
Exploiting and Defending Against the Approximate Linearity of Apple’s NEURALHASH

Jagdeep S Bhatia^{1*} Kevin Meng^{1*}

Abstract

Perceptual hashes map images with identical semantic content to the same n -bit hash value, while mapping semantically-different images to different hashes. These algorithms carry important applications in cybersecurity such as copyright infringement detection, content fingerprinting, and surveillance. Apple’s NEURALHASH is one such system that aims to detect the presence of illegal content on users’ devices without compromising consumer privacy. We make the surprising discovery that NEURALHASH is *approximately linear*, which inspires the development of novel black-box attacks that can (i) evade detection of “illegal” images, (ii) generate near-collisions, and (iii) leak information about hashed images, all without access to model parameters. These vulnerabilities pose serious threats to NEURALHASH’s security goals; to address them, we propose a simple fix using classical cryptographic standards.

1. Introduction

In 2021, Apple unveiled NEURALHASH, an algorithm aimed at detecting Child Sexual Abuse Material (CSAM) in images that are uploaded to iCloud (Apple, 2021).

NEURALHASH is a perceptual hash that, in general terms, aims to map images containing the same semantic information to the same hash, while mapping images with differing semantic content to different, random-looking hashes. If executed correctly, such algorithms could potentially enable organizations to prevent illegal activity (e.g. copyright infringement, storing illicit audiovisual content) *without compromising user privacy*. Thus, NEURALHASH is of great interest to the cybersecurity community.

*Equal contribution ¹Massachusetts Institute of Technology. Correspondence to: Jagdeep Singh Bhatia <jagdeep@mit.edu>, Kevin Meng <mengk@mit.edu>.

Unfortunately, NEURALHASH has been controversial since its inception, due to concerns about efficacy and user privacy (Abelson et al., 2021). Recently, these concerns have been made concrete, with studies identifying problems with NEURALHASH’s collision resistance, privacy guarantees, and false positive/negative rates under gradient-based adversarial attacks (Struppek et al., 2021; Ygvar, 2021; Roboflow, 2021; Dwyer, 2021).

NEURALHASH relies on deep learning to hash images, in part because semantic image matching is difficult with rule-based or classical learning algorithms (Di Martino & Sessa, 2013; Chum et al., 2008). Despite their advantages, however, neural networks are known to be sensitive to edge cases, susceptible to adversarial attacks, and prone to unexplainable behavior (Goodfellow et al., 2014).

White-box adversarial attacks on neural networks are well studied in literature; as a result, numerous gradient-based attacks have been proposed against NEURALHASH (Athalye, 2021; Struppek et al., 2021). We ask whether the same attacks can be achieved *without* access to model weights, as the ability to construct such attacks in a gradient-free setting would imply a deeper understanding of NEURALHASH’s inner workings. In particular, it is generally accepted that gradient-based attacks are an unavoidable phenomenon in deep neural networks (Ilyas et al., 2019); the presence of strong *black-box* attacks may suggest deeper structural flaws with NEURALHASH.

In this paper, we make the striking discovery that NEURALHASH is approximately linear in its inputs (Section 3), which not only violates NEURALHASH’s privacy guarantees (Section 2.1), but also provides a framework under which we develop several strong black-box attacks (Section 4). In Section 5, we suggest a defense mechanism against the proposed attacks and demonstrate its effectiveness. To the best of our knowledge, we are the first to study the approximately linear nature of NEURALHASH’s hashing mechanism.

2. Preliminaries

2.1. NEURALHASH Security and Privacy Goals

Definition 2.1. Semantic Similarity: Two images $x_1, x_2 \in X$ are semantically similar if, given some metric function




Figure 1. **NEURALHASH Pipeline.** The pre-processed image is first embedded with a contrastively-trained MobileNet (Howard et al., 2017) convolutional neural network. The extracted features are then transformed with a randomized hashing matrix, whose outputs are passed through a step function to generate the final binary hash (Struppek et al., 2021; Apple, 2021).

$d : X \times X \rightarrow \mathbb{R}$ and some threshold ϵ , we have that $d(x_1, x_2) \leq \epsilon$. d can be implemented in various ways, depending on the evaluation context.

The primary difference between NEURALHASH and classical cryptographic hashing is that NEURALHASH aims to remain invariant under transformations that preserve semantic similarity. This gives rise to the first desired property of NEURALHASH:

Property 2.2. Neighborhood Consistency: Let $x_1, x_2 \in X$ be any semantically similar images and $\mathcal{N}(x) : X \rightarrow \mathcal{H}$ be the NEURALHASH perceptual hashing algorithm. It should be the case that $\mathcal{N}(x_1) = \mathcal{N}(x_2)$. Note that, in NEURALHASH, $\mathcal{H} = \{0, 1\}^{96}$.

In analyzing the behavior of NEURALHASH, we adopt a variant of the Random Oracle Model that accounts for semantic similarity. In our setting, a Perceptual Random Oracle (PRO) should return a random hash for any previously-unseen query x , and the same hash value as a previous input x' if x and x' are semantically similar:

$$\mathcal{N}(x) = \begin{cases} \mathcal{N}(x') & \text{if } \exists x'. d(x, x') \leq \epsilon \\ h \sim \mathcal{H} & \text{otherwise.} \end{cases} \quad (1)$$

Because we expect NEURALHASH to behave as an idealized PRO, it should satisfy the following properties:¹

Property 2.3. Uniform Hashing: Let $x_1, x_2 \in X$ be any images that are *not* semantically similar. It should be the case that (i) $\mathcal{N}(x_1) \neq \mathcal{N}(x_2)$ and (ii) the distributions over $\mathcal{N}(x_1)$ and $\mathcal{N}(x_2)$ look indistinguishable to a computationally bounded adversary. $\mathcal{N}(x_1)$ and $\mathcal{N}(x_2)$ should appear to be random elements of the hash space, and therefore,

$$\mathbb{E}_{x_1, x_2 \in X, d(x_1, x_2) > \epsilon} [\text{sim}(\mathcal{N}(x_1), \mathcal{N}(x_2))] \leq 0.5 + \text{negl}, \quad (2)$$

where, treating hashes as binary vectors,

$$\text{sim}(h_1, h_2) = \frac{h_1 \cdot h_2}{|h_1| |h_2|}. \quad (3)$$

¹Assume that adversaries are *computationally bounded* and have unrestricted access to the hashing oracle $\mathcal{N}(\cdot)$.

Note that $\text{sim}(h_1, h_2)$ is equal to one minus the normalized Hamming distance; call it *Hamming similarity*.

Property 2.4. Weak Non-Invertibility: Given target hash $h^* \in \mathcal{H}$, it should be impossible for an adversary to generate any image x such that $\mathcal{N}(x) = h^*$.

Property 2.5. Strong Non-Invertibility: Given any hash $h^* \in \mathcal{H}$ and n images x_1, \dots, x_n sampled uniformly at random from X , let s be the expected value of the highest Hamming Similarity score between h^* and the hashes of all sampled images:

$$s = \mathbb{E}_{\{x_1, \dots, x_n\} \sim X} \left[\max_{x_i} \text{sim}(h^*, \mathcal{N}(x_i)) \right]. \quad (4)$$

Given n queries to the hashing oracle, it should be hard for an adversary to consistently generate near collisions:

$$\mathbb{E}_x [\text{sim}(\mathcal{N}(x), h^*)] \leq s + \text{negl}. \quad (5)$$

The ability for an adversary to do so would suggest that $\mathcal{N}(x)$ leaks information about x .

Note that our security goals for NEURALHASH are considerably stricter than those of other perceptual hashes (Microsoft, 2022; Facebook, 2020), particularly in terms of privacy (Properties 2.3, 2.5). This is because NEURALHASH is used in a *client-side scanning* (CSS) scheme (Bhowmick et al., 2021) on users’ devices, where potential false positives or privacy leaks would significantly damage trust between consumers and providers.

2.2. Functional Description of NEURALHASH

NEURALHASH hashes are generated in two steps. First, a MobileNet-based (Howard et al., 2017) convolutional neural network generates a feature vector describing the image using 128 floating-point numbers. This vector is then multiplied by a randomized hashing matrix; the result is then thresholded using a step function to produce the final 96-bit binary hash. See Figure 1 for a visual depiction.

Note that the hashing matrix can be interpreted as containing ninety-six random 128-dimensional hyperplanes. To




Figure 2. Images produced by interpolation of a face and a tree, with parameter α denoting the proportion of the first image.

generate the hash, the 128-dim feature vector is compared with each of these hyperplanes; if the feature vector is on the positive side of the hyperplane, the bit is 1, else 0.

2.3. Related Work in Attacking NEURALHASH

In the realm of intervention-free attacks, Dwyer (2021); Roboflow (2021) discovered naturally-occurring NEURALHASH collisions by mining publicly-available images that were not artificially modified in any way. This demonstrates a failure of Property 2.3.

Because NEURALHASH relies on a neural network whose weights have been extracted and made public (Ygvar, 2021), a variety of gradient-based attacks have also been proposed (Struppek et al., 2021). These white-box attacks rely on the susceptibility of neural networks to small adversarial perturbations (Goodfellow et al., 2014), allowing malicious actors to both evade detection (violating Property 2.2) and generate collisions (violating Properties 2.4, 2.5) using small amounts of adversarial noise.

Even more interesting, however, would be the existence of strong *black-box* attacks; it is generally accepted that adversarial examples are an unavoidable phenomenon in neural networks. Struppek et al. (2021) makes initial headway into this line of inquiry, showing that basic image transformations including translation, rotation, flipping, cropping, and brightness/contrast changes can help evade detection (violating Property 2.2). The other security properties remain underexplored in the black-box setting.

In summary, the existing literature around gradient-based white-box attacks is fairly well-developed, but we lack a sufficiently deep understanding of NEURALHASH’s behavior to develop strong black-box attacks.

3. On NEURALHASH’s Approximate Linearity

Imagine that we have two images x_1, x_2 where $d(x_1, x_2) > \epsilon$. By Property 2.3, they should initially have random-looking hashes on average, i.e. $\mathbb{E}[\text{sim}(\mathcal{N}(x_1), \mathcal{N}(x_2))] \leq 0.5 + \text{negl}$. But what happens as we gradually move x_2 towards x_1 ? In an idealized NEURALHASH, these hashes

should remain random-looking until $d(x_1, x_2') \leq \epsilon$, at which point the hashes should match exactly.

A natural strategy for analyzing the gradual movement of images toward each other is *image interpolation*. To the best of our knowledge, this has not yet been studied in the context of NEURALHASH.

More formally, consider taking two source images x_1 and x_2 and linearly interpolating them with respect to some parameter α . The resulting image is $I_\alpha = \alpha x_1 + (1 - \alpha)x_2$. Figure 2 provides a visualization of interpolation where α is varied from 0 to 1 between an image of a French horn (x_1) and an image of a church (x_2).

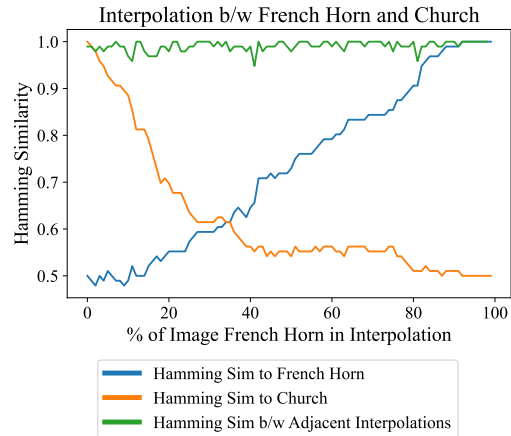


Figure 3. **Approximate Piecewise Linearity of NEURALHASH.** When interpolating between two images, we observe violations of NEURALHASH’s security criteria (Property 2.3). Hashes of semantically-unrelated images do *not* look random.

To quantitatively analyze the behavior of NEURALHASH as x_1, x_2 are interpolated, we plot the Hamming similarity between I_α and both x_1, x_2 for all α ; Figure 3 displays the results. Strikingly, as α moves from 0 to 1, the Hamming similarity between I_α and the original images x_1, x_2 changes approximately *piecewise linearly*.

This provides strong evidence that NEURALHASH fails to meet its privacy requirements. Specifically, by Property 2.3, we would expect (i) the orange and blue lines to remain at

roughly 0.5 in the middle region, where I_α is semantically-identical to neither x_1 nor x_2 , and (ii) the green line to oscillate between 1 and 0.5. Empirically, we observe failures of both expectations.

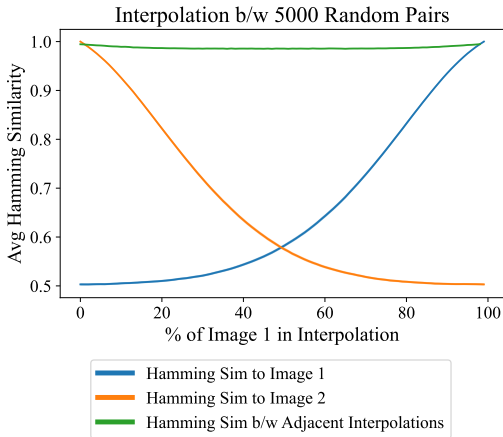


Figure 4. **Averaged Interpolation Behavior.** We average interpolation similarities over a large set of randomly sampled ImageNette pairs, plotted with 99% confidence intervals. The approximate linearity of NEURALHASH is systematic.

Furthermore, this approximate piecewise linearity is *systematic*. In Figure 4, we average Hamming similarities over a large set of ImageNette² (FastAI, 2019) pairs, finding the same pattern with tight 99% confidence intervals.

These results are significant not only because they demonstrate the violation of several security properties. Perhaps more importantly, they clarify a key operational behavior of NEURALHASH, approximate linearity, which enables us to make fairly reliable inferences about the hashes of the interpolations between a set of images.

4. Attacks

The approximate linearity of NEURALHASH lends itself naturally to several attacks. Firstly, the interpolation graphs (Figures 3, 4) suggest the possibility of an **evasion attack** where, given a target image x , we can compute a semantically similar image x' by adding a small percentage of another image, such that $\mathcal{N}(x) \neq \mathcal{N}(x')$. This would break the neighborhood property.

The predictability of NEURALHASH’s hashing behavior also suggests that we might be able to **generate near-collisions** where, given any target image x , we can *on average* generate an image x' that nearly collides with but is not semantically similar to x : $\mathbb{E}_x[\text{sim}(\mathcal{N}(x), \mathcal{N}(x'))] \gg 0.5 + \text{negl}$. This

²ImageNette is a subset of ImageNet containing roughly 1000 examples from 10 classes. Note that since NEURALHASH is meant to be a generalized perceptual hashing algorithm, it is reasonable to use any dataset for our experiments.

breaks uniform hashing and strong non-invertibility.

Finally, we hypothesize that NEURALHASH **leaks information**. Given an image x which belongs to one of n known image classes $\{C_1, C_2, \dots, C_n\}$, we can predict target class C_t with probability greater than $\frac{1}{n} + \text{negl}$. This would break the uniform hashing property.

In the following subsections, we demonstrate that all attacks can be realized using interpolation-based algorithms.

4.1. Evasion Attack

Attack Description. Given a source image x and an arbitrary semantically-unrelated image x_0 , we produce a semantically-similar image $x' = \alpha^*x + (1 - \alpha^*)x_0$ such that $\mathcal{N}(x) \neq \mathcal{N}(x')$. α^* is chosen as the largest α for which the hashes differ:

$$\alpha^* = \max \left\{ \alpha \mid \mathcal{N}(x) \neq \mathcal{N}(\alpha x + (1 - \alpha)x_0) \right\} \quad (6)$$

Intuitively, this selects for the interpolation that contains the maximum possible percentage of x .

Attack Results. We consider an attack “successful” iff (i) $\mathcal{N}(x) \neq \mathcal{N}(x')$ and (ii) x and x' contain the same semantic information. To formalize the latter concept, we adopt the Structural Similarity Index Measure (SSIM) metric, as used in Struppek et al. (2021):

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)}, \quad (7)$$

where x and y are images with means μ_x, μ_y , variances σ_x^2, σ_y^2 , and covariance σ_{xy} . Moreover, $c_1 = 0.01, c_2 = 0.03$ for numerical stability. The intuition is that, by accounting for mean, variance, and covariance in the pixel space, SSIM focuses on *structure* better than Mean-Squared Error (MSE), for example, which can be distracted by noise.³




Figure 5. **Qualitative Example of Evasion.** Interpolating between the gas pump and cassette player ($\alpha^* = 0.95$, as computed using Eqn. 6) results in an image that is semantically indistinguishable from the original gas pump, but has a different hash.

We carry out a systematic evaluation by randomly sampling 10,000 pairs of images from ImageNette, designating one

³Note that SSIM is *not* robust under various transformations such as rotation or translation, but this is not a concern because the interpolation attack is executed with a pixelwise sum.

of each pair to be x and the other to be x_0 . On average, the attack achieves **100% efficacy** with an **SSIM of 0.981** ($\pm 5 \times 10^{-4}$ for a 95% confidence interval). Figure 5 shows a representative example of the attack. The interpolated image is hardly distinguishable from the original gas pump picture, but its hash is different in multiple bit positions.

Our interpolation-based evasion attack has several advantages over existing methods. In particular, transformations such as rotation, cropping, translation, reflection, and brightness/contrast adjustment (Struppek et al., 2021) *do* preserve semantic content but nonetheless result in significant changes to the appearance of images. By comparison, our interpolation attacks cause minimal visual impact and involve less guesswork than performing arbitrary transformations. A limitation, however, is that the efficacy of the interpolation attack is somewhat dependent on a good choice of x_0 . We show that, on average, the method is highly robust, but there are a few outlier cases where a poor choice of x_0 causes SSIM to drop below 0.7.

4.2. Near-Collision Generation Attack

Attack Description. Consider a database of images $X = \{x_1, x_2, \dots, x_k\}$. Let P be a parameter space describing all possible interpolations of images in X . Thus, each element $p \in P$ represents an interpolated image I_p where

$$I_p = \sum_{i=1}^k p_i x_i, \quad (8)$$

and $p_i \in [-1, 1]$ and $\sum_{i=1}^k |p_i| = 1$. Then, given a target hash h^* , we employ a genetic algorithm (Whitley, 1994) to search P for an interpolated image whose hash maximizes Hamming similarity with h^* . Formally, we optimize:

$$\max_{p \in P} \text{sim}(\mathcal{N}(I_p), h^*) \quad (9)$$

using the Hamming similarity score defined in Eqn. 3 and with I_p as defined in Eqn. 8.

Crucially, our algorithm employs two genetic operators that rely heavily on the properties of interpolated images explored in Section 3.

Crossover Operator Given $p, q \in P$ and $\alpha \in [0, 1]$. Compute a new crossed-over parameter $r = \alpha p + (1 - \alpha)q$. Note that r is normalized after interpolation such that $I_r \approx \alpha I_p + (1 - \alpha)I_q$. By the interpolation property, with high probability, $\text{sim}(\mathcal{N}(I_r), h^*)$ will be between $\text{sim}(\mathcal{N}(I_p), h^*)$ and $\text{sim}(\mathcal{N}(I_q), h^*)$, and with some nontrivial probability it will be greater.

Mutation Operator Given $p \in P$, index $m \in \{1, \dots, k\}$ and $\alpha \in [-0.05, 0.05]$. Compute a new mutated parameter r by mutating p_m by α . More formally,

$\forall i \neq m. r_i = p_i$, and $r_m = p_m + \alpha$. By our interpolation property, such a mutation should only slightly change the Hamming similarity. Additionally, with some nontrivial probability, $\text{sim}(\mathcal{N}(I_r), h^*) > \text{sim}(\mathcal{N}(I_p), h^*)$. As with the crossover operator, note that r is normalized after mutation.

We use a population size of 100 decreasing exponentially to 10 at a rate of 0.97. Our algorithm runs for 50 iterations, each time generating 20 children. Each child is equally likely to be formed as a product of a Crossover or Mutation operator, and all other parameters (α, m, p, q) are selected uniformly at random from the appropriate ranges.

Our database X consists of 250 images, 25 training images from each of the 10 ImageNette classes (FastAI, 2019). We test our algorithm on the hashes of 150 images, 15 validation images from each of the same 10 classes. To improve performance, our hashes are encoded as $\{-1, 1\}^{96}$ vectors rather than the standard $\{0, 1\}^{96}$.⁴

Note that the approximate linearity of NEURALHASH is crucial in enabling this attack, because it narrows the search space from all possible images (which is unimaginably large) to P , the set of interpolated images within a fixed dataset. This mitigates the combinatorial explosion and makes the problem tractable.

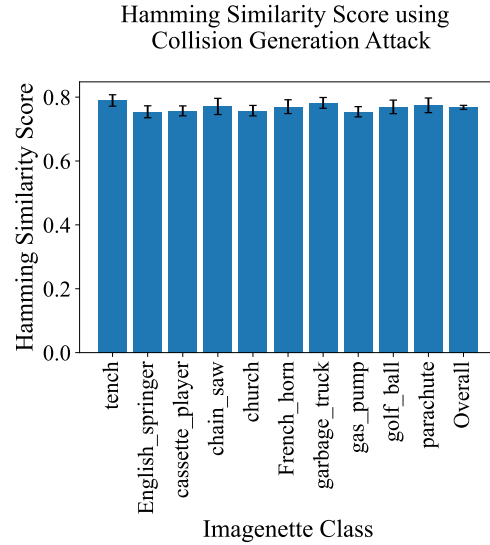


Figure 6. **Near Collision Generation with Genetic Algorithm.** Hamming similarity score of images generated by the genetic algo, averaged across 150 target hashes: 15 per class \times 10 classes.

Attack Results.

Our results are shown in Figure 6. While we are not directly able to find perfectly colliding images, the average

⁴Linear combinations of the value 0 are not meaningful.



Figure 7. **Near Collisions Created by the Genetic Algorithm.** The first row contains the source image and its hash (the target hash). The second row contains the interpolated image optimized by the genetic algorithm, its hash, and the Hamming similarity between the interpolated image’s hash and the target hash.

Hamming similarities of generated images are about *six* standard deviations above the mean of 0.5,⁵ indicating that the genetic algorithm performs much better than random generation and thus violates strong non-invertibility (Property 2.5). This conclusion is further reinforced in Section 5 where we perform the same attack after improving the security of the NEURALHASH system (see Figure 11).

Qualitatively, it is also interesting to see that the genetic algorithm’s generations often share semantic similarities in color and shape with the source image, from which the target hash is derived (Figure 7).

4.3. Information Extraction Attack

Recall the goal of information extraction from Section 4, where the goal is to predict the class of an image x from a set of n classes with probability greater than $\frac{1}{n} + \text{negl}$ given only its hash, h^* . We also assume access to a database of images $\{x_1, x_2, \dots, x_k\}$ with known hashes $\{h_1, h_2, \dots, h_k\}$ and known class labels $\{c_1, c_2, \dots, c_k\}$.

Attack Description. Let $\{p_1, p_2, \dots, p_k\} \in [-1, 1]^k$ be parameters such that $\sum_{i=1}^k |p_i| = 1$. Our attack is simple: first, minimize the following objective using gradient decent over parameters $\{p_1, p_2, \dots, p_k\}$

$$L(p_1, p_2, \dots, p_k, h^*) = \left(\sum_{i=1}^k p_i h_i - h^* \right)^2 - \sum_{i=1}^k p_i \log(p_i). \quad (10)$$

⁵Viewing hashes as a binomial RV, $\sigma = \sqrt{npq} \approx 0.05$.

Then, compute the class C with the most support:

$$\max_C \sum_{i \text{ s.t. } c_i = C} |p_i| \quad (11)$$

Note that this attack does not require model access except when labeling the dataset, as it only utilizes the hashes of the dataset images h_i , and class labels c_i .

To understand why this attack is reasonable, consider the objective function $L(\cdot)$ from Eqn. 10. Imagine if we could find a setting of parameters $\{p_1, p_2, \dots, p_k\}$ such that $\sum_{i=1}^k p_i h_i = h^*$. The interpolation property of the NEURALHASH model would imply that $\mathcal{N}\left(\sum_{i=1}^k p_i x_i\right) \approx h^*$. This motivates the first term of $L(\cdot)$.⁶ The second term simply aims to reduce entropy or, in other words, use the fewest number of parameters possible to achieve the objective.

Some additional information about our approach is as follows. Our database X consists of 500 images, 50 training images from each of the 10 ImageNette classes (FastAI, 2019). We tested our algorithm on the hashes of 1000 images, 100 validation images from each of the same 10 classes. Our algorithm runs for 25 epochs, 100 steps per epoch, and with a learning rate of $2 \cdot 10^{-5}$. To improve performance, our hashes are encoded as $\{-1, 1\}^{96}$ vectors

⁶One might wonder why we chose not to use a similar objective in Section 4.2 while generating near collisions. We find that, because the space of solutions to Eqn. 10 is fairly large, it is difficult to consistently generate reasonable-looking in-distribution images that resulted in hash collisions. However, we find that it *does* leak a significant amount of class information.

rather than the standard $\{0, 1\}^{96}$.




Figure 8. Classification Accuracy Shows Significant Information Leakage. Classification accuracy of ImageNette images from just their hash using an Interpolation Information Extraction Attack. Results averaged across 1000 target hashes: 100 per class over 10 classes.

Attack Results. Our results, seen in Figure 8 show that NEURALHASH hashes leak a significant amount of information about the semantic contents of images they are derived from. This is a clear violation of the Uniform Hashing Property, Property 2.3. It is interesting that classification accuracy (and therefore information leakage) seems to be quite class dependent. As one might expect, classes containing images with less clutter and clearer subjects tend to have more information leakage.

5. Improving NEURALHASH: SHA-at-the-End

All black-box interpolation attacks presented in Section 4 all rely on the approximate linearity of $\mathcal{N}(x)$. We aim to nullify this property.

Method Description. To address the undesirable linearity of NEURALHASH, we propose a simple yet effective extension: adding a SHA-256 block at the very end of the computation. Revisiting the security definitions from Section 2, we find that this strengthens the uniform hashing and noninvertibility properties while maintaining the neighborhood property: if $\mathcal{N}(x) = \mathcal{N}(x')$, then it must be true that $\text{SHA}(\mathcal{N}(x)) = \text{SHA}(\mathcal{N}(x'))$. Moreover, if $\mathcal{N}(x) \neq \mathcal{N}(x')$, then under the Random Oracle Model (ROM) their hashes will (i) look uniformly random and (ii) be non-invertible.

Method Results. First, we check whether the interpolation graph behaves as expected. As seen in Figure 9, the Ham-

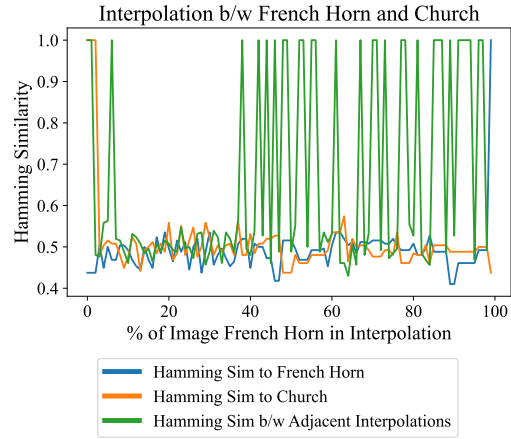


Figure 9. Interpolation Results using SHA-at-the-End. As expected, SHA breaks the linearity of NEURALHASH (compare to Figure 3).




Figure 10. Averaged Interpolation Under SHA. Here we demonstrate that the impact of SHA is systematic, and the violation of security properties is alleviated (compare to Figure 4).

ming similarities hover around 0.5 (except at the borders), and the similarity between adjacent interpolations oscillates between 1 and ≈ 0.5 . Figure 10 demonstrates that these trends are systematic. This suggests that SHA-at-the-end successfully breaks approximate linearity.

We further these conclusions experimentally by re-attempting the same interpolation attacks for collision generation (Section 4.2) and information extraction (Section 4.3) with the new system architecture.⁷ From Figure 11, we see that the ability to generate near-collisions is significantly hindered by the SHA block. The algorithm still consistently

⁷Note that we do not attempt to defend against the evasion attack (Section 4.1), as it will clearly work just as well. While the evasion attack does rely on the interpolation property, it is ultimately determined by NEURALHASH’s neighborhood property, which our proposed change has no impact on.




Figure 11. **SHA Hinders Collision Generation with Genetic Algorithm.** Hamming similarity score of images generated by the genetic algorithm averaged across 150 SHA-ed target hashes: 15 per class over 10 classes (compare to Figure 6).

produces images with Hamming similarities greater than 0.5, but this is reasonable; the expected similarity of the *maximum* of many randomly-sampled hashes is greater than the expectation of just one hash. From Figure 12, we also see that SHA-at-the-end completely disables information leakage, and the overall classification accuracy is approximately 10%, as expected. We conclude that SHA provides a simple defense mechanism against both these attacks.

6. Conclusion

In this paper, we have proposed and analyzed a variety of black-box attacks that expose vulnerabilities of Apple’s NEURALHASH algorithm. In particular, our black box attacks suggest that NEURALHASH is fundamentally flawed *aside* from the existence of gradient-based adversarial examples; this is a much stronger statement than the white-box case. At the core of these attacks is our discovery that NEURALHASH’s outputs change approximately linearly with respect to its inputs. Furthermore, we suggest a simple modification and demonstrate that it nullifies the piecewise linear property while successfully safeguarding against some of our proposed attacks.

This goes to show the immense difficulty of designing perceptual hashing algorithms that are both robustly effective and private. There remains much work to be done before such systems can be integrated into client-side scanning for real-world use.

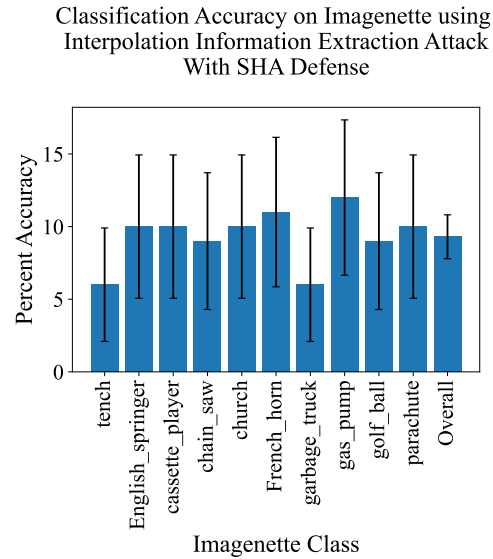


Figure 12. **SHA Disables Information Leakage.** Imagenette classification accuracy from SHA-ed hashes using an Interpolation Information Extraction Attack. Results averaged across 1000 target hashes: 100 per class over 10 classes (compare to Figure 8).

Acknowledgements. We thank Ronald Rivest, Yael Kalai, Andres Fabrega, Kyle Hogan, Nithya Attaluri, and Albert Xing for their valuable assistance and feedback.

References

- Abelson, H., Anderson, R., 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., and Troncoso, C. Bugs in our pockets: The risks of client-side scanning, 2021.
- Apple. Child sexual assault material detection, 2021. URL https://www.apple.com/child-safety/pdf/CSAM_Detection_Technical_Summary.pdf.
- Athalye, A. Neural hash collider, 2021. URL <https://github.com/anishathalye/neural-hash-collider>.
- Bhowmick, A., Boneh, D., and Myers, S. The apple psi system, 2021. URL https://www.apple.com/child-safety/pdf/Apple_PSI_System_Security_Protocol_and_Analysis.pdf.
- Chum, O., Philbin, J., and Zisserman, A. Near duplicate image detection: min-hash and tf-idf weighting. *BMVC 2008 - Proceedings of the British Machine Vision Conference 2008*, 01 2008. doi: 10.5244/C.22.50.

- Di Martino, F. and Sessa, S. Image matching by using fuzzy transforms. *Advances in Fuzzy Systems*, 2013, 2013. doi: 10.1155/2013/760704.
- Dwyer, B. Imagenet contains naturally occurring neuralhash collisions, 2021. URL <https://blog.roboflow.com/neuralhash-collision/>.
- Facebook. The pdq perceptual hash, 2020. URL <https://github.com/facebook/ThreatExchange/blob/main/hashing/hashing.pdf>.
- FastAI. Imagenette: A subset of imagenet, 2019. URL <https://github.com/fastai/imagenette>.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017.
- Ilyas, A., Santurkar, S., Tsipras, D., Engstrom, L., Tran, B., and Madry, A. Adversarial examples are not bugs, they are features. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/e2c420d928d4bf8ce0ff2ec19b371514-Paper.pdf>.
- Microsoft. Photodna, 2022. URL <https://www.microsoft.com/en-us/photodna>.
- Roboflow. Neuralhash collisions, 2021. URL <https://github.com/roboflow-ai/neuralhash-collisions>.
- Struppek, L., Hintersdorf, D., Neider, D., and Kersting, K. Learning to break deep perceptual hashing: The use case neuralhash. *CoRR*, abs/2111.06628, 2021. URL <https://arxiv.org/abs/2111.06628>.
- Whitley, D. A genetic algorithm tutorial. *Statistics and computing*, 4(2):65–85, 1994.
- Ygvar, A. Apple neural hash to onnx, 2021. URL <https://github.com/AsuharietYgvar/AppleNeuralHash2ONNX>.