^2 \\ &= (p_n^2 + (1-p_n)^2) \sum_{h=H_{i,n-1}^n} \mathbb{P}(a_{n-1}^n = h)^2\,.\end{aligned}$$

The result follows by recurrence over the remaining $n-1$ bits. □

Proposition 2. *Denote with* $\mathbb{P}(E_p^n)$ *the probability that two* n*-bit hash values with distribution* p *are equal and define* $M = 2^n$*. The expected number of hash values to compute before the first collision* $\mathbb{E}(D_p^n)$ *is equal to:*

$$\mathbb{E}(D_p^n) = 2 + (M+1) \times y^{\frac{M(M+1)}{2}} + \sum_{x=1}^{M-1} y^{\frac{x(x+1)}{2}} \text{ with } y = 1 - \mathbb{P}(E_p^n)\,.$$

Proof. $\mathbb{P}(E_p^n)$ denotes the probability of a collision between two hash values, and D_p^n denotes the random variable representing the number of hash evaluations required before encountering the first collision. The probability that the first collision occurs after x hash function evaluations is denoted by $\mathbb{P}[D_p^n = x]$.

It is worth noting that $\mathbb{P}[D_p^n < 2] = 0$ and $\mathbb{P}[D_p^n > M+1] = 0$, as there are a maximum of $M = 2^n$ different hash values.

By definition:

$$\sum_{i=1}^{M+1} \mathbb{P}[D_p^n = i] = 1 \texttt{ and: } \mathbb{E}(D_p^n) = \sum_{x=2}^{M+1} x \times \mathbb{P}[D_p^n = x]\,.$$

We first define $\mathbb{P}[D_p^n = n]$. The probability $\mathbb{P}[D_p^n = 2]$ is the probability that a collision occurs after 2 hash evaluations, which is simply: $\mathbb{P}[D_p^n = 2] = \mathbb{P}(E_p^n)$. The probability $\mathbb{P}[D_p^n = 3]$ is the probability that a collision occurs after exactly

3 hash evaluations. This is the probability that the first two hash values are different, multiplied by the probability that the third hash value matches one of the first two. Therefore:

$$\mathbb{P}[D_p^n = 3] = (1 - \mathbb{P}(E_p^n)) \times (1 - (1 - \mathbb{P}(E_p^n))^2)\,.$$

Similarly, the probability $\mathbb{P}[D_p^n = 4]$ equals:

$$\mathbb{P}[D_p^n = 4] = (1 - \mathbb{P}(E_p^n))^3 \times (1 - (1 - \mathbb{P}(E_p^n))^3)\,.$$

This approach generalizes to obtain $\mathbb{P}[D_p^n = n]$, the probability that a collision occurs after exactly n hashes with $(3 \leq n \leq M + 1)$:

$$\mathbb{P}[D_p^n = n] = \left(\prod_{i=1}^{n-2}(1 - \mathbb{P}(E_p^n))^i\right) \times (1 - (1 - \mathbb{P}(E_p^n))^{n-1})\,,$$

which simplifies to:

$$\mathbb{P}[D_p^n = n] = (1 - \mathbb{P}(E_p^n))^{\sum_{i=1}^{n-2} i} \times (1 - (1 - \mathbb{P}(E_p^n))^{n-1})\,.$$

Expanding the sums, we obtain:

$$\mathbb{P}[D_p^n = n] = (1 - \mathbb{P}(E_p^n))^{\frac{(n-2)(n-1)}{2}} - (1 - \mathbb{P}(E_p^n))^{\frac{(n-1)n}{2}}\,.$$

The expected number of hash values to compute before the first collision is given by $\mathbb{E}(D_p^n)$, which is expressed as:

$$\mathbb{E}(D_p^n) = \sum_{x=2}^{N+1} x \times \mathbb{P}[D_p^n = x]\,,$$

which is equivalent to:

$$\mathbb{E}(D_p^n) = 2 \times \mathbb{P}(E_p^n) + \sum_{x=3}^{N+1} x \times \mathbb{P}[D_p^n = x]\,.$$

Substituting the expression for $\mathbb{P}[D_p^n = x]$ and expanding the sum, we obtain:

$$\begin{aligned}\mathbb{E}(D_p^n) = 2 \times \mathbb{P}(E_p^n) &+ \left(3 \times (1 - \mathbb{P}(E_p^n)) - 3 \times (1 - \mathbb{P}(E_p^n))^3\right) \\ &+ \left(4 \times (1 - \mathbb{P}(E_p^n))^3 - 4 \times (1 - \mathbb{P}(E_p^n))^6\right) \\ &+ \left(5 \times (1 - \mathbb{P}(E_p^n))^6 - 5 \times (1 - \mathbb{P}(E_p^n))^{10}\right) + \ldots \\ &+ \left((M+1) \times (1 - \mathbb{P}(E_p^n))^{\frac{(M-1)M}{2}} - (M+1) \times (1 - \mathbb{P}(E_p^n))^{\frac{M(M+1)}{2}}\right)\,.\end{aligned}$$

Reordering the terms, we then obtain:

$$\mathbb{E}(D_p^n) = 2 \times \mathbb{P}(E_p^n) + 2 \times (1 - \mathbb{P}(E_p^n)) + (1 - \mathbb{P}(E_p^n))$$

$$+(M+1)\times(1-\mathbb{P}(E_p^n))^{\frac{M(M+1)}{2}}$$
$$+\sum_{x=4}^{M+1}\left((x+1)\times(1-\mathbb{P}(E_p^n))^{\frac{(x-2)(x-1)}{2}}-x\times(1-\mathbb{P}(E_p^n))^{\frac{(x-2)(x-1)}{2}}\right).$$

Simplifying the terms of the sum and shifting the bounds by changing the variable x to $x-2$, we include the term $(1-\mathbb{P}(E_p^n))$ as the first term of the sum:

$$\mathbb{E}(D_p^n)=2+(M+1)\times(1-\mathbb{P}(E_p^n))^{\frac{M(M+1)}{2}}+\sum_{x=1}^{M-1}(1-\mathbb{P}(E_p^n))^{\frac{x(x+1)}{2}}.$$

With $y=1-\mathbb{P}(E_p^n)$, we conclude that

$$\mathbb{E}(D_p^n)=2+(M+1)\times y^{\frac{M(M+1)}{2}}+\sum_{x=1}^{M-1}y^{\frac{x(x+1)}{2}}.$$

□

The computation of the exact value of $\mathbb{E}(D_p^n)$ is infeasible given the sum over M elements (in the NeuralHash case $M=2^{96}$ and in case of PhotoDNA $M=2^{1152}$). Proposition 3 provides a bound for $\mathbb{E}(D_p^n)$ that is easy to compute.

Proposition 3. *Denote with $\mathbb{P}(E_p^n)$ the probability that two n-bit hash values with distribution p are equal and define $M=2^n$. The expected number of hash values to compute before the first collision $\mathbb{E}(D_p^n)$ is upper bounded by:*

$$\mathbb{E}(D_p^n)\leq 1+(M+1)\times y^{\frac{M(M+1)}{2}}+\frac{\theta_2(0;y^{\frac{1}{2}})}{2\times y^{\frac{1}{8}}},$$

with $\theta_2(0;y^{\frac{1}{2}})$ the Jacobi theta function $\theta_2(z;q)$ with $z=0$ and $q=y^{\frac{1}{2}}$ and with

$$y=1-\mathbb{P}(E_p^n).$$

Proof. From Proposition 2 we have:

$$\mathbb{E}(D_p^n)=2+(M+1)\times y^{\frac{M(M+1)}{2}}+\sum_{x=1}^{M-1}y^{\frac{x(x+1)}{2}}.$$

Note that $\sum_{n=1}^{M-1}y^{\frac{n(n+1)}{2}}=\sum_{n=0}^{M-1}y^{\frac{n(n+1)}{2}}-1$. As $\sum_{n=0}^{M-1}y^{\frac{n(n+1)}{2}}\leq\sum_{n=0}^{\infty}y^{\frac{n(n+1)}{2}}$ we thus have:

$$\sum_{n=1}^{M-1}y^{\frac{n(n+1)}{2}}\leq -1+\sum_{n=0}^{\infty}y^{\frac{n(n+1)}{2}}. \tag{1}$$

The Jacobi theta function $\theta_2(z;q)$ is defined as follows:

$$\theta_2(z;q)=2\times q^{\frac{1}{4}}\times\left(\sum_{n=0}^{\infty}q^{n(n+1)}\cos((2n+1)z)\right).$$

Using $z = 0$ and $q = y^{\frac{1}{2}}$ we obtain:

$$\theta_2(0; y^{\frac{1}{2}}) = 2 \times y^{\frac{1}{8}} \times \left(\sum_{n=0}^{\infty} y^{\frac{n(n+1)}{2}} \right) .$$

By introducing the Jacobi function in the right hand side of Equation (1), we conclude that $\mathbb{E}(D_p^n)$ satisfies

$$\mathbb{E}(D_p^n) \leq 1 + (M+1) \times y^{\frac{M(M+1)}{2}} + \frac{\theta_2(0; y^{\frac{1}{2}})}{2 \times y^{\frac{1}{8}}} .$$

□

Table 3.2 is derived from Proposition 3 and the distributions of NeuralHash presented in Figure 4. It provides the expected number of hash values to compute before reaching the first illegitimate collision for each image type.

Table 1. Expected number of hash operations before first illegitimate collision using NeuralHash

Type	Non-human	Non BF	Light BF	BF only	Medium BF	High BF
$\mathbb{E}(D_p^n)$	$2^{47.8}$	$2^{43.5}$	$2^{41.1}$	$2^{38.9}$	$2^{37.5}$	$2^{34.1}$

Conclusion on Distribution. From Table 3.2, we conclude that, except for the non-human images, the expected number of hash values to compute before the first NeuralHash illegitimate collision is significantly lower than the theoretical value 2^{48}. For other image types, the expected number of hash values to compute before an illegitimate collision varies between $2^{43.5}$ (non-blurred face) and $2^{34.1}$ (highly blurred face).

It is worth noting that these values assume each bit of the hash is independent of the others. If this independence does not hold, the expected number of hash values before reaching an illegitimate collision decreases substantially.

Independence of Bits. In the previous section, we assumed that the bits of the hash values are independent of each other. In this section, we test this independence hypothesis. If the bits are independent, the value of any bit j of a hash value should not provide any information about the value of bit i with $j \neq i$.

Simple Matching Coefficient. The simple matching coefficient (*SMC*) is a statistic ranging from 0 to 1 used to compare the similarity of symmetric binary sample sets. When two samples are identical, the *SMC* is equal to 1. When two samples have no common values (i.e., they are completely different), the *SMC* is equal to 0. The expected *SMC* of two independent samples is 0.5. It is worth

noting that while the *SMC* can be used to verify non-independence, it cannot be used to confirm independence.

For two bits from the same hash value h, Definition 1 provides the value $m_{i,j}(h)$, corresponding to the simple matching value of bits i and j denoted as $h(i)$ and $h(j)$ respectively.

Definition 1. *For a hash value h, $m_{i,j}(h)$ is defined as:*

$$m_{i,j}(h) = \begin{cases} 1 \text{ if } h(i) = h(j) \\ 0 \text{ if } h(i) \neq h(j)\,. \end{cases}$$

Using Definition 1, Definition 2 provides the *SMC* for two bits i and j of a sample of N hash values of n bits.

Definition 2. *Consider a set D_p^n of ℓ ordered hash values of length n. Given h_k, the k-th element of D_p^n, the SMC of bits i and j is defined as:*

$$SMC_{i,j}(D_p^n) = \frac{1}{N} \cdot \sum_{k=1}^{\ell} m_{i,j}(h_k)\,.$$

By definition, $\forall i, j$, $SMC_{i,j}(D_p^n) = SMC_{j,i}(D_p^n)$ and $SMC_{i,j}(D_p^n) = 1$ when $i = j$.

Using Definition 2, the simple matching matrix (*SMM*) of a set D_p^n is built by computing the *SMC* for every pair of bits. We thus use the *SMM* to compute the *similarity* between every pair of bits of hash values from each of the six image types.

Bits Sample Similarity. For each image type, the *SMM* of the set is computed using Definition 2 and the hash values of 30 000 random images per set. A color visualization of the NeuralHash *SMM* for the Non-human and Non-blurred face sets is presented in Figure 11 in Appendix.

Table 2 summarizes, for each set, the values of the NeuralHash *SMM* that fall within different ranges, from $0.0-0.1$ (indicating that the pair of bits almost always have opposite values) to $0.9-1.0$ (indicating that the pair of bits almost always have equal values). For hash values with independent bits, 100% of these values should fall within the $0.4-0.6$ ranges.

Conclusion on Independence. Table 2 shows that for the Non-human type image, only 0.22% of the NeuralHash values fall outside the $0.4-0.6$ range. This corresponds to 20 pairs of bits over the 9120 possible pairs. These 20 values are all very close to either 0.4 or 0.6. These values are not sufficient to conclude that bits of non-human images are dependent.

However, for all other types of images, a significant number of NeuralHash values fall outside the $0.4-0.6$ range. For the non-blurred face type, 13.97% of the pairs are out of range. For the blurred face types, the percentages of values

Table 2. Similarity score, in percentage, for each image type

Type \ SMC	$0-0.1$	$0.1-0.2$	$0.2-0.3$	$0.3-0.4$	$0.4-0.5$	$0.5-0.6$	$0.6-0.7$	$0.7-0.8$	$0.8-0.9$	$0.9-1$
Non-Human	0.0	0.0	0.0	0.18	48.6	51.18	0.04	0.0	0.0	0.0
Non BF	0.0	0.0	0.29	6.97	43.05	42.98	6.56	0.13	0.02	0.0
Light BF	0.0	0.02	1.1	10.72	38.62	38.07	10.48	0.96	0.02	0.0
BF Only	0.0	0.2	2.7	11.64	33.93	35.68	12.7	2.89	0.26	0.0
Medium BF	0.0	0.33	3.53	15.7	31.47	29.69	14.76	3.88	0.64	0.0
High BF	0.2	1.69	6.67	15.75	26.45	29.99	13.64	7.61	2.08	0.11

outside the range goes from 23.3% for the lightly blurred face type to 47.75% for the highly blurred face type. For the blurred face only type, 30.39% of the values are outside the range.

These results indicate that, except for the non-human images, the bits of the NeuralHash value are not independent. Consequently, Table 3.2 substantially overestimates the expected number of hash values required to obtain an illegitimate collision.

4 NeuralHash Collision Search

The results from Section 3 indicate that, for all types of images except the non-human ones, the expected number of NeuralHash evaluations before reaching a collision is significantly lower than the theoretical value of 2^{48}.

Therefore, we hashed all 202 599 images of each set, searching for illegitimate collisions between hash values. For each identified collision, we classified it as *legitimate* when the colliding hash values correspond to images identical or perceptually similar, and *illegitimate* when the colliding hash values correspond to images that are perceptually different. In the few cases of colliding hash between images that cannot clearly be identified as perceptually similar or not, we classified the collision as legitimate. Next we determined the number of illegitimate collisions (false positives) and illegitimate non-collisions (false negatives).

4.1 False Positives

The number of illegitimate collisions for each image type and the number of hash evaluations before the first collision are presented in Table 3. Since some collisions involve more than two images (multiple images sharing the same hash), we count the number of illegitimate collisions as the number of images sharing their hash values with at least one perceptually different image.

Table 3 shows that out of the 202 599 hash values computed for each image type, the non-human type is the only one without any illegitimate collisions. All other image types present several illegitimate collisions. The number of these collisions, given the number of image hashed, is much higher than expected. This number of false positive for sets of only 202 599 images implies that the illegitimate collision rate will be much higher when millions or billions of images will

be hashed. Section 6 presents estimates for the number of illegitimate collisions for large-scale use.

Figure 5 presents an example of a collision obtained for each image type. It is worth noting that for all blurred face images (light, medium, and high), there are hash values shared by three or more perceptually different images. An example of three images sharing a hash value is given for the light blurred type. For the light blurred face set, three images share the same hash. For the medium blurred face set, eight different hash values are shared by three different images each, and four hash values are shared by four different images each. For the highly blurred face set, instances where three images share the same hash value are very frequent, and two hash values are shared by five different images each.

Table 3. False positives for each image type

Type	Non Human	Non BF	Light BF	BF only	Medium BF	High BF
# illegitimate collisions	0	24	25	12	260	588
# hash before collision	–	$2^{16.1}$	$2^{15.2}$	$2^{16.1}$	$2^{12.6}$	$2^{11.6}$

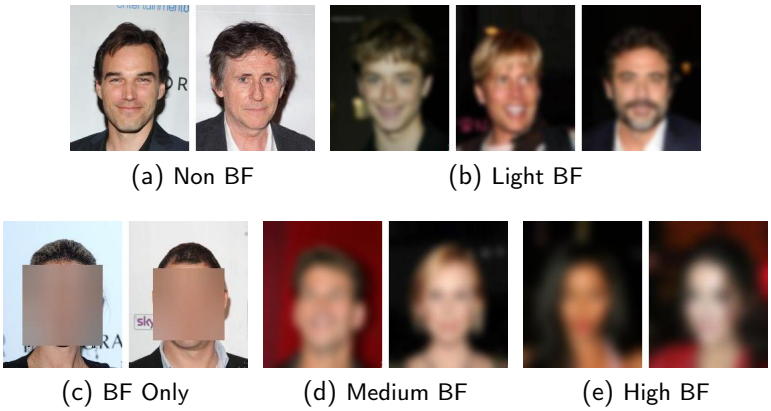

Fig. 5. Collision examples for each image type

4.2 False Negatives

As mentioned in Section 3.1, at least 4425 images out of the 202, 599 should result in legitimate collisions. The minimum false negative rate can thus be computed as follows:

Given m legitimate collisions, the false negative rate in % is at least

$$100 \cdot (4425 - m)/4425.$$

The total number of legitimate collisions and the corresponding false negative rate is provided in Table 4.

Table 4. Number of false negatives

Type	Non BF	Light BF	BF only	Medium BF	High BF
# legitimate collisions	1559	1513	1005	1864	2333
False negative rate (%)	64.7	65.9	77.3	57.9	47.3

It is worth noting that the 4425 images expected to legitimately collide is a lower bound. For the medium and highly blurred face sets, as the blur degrades the images, some images that are not perceptually the same before blurring become perceptually the same after blurring, thus increasing the theoretical number of legitimate collisions. Despite this higher number of theoretical legitimate collisions, a significant number of images are still not detected as colliding images, even though they should be.

For the non-blurred face, lightly blurred face, and blurred face only image types, a significant number of perceptually identical images do not have the same hash value. For the blurred face only set, this rate reaches 77.3%. These results indicate that the function does not properly detect perceptually identical images. This ratio should be considered alongside the false positive rate. For the blurred face only set, for instance, the number of illegitimate collisions is lower than in other sets, but the number of legitimate collisions is also lower. This indicates that the function is particularly inaccurate when only faces are blurred.

5 PhotoDNA Collision Search

To support our hypothesis that the vulnerabilities identified in NeuralHash are inherent to the current design of perceptual hashing, we conducted similar black-box experiments on Microsoft's PhotoDNA. Similar to NeuralHash, PhotoDNA is widely used for content moderation. Particularly PhotoDNA is the function the most commonly used by NCMEC for detecting CSAM.

5.1 Experimental Setup

PhotoDNA (described in Section 2.6) produces a hash vector of 144 bytes (1152 bits), which implies that perceptually similar images do not have the same hash values, but hash values that are close to each other (near-collisions [39]) according to some distance measure. The common distance measure are L1 and L2.

Definition 3. *Given two hash vectors $x, y \in \mathbb{R}^n$, the L1 distance is defined as:*

$$L_1(x, y) = \sum_{i=1}^{n} |x_i - y_i| ,$$

Definition 4. *Given two hash vectors $x, y \in \mathbb{R}^n$, the L2 distance is defined as:*

$$L_2(x, y) = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2}.$$

To break PhotoDNA, one does not need to search for collisions but rather for perceptually different images with hash values that have a distance below a certain threshold. Microsoft has not disclosed the specific distance metric used for PhotoDNA matching, nor the precise threshold values they apply. However, prior research [52] has analyzed the function and recommended using the L2 distance metric with a threshold between 150 and 175. According to this study, these values prevent false positives while ensuring more than 99% true positive detection. Additionally, [51] employs the L1 distance metric and suggests that setting an L1 threshold of 1800 allows to avoid all false positives.

In many practical scenarios, assuming that the differences per coordinate are roughly of equal magnitude, the L1 and L2 distances are related [18] by the approximation given in Equation (2).

$$L_2(x, y) \approx \frac{L_1(x, y)}{\sqrt{n}}. \tag{2}$$

Converting the 1800 L1 distance given in [51] to an L2 approximation using Equation (2), with $n = 144$ (the size of the hash vector), yields an L2 threshold of approximately 150, which aligns with the findings of [52].

Based on these observations, we conducted experiments using the CelebA dataset. Specifically, we extracted PhotoDNA hash values for all non-blurred face (Non-BF) images using the implementation provided in [34]. We then computed the L2 distances between all pairs of images, counting both legitimate collisions (cases where perceptually identical images yield similar hashes) and illegitimate collisions (cases where perceptually distinct images yield similar hashes). This methodology mirrors the approach we employed for NeuralHash, enabling a direct comparison of results.

5.2 Collision Search

Figure 6 presents the number of legitimate and illegitimate collisions as a function of the L2 distance. Each bar represents the number of collisions (either legitimate or illegitimate) that occur with an L2 distance below a given threshold. For example, 745 perceptually similar images exhibit an L2 distance between 0 and 25, while 1060 images have an L2 distance between 25 and 50, and so on.

Previous work [52] reports that the first observed illegitimate collision occurs at an L2 distance between 200 and 225. In contrast, our experiments identify the first illegitimate collision at a significantly lower distance of 144. Moreover, while [52] claims to observe no illegitimate collisions for L2 distances below 200 and only one illegitimate collision between 200 and 225, we observe 342 illegitimate collisions with an L2 distance below 200 and more than 1800 with an L2 distance below 220. Although [52] extends the analysis up to a distance of 500, we limited our study to 220 as beyond this threshold the number of illegitimate collisions increases dramatically.

In [52], the author concludes that setting the threshold between 150 and 175 ensures no false positives while maintaining a false negative rate below 1%. This assessment is corroborated by [51]. However, our results indicate that 1117 legitimate collisions occur at L2 distances greater than 150. This suggests that for a threshold of 150, the false negative rate is not below 1% but instead exceeds 24%. For a threshold of 175, the false negative rate is 18%. While a threshold of 150 may yield an acceptable false positive rate, it does not guarantee the absence of collisions; similarly, a threshold of 175, as proposed in [52], actually leads to both a high false positive rate and a high false negative rate. This discrepancy significantly diminishes the effectiveness of PhotoDNA compared to the claims made by Microsoft and previous studies.

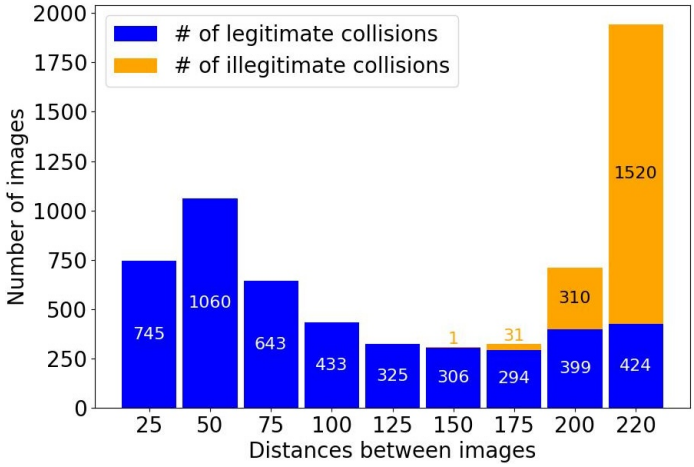

Fig. 6. PhotoDNA collision distribution according to L2 distance

In addition to these observations on near-collisions, we also note—as highlighted in [52]—that the distribution of hash values is not uniform. On average, each hash tuple contains 88 out of 144 values that are below 64, with an overall mean of 62, even though the possible range spans from 0 to 255. Moreover, while the CelebA dataset consists of 58% female and 42% male subjects, false positives

disproportionately affect female subjects, accounting for 62% of cases compared to 38% for males.

6 Estimation for Large-Scale Applications

We use results presented in the previous sections to estimate the implications of the use NeuralHash in the context of large scale CSAM detection. We also evaluate the impact of a 2023 NeuralHash design update. Subsequently we present the consequences of large-scale usage of PhotoDNA.

6.1 Model

This section explains how we estimate the number of illegitimate collisions when scaling up the number of images hashed.

Approximation. To estimate the number of illegitimate collisions when hashing a larger number of images than in our current sets, we reduce the problem to a uniformly distributed hash values with independent bits problem. This approach provides a worst-case scenario, ensuring that our estimates are conservative.

Given the birthday paradox for a uniform distribution with independent bits, we can estimate the number of images that share the same hash values. Let q denote the number of images that illegitimately share their hash values with at least one other image. For a set of uniformly distributed hash values, where each bit is independent of the others, Equation (3), derived from the birthday paradox, estimates q for a set of N images:

$$q = N \left(1 - \left(\frac{2^n - 1}{2^n} \right)^{N-1} \right) . \tag{3}$$

To address the non-uniform distribution and interdependence of bits in NeuralHash case, we reduce the case of NeuralHash to the case of an 'ideal' hash value of $n' < n = 96$ bits, that are uniformly distributed and independent.

To obtain n' from our experiments, we use Equation (3) with q taken from Table 3 and $N = 198\ 174$ unique images[3] in the set. By transforming Equation (3) into Equation (4), we obtain an approximation for the value n':

$$n' \approx -\log_2 \left(1 - \left(1 - \frac{q}{N} \right)^{\frac{1}{N-1}} \right) . \tag{4}$$

Table 5 provides the corresponding n' for each set.

Given a hash size n', we apply the usual birthday paradox (Equation (3)), for any number N of hash values. This provides an approximation of the number of illegitimate collisions for any number of hashes for each image type.

Next we test our approximation on our set of images to verify that the approximation matches the reality.

[3] To ensure a conservative estimate, since at least 4425 images appear more than once, we use $202\,599 - 4425 = 198\ 174$ unique images instead of the full set of 202 599.

Table 5. Size of equivalent uniformly distributed and independent hashes

Type	Non BF	Light BF	BF only	Medium BF	High BF
Size of equivalent hash	30.7	30.6	31.9	27.2	26.0

Verification on 56 Bits of NeuralHash. For the medium blurred and highly blurred types, the number of illegitimate collisions for the entire 96-bit hash value is sufficiently high to verify that the approximation given by Equation (4) fits our experiments. For the other three sets, we randomly selected 56 bits of the NeuralHash hash values and then computed the approximate number of illegitimate collisions for these 56 bits using Equation (4), with $N = 198\,174$.

By varying N for n' obtained from Equation (4) we compute the number of observed illegitimate collisions (orange curve) alongside the approximate number of illegitimate collisions predicted by Equation (3) (blue curve). The number of observed illegitimate collisions represent the number of images in our non-blurred face, light blurred face, and blurred face only sets that have a hash value with at least one other image sharing the selected 56 bits. The results are presented in Figure 7. The observed number of illegitimate collisions closely matches the theoretical values when using the approximation of a well-distributed and independent hash value. Therefore, we apply this approximation to the entire 96-bit hash result and for every image type in the following section.

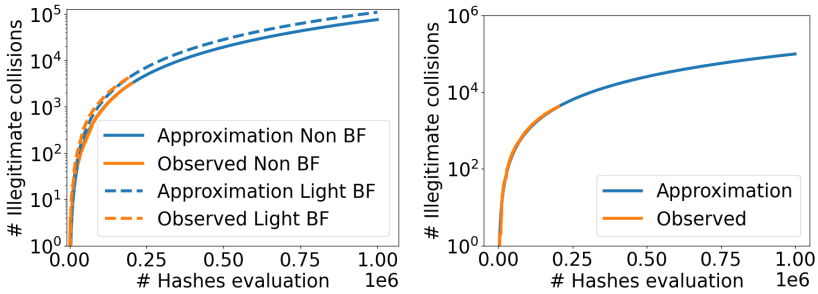

(a) Non BF and Light BF (b) BF only

Fig. 7. Approximated and observed number of illegitimate collisions on 56 bits

6.2 False Positives for NeuralHash

To estimate the potential for false positives in a large-scale context, we extrapolated the number of illegitimate collisions for each image type. Figure 8 shows the approximated (blue curve) and observed (orange curve) number of collisions for medium and highly blurred images. For both types, the actual number of false positives observed in our set closely matches the approximation. For other image types, the number of false positives is lower, making the comparison on 96 bits between the approximation and observed collisions less relevant. However,

the approximation has been validated in the previous section on 56 bits. We nevertheless provide Figure 12 in the Appendix A, which compares the approximation with the observed collisions on 96 bits for the remaining three image types.

The close alignment between the approximated and observed false positives across image types supports the reliability of our method for estimating false positive rates, even at larger scales. The approximation results, derived from Equation (3) and the values of n for each image type (see Table 5), are summarized in Table 6, which shows the theoretical percentage of collisions for 1 million and 10 million images.

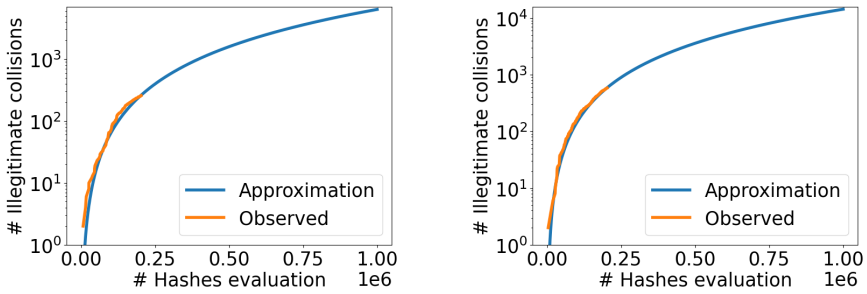

Fig. 8. Approximated and observed number of illegitimate collisions on 96 bits

Table 6. Estimated number of false positive rate (in percentage)

Type	Non BF	light BF	BF only	medium BF	high BF
Real 202 599 hashes	0.01	0.01	0.005	0.13	0.29
Extrapolation for 1M hashes	0.06	0.06	0.03	0.65	1.48
Extrapolation for 10M hashes	0.55	0.63	0.25	6.36	13.84

From Table 6, we observe that the percentage of collisions ranges from 0.03% to 1.2% for 1 million hash values. For 10 million hash values, the percentage of illegitimate collisions increases significantly, ranging from 0.25% to 13.84% according to the image type.

These findings suggest a concerning trend: as the number of hash values increases, the likelihood of false positives also rises substantially. This has significant implications for the deployment of perceptual hash functions such as NeuralHash in large-scale CSAM detection systems. The increasing rate of false positives could lead to numerous legitimate users being incorrectly flagged, thereby undermining the effectiveness and reliability of the detection system. The impact of our results on CSAM detection in a realistic scenario is provided in Section 7.

6.3 False Positives for the 2023 NeuralHash Design

The NeuralHash files deployed on macOS devices have been updated around the end of 2023. The new model uses seeds in fp16 and is trained with fp16 precision (instead of fp32 for the previous model). Apple has not provided any official communication regarding this change. This update is surprising as Apple announced in December 2022 that it canceled its plan to scan photos on Apple devices for CSAM.

Previous attacks on NeuralHash were conducted in a white-box setting, meaning the attackers had full knowledge of the model. With this new model, which has not yet been reverse-engineered, it is currently impossible to fully understand or manipulate the inner workings of the algorithm, especially the CNN involved.

However, it is still possible to run NeuralHash using the files present on a device through code available on the GitHub project [56], without detailed knowledge of the internal processes. We applied our collision attacks using this method, and the results in Table 7 are very similar to those obtained in the previous section.

Table 7. False positives for each image type for the fp16 model of NeuralHash

Type	Non-Human	Non BF	Light BF	BF only	Medium BF	High BF
# illegitimate collisions	0	12	50	90	> 500	> 1000

The results indicate that, except for non-blurred face images where the rate of illegitimate collisions decreased, all other types showed a significant increase in illegitimate collision rates, particularly for medium and highly blurred face images. Therefore, the use of this new model does not alter the conclusions. On the contrary, it tends to exacerbate the issues highlighted.

6.4 False Positives for PhotoDNA

We applied Equations (3) and (4) to our experimental results on PhotoDNA for the Non-Blur Face type, considering threshold distances of 150 and 175. The results are presented in Figure 9.The effective size of uniformly distributed and independent hashes for PhotoDNA is evaluated to be approximately 34 bits for a threshold of 150 and approximately 29 bits for a threshold of 175. It is worth noting that [39] proposes a very similar method for evaluating the equivalent size of uniformly distributed and independent hashes with near-collisions; the main difference is that [39] uses the Hamming distance to find near-collisions, whereas we use the L2 distance.

In Appendix B, we provide a comparison between the estimated number of illegitimate collisions and the observed illegitimate collisions for thresholds of 175 and 220. Since there is only one collision for a threshold of 150, the comparison is not relevant in that case. Nevertheless, the estimates for thresholds of 175 and

220 match our observations, which gives us confidence in both the effective hash size approximation and the estimation of false positives at large scale.

When the threshold is set to 150, the near-collision rate remains relatively low; however, as mentioned above, this comes at the cost of a high number of false negatives. Conversely, when the threshold is increased to 175, as suggested in [52], the false positive rate becomes higher than that observed for NeuralHash, despite PhotoDNA operating having a result of 1152 bits compared to the 96-bit output of NeuralHash.

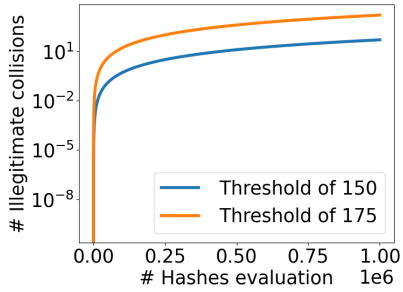

Size of equivalent uniformly distributed and independent hashes

Threshold	150	175
Size of hash	34.2	29.2

Percentage of false positives

Threshold	150	175
Real 202 599 hashes	0.001	0.03
Approx. for 1M hashes	0.005	0.16
Approx. for 10M hashes	0.05	1.6

Fig. 9. Approximated number of PhotoDNA false positives for Non-Blur type images and according to the threshold used

Our analysis of PhotoDNA reinforces the conclusions drawn from NeuralHash: perceptual hash functions currently deployed for content moderation are highly vulnerable to false positives and false negatives, especially when applied to human face images.

7 Collisions with CSAM Databases

This section applies our results on NeuralHash and PhotoDNA to estimate the expected number of false positives when CSAM content would be detected at scale using the current CSAM databases.

CSAM Hash Database Scale. According to the NCMEC report [46], *"As of December 31, 2023, NCMEC has added 7,705,865 hashes to the Non-Governmental Apparent Child Pornography Hash-Sharing Initiative. Other non-governmental organizations have submitted an additional 9,802,435 hashes."* Furthermore, the report [57] from the organization Thorn states, *"Safer, our all-in-one solution for CSAM detection, uses hashing and matching as a core part of its detection technology. With the largest database of verified hashes (29 million hashes) to match against, Safer can cast the widest net to detect known CSAM."*

From these reports, we can deduce that at least these two organizations have access to databases containing 17.5 million and 29 million hash values,

respectively. This indicates that CSAM hash databases contain between 2^{24} and $2^{24.8}$ hash values; in order to be on the safe side, we will use the value 2^{24} in our estimates. Unfortunately, it can be expected that these databases will become larger in the future. For our research, we did not have access to this database of highly sensitive data. Table 5 shows that for images containing human faces, the maximum effective size of a NeuralHash string, assuming uniformly distributed and independent bits, is reached for the blurred face only type. Therefore, to remain conservative, we will assume that images with only the face blurred are representative of the images in these databases.

Probability of Collision with CSAM Databases. Let p denote the probability that the hash value of a given image matches one of the stored CSAM hash values. Given that we consider the maximum effective size of a NeuralHash string (32 bits), we have $p < \frac{2^{24}}{2^{32}}$. For PhotoDNA, from Figure 9, we have $p < \frac{2^{24}}{2^{34.2}}$ for threshold of 150 and $p < \frac{2^{24}}{2^{29.2}}$ threshold of 175.

Under the assumption that N images are "independent" (which is not completely realistic as many users take several variants of the same image), the probability $\mathbb{P}(X = k)$ of exactly k collisions between user images and the CSAM database follows a binomial distribution:

$$\mathbb{P}(X = k) = \binom{N}{k} p^k (1-p)^{N-k}\,.$$

The EU Council draft proposal of July 2024 [23] and Apple's reports [2, 9] suggest increasing the number of required flagged images to 2-3 before reporting a user. Given this, the probability of illegitimately reporting a user based on k required flagged images can be expressed as $1 - \sum_{i=0}^{k} \mathbb{P}(X = i)$.

In 2019, European citizens took an average of $N = 597$ selfies per year [30]. Table 8 provides an estimate of the number of EU citizens that would be illegitimately reported in the first year of using NeuralHash or PhotoDNA, based on k, the number of flagged images required before reporting a user. It is important to note that these estimates account only for selfies and apply only to the first year. As users generally accumulate more images over time, the number of flagged citizens would rise significantly in subsequent years.

Our results demonstrate that regardless of the hash function used, hashing face images leads to a huge number of false positives. The false positive rate and the expected number of matches are so high that even raising the threshold to require two or three flagged images would not effectively prevent legitimate users from being mistakenly reported as CSAM owners. Even with this threshold, hundreds of millions of EU citizens would still be falsely reported. Moreover, increasing the hash size to 128 bits is insufficient to reduce the probability to an acceptable level.

8 Conclusion

By using NeuralHash and PhotoDNA as examples of real-world applications, our work demonstrated significant flaws in the use of perceptual hashing for

Table 8. Number of EU citizens illegitimately reported

NeuralHash

# images to flag before reporting	1	2	3	5	8	10
# reported EU citizens	406M	304M	185M	39M	1.2M	69K

PhotoDNA

# images to flag before reporting	1	2	3	5	8	10
# reported EU citizens (Threshold of 150)	179M	41M	6.8M	84K	31	0
# reported EU citizens (Threshold of 175)	Every Citizens				445M	433M

M stands for million K stands for thousand

CSAM detection. Specifically, we found that applying NeuralHash to human faces (whether blurred or not) results in a very high false positive rate. We observed a similar result when using PhotoDNA on non-blurred human faces. Given that CSAM content often contains faces, we expect the false positive rate in this context to be equally high, leading to many users being incorrectly flagged as CSAM owners. Furthermore, NeuralHash was also found to be inaccurate in detecting the same image of a face, exhibiting a very high false negative rate, even for images without any blurring. These false negative rates suggest that NeuralHash is not reliable for its intended purpose either. Our analysis of PhotoDNA has demonstrated similar flaws.

Our work leads to the conclusion that designs of current perceptual hash functions (including the widely used NeuralHash and PhotoDNA) are not suitable for CSAM detection based on Client-Side Scanning. Our black-box attacks show that these functions fail to provide the necessary accuracy and reliability to be deployed for such a sensitive task. Indeed, large-scale deployment of these functions would mean that a large number of innocent citizens would be flagged while the mechanism would be easy to avoid for those who share illegal content. In view of this, we strongly recommend against using NeuralHash, PhotoDNA or similar perceptual hash functions for this application: the potential for harm and the infringement on privacy far outweigh the intended benefits of such systems.

It remains an open problem whether new perceptual hash functions with longer hash values can be designed that provide a better distribution of outputs for relevant inputs such that they resist black-box attacks. Note that increasing the output size means that more information is provided about the inputs, which increases the risk for preimage attacks.

The design of perceptual hash functions that can resist white-box attacks is a much harder open problem, in particular because Client-Side Scanning means that the design cannot be kept secret. Recent research has shown that for the current designs, it is easy to construct false positives and false negatives in a white-box setting.

Even if the accurate and reliable perceptual hash functions were available, the use of Client-Side Scanning remains highly problematic due to the risk of function creep and abuse.

References

1. Abelson, H., Anderson, R.J., Bellovin, S.M., Benaloh, J., Blaze, M., Callas, J., Diffie, W., Landau, S., Neumann, P.G., Rivest, R.L., Schiller, J.I., Schneier, B., Teague, V., Troncoso, C.: Bugs in our pockets: The risks of client-side scanning. J. Cybersecur. **10**(1) (2024). https://doi.org/10.1093/cybsec/tyad020
2. Apple: Csam detection – technical summary 2021. Tech. rep., Apple (2021), `https://www.apple.com/child-safety/pdf/CSAM%20Detection%20Technical%20Summary.pdf`, accessed on July 8, 2024
3. Apple: Expanded protections for children, frequently asked questions (2021), `https://www.apple.com/child-safety/pdf/Expanded_Protections_for_Children_Frequently_Asked_Questions.pdf`, accessed: July 8, 2024
4. Asano, Y.M., Rupprecht, C., Zisserman, A., Vedaldi, A.: PASS: An ImageNet replacement for self-supervised pretraining without humans. NeurIPS Track on Datasets and Benchmarks (2021)
5. Athalye, A.: Inverting PhotoDNA. `https://www.anishathalye.com/2021/12/20/inverting-photodna/` (Dec 2021)
6. Athalye, A.: NeuralHash Collider (2021), `https://github.com/anishathalye/neural-hash-collider`, accessed: July 7, 2024
7. Bai, J., Chen, B., Li, Y., Wu, D., Guo, W., Xia, S., Yang, E.: Targeted attack for deep hashing based retrieval. In: Computer Vision – ECCV 2020 – 16th European Conference, Glasgow, UK, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12346, pp. 618–634. Springer (2020). https://doi.org/10.1007/978-3-030-58452-8_36
8. Bartusek, J., Garg, S., Jain, A., Policharla, G.: End-to-end secure messaging with traceability only for illegal content. In: Advances in Cryptology – EUROCRYPT. Lecture Notes in Computer Science, vol. 14008, pp. 35–66. Springer (2023). https://doi.org/10.1007/978-3-031-30589-4_2
9. Bhowmick, A., Boneh, D., Myers, S., Talwar, K., Tarbe, K.: The apple psi system. Tech. rep., Apple, Inc. (2021), `https://www.apple.com/child-safety/pdf/Apple_PSI_System_Security_Protocol_and_Analysis.pdf`
10. Brunet, D., Vrscay, E.R., Wang, Z.: On the mathematical properties of the structural similarity index. IEEE Trans. Image Process. **21**(4), 1488–1499 (2012). https://doi.org/10.1109/TIP.2011.2173206
11. Bursztein, E., Clarke, E., DeLaune, M., Elifff, D.M., Hsu, N., Olson, L., Shehan, J., Thakur, M., Thomas, K., Bright, T.: Rethinking the detection of child sexual abuse imagery on the internet. In: The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13–17, 2019. pp. 2601–2607. ACM (2019). https://doi.org/10.1145/3308558.3313482
12. Carlini, N., Wagner, D.A.: Towards evaluating the robustness of neural networks. In: 2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22–26, 2017. pp. 39–57. IEEE Computer Society (2017). https://doi.org/10.1109/SP.2017.49
13. Chandler, D.M., Hemami, S.S.: VSNR: A wavelet-based visual signal-to-noise ratio for natural images. IEEE Trans. Image Process. **16**(9), 2284–2298 (2007). https://doi.org/10.1109/TIP.2007.901820
14. Corporation, M.: New technology fights child porn by tracking its "PhotoDNA" (December 15, 2009), `https://news.microsoft.com/2009/12/15/new-technology-fights-child-porn-by-tracking-its-photodna/`, accessed: July 8, 2024

15. Counter Extremism Project: How CEP's eGLYPH Technology Works (Dec 08, 2016), https://www.counterextremism.com/video/how-ceps-eglyph-technology-works, accessed: July 8, 2024
16. Croce, F., Hein, M.: Minimally distorted adversarial examples with a fast adaptive boundary attack. In: Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event. Proceedings of Machine Learning Research, vol. 119, pp. 2196–2205. PMLR (2020)
17. Dalins, J., Wilson, C., Boudry, D.: PDQ & TMK + PDQF - a test drive of facebook's perceptual hashing algorithms. CoRR **abs/1912.07745** (2019)
18. Datar, M., Immorlica, N., Indyk, P., Mirrokni, V.S.: Locality-sensitive hashing scheme based on p-stable distributions. In: Proceedings of the Twentieth Annual Symposium on Computational Geometry. pp. 253–262. ACM (2004)
19. Dong, Y., Liao, F., Pang, T., Su, H., Zhu, J., Hu, X., Li, J.: Boosting adversarial attacks with momentum. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018. pp. 9185–9193. Computer Vision Foundation / IEEE Computer Society (2018). https://doi.org/10.1109/CVPR.2018.00957
20. Drmic, A., Silic, M., Delac, G., Vladimir, K., Kurdija, A.S.: Evaluating robustness of perceptual image hashing algorithms. In: 40th International Convention on Information and Communication Technology, Electronics and Microelectronics, MIPRO, Opatija, Croatia. pp. 995–1000. IEEE (2017). https://doi.org/10.23919/MIPRO.2017.7973569
21. Du, L., Ho, A.T.S., Cong, R.: Perceptual hashing for image authentication: A survey. Signal Process. Image Commun. **81** (2020). https://doi.org/10.1016/j.image.2019.115713
22. Cross-party letter of members of the european parliament against general monitoring (2021), https://www.patrick-breyer.de/wp-content/uploads/2021/11/20211020_Letter_General_Monitoring.pdf, patrick Breyer, Alviina Alametsä, Rosa D'Amato, Pernando Barrena, Saskia Bricmont, Antoni Comín, Gwendoline Delbos-Corfield, Francesca Donato, Cornelia Ernst, Claudia Gamon, Markéta Gregorová, Francisco Guerreiro, Svenja Hahn, Irena Joveva, Petra Kammerevert, Marcel Kolaja, Moritz Körner, Karen Melchior, Clara Ponsatí, and Mikuláš Peksa; accessed: July 7, 2024
23. Proposal for a regulation of the european parliament and of the council laying down rules to prevent and combat child sexual abuse (2024), https://www.patrick-breyer.de/wp-content/uploads/2024/04/2024-03-28-conseil-csam-compromis-presidence-belge.pdf, accessed on July 6, 2024
24. Facebook: The TMK+PDQF video-hashing algorithm and the PDQ image-hashing algorithm (2020), https://github.com/facebook/ThreatExchange/blob/main/hashing/hashing.pdf, accessed: July 8, 2024
25. Farid, H.: An overview of perceptual hashing. Journal of Online Trust and Safety **1**(1) (Oct 2021). https://doi.org/10.54501/jots.v1i1.24
26. Gionis, A., Indyk, P., Motwani, R.: Similarity search in high dimensions via hashing. In: VLDB'99, Proceedings of 25th International Conference on Very Large Data Bases, September 7–10, 1999, Edinburgh, Scotland, UK. pp. 518–529. Morgan Kaufmann (1999)
27. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings (2015)
28. Google Inc.: How Content ID works (2007), https://support.google.com/youtube/answer/2797370?hl=en, accessed: July 8, 2024

29. Hao, Q., Luo, L., Jan, S.T.K., Wang, G.: It's not what it looks like: Manipulating perceptual hashing based applications. In: CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15–19, 2021. pp. 69–85. ACM (2021). https://doi.org/10.1145/3460120.3484559
30. for HONOR, G.: Honor 9x research reveals europeans now post a staggering 597 photos of themselves every year (December 9, 2019), \url{https://www.honor.com/global/news/honor-9x-research-reveals-europeans-now-post-a-staggering-597-photos-of-themselves-every-year/}, research carried out by GingerComms on behalf of HONOR on 2,053 UK adults and 750 adults in France, Germany, Spain, Italy and the Netherlands between 14–18 November 2019.
31. Indyk, P., Motwani, R.: Approximate nearest neighbors: Towards removing the curse of dimensionality. In: Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23–26, 1998. pp. 604–613. ACM (1998). https://doi.org/10.1145/276698.276876
32. Jain, S., Cretu, A., Cully, A., de Montjoye, Y.: Deep perceptual hashing algorithms with hidden dual purpose: When client-side scanning does facial recognition. In: 44th IEEE Symposium on Security and Privacy, SP 2023, San Francisco, CA, USA, May 21–25, 2023. pp. 234–252. IEEE (2023). https://doi.org/10.1109/SP46215.2023.10179310
33. Jain, S., Cretu, A., de Montjoye, Y.: Adversarial detection avoidance attacks: Evaluating the robustness of perceptual hashing-based client-side scanning. In: 31st USENIX Security Symposium, USENIX Security 2022, Boston, MA, USA, August 10–12, 2022. pp. 2317–2334. USENIX Association (2022)
34. Kaiser, J.: pyphotodna (2023), https://github.com/jankais3r/pyPhotoDNA, accessed: Feb 12, 2025
35. Kilcher, Y.: Neural hash collision creator (2021), https://github.com/yk/neural_hash_collision, accessed: July 7, 2024
36. Klinger, E., Starkweather, D.: pHash: The open source perceptual hash library (2010), https://www.phash.org/, accessed on July 22, 2024
37. Kulis, B., Darrell, T.: Learning to hash with binary reconstructive embeddings. In: Advances in Neural Information Processing Systems 22: 23rd Annual Conference on Neural Information Processing Systems 2009. Vancouver, British Columbia, Canada. pp. 1042–1050. Curran Associates, Inc. (2009)
38. Kulshrestha, A., Mayer, J.R.: Identifying harmful media in end-to-end encrypted communication: Efficient private membership computation. In: 30th USENIX Security Symposium, USENIX Security 2021, August 11–13, 2021. pp. 893–910. USENIX Association (2021)
39. Lamberger, M., Teufl, E.: Memoryless near-collisions, revisited. Information Processing Letters **113**(3), 60–66 (2013). https://doi.org/10.1016/j.ipl.2012.11.004
40. Levy, I., Robinson, C.: Thoughts on child safety on commodity platforms (2022)
41. Liong, V.E., Lu, J., Wang, G., Moulin, P., Zhou, J.: Deep hashing for compact binary codes learning. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7–12, 2015. pp. 2475–2483. IEEE Computer Society (2015). https://doi.org/10.1109/CVPR.2015.7298862
42. Liu, H., Wang, R., Shan, S., Chen, X.: Deep supervised hashing for fast image retrieval. Int. J. Comput. Vis. **127**(9), 1217–1234 (2019). https://doi.org/10.1007/s11263-019-01174-4
43. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: Proceedings of International Conference on Computer Vision (ICCV) (December 2015)

44. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 – May 3, 2018, Conference Track Proceedings. OpenReview.net (2018)
45. NCMEC: Ncmec's statement regarding end-to-end encryption. Tech. rep., National Center for Missing and Exploited Children (2019), https://www.missingkids.org/blog/2019/post-update/end-to-end-encryption
46. NCMEC: OJJDP-NCMEC-Transparency-CY-2023-Report. Tech. rep., NCMEC (2024), https://www.missingkids.org/content/dam/missingkids/pdfs/OJJDP-NCMEC-Transparency-CY-2023-Report.pdf, accessed: July 22, 2024
47. Joint statement of scientists and researchers on EU's proposed child sexual abuse regulation (2023), https://edri.org/wp-content/uploads/2023/07/Open-Letter-CSA-Scientific-community.pdf, accessed: July 7, 2024
48. An open letter against Apple's privacy-invasive content scanning technology (2021), https://appleprivacyletter.com/, accessed: July 23, 2024
49. Joint statement of scientists and researchers on EU's new proposal for the child sexual abuse regulation (2024), https://edri.org/wp-content/uploads/2023/07/Open-Letter-CSA-Scientific-community.pdfhttps://nce.mpi-sp.org/index.php/s/eqjiKaAw9yYQF87, accessed: July 7, 2024
50. Papernot, N., McDaniel, P.D., Jha, S., Fredrikson, M., Celik, Z.B., Swami, A.: The limitations of deep learning in adversarial settings. In: IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21–24, 2016. pp. 372–387. IEEE (2016). https://doi.org/10.1109/EUROSP.2016.36
51. Prokos, J., Fendley, N., Green, M., Schuster, R., Tromer, E., Jois, T., Cao, Y.: Squint hard enough: Attacking perceptual hashing with adversarial machine learning. In: 32nd USENIX Security Symposium (USENIX Security 23). pp. 211–228. USENIX Association, Anaheim, CA (2023)
52. Steinebach, M.: An analysis of photodna. In: Proceedings of the 18th International Conference on Availability, Reliability and Security, ARES 2023, Benevento, Italy, 29 August 2023–1 September 2023. pp. 44:1–44:8. ACM (2023). https://doi.org/10.1145/3600160.3605048
53. Struppek, L., Hintersdorf, D., Neider, D., Kersting, K.: Learning to break deep perceptual hashing: The use case neuralhash. In: FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, Seoul, Republic of Korea, June 21–24, 2022. pp. 58–69. ACM (2022). https://doi.org/10.1145/3531146.3533073
54. Su, J., Vargas, D.V., Sakurai, K.: One pixel attack for fooling deep neural networks. IEEE Trans. Evol. Comput. **23**(5), 828–841 (2019). https://doi.org/10.1109/TEVC.2019.2890858
55. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I.J., Fergus, R.: Intriguing properties of neural networks. In: 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings (2014)
56. Tian, K., Hall, M.: nhcalc (2021), https://github.com/KhaosT/nhcalc, accessed: July 23, 2024
57. Torn: How hashing and matching can help prevent revictimization (2023), https://www.thorn.org/blog/hashing-detect-child-sex-abuse-imagery, accessed: July 22, 2024
58. UK Department for Science, Innovation & Technology: Online safety act: Explainer (8 May 2024), https://www.gov.uk/government/publications/online-safety-act-explainer/online-safety-act-explainer

59. Wang, X., Zhang, Z., Lu, G., Xu, Y.: Targeted attack and defense for deep hashing. In: SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada. pp. 2298–2302. ACM (2021). https://doi.org/10.1145/3404835.3463233
60. Wang, X., Zhang, Z., Wu, B., Shen, F., Lu, G.: Prototype-supervised adversarial network for targeted attack of deep hashing. In: IEEE CVPR. pp. 16357–16366. Computer Vision Foundation / IEEE (2021). https://doi.org/10.1109/CVPR46437.2021.01609
61. Wang, Y., Palangi, H., Wang, Z.J., Wang, H.: Revhashnet: Perceptually de-hashing real-valued image hashes for similarity retrieval. Signal Process. Image Commun. **68**, 68–75 (2018). https://doi.org/10.1016/J.IMAGE.2018.06.018
62. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. **13**(4), 600–612 (2004). https://doi.org/10.1109/TIP.2003.819861
63. Wu, D., Lin, Z., Li, B., Ye, M., Wang, W.: Deep supervised hashing for multi-label and large-scale image retrieval. In: Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval, ICMR 2017, Bucharest, Romania, June 6–9, 2017. pp. 150–158. ACM (2017). https://doi.org/10.1145/3078971.3078989
64. Yang, E., Liu, T., Deng, C., Tao, D.: Adversarial examples for hamming space search. IEEE Trans. Cybern. **50**(4), 1473–1484 (2020). https://doi.org/10.1109/TCYB.2018.2882908
65. Ygvar, A.: AppleNeuralHash2ONNX (Aug 2021), `https://github.com/AsuharietYgvar/AppleNeuralHash2ONNX`, accessed: July 12, 2024
66. Zauner, C.: Implementation and benchmarking of perceptual image hash functions, MsC Thesis, University of Applied Sciences, Hagenberg, Austria (2010)
67. Zhao, F., Huang, Y., Wang, L., Tan, T.: Deep semantic ranking based hashing for multi-label image retrieval. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7–12, 2015. pp. 1556–1564. IEEE Computer Society (2015). https://doi.org/10.1109/CVPR.2015.7298763
68. Zhu, M., Chen, T., Wang, Z.: Sparse and imperceptible adversarial attack via a homotopy algorithm. In: Proceedings of the 38th International Conference on Machine Learning, ICML, Virtual Event. Proceedings of Machine Learning Research, vol. 139, pp. 12868–12877. PMLR (2021)

Appendix A

Distribution of Bit Values

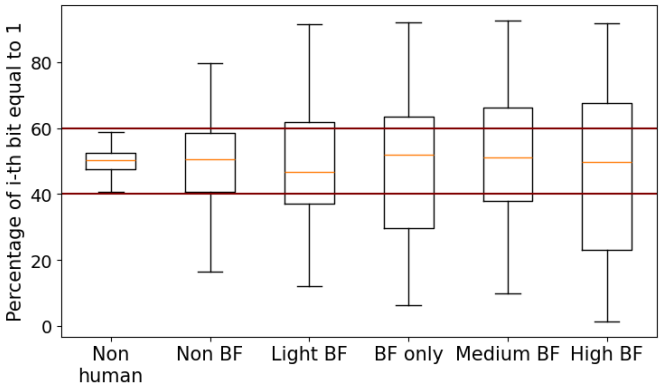

Fig. 10. Disparity of bit values for each type of image

SMM Color Visualization

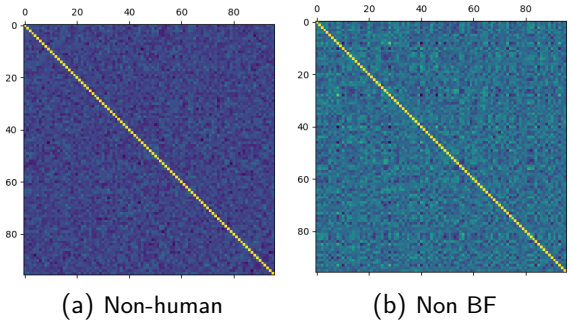

Fig. 11. *SMM* for Non-human and Non BF sets

Approximation on 96 bits

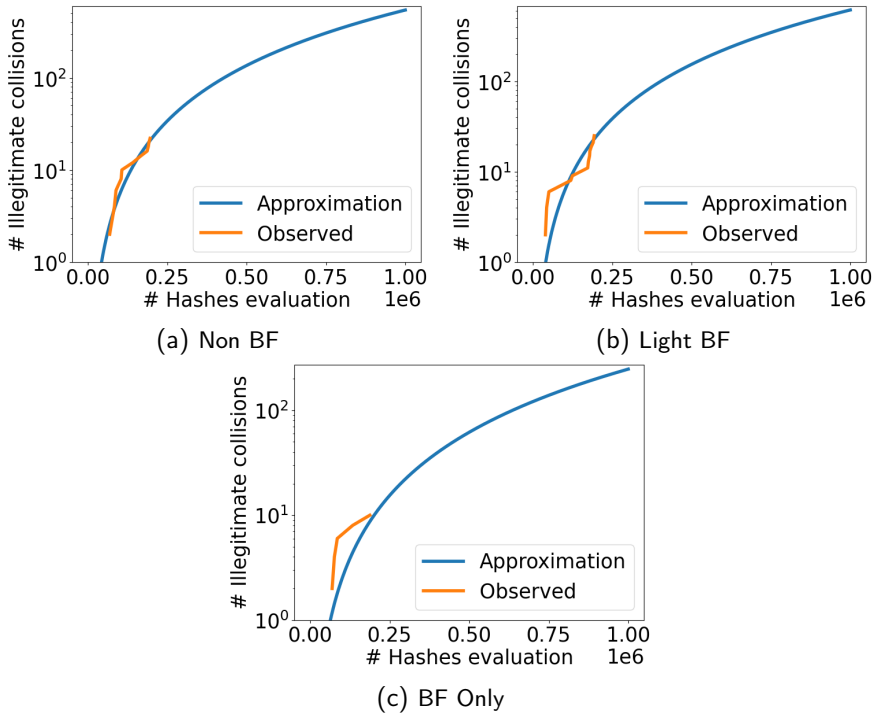

Fig. 12. Approximated and observed number of illegitimate collisions on 96 bits

Appendix B

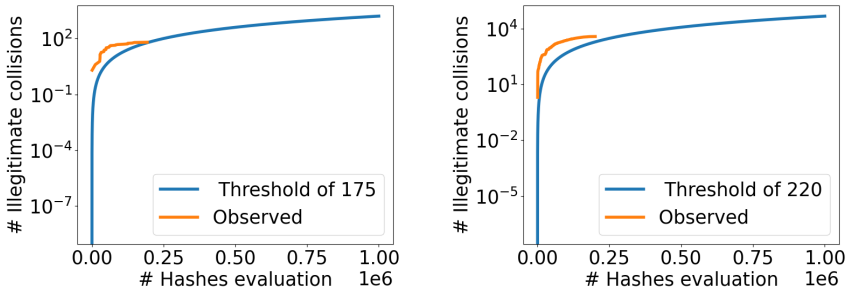

(a) Threshold of 175 (b) Threshold of 220

Fig. 13. Approximated and observed number of illegitimate collisions for PhotoDNA with threshold of 175 and threshold of 220