# Fuzzy Hashing as Perturbation-Consistent Adversarial Kernel Embedding

**Ari Azarafrooz** and **John Brock**
Cylance, Department of Research and Intelligence.
Irvine, CA, USA

## Abstract

Measuring the similarity of two files is an important task in malware analysis, with fuzzy hash functions being a popular approach. Traditional fuzzy hash functions are data agnostic: they do not learn from a particular dataset how to determine similarity; their behavior is fixed across all datasets. In this paper, we demonstrate that fuzzy hash functions can be learned in a novel minimax training framework, and that these learned fuzzy hash functions outperform traditional fuzzy hash functions at the file similarity task for Portable Executable files. In our approach, hash digests can be extracted from the kernel embeddings of two kernel networks, trained in a minimax framework, where the roles of players during training (i.e adversary versus generator) alternate along with the input data. We refer to this new minimax architecture as *perturbation-consistent*. The similarity score for a pair of files is the utility of the minimax game in equilibrium. Our experiments show that learned fuzzy hash functions generalize well, capable of determining that two files are similar even when one of those files was generated using *insertion* and *deletion* operations.

## Introduction

### File Similarity in Malware Analysis

The rapid proliferation of malware poses a substantial set of technical and organizational challenges for malware analysts. For example, for the year 2017, there were more than 121.6 million instances of new malware recorded, i.e., on average there were about 3.9 new malware instances being created every second (AV-TEST 2018). While these numbers may seem large, it is important to note that most new malware is not written from scratch. Instead, most new malware files are only slightly modified versions of previously known malware. New malware files will often be near-exact duplicates of previously seen malware. There are several reasons for this near-duplication, including malware authors making small changes to bypass new malware defenses, as well as code reuse and code sharing among malware authors (Walenstein and Lakhotia 2007).

Given some unknown file, two important questions a malware analyst might ask are: Is this new file benign or malicious; and, if the file is malicious, what kind of malware is

---

AAAI-19 Workshop on Artificial Intelligence for Cyber Security.

it (e.g., ransomware, Trojan, credential stealer, etc.)? Malware analysts can significantly reduce the burden of answering these two questions by exploiting the fact that most new malware is very similar to previously known malware. If we already had a database of known malware, we could essentially do a nearest neighbor search using our file similarity metric to query for known files that are similar to our new unknown file. If, for example, the most similar files in the database are malicious, we may assume that the new file is also malicious. And if many of the similar files in the database are a particular type of malware, we may also assume that the new file is that same type of malware.

### Fuzzy Hashing

In the computer security community, file similarity is commonly measured using fuzzy hashing functions. A typical fuzzy hashing algorithm has two parts: (1) a hashing algorithm to generate a hash digest (e.g., the hash value representing a file), and (2) a comparison algorithm to compute a similarity metric from a pair of input hash digests (Bass et al. 2012).

Fuzzy hashing algorithms differ from cryptographic hash functions (e.g. MD5, SHA256, etc.) in that the output of a fuzzy hash function is insensitive to small changes in the input. For example, changing a single bit in the input to SHA256 will change the output drastically. However, with typical fuzzy hash functions, changing a single bit will not affect the output. Ideally, fuzzy hash functions should be as robust as possible to adversarial and random perturbations.

A well-known fuzzy hash function is ssdeep (Kornblum 2006). Given a file, ssdeep first breaks it up into several pieces, and then uses a rolling hash function to compute the digest for each piece. These individual digests are concatenated to produce a final *fixed-size* similarity digest for the whole file. Given two hash digests, ssdeep will return a score of between 0 and 100 that expresses a level of confidence that two files are similar.

Despite ssdeep's popularity, it has several known weaknesses. For example, ssdeep is sensitive to random bit flips made across an entire file. In fact, if more than 70 bytes of a file are modified, ssdeep will always return a similarity score of zero when given the original and modified files as inputs (Breitinger 2011).

sdhash (Roussev 2010) is another widely used fuzzy hash

function with increased robustness to random modifications of the input compared to ssdeep: sdhash can measure that a file is similar to a modified version of that file even when up to 1.1% of the file has been modified at random. Unlike ssdeep, sdhash generates varying length digests (2-3% of the length of the input). sdhash uses normalized Shannon entropy to extract 64-byte sequence features. The extracted features are hashed and then put into a Bloom filter (Bloom 1970). For the comparison of two files, Bloom filters are compared using a Hamming distance measure. The result is the estimated fraction of Bloom filter features that the two filters have in common that are not due to chance.

The traditional fuzzy hash function that is most similar to our work is sahash (Breitinger 2014). The simplistic design of sahash allows for its similarity score to be a lower bound on the Levenshtein distance. However, sahash is constrained to only fixed-sized inputs and cannot measure the similarity digests of files with different sizes. sahash works by concatenating the following sub-hashes: 1) the hash of the frequency of byte-arrays of files, 2) the hash of the frequency of 4-bit circularly shifted byte-arrays of files 3) a so-called `uneva` function. `uneva` not only counts the frequency of occurrences in the byte-arrays, but it also counts the frequency of the repeated occurrences, as a measure of how sporadic the occurrences of a specific byte are. We refer the reader to (Breitinger et. al 2014) for a more detailed description. Since we train our model using sahash features, we describe sahash features more in the feature preparation section.

We don't intend to provide a panacea to the fuzzy hash problem. The main novelty of our work is to propose a *minimax training* framework, where fuzzy hash functions can be *learned*. We show that learned fuzzy hash functions are more robust than popular traditional fuzzy hash functions to modifications of the inputs. We hypothesize that this robustness is due to traditional fuzzy hash functions being data agnostic, i.e., they don't learn from data, whereas our fuzzy hashing function can exploit learned statistical relationships in the underlying dataset. We also show that the tolerance of the fuzzy hashing function to file modifications can be implicitly learned from the corpus of data on which the fuzzy hashing function is trained.

Note that we focus on fixed-size digests to have a *differentiable end-to-end* mechanism for learning. This choice of fixed-sized digest helps us to do a fair comparison between ssdeep, sahash, and our method.

For our dataset, we start with a collection of binary files, and then modify the bytes of these files at random using a fuzzer, generating a new collection of perturbed files. Note that, while we use the term *perturbed* in our paper to refer to these fuzzed examples in the training data, the term 'perturbed' is relative, and the original files can be also considered to be the perturbed versions of the fuzzed files. We'll describe the dataset later on in the paper in our Experiments section.

Note that while we focus on applications of fuzzy hashing to file similarity for the purposes of malware analysis, fuzzy hashing can be applied to many kinds of data, including images and text.

## Fuzzy Hash Functions as Kernel Networks

We view fuzzy hash functions as kernel networks. Inspired by (Goodfellow 2014) (Bo 2014) (Azarafrooz 2017), we incorporate kernel networks into a novel minimax framework. Our proposed minimax framework shares some similarities with cycle-consistent minimax frameworks (Zhu 2017) (Yi 2017). However, our approach is different in two ways: 1) The players of minimax games are kernel networks rather than neural architectures, and 2) Consistency is expressed in terms of the roles of the kernel networks w.r.t to the perturbed data. We refer to this as 'perturbation consistency' because networks alternate their input data in each stage of learning, and whichever network uses the perturbed data as input will play the role of an adversary. We found such consistency to be crucial for satisfying the end-to-end symmetrical properties of a learned fuzzy hash function (we explain this further in the Learning Fuzzy Functions section).

## Feature Preparation

Our proposed methodology works at the byte level and, therefore, data preparation requires no special parser. Given a file with byte-array `f`, we generate a perturbed file with byte-array `f'` using random $\rho$-bit *substitutions* s.t. $0 < $ `f` $\oplus$ `f'` $< \rho$, without changing the physical size or the run-time behavior of files, where $\oplus$ is `XOR` operator. This is not to say that a file is randomly selected from Hamming sphere of radius $\rho$, since some points in the Hamming sphere do *not* remain truthful to the functionality of `f`. In other words, we avoid flipping bits that are essential to the functionality of the underlying executable files. The choice of using a bit substitution operation is to guarantee that the functionality of the perturbed version remains intact. Enabling perturbations with other edit operators, *deletion* and *insertion*, requires special care to avoid changing an executable file in such a way that it is no longer able to be executed. As explained later, we prepare test data in which files are fuzzed using deletion and insertion operations.

Similar to sahash, for each byte-array `f`, we prepare features `x` by concatenating the histogram of byte counts with the histogram of byte counts after a circular shift:

$$x = \text{hist}(f) \mid \text{hist}(f >> 4) \qquad (1)$$

, where `f >> n` denotes `n`-bit circular shift to the right, `hist(f)` is the histogram of byte counts and `|` is the concatenation operator. Since each byte can be represented as a value between `0` and `255`, `x` is a `512` dimensional vector. For example, the zero-based indexes of `255` and `511` represent the counted numbers of byte occurrences of `0xff` in the original bytes and circularly shifted bytes, respectively. A more aggressive feature selection is also possible, where histograms of byte counts for other shift values `n : n` mod $8 \neq 0$ are incorporated into the features. Similar feature selection is done for the perturbed byte-array `f'` to get `x'`. The pairs $(x, x') \in X, X'$ are then used to learn fuzzy hash functions.

## Learning Fuzzy Functions

sahash (Breitinger et. al 2014) digest gets calculated as:

$$\text{x} \mod \text{m} \quad \text{where} \quad \text{m} = 2^{\max(8, \lceil \frac{\log_2 1}{2} \rceil)} \tag{2}$$

, where $1$ is the length of $\text{f}$. The similarity measure for the sahash digests is then related to the maximum number of byte frequencies that are different among both the circulated and the original byte-array. This simply provides a lower bound for the edit distance between the files. Note that edit distance between two files is defined as the minimum number of substitutions, insertions, and deletions that are required to transform one file to another. The value of modulus $\text{m}$ is selected to be file-size dependent for a suitable balance between compressibility and robustness. This, however limits its applicability to different file sizes. Instead of file-size dependent modulo operations, we learn a high-order non-linear function from the training data pairs $\text{X}, \text{X}'$. By doing so, similarity measure between digests finds intuitive mathematical interpretation in terms of maximum mean discrepancy (MMD) measure (Gretton 2012).

To incorporate learning mechanisms in fuzzy hash functions, we need to extend the fuzzy hash function properties to include the following:

- *Differentiability*: The desire to learn from the structure of the data $\text{X}, \text{X}' = (\text{x}_1, \text{x}'_1), (\text{x}_2, \text{x}'_2), ..., (\text{x}_n, \text{x}'_n)$ asks for a differentiable network. For example, we cannot use the modulo operators used by sahash.

- *End-to-end symmetrical property*: We use a minimax framework with two different functions $(G, D)$ to compute the similarity measure $\delta$. This, however, requires an end-to-end symmetrical behavior, formalized below:

$$\delta(G(\text{x}), D(\text{x}')) = \delta(G(\text{x}'), D(\text{x})) \tag{3}$$

Note this is more specific than the symmetrical properties of similarity/distance measures. The MMD similarity measure used in our paper is symmetric by default.

To achieve the introduced properties above and those introduced in the introduction section, we model the fuzzy hash function selection using kernel minimax frameworks. This leads to differentiable non-linear embeddings. We view these embeddings as fuzzy hash digests which turn out to be robust against perturbations, thanks to the qualities of MMD distance measure and the minimax training. In order to achieve the end-to-end symmetrical property, a perturbation-consistent version of the minimax framework is presented.

## Kernel embedding in a Minimax Framework

Let $(\text{h}, \text{h}')$ denotes the digest of pairs $(\text{x}, \text{x}')$ with similarity measures of $\delta(\text{h}, \text{h}')$. We view the digests in our framework as the kernel embeddings of $(\text{x}, \text{x}')$. This view establishes the connection between the required similarity measures of the fuzzy hash functions and maximum mean discrepancy (MMD) measures (albeit a stochastic version). The key idea behind MMD measure is the fact that one can measure the difference between distributions $p(\text{X})$ and $p(\text{X}')$ by linear witness in Hilbert space such as:

$$\sup_{f \in \mathcal{H}} \frac{\mathbb{E}[\psi(\text{X})] - \mathbb{E}[\psi(\text{X}')]}{\delta(\text{h}, \text{h}')} \tag{4}$$

, without using density estimators over the original domains $(\text{X}, \text{X}')$, where $\psi$ is a function living in a *reproducing kernel Hilbert space* (RKHS) $H$. Conveniently, in the process, we get the kernel embedding as the digest of the fuzzy hash functions. This provides digests and similarity measures in one shot. We further achieve *learnable parameters* and robustness by incorporating the kernel networks into minimax frameworks followed by a simple mathematical trick introduced in (Bo 2014) (Azarafrooz 2017): take the derivative of Eq. 4 to get gradients and then reconstruct a new network with fresh learnable parameters.

Before we continue with demonstrating the derivation of the new network with learnable parameters, we have to touch on a few definitions and one main theorem.

**Definition 1** (kernel mean embedding). Due to the *reproducing property* of $H$, the expectation of any function $\psi$ in RKHS $H$ with respect to random variable $\text{X}$ can be computed as an inner product with its so called *kernel mean embedding* $\mathbb{E}[k(X, .)]$:

$$\mathbb{E}[\psi(\text{X})] = \langle \psi, \mathbb{E}[k(\text{X}, .)] \rangle \tag{5}$$

where $k$ is the kernel associated with RKHS $H$.

**Definition 2** A kernel $k(\text{x}, \text{x}') : \text{X} \times \text{X} \to \mathbb{R}$ is positive definite (PD), when for all $n > 1$ and $x_1, x_2, .., x_n \in \text{X}$ and $c_1, ..., c_n \in \mathbb{R}$, we have $\sum_{i,j} c_i c_j k(\text{x}_i, \text{x}_j) \geq 0$.

**Theorem 1** *Duality between Kernels and Random Processes* (Devinatz 1953): If $k(\text{x}, \text{x}')$ is a positive definite kernel, then there exits a set $\Omega$, a measure $\mathbb{P}$ on $\Omega$, and random function $\phi_{\text{w}}(\text{x}) : \text{X} \to \mathbb{R}$ from $L_2(\Omega, \mathbb{P})$, such that $k(\text{x}, \text{x}') = \int_\Omega \phi_{\text{w}}(x) \phi_{\text{w}}(x') d\mathbb{P}(\text{w})$.

**Reconstructing a minimax network** Using Eq. 4 and Eq. 5, we note that gradient of the similarity measure is:

$$\frac{\partial \delta}{\partial \psi} = \mathbb{E}[k(\text{X}, .)] - \mathbb{E}[k(\text{X}', .)] \tag{6}$$

Applying Theorem 1 to Eq. 6, followed by Monte Carlo approximation analogous to random Laplace feature maps (Rahimi 2007), we arrive at the following gradient terms:

$$\tilde{\phi}(\text{x}) - \tilde{\phi}(\text{x}'), \quad \tilde{\phi}(\text{x}) \equiv \frac{\exp(-\text{x}^T \text{w}_1)...exp(-\text{x}^T \text{w}_\text{s})}{\text{s}} \tag{7}$$

where $\text{w}_j$ are drawn from a proper distribution. In our paper, it is drawn from Gaussian $\mathbb{P}(\text{w}) = \frac{\exp(-\|\text{w}\|_2^2/2)}{(2\pi)^{d/2}}$. We also investigated Laplace features maps (more on this can be found in the Appendix).

Now we are in the position to construct a new parametrized learnable framework by the linear combination of the first-order gradient terms in Eq. 7. To implement this, we treat the gradient terms as the activation of the first layer in our network (note the non-linearity of the kernel embedding) followed by one more layer with softmax activation. We use batch normalization immediately after the first layer followed by dropout during training. This completes the design description for both involved networks in the minimax training framework, the generator $G$ and adversary $D$.

In the well-known minimax process used by GAN (Goodfellow 2014), a generator tries to sample from a desired

probability space, given random noise as input. An adversary then evaluates the generated sample by comparing it against the real data, according to a suitable distance measure, such as Jenson-Shannon, Wasserstein, MMD, etc. However, in our case, the inputs for both generator and adversary belong to the space of training data $X, X'$. Parameters $\theta$ and $\phi$ get adjusted by the gradient flowing through both networks, in every stage of the game.
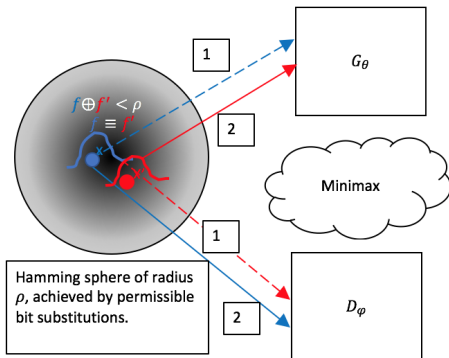


Figure 1: Scheme of the proposed perturbation-consistent minimax training: Every stage of learning includes 2 rounds. Round $\boxed{1}$: $G_\theta$ uses $\texttt{x}$ as input data, $D_\phi$ uses $\texttt{x'}$. $G_\theta$ plays

the role of the generator and $D_\phi$ plays the role of adversary. Round $\boxed{2}$: $D_\phi$ uses $\texttt{x}$ as input data, $G_\theta$ uses $\texttt{x'}$. $D_\phi$ plays

the role of the generator and $G_\theta$ plays the role of adversary. Both $G_\theta$ and $D_\phi$ backpropagate gradients to adjust their parameters, $\theta$ and $\phi$. Learning continues similarly with the next batch of data until convergence is achieved.

$$\min_{\theta \in \Theta} \delta(G_\theta(\texttt{x}), D_\phi(\texttt{x'})) \tag{8}$$

If the minimax characteristic of the framework is not clear from Eq. 8 please refer to the definition of MMD distance in Eq. 4 (find sup/max in Eq. 4 and min in Eq. 8). The embedded features of either $G_\theta$ or $D_\phi$ can be considered the hash digest of resulting fuzzy hash function, with the similarity measure being the utility of the minimax game in *equilibrium*.

To induce end-to-end symmetry, we train in a *perturbation-consistent* way, where the role of players in the minimax training (i.e discriminator versus generator) *alternates* along with their input data. In other words, during each state of training, we alternate input data to both networks, and whoever has access to the perturbed data $X'$ in that round, plays the role of an adversary. This is visually illustrated in Fig. 1. We refer to our proposed method as *adversarial kernel hash (akash)*.

## Experiments

**Data Preparation** To show that learning from data can result in an improved fuzzy hash function relative to traditional

fuzzy hashing, we base our study on a set of Portable Executable (PE) files. The PE file format describes a data structure that encapsulates the information necessary for the Windows OS loader to manage the wrapped executable code. We curated a training set of $\texttt{80}$ thousand malicious PEs, fuzzed using the aforementioned $\rho$-bit substitutions mechanism with $\rho$ selected uniformly random from range $\texttt{1}$ to $\texttt{500}$. We generate two types of test data. One set of test data is generated using the explained flip substitution, while the other set is generated using LIEF (Quarkslab 2017). LIEF (Library to Instrument Executable Formats) is an open source library capable of parsing and manipulating PE files using advanced edit operations.

We avoided using LIEF for preparing the training data set for two reasons: 1) In order to show that even training over the Hamming sphere of permissible bit-substitutions generalizes to deletion and insertion edit operations. 2) Advanced edit operations using tools like LIEF can still sometimes generate a fuzzed PE file that is corrupted in such a way as to be unexecutable. However, our bit substitution mechanism is guaranteed to keep the functionality of the PE intact.

**Hyperparameter tuning** The best results were achieved by setting batch size to $\texttt{1000}$, dropout to $\texttt{0.75}$ and learning rates to $\texttt{5e-4}$, for both generator and adversary networks. We found our experiments to be sensitive to the batch size. Also, we set the embedding size to $\texttt{512}$, without performing any tuning, in order to keep the digest size consistent with sahash. For optimization, we used (adaptive moment estimation) Adam stochastic optimization. We trained the model for $\texttt{5000}$ epochs.

**Thresholds** sahash uses the absolute value of the difference between the $\texttt{uneva}$ of the files in order to reduce false positives. While it is possible to incorporate the $\texttt{uneva}$ arrays into input data $\texttt{x}$, we chose not to do so, in order to have a fair comparison between akash and sshash. Similar to sahash, we call two files similar if they fall below a certain threshold for $\delta$ (Eq. 4) *and* above a threshold for $\texttt{uneva}$. While sahash selected the threshold to optimal value of $\texttt{97}$, we achieved zero false positives by smaller value of threshold $\texttt{80}$.

**Visual Comparisons of Robustness** To compare the robustness of akash with other fuzzy hash techniques, we plotted the effect of various types of perturbations on the distance measures in Fig. 2 and Fig. 3. In all these plots, the $\texttt{y}$ axis in ssdeep and sdhash are ($\texttt{100}$-confidence score of finding files similar) and ($\texttt{100}$-estimate of the fraction of common features between files) respectively, as a proxy for their distance measures. The $\texttt{y}$ axis in sahash provides a lower bound estimate on the edit distance. The $\texttt{y}$ axis in akash is representative of distance measure $\delta$, described in Eq. 4 as the utility of the minimax game at equilibrium. The $\texttt{x}$ axis in Fig. 2 represents the number of bit-substitutions performed in randomly selected files. However, the modified file has the same physical size and functionality as the unmodified file. Disregarding the plot scales, the large spikes in the figures correspond to the events where fuzzy hash functions are being bypassed, i.e., where the fuzzy hash functions fail to consider two similar files to be similar. Lower numbers of spikes implies more robustness. It can be observed that

akash is consistently more robust than the other fuzzy hash functions. Despite training on only $\rho$-bit substitutions within $0 < \rho < 500$, Fig. 2 demonstrates that akash's performance generalizes to far more than 500 modifications.
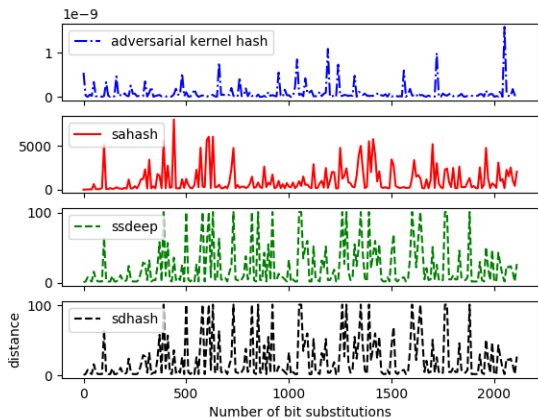


Figure 2: The effect of increase in the number of (permissible) bits substitutions on the distance measures. The smoother the trends, the more robust the fuzzy hasher.
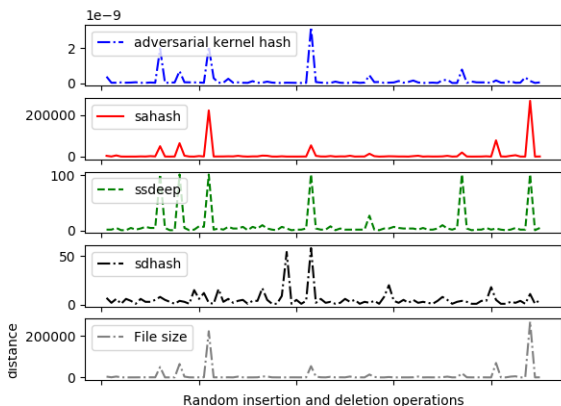


Figure 3: The effect of random insertion/deletion operations on the distance measures. The smoother the trends, the more robust the fuzzy hasher. Notably, training over bit substitutions results in a fuzzy hasher that generalizes well to insertions and deletions.

More interestingly, Fig. 3 shows that akash generalizes well to even more advanced edit operations such as insertions and deletions. The x axis in Fig. 3 represent the events with random action selected from the insertion and deletion actions available to LIEF. The insertion operations include the following LIEF actions: 'append overlay', 'append import, 'add section' and 'append section'. The deletion operations include the following LIEF actions: 'remove the signature' and 'remove debug'.

Table 1: True Negative Ratio: **bit-substitions**

| Fuzzy Hash | akash | ssdeep | sahash | sdhash |
|---|---|---|---|---|
| TN % | 97.9 | 86.3 | 31 | 93.4 |

Table 2: True Negative Ratio: **Insertion and deletion**

| Fuzzy Hash | akash | ssdeep | sahash | sdhash |
|---|---|---|---|---|
| Overall % | 85 | 96.3 | 28 | 90.1 |
| Insertion % | 88 | 96 | 42 | 93.4 |
| Deletion % | 75 | 99 | 36 | 77 |

We also plotted the byte size of the modifications associated with the performed edit actions. Fig. 3 shows that akash is robust even to the cases where files are generated via deletion and insertion operators rather than bit substitutions.

**Performance** We note that sdhash (unlike akash, ssdeep, and sahash) generates varying length digests (2-3% of the length of the input). Therefore, a fair comparison between sdhash and the rest of the mentioned fuzzy hashers is not possible, but we report the sdhash results to be comprehensive. Table 1 and Table 2 shows the result of comparison over a randomly selected set of 10 thousand benign and malicious hashes. We use the recommended threshold of `21` for sdhash, `50` for ssdeep and the exact thresholds of sahash reported in (Breitinger et. al 2014). *Table 1 shows that akash outperforms the traditional fuzzy hashing functions, when perturbations are induced using bit-substitution operations*. Table 2 shows that, despite being trained on the space of permissible bit-substitutions, akash generalizes well even to the cases where perturbations are induced using insertions and deletions.

We also experimented with more advanced edit operation such as compression. Specifically, we used the UPX pack/unpack compression action in LIEF when conducting this new experiment, and found zero detections in this case, as is reported in Table 3. This result is expected, since compression changes the structure of the original files in much more substantial way than substitution, insertion, and deletion operations.

**False Positives** In order to measure the false positive rate, we selected a random set of `n=10000` malicious hashes. We then selected a random selection of `10000` pairs from the possible set of combinations `n × n-1/2`. The results are reported in Table 4. We performed similar experiments for a set of benign hashes. The numbers are consistent with Table 4 as well. We found the true positive ratio for both sahash (using the same exact settings in the original paper) and ssdeep to be the same: 98.3%, which is lower than the perfect

Table 3: True Negative Ratio for Compression Operations

| Fuzzy Hash | akash | ssdeep | sahash | sdhash |
|---|---|---|---|---|
| UPX % | 0 | 20 | 20 | 14.2 |

Table 4: True Positive Ratio

| Fuzzy Hash | akash | ssdeep | sahash | sdhash |
|------------|-------|--------|--------|--------|
| TP % | 100 | 98.3 | 98.3 | 100 |

100% performance of akash and sdhash. Results in Table 4 shows that akash performs as good as sdhash, in terms of ratio of false positives, despite having fixed-size digest.

## Conclusion and Future Work

We introduced a novel minimax training framework capable of learning to perform fuzzy hashing for the purposes of file similarity. The kernel embeddings of the minimax players can be used as hash digests, and the utility of the minimax game in equilibrium provides the similarity score. As shown in our experiments, the trained fuzzy hash function outperforms widely used traditional fuzzy hash functions when perturbations are induced using bit substitutions. Also, despite being trained on the space of permissible bit substitutions, the trained fuzzy hash function generalizes well even to the cases where perturbations are induced using insertions and deletions.

We ran our experiments with PE files, but we suspect our approach can be fruitfully applied to other kinds of data, such as images or audio. For example, one could modify a collection of images using cropping, color changes, and geometric transformations, and then use our technique to measure the similarity of the modified and original images.

For files in particular, an interesting future direction is to investigate to what extent a trained model is generalizable to perturbations involving more advanced edit operations, such as compression, instead of simple bit substitutions.

## References

AV-TEST Security Report 2017/2018. 2018. Magdeburg, Germany. AV-TEST GmbH.

Azarafrooz, M. 2017. Doubly Stochastic Adversarial Autoencoder. *Second Workshop on Bayesian Deep Learning, NIPS 2017*: Long Beach, Calif.

Baier, H. and Breitinger, F. 2011. Security Aspects of Piecewise Hashing in Computer Forensics. In *Proceedings of the Sixth International Conference on IT Security Incident Management and IT Forensics*, 21-36. IEEE.

Bass, L., Brown, N., Cahill, G., Casey, W., Chaki, S., Cohen, C., Niz, D. de, French, D., Gurfinkel, A., Kazman, R., Morris, E., Myers, B., Nichols, W., Nord, R.L., Ozkaya, I., Sangwan, R.S., Simanta, S., Strichman, O., Valetto, P., 2012. Results of SEI Line-Funded Exploratory New Starts Projects, Technical Report. CMU/SEI-2012-TR-004, ESC-TR-2012-004, Software Engineering Institute, Carnegie Mellon Univ., Pittsburgh, Penn.

Bloom, B.H.. 1970. Space/Time Trade-Offs in Hash Coding with Allowable Errors. *Communications of the ACM* 13(7): 422-426.

Breitinger, F., Ziroff G., Lange, S., Baier, H. 2014. Similarity Hashing Based on Levenshtein Distances. In *Advances in Digital Forensics X. DigitalForensics 2014. IFIP Advances in Information and Communication Technology*, vol. 433. Berlin, Heidelberg: Springer.

Dai, B., Xie, B., He, N., Liang, Y., Raj, A., Balcan, M.F.F. and Song, L.. 2014. Scalable Kernel Methods via Doubly Stochastic Gradients. In *Advances in Neural Information Processing Systems* 27: 3041-3049.

Devinatz, A. 1953. Integral Representation of Positive Definite Functions. *Transactions of the American Mathematical Society* 74(1) 56-77.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*: 2672-2680.

Gretton, A., Borgwardt, K.M., Rasch, M.J., Scholkopf, B. and Smola, A.. 2012. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13 (March 2012): 723-773.

Kornblum, J.. 2006. Identifying Almost Identical Files Using Context Triggered Piecewise Hashing. *Digital Investigation* 3: 91-97.

Quarkslab. 2017-2018. LIEF: Library for Instrumenting Executable Files. https://lief.quarkslab.com. Last accessed: Sept. 5, 2018.

Rahimi, A. and Recht, B.. 2008. Random Features for Large-scale Kernel Machines. In *Advances in Neural Information Processing Systems*: 1177-1184.

Roussev, V. Data Fingerprinting with Similarity Digests. In *Advances in Digital Forensics VI*, K. Chow and S. Shenoi, Eds., *Springer*, Heidelberg, Germany, 207–226, 2010.

Walenstein, A., and Lakhotia, A. 2007. The software similarity problem in malware analysis. In *Dagstuhl Seminar Proceedings*. Dagstuhl, Germany: Schloss Dagstuhl-Leibniz-Zentrum for Informatik.

Yi, Z., Hao (Richard) Zhang, Tan, P. and Gong, M.. 2017. Unsupervised Dual Learning for Image-to-Image Translation. In *International Conference on Computer Vision (ICCV)*.

Zhu, J.Y., Park, T., Isola, P. and Efros, A.A.. 2017. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *International Conference on Computer Vision (ICCV)*.

## Appendix

We also experimented with Laplace feature maps with samples drawn Exponential and Levi. But we found the Fourier features maps to perform much better than Laplace feature maps. The Laplace feature maps led to fuzzy hashers that have an impractical FP ratio. This is an interesting observation since semigroup kernels associated with Laplace feature maps provide a much more intuitive interpretation than Fourier Features maps. This is because one can consider histograms X as groups closed over the + operation. Since byte counts can only get positive values. We did however, find that the convergence for Laplace Feature maps with samples drawn from exponential to be much smoother.

**Definition 3** (Abelian semigroup). A semigroup $(X, \circ)$ is a nonempty set $X$ equipped with an associative composition $\circ$, i.e for any $x, x', x'' \in X : x \circ (x' \circ x'') = (x \circ x') \circ x''$ and a identity element $e$, i.e, for any $X \in X : x \circ e = x$. For an abelian semigropu, the composition is commutative, i.e. for any $x, x' \in X : x \circ x' = x' \circ x$.

**Definition 4** (Kernels on Abelian semigroups). A function $k : X \times X \rightarrow \mathbb{R}$ is a positive definite (PD) kernel function on an abelian semigroup $(X, \circ)$ if $k(x, x') = \Phi(x \circ x')$ where $\Phi : x \rightarrow \mathbb{R}$ is a PD function, i.e for any $x_1, ..., x_n \in X$, any real-valued scalars $c_1, ..., c_n$, the following holds: $\sum_{i,j=1}^{n} c_i c_j \Phi(x_i, x_j) \geq 0$.

The following theorem establishes a one-to-one correspondence between semigroup kernels and probability densities on $\mathbb{R}^d$, via the Laplace transform.

**Theorem 1** (Duality between semigroup kernels and random processes): For every bounded continuous kernel function $k(x, x')$ on the Abelian Semigroup $(\mathbb{R}_+^d, +)$ there exits a non-negative measure $w$ with associated *random Laplace* construction $\phi_W$ s.t that $k(x, x') = \int_{\mathbb{R}_+^D} \phi_{\mathcal{W}}(x) \phi_{\mathcal{W}}(x') d\mathbb{P}(\mathcal{W})$.