



# ENCOD: Distinguishing Compressed and Encrypted File Fragments

Fabio De Gaspari<sup>1</sup>, Dorjan Hitaj<sup>1(✉)</sup>, Giulio Pagnotta<sup>1</sup>, Lorenzo De Carli<sup>2</sup>,  
and Luigi V. Mancini<sup>1</sup>

<sup>1</sup> Dipartimento di Informatica, Sapienza Università di Roma, Rome, Italy  
{degaspari,hitaj.d,pagnotta,mancini}@di.uniroma1.it

<sup>2</sup> Department of Computer Science, Worcester Polytechnic Institute,  
Worcester, MA, USA  
ldecarli@wpi.edu

**Abstract.** Reliable identification of encrypted file fragments is a requirement for several security applications, including ransomware detection, digital forensics, and traffic analysis. A popular approach consists of estimating high entropy as a proxy for randomness. However, many modern content types (e.g. office documents, media files, etc.) are highly compressed for storage and transmission efficiency. Compression algorithms also output high-entropy data, thus reducing the accuracy of entropy-based encryption detectors.

Over the years, a variety of approaches have been proposed to distinguish encrypted file fragments from high-entropy compressed fragments. However, these approaches are typically only evaluated over a few, selected data types and fragment sizes, which makes a fair assessment of their practical applicability impossible. This paper aims to close this gap by comparing existing statistical tests on a large, standardized dataset. Our results show that current approaches cannot reliably tell apart encryption and compression, *even for large fragment sizes*. To address this issue, we design ENCOD, a learning-based classifier which can reliably distinguish compressed and encrypted data, starting with fragments as small as 512 bytes. We evaluate ENCOD against current approaches over a large dataset of different data types, showing that it outperforms current state-of-the-art for most considered fragment sizes and data types.

## 1 Introduction

Reliable detection of encrypted data fragments is an important primitive with many applications to security and digital forensics. For instance, ransomware detection algorithms use estimates of write-operations' data randomness to quickly identify evidence of malicious encryption processes [17, 25, 26, 33]. When performing digital forensic analysis of hard drives and phones, it is oftentimes important to identify encrypted archives [16]. Finally, encryption detection is widely used in network protocol analysis [18, 20].

A popular approach to address this problem is to estimate the Shannon entropy of the sequence of interest using the Maximum Likelihood Estimator (MLE):  $\hat{H}_{MLE}$ . This approach leverages the observation that the distribution of byte values in an encrypted stream closely follows a uniform distribution; therefore, high entropy is used as a proxy for randomness. This estimator has the advantage of being simple and computationally efficient. As non-encrypted digital data is assumed to have low byte-level entropy, the estimator is expected to easily differentiate non-encrypted and encrypted content.

While this approach remains widely used (e.g., [17,25,26,33]), a number of works have highlighted its limitations. Modern applications tend to compress data prior to both storage and transmission. Popular examples include the zip compressed file format, and HTTP compression [21] (both using the DEFLATE algorithm). As compression removes recurring patterns in data, compressed streams tend to exhibit high Shannon entropy. As a result, compressed data exhibit values of  $\hat{H}_{MLE}$  that are close and oftentimes overlapping with those obtained by encryption. In principle, compressed content can be identified by using appropriate parsers. However, many security-related applications, such as ransomware detection, traffic analysis and digital forensics, generally do not have access to whole-file information, but rather work at the level of *fragments* of data. In these settings, the metadata that is required by parsers is not present or is incomplete [35]. Given this issue, a number of works have been looking at alternative tests to distinguish between encrypted and compressed content [12,14,23,30,32,34,41]. While these works have the potential to be useful, there has been limited evaluation of their performance on a standardized dataset. Consequently, there is no clear understanding of how these approaches: (i) fare on a variety of compressed file formats and sizes, and (ii) compare to each other. The potential negative implications are significant: the use of ineffective techniques for identifying encrypted content can hinder the effectiveness of ransomware detectors [19], and significantly limit the capability of forensic tools.

Our work compares state-of-the-art approaches on a large dataset of different data types and fragment sizes. We find that, while more useful than entropy estimates, current approaches fail to achieve consistently high accuracy. To address this, we propose ENCOD (**E**ncryption/**C**ompression **D**istinguisher), a novel neural network-based approach. Our evaluation shows that ENCOD outperforms existing approaches for most considered file types, over all considered fragment sizes. ENCOD can classify data fragments as small as 512B with 86% accuracy, increasing to up to 94% for purely compressed data (i.e., zip, gzip) and up to 100% for compressed application data fragments (e.g., pdf, jpeg, mp3) for 8KB fragment sizes. Overall, this paper makes the following contributions:

- We review and categorize existing literature on the topic of distinguishing compressed and encrypted data fragments.
- We systematically evaluate and compare state-of-the-art approaches on a large, standardized dataset including a variety of fragment formats and sizes.

- We propose a new neural-network based approach, which outperforms current state-of-the-art tests in distinguishing encrypted from compressed content for most considered formats, over all considered fragment sizes.
- We propose a new multi-class classifier that can label a fragment with high accuracy as encrypted data, general-purpose compressed data (zip/gzip/rar), or one of multiple application-specific compressed data (png, jpeg, pdf, mp3).
- We thoroughly discuss the implications of our findings, in terms of the applicability of the evaluated approaches.

The rest of this paper is structured as follows: Section 2 provides background on entropy estimation and its applications. Section 3 reviews existing approaches to the problem. Section 4 presents and evaluates a novel approach to the problem, based on deep learning. Section 5 estimates the performance of the considered approaches, discussing their strengths and limitations. Section 6 discusses the implications of our findings. Section 7 discusses related work and Sect. 8 concludes the paper.

## 2 Background

Determining the format of a particular data object (e.g. a file in permanent storage, or an HTTP object) is an extremely common operation. Under normal circumstances, it can be accomplished by looking at content metadata or by parsing the object. Things get more complicated, however, when no metadata is available and the data object is corrupted or partly missing. In this paper, we focus on detection of *encrypted content* and, in particular, on distinguishing between encrypted and compressed data fragments. We begin by examining relevant applications of encryption detection primitives.

### 2.1 Ransomware Detection

*Ransomware* encrypts user files with the aim of making them unusable for the user. It then presents a prompt asking the user to pay a ransom in order to receive the decryption key. Ransomware attacks can cause significant financial damage to organizations [3, 4, 10].

Mitigating a ransomware infection requires rapid detection and termination of all ransomware processes. A number of approaches based on *behavioral process analysis* have been proposed for this purpose [17, 25, 26, 33]. These approaches typically rely on a classifier trained on various process-related features to distinguish benign and ransomware processes. Virtually all proposed behavioral detectors use entropy of file write operations as one of the key features, based on the insight that frequently writing encrypted content is a characteristic behavioral fingerprint of ransomware. Entropy is typically estimated using  $\hat{H}_{MLE}$ . In several approaches entropy is estimated on the content of individual file writes [17, 25, 26], therefore the estimation procedure has only access to partial file fragments.

## 2.2 Forensics

Digital forensics oftentimes involves analysis of phone [39] or PC [35] storage that has been corrupted, or uses an unknown format. Therefore, forensic techniques attempt to recover data of interest (contacts, pictures, etc.) by searching for blocks with recognizable structure. These techniques typically only have access to data fragments, rather than whole files.

Encrypted and compressed data represent a corner case, as they exhibit a complete lack of structure. Still, detecting such content may be important in data recovery operations (e.g., if sensitive data is known to have been encrypted). Distinguishing between compressed and encrypted blocks is notoriously difficult, and some forensic approaches label data as “*compressed or encrypted*”, without attempting to pinpoint which one of the two it is [16].

## 2.3 Network Traffic Analysis

Network traffic analysis examines flows in/out of a network to identify security issues. Regulations (e.g. HIPAA in the U.S.) and best practices expect sensitive data to be encrypted in transit; therefore, entropy-based analyzers have been proposed to ensure that all traffic leaving a monitored network is encrypted [20]. Another application is reverse-engineering of network protocols used by malware. It has been observed [18] that malware protocols may mix encrypted and non-encrypted content within the same message. Encryption detection primitives can be applied to break messages into encrypted and non-encrypted fields.

In both cases above, encryption detectors have partial visibility on the data stream and can only access fragments of data (e.g., an encrypted stream broken into individual packets), rather than whole data objects.

## 2.4 Challenges

In the three domains above, the use of Shannon entropy has been proposed in order to identify encrypted content. Entropy is used to measure the *information content* of a byte sequence; highly structured data exhibit low entropy, while unstructured data—such as a randomly distributed sequence—have high entropy. Therefore, an entropy estimate can be used as a proxy for how close a sequence of bytes is to being randomly distributed. Most encryption algorithms output ciphertexts whose byte-value distributions tend to follow a uniform distribution. As a result, an encrypted bytestream will almost invariably exhibit high entropy.

One of the most common approaches to entropy estimation is the Maximum Likelihood Estimator  $\hat{H}_{MLE} = -\sum_{i=0}^{255} f_i \log_2(f_i)$ , where  $f_i$  is the frequency of byte value  $i$  in the sequence. The entropy range is  $[0 - 8]$ . The frequency  $f_i$  of byte value  $i$ , which is measurable, is used in place of the probability  $P(i)$  of that value occurring, which is unknown. This approach is commonly used in some of the applications above (e.g., [20, 25]), due to its simplicity and efficiency.

This reasoning assumes that, while encrypted data has high entropy, non-encrypted data does not. This appears reasonable, as most relevant data types (e.g., text, images, audio) are information-rich and highly structured. However, this assumption does not hold true in modern computing. Modern CPUs can efficiently decompress data for processing, and compress it back for storage or transmission; this is oftentimes performed in real time and transparent to the user. As a result, most formats tend to apply compression [2, 38]. Informally, a good compression algorithm works by identifying and removing recognizable structures from the data stream; as a result, compressed data tend to exhibit high entropy. In practice, this fact compromises the ability of entropy-based detectors to distinguish encrypted and non-encrypted, compressed content.

**Entropy Estimates for Common Data Formats.** In order to substantiate the claim above, we computed entropy estimates using a dataset consisting of **10,000** file fragments. The dataset covers various popular file formats and AES-256-encrypted data. We considered multiple fragment sizes, from 512B to 8 KB (details in Sect. 4). Figure 1 summarizes the distribution of estimated entropy values for eight different formats with block size 2048 (some ranges truncated for clarity). Results for other block sizes were qualitatively similar; full results are tabulated in Appendix A. As illustrated in Fig. 1, both general-purpose (zip, rar) and domain-specific (jpeg, mp3) compression algorithms result in data which exhibits entropy whose ranges are overlapping with that of encrypted content (enc). The only format that can be unambiguously distinguished is png. Even so, png still overlaps with various other formats. Interestingly, utilities that create and modify data in zip, gzip and png format internally all use the DEFLATE algorithm for compression: the differences in entropy are likely due to differences in file structure and algorithm implementation.

Due to the limits of entropy estimation, the attention of the community has been increasingly focusing on alternative measures that can more precisely estimate whether data follow a random distribution. However, no comprehensive review of such approaches exists. In the next section, we review state-of-art approaches, while we evaluate and compare them in Sect. 5.

### 3 Review of Existing Techniques

This section reviews three state-of-the-art approaches to distinguishing encrypted and compressed content: the NIST suite,  $\chi^2$  and HEDGE [12]. Strictly speaking, these approaches test the *randomness* of a string of bytes, and make no attempt to determine its type. However, due to their high precision they can be used to distinguish true pseudorandom (encrypted) sequences and compressed ones which, while approximating a randomly generated stream, maintain structure.

The NIST suite and  $\chi^2$  are standard statistical tests for identifying randomly-distributed data. HEDGE is a recently proposed statistical approach which shows promising results. HEDGE is a combination of a subset of the NIST tests and two forms of  $\chi^2$  tests. Note that, despite the inclusion of HEDGE, we decided

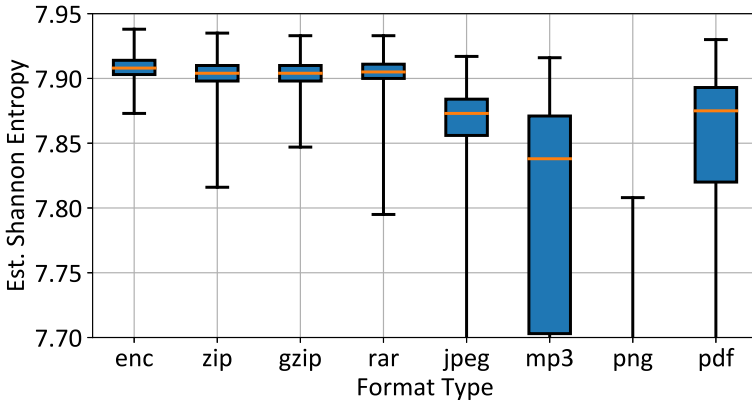


Fig. 1. Entropy ranges for common formats (2048B blocks)

to also report separate results for NIST and  $\chi^2$  due to the fact that those are designed to be, and oftentimes are, used as standalone tests.

### 3.1 NIST SP800-22

The NIST SP800-22 specification [36] describes a suite of tests whose intended use is to evaluate the quality of random number generators. The suite consists of 15 distinct tests, which analyze various structural aspects of a byte sequence. These tests are commonly employed as a benchmark for distinguishing compressed and encrypted content (e.g., [12, 14]). Each test analyzes a particular property of the sequence, and subsequently applies a test-specific decision rule to determine whether the result of the analysis suggests randomness or not. When using the NIST suite for discriminating random and non-random sequences, an important question concerns aggregation of the results of individual tests. Analysis of the tests [36] suggests that they are largely independent. Given this observation, and the intrinsic complexity of *a priori* defining a ranking between the tests, we use a *majority voting* approach. In other words, we consider a fragment to be random (and therefore encrypted) when the majority of tests considers it so. Since some of the tests require a block length much bigger than the ones we use for our smaller fragment sizes, we did not consider in the voting the tests that cannot be executed.

### 3.2 $\chi^2$ Test

The  $\chi^2$  test is a simple statistical test to measure goodness of fit. It has been widely applied to distinguish compressed and encrypted content [12, 30, 34]. Given a set of samples, it measures how well the distribution of such samples

follows a given distribution. Mathematically, the test is defined as:

$$\chi^2 = \sum_{i=0}^{255} \frac{(N_i - E_i)^2}{E_i}$$

where  $N_i$  is the actual number of samples assuming value  $i$ , and  $E_i$  is the expected number of samples assuming value  $i$  according to the known distribution of interest. Since the distribution being evaluated for goodness of fit is the discrete uniform distribution,  $\forall i E_i = L/256$ , where  $L$  is the particular fragment length being considered. The results of the test can be interpreted using either a fixed threshold, or a confidence interval [12].

### 3.3 HEDGE

HEDGE [12] simultaneously incorporates three methods to distinguish between compressed and encrypted fragments:  $\chi^2$  test with absolute value,  $\chi^2$  with confidence interval and a subset of NIST SP800-22 test suite. Out of the NIST SP800-22 test suite HEDGE incorporates 3 tests: *frequency within block test*, *cumulative sums test*, and *approximate entropy test*. These tests were selected due to (i) their ability to operate on short byte sequences, and (ii) their reliable performance on a large and representative dataset. In the HEDGE detector the threshold of the number of the above-mentioned NIST SP800-22 tests failed is set to 0. For the  $\chi^2$  with absolute value test, the thresholds are pre-computed for each of the considered packet sizes, by considering the average and its standard deviation. For  $\chi^2$  with confidence interval, the  $\chi\%$  interval is ( $\chi\% > 99\% || \chi\% < 1\%$ ). For classifying the content of a packet, HEDGE applies the three randomness tests to the input data. Data is considered random only if it passes all tests.

## 4 ENCoD: A Learning-Based Approach

Past work and our own evaluation suggest that tests based on byte-value distribution, such as  $\chi^2$ , can distinguish some encrypted and compressed content, but have accuracy issues (ref. Sect. 5). Such tests, in a sense, “collapse” the entire distribution to a single scalar value, losing information concerning the shape of the distribution. It is therefore natural to ask if Deep Neural Networks (DNNs) can improve such results. DNNs can consider the entire discrete distribution (modeled as a feature vector), and can learn to recognize complex distributions [29].

In order to evaluate the potential of DNNs we designed ENCoD, a set of two distinct neural network-based approaches for distinguishing encryption and compression.

### 4.1 Model Architecture #1: Binary Classifiers

Our first model is a binary classifier trained to distinguish a single specific compressed format from encrypted content. It may be used in cases where only one

compressed format is known to exist in the dataset (e.g., detecting writes of encrypted data performed by a potential ransomware on image files vs legitimate writes of JPEG-compressed data). We explored several alternative architectures for this application, and we found that the structure depicted in Fig. 2a provides the best performance. The binary-classifier architecture consists of 4 fully-connected layers with dimensions as shown in the figure. We initialize the model weights using Glorot uniform [22]. The activation function is ReLU for the first 3 layers, followed by a softmax on the output layer. We used a batch size of 64 for training our model. Each hyperparameter has been chosen using grid search. We used the same procedure also for the model described in Sect. 4.2.

## 4.2 Model Architecture #2: Content-Type Detector

In many applications, a classifier may encounter more than one type of compressed data. Furthermore, it may be important to determine the specific type being encountered. To support these use cases, we design a *content-type detector*: a multi-class classifier that can determine whether a given fragment is encrypted, or belongs to one of multiple known compressed formats. We explored several designs for the neural network, converging to the model depicted in Fig. 2b. Its architecture consists of 5 fully-connected layers with dimensions as shown in the figure. We initialize model weights using LeCun normal [28]. Differently from the binary models, this multi-class classifier seemed prone to the dying neuron problem associated with the ReLU activation function [37]. We therefore opted for the SelU activation function [27] for the first 4 layers, followed by a softmax on the output layer. We used a batch size of 64 instances for training.

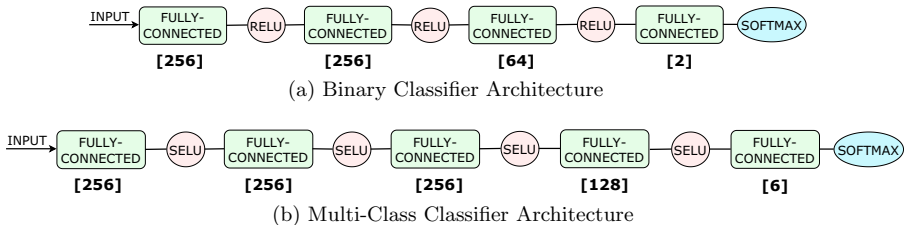


Fig. 2. Neural network architectures

## 4.3 Fragment Dataset

We built a dataset of **200M encrypted and compressed fragments**. For the compressed data, we selected a set of formats covering common, popular content types. To generate the encrypted data fragments, we used the AES cipher in CBC mode implemented by the PyCryptodome library [1]. We chose AES because it



is the most widely used and well known symmetric cipher, representative of modern ciphers which result in byte streams consistently close to random.

In constructing the dataset, we focused on ensuring a diversity of compressed formats, rather than compression algorithms. While algorithms such as DEFLATE are used in multiple compression formats, they are generally used with different parameters and/or embed compressed data in different ways within the compressed archive. Consequently, compressed archives created with different formats tend to differ considerably from each other even when using the same underlying compression algorithm. This observation is empirically confirmed by our evaluation in Sect. 5. Finally, our dataset does not include data which is both compressed and encrypted, and we ensured such data is not present in the dataset. The dataset is comprised of the following data types:

1. **AES encrypted data (enc).** We used the AES implementation provided by the Cryptodome Python library. AES was configured to use CBC mode with 256-bit keys, with a random IV generated before encrypting each file.
2. **DEFLATE- and rar-compressed data (zip/gzip and rar):** both DEFLATE and rar are de-facto standards for generic file compression. DEFLATE is also widely used for documents (such as in the MS OFFICE file formats), and network applications (e.g., HTTP header compression).
3. **png and jpeg images:** png is used for lossless image compression; it internally uses DEFLATE, but png files present a structure that is different from that of zip files. jpeg uses DCT-based lossy compression.
4. **mp3 audio files:** MP3 compressors use a psychoacoustic model to remove inaudible frequencies from audio data, and compress the resulting data using a lossy algorithm based on the modified-DCT transform.
5. **pdf documents:** PDF is an office format used for document exchange and form filling. Internally, PDF files consist of a tree of objects that can be compressed using a variety of techniques. In practice, most PDF documents contain a large amount of compressed content, such as embedded images.

**Fragment Generation Process.** We generate fragments from a dataset of files:

- **zip/gzip/rar/enc:** we used various textual documents obtained from a 2020 English Wikipedia dump [8]. We created four copies of each file, each of which was either compressed using one of zip, gzip, rar utilities (with default parameters), or encrypted using AES-256.
- **png:** we crawled 5000 png images from the web and various repositories [5].
- **jpeg:** we downloaded 10,000 images from the Open Images Dataset v5 [7].
- **mp3:** we used the FMA dataset [6], which contains 8000 mp3 files.
- **pdf:** we crawled 1,000 randomly-selected papers from arXiv [9].

We split each file into fragments of 512B, 1 KB, 2 KB, 4 KB, and 8 KB. We then selected 5M fragments for each fragment size/data type combination (this ensures that the dataset remains balanced).

#### 4.4 Dataset Analysis Methodology

*Statistical tests (NIST,  $\chi^2$ , HEDGE):* For each fragment size, we randomly selected 10,000 compressed fragments (evenly distributed across the different compressed data types) and 10,000 encrypted fragments. We then executed the tests directly on these fragments.

*ENCOD/Binary Classifiers:* We separately trained and evaluated classifiers for each fragment size. The features that are fed to our models for training/classification are derived from the histograms of the byte values for the observed fragment size. Each feature is the value of the probability density function at a given bin, normalized such that the integral over the range is 1.

We trained the binary classifiers by randomly selecting 3M vectors from the encrypted class and 3M vectors from the data type that we aim to distinguish. We partitioned this dataset into 85% training, 5% development and 10% test. Before fitting the data to the model for training, we applied a MinMax scaler to scale the dataset from the range  $[0, 1]$  to the range  $[0, 2]$  (range selected via grid search). Scaling helps the ML model to more easily capture minute differences in the inputs, allowing to better distinguish among the classes and converge faster.

*ENCOD/Content-Type Detector:* To train the content-type detectors, for each fragment size we randomly sampled 6M feature vectors consisting of a mix of the considered file types. This dataset was partitioned into training, development and test sets in the same ratios used for the binary classifiers. We also scaled the dataset using the MinMax scaler with the same parameters used above.

## 5 Evaluation

In this section, we comprehensively evaluate existing approaches discussed in Sect. 3, in addition to ENCOD, our novel neural network-based approach (see Sect. 4). We frame the evaluation in terms of the following comparisons:

1. **Binary classification: all formats.** In Sect. 5.2, we consider the ability of different detectors to discriminate encrypted and compressed data, regardless of the specific compressed format. Results show that our classifier heavily outperforms NIST,  $\chi^2$ -test and HEDGE for all fragment sizes.
2. **Binary classification by format.** In Sect. 5.3, we break down the performance of  $\chi^2$ , NIST and HEDGE by compressed format. We also report the performance of our per-format binary classifiers (see Sect. 4.1). The latter perform comparably or better than other tests on all formats but one.
3. **Format fingerprinting.** In Sect. 5.4, we evaluate the accuracy of our multi-class classifier in labeling unknown fragments as the correct compressed format (or as encrypted). Results show that our classifier is able to distinguish the file type with an overall accuracy of 90% for the 2048 byte chunk size. It also achieves high precision, especially on png, jpeg, mp3.

## 5.1 Implementation

We implemented the classifier described in Sect. 4 using the Keras Library [13] for machine learning. For the NIST tests, we used the official implementation [15]. In order to aggregate the NIST tests results, we use the majority voting approach described in Sect. 3.1. In order to label fragments as compressed or encrypted based on  $\chi^2$  results, we used the thresholds suggested in the HEDGE paper [12], as the analysis in HEDGE is specifically aimed at producing a dataset-independent threshold for general use. We implemented HEDGE according to the published description [12]. Finally, all experiments were conducted using the dataset described in Sect. 4.

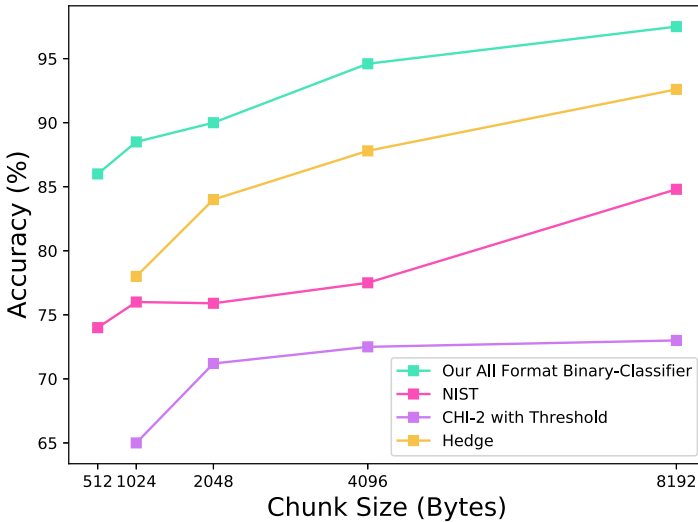


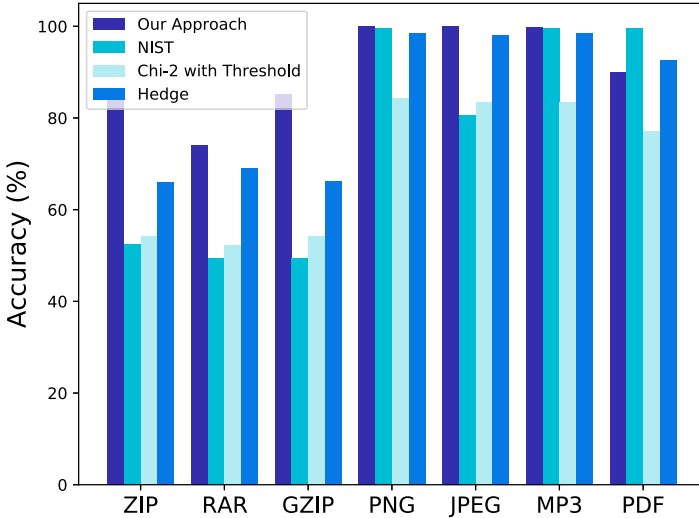
Fig. 3. Performance comparison (binary classification: all formats)

## 5.2 Binary Classification: All Formats

In the first part of our evaluation, we consider the binary classification problem of determining whether a given high-entropy data fragment is compressed or encrypted. Given a fragment, the  $\chi^2$  test, HEDGE, and the NIST test suite return whether the fragment’s content appears random or not. Therefore, a binary classifier can be derived simply by labeling random content as encrypted. Our binary classifier used for this evaluation is based on our multi-class classifier. The multi-class classifier labels each fragment either as encrypted, or as one of the seven supported compressed formats. Since in this experiment we are only interested in distinguishing encryption and compression, regardless of the type, we combine all compressed type labels into one (we look at content

fingerprinting accuracy in Sect. 5.4). Effectively, we consider classification in two labels: (1) a macro-label “compressed”, which is comprised of the labels  $\{zip, rar, gzip, png, jpeg, mp3, pdf\}$  and (2) the label “encrypted”.

The results of this evaluation are depicted in Fig. 3. All classifiers tend to improve as fragment size increases; we discuss this phenomenon in Sect. 6. Our neural network-based approach heavily outperforms all the other approaches on all block sizes. The  $\chi^2$  accuracy remains consistently low across the range of block sizes. Results suggest that this test has an intrinsic difficulty in discriminating non-random content which closely approaches a uniform random distribution.

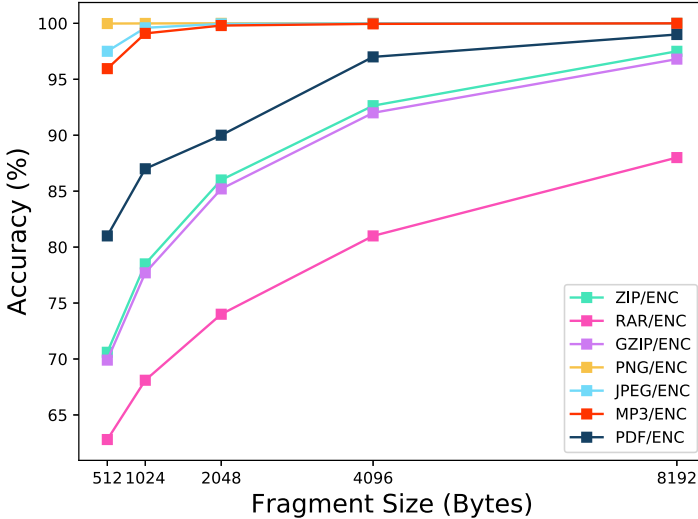


**Fig. 4.** Performance comparison between our binary-classifier approach, NIST,  $\chi^2$  with Threshold and HEDGE (2048B blocks)

### 5.3 Binary Classification by Format

In the second experiment, we consider the question of whether some compressed formats are harder than others to distinguish from encrypted content. Such phenomenon may arise due to (i) differences in effectiveness between compression algorithms in removing redundancy (and thus structure) from the uncompressed data; and (ii) presence (or absence) of metadata, or other structured information interleaved with compressed data.

In order to answer this question, we break down results for the  $\chi^2$ -test, NIST suite, and HEDGE test by format. We do not evaluate our multi-class classifier in this experiment. This is due to the fact that this classifier can generate two different types of classification errors for compressed formats: (1) mislabeling a compressed fragment as an encrypted one; and (2) mislabeling a compressed



**Fig. 5.** Performance of our binary classifiers (all block sizes)

fragment of a given type for a compressed fragment of a different type. As  $\chi^2$ -test, NIST and HEDGE can only generate errors of the former type, a direct comparison is not possible.

Instead, in this experiment we evaluate multiple binary neural network-based classifiers (see Sect. 4.1). With this approach, we train one binary classifier per compressed format. Each classifier is trained to distinguish content in that format from encrypted content. It is important to note that, while each of these classifiers is trained specifically on one format, the other tests ( $\chi^2$ , NIST and HEDGE) work the same regardless of the format. Despite this limitation, we believe this to be an informative analysis of the potential of learning-based approaches.

Figure 4 shows the comparison between the three approaches on 2048-byte blocks. Overall, neural network-based classifiers tend to fare better than the other tests, particularly on challenging formats such as zip/gzip and rar. PDF is the only format on which the NIST and HEDGE tests outperform the neural network classifier. Interestingly, the  $\chi^2$  fares slightly better than NIST on most formats, but its accuracy is significantly worse on formats that are typically easy to distinguish, such as PNG. We believe this to be due to the fact that the NIST tests look at a richer set of properties beyond byte value distribution, such as a presence of runs and repeated sequences. HEDGE test outperform  $\chi^2$  on all file types, while outperforming NIST on most formats, beside PDF, and have similar performance on PNG and MP3 formats. Finally, Fig. 5 presents the performance of all our binary classifiers across the range of fragment sizes. These results again show that accuracy increases significantly as block size increases. For 8 KB-blocks, accuracy is above 85% for all types.

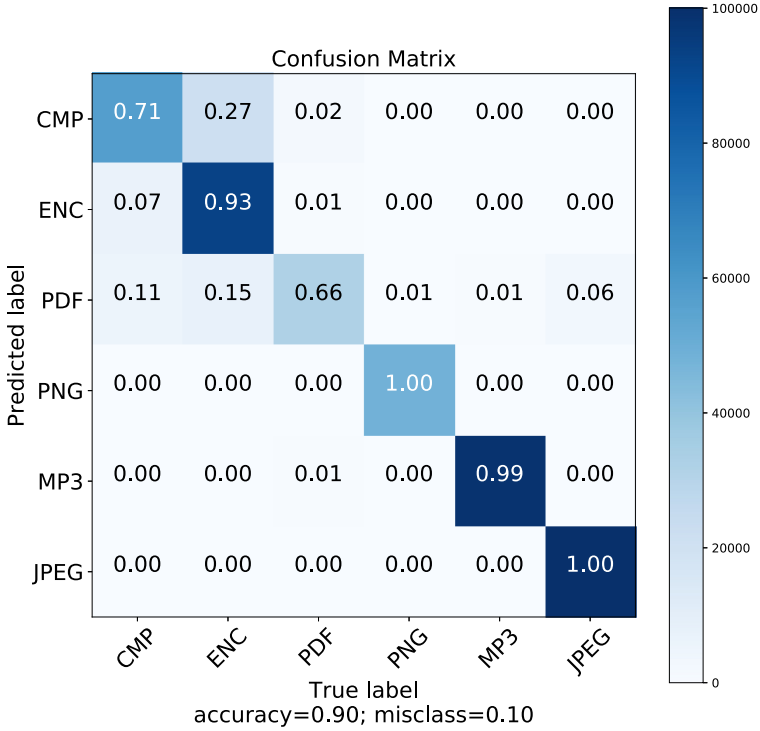


Fig. 6. Confusion matrix for the content-type classifier

#### 5.4 Format Fingerprinting

Our multiclass classifier has the ability to (1) distinguish encrypted and compressed data, and (2) pinpoint the specific format compressed data belong to. This is a significant improvement over the functionality of existing tests, that can only distinguish encryption and compression. In this section, we evaluate the effectiveness of our multi-class classifier in fingerprinting the correct type of compressed content. Figure 6 shows the confusion matrix for the multi-class classifier. Results indicate that our classifier is able to pinpoint the file type with consistently high precision for most formats, especially png, mp3, and jpeg. It performs fairly well on the other considered compressed formats such as cmp (which contains a mixture of zip, rar, and gzip feature vectors) but with a slightly higher rate of misclassified instances between enc and cmp. This can be explained by the fact that their distributions are very close, and intrinsically hard to distinguish.

#### 5.5 Overhead

In the final part of our evaluation, we analyze the practical applicability of the three approaches, comparing their runtime in order to understand if they can be deployed in time-critical applications. For this test, we used a small dataset

comprised of 1000 randomly-selected compressed or encrypted samples. We ran all three approaches (NIST, HEDGE and our binary ML model) on each sample, taking individual runtime and repeating the experiment 1000 times. Table 1 presents the results of our evaluation. As we can see, while both mean and median runtime for NIST tests are faster than HEDGE, our proposed binary classifier is considerably faster than both. Both mean and median runtime for the ML model are three orders of magnitude faster than both NIST and HEDGE, making it easily applicable to scenarios that require fast classification results such as ransomware detection. It is worth noting that the evaluation of our ML model was carried out by measuring the time required to predict a single sample, rather than a batch of samples. However, our model can easily classify multiple samples in parallel by exploiting the heavy parallelism of GPUs, further decreasing the runtime required per individual sample.

**Table 1.** Time required by each approach to classify one sample, in seconds.

Approach	Mean	Median	Std.dev
NIST	0.1	0.1	0.004
HEDGE	0.44	0.43	0.008
Binary classifier	0.00046	0.00044	0.00012

## 6 Discussion of Findings

Results shown in Sect. 5 highlight the difficulty of discriminating compressed and encrypted fragments. State-of-the-art statistical tests tend to fare better than entropy measures (ref. Sect. 2), but their performance varies significantly depending on the specifics of the compressed format and fragment size. Moreover, such approaches can only determine whether a given fragment is encrypted with a certain confidence, but cannot distinguish between different compressed formats. ENCOD, the learning-based approach introduced in Sect. 4, tackles both these limitations. Both per-format and multi-class classifiers outperform existing tests on all considered file types/block sizes. Moreover, our multi-class classifier can be used to determine the format of a given unknown fragment, even in the complete absence of any context or information on its type.

Results show that accuracy improves consistently with increasing fragment size. This is in a sense to be expected; all approaches considered in this paper leverage differences between the byte value distribution of random data (which is uniform) and that of compressed data. Perfectly estimating the byte value distribution of a short data stream is generally not possible. As sequences get shorter, the probability that the estimated distribution may not reflect the typical distribution for their content type increases. However, as the size of the sample increases, the estimated empirical distribution approaches the underlying data distribution, allowing us to capture any deviation from the uniform distribution.

For modern compression algorithms, these deviations are quite minor, and a 512-byte block gives even accurate tests very little data to work with. However, when enough data is available, it is possible to identify the class of data with high accuracy; our learning-based classifier exceeds 90% accuracy already for 2048-byte blocks. In general, we recommend against using any one approach as the sole guidance for automated security decisions (e.g. dropping/allowing flows, terminating processes, etc.). However, when integrated as part of a more complete set of features in a larger system, our proposed classifiers can provide an additional robust feature to use in the decision-making process.

Given the discussion above, we suspect an intrinsic bound on the accuracy reachable by any classifier which looks purely at byte value distributions. However, approaches attempting to parse fragments or identify recognizable structures are likely to incur an impractical computational cost. Moreover, it is not apparent that any such structure is preserved for very short fragment sizes.

## 7 Related Work

### 7.1 Entropy-Based Encryption Detection

Use of entropy estimation to detect encrypted content is common in ransomware detection. Proposals such as RWGuard [33], UNVEIL [26], Redemption [25] and ShieldFS [17] use entropy of written content either directly as a feature, or as part of feature calculation. It should be noted that none of these detectors use entropy as the sole feature for detection. However, evidence from Sect. 2 suggests that they may be better ignoring entropy altogether. In the realm of digital forensics, entropy estimation has been used to determine the type of unknown disk data fragments. One of the most complete approaches is that of Conti et al. [16]. However, the same authors found that such estimates have limited discerning power in distinguishing encrypted and compressed content, and aggregated the two types under a single label.

Entropy estimation has also been applied to the real-time analysis of network traffic. Dorfinger’s Master thesis [20] proposes a system for discriminating encrypted and non-encrypted traffic, to ensure that all communications from a target network are encrypted. Similar approaches were also proposed by Mamun et al. [31] and Malhotra [30]. Zhang et al. proposed an entropy-based classifier for the identification of botnet traffic [42]. All these approaches also suffer from the limitations of using high entropy as a fingerprint of encryption. Wang et al. [41] report positive results in using an SVM classifier to discriminate between various data types using entropy estimates. Their application scenario is different from ours, as they consider both low-entropy (non-compressed) and high-entropy (compressed or encrypted) formats. We only consider high-entropy formats, which are difficult to distinguish using entropy alone.

Finally, MovieStealer [40] aims at identifying encrypted and decrypted-but-compressed media buffers in order to break DRM. It uses an entropy test to single out encrypted and compressed buffers from other data, and the  $\chi^2$ -test to distinguish them. It requires 800 KB of data to reliably identify random data, which is far beyond the fragment size in the scenarios that we consider.



## 7.2 Non-entropy-based Approaches

HEDGE, by Casino et al. [12], evaluates a combination of  $\chi^2$ -test and a subset of NIST SP800-22 [36] to discriminate encrypted and compressed traffic. They use a dataset which is significantly smaller than ours, and do not discuss learning-based approaches. A limitation of this class of approaches is the fairly low accuracy, especially for small block sizes (ref. Sect. 5). Also, this and other similar approaches based on statistical randomness tests (e.g., [14, 34]) cannot distinguish between different types of compressed archives. Mbol et al. [32] investigate the use of the Kullback-Leibler divergence (relative entropy) to differentiate encrypted files from JPEG images. Their analysis does not investigate other formats, and assumes the availability of blocks of significant size (128 to 512 KB) from the beginning of each file. Especially in forensic and networking applications, uninterrupted blocks of such size are difficult to obtain.

While the application of neural networks to the problem at hand is fairly new, there exist some preliminary work. Ameeno et al. [11] show promising preliminary results, however the analysis is limited in scope: it only attempts to distinguish zip archives from rc4-encrypted data, and considers whole files (not fragments). Hahn et al. [24] perform an exploratory analysis of machine learning models. Their dataset is order of magnitudes smaller than ours, and they lack a comparative analysis of statistical approaches.

## 8 Conclusions

Discriminating encrypted from non-encrypted content is important for a variety of security applications, and oftentimes tackled via entropy estimation. We comprehensively highlighted the limits of this technique and reviewed the effectiveness of the leading alternative approaches:  $\chi^2$ -test, NIST SP800-22 test suite, and HEDGE. In addition, we proposed ENCOD, a novel neural network classifier of our own design. In order to ensure generality of results, we created a dataset of 200M fragments covering 5 different sizes and 8 data formats.

Results show that previous state-of-the-art methods have blind spots which result in low accuracy for certain fragment sizes/data types. However, our neural network-based approach appears promising. Besides being able to discriminate between compressed formats (which traditional statistical tests cannot), it exceeds 90% accuracy already on fragments of only 2KB. This suggests that systems incorporating encrypted content detection (e.g., ransomware detectors) would be better served by learning-based, rather than hand-crafted statistical approaches. This finding also suggests that learning may have useful applications to other problems in content type inference. Overall, we believe this work is an important step forward towards reliable encryption detection.

**Acknowledgments.** We would like to thank Daniele Venturi and Guinevere Gilman for their useful insights and comments. This work was supported by Gen4olive, a project that has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 101000427, and in part by the Italian

MIUR through the Dipartimento di Informatica, Sapienza University of Rome, under Grant Dipartimenti di eccellenza 2018–2022.

## Appendix

### A Entropy Analysis Results

Full results for the entropy analysis discussed in Sect. 2.4:

Chunk size: 512B					
Format	Min	Q1	Median	Q3	Max
enc	7.427	7.569	7.591	7.613	7.709
zip	7.163	7.560	7.584	7.607	7.695
gzip	7.154	7.560	7.585	7.607	7.703
rar	7.381	7.563	7.587	7.610	7.692
jpeg	3.820	7.512	7.548	7.576	7.676
mp3	0.000	7.451	7.527	7.565	7.680
png	0.000	1.070	2.605	4.549	7.572
pdf	0.000	7.453	7.534	7.574	7.676
Chunk size: 2048B					
Format	Min	Q1	Median	Q3	Max
enc	7.873	7.903	7.908	7.914	7.938
zip	7.816	7.898	7.904	7.910	7.935
gzip	7.847	7.898	7.904	7.910	7.933
rar	7.795	7.900	7.905	7.911	7.933
jpeg	5.123	7.856	7.873	7.884	7.917
mp3	0.379	7.703	7.838	7.871	7.916
png	0.000	1.312	2.815	4.752	7.808
pdf	0.000	7.820	7.875	7.893	7.930
Chunk size: 8192B					
Format	Min	Q1	Median	Q3	Max
enc	7.969	7.976	7.978	7.979	7.984
zip	7.955	7.973	7.975	7.976	7.983
gzip	7.955	7.973	7.975	7.976	7.983
rar	7.960	7.974	7.976	7.977	7.983
jpeg	5.646	7.930	7.945	7.952	7.967
mp3	0.497	7.789	7.918	7.942	7.971
png	0.014	1.451	2.963	4.852	7.914
pdf	0.010	7.903	7.953	7.968	7.981

## References

1. Pycryptodome library. <https://pycryptodome.readthedocs.io/en/latest/src/introduction.html>
2. DOCX Transitional (Office Open XML), January 2017. <https://www.loc.gov/preservation/digital/formats/fdd/fdd000397.shtml>
3. Atlanta spent \$2.6m to recover from a \$52,000 ransomware scare (2018). <https://www.wired.com/story/atlanta-spent-26m-recover-from-ransomware-scare/>
4. Wannacry cyber attack cost the NHS £92m as 19,000 appointments cancelled (2018). <https://www.telegraph.co.uk/technology/2018/10/11/wannacry-cyber-attack-cost-nhs-92m-19000-appointments-cancelled/>
5. Evolvingai: Deep neural networks are easily fooled: High confidence predictions for unrecognizable images, December 2019. <http://www.evolvingai.org/fooling>
6. FMA: A dataset for music analysis, December 2019. <https://github.com/mdeff/fma>
7. Open images dataset v5, December 2019. <https://www.figure-eight.com/dataset/open-images-annotated-with-bounding-boxes/>
8. Wikipedia: database download, December 2019. <https://dumps.wikimedia.org/enwiki/>
9. arXiv.org e-Print archive, February 2020. <https://arxiv.org/>
10. Ransomware attacks grow, crippling cities and businesses (2020). <https://www.nytimes.com/2020/02/09/technology/ransomware-attacks.html>
11. Ameeno, N., Sherry, K., Gagneja, K.: Using machine learning to detect the file compression or encryption. *Amity J. Comput. Sci.* **3**(1), 6 (2019)
12. Casino, F., Choo, K.K.R., Patsakis, C.: HEDGE: efficient traffic classification of encrypted and compressed packets. *IEEE Trans. Inf. Forensics Secur.* **14**(11), 2916–2926 (2019)
13. Chollet, F., et al.: Keras (2015). <https://keras.io>
14. Choudhury, P., Kumar, K.R.P., Nandi, S., Athithan, G.: An empirical approach towards characterization of encrypted and unencrypted VoIP traffic. *Multimedia Tools Appl.* **79**(1–2), 603–631 (2020)
15. Computer Security Division, I.T.L.: NIST SP 800-22: Documentation and Software, May 2016. <https://csrc.nist.gov/projects/random-bit-generation/documentation-and-software>
16. Conti, G., et al.: Automated mapping of large binary objects using primitive fragment type classification. *Digital Invest.* **7**, S3–S12 (2010)
17. Continella, A., et al.: Shieldfs: a self-healing, ransomware-aware filesystem. In: ACSAC (2016)
18. De Carli, L., Torres, R., Modelo-Howard, G., Tongaonkar, A., Jha, S.: Botnet protocol inference in the presence of encrypted traffic. In: INFOCOM (2017)
19. De Gaspari, F., Hitaj, D., Pagnotta, G., De Carli, L., Mancini, L.V.: The naked sun: malicious cooperation between benign-looking processes. In: 18th International Conference on Applied Cryptography and Network Security. ACNS (2020)
20. Dorfinger, P., Panholzer, G., John, W.: Entropy estimation for real-time encrypted traffic identification. In: Traffic Monitoring and Analysis (2011)
21. Fielding, R., et al.: RFC 2616, hypertext transfer protocol - HTTP/1.1 (1999). <http://www.rfc.net/rfc2616.html>
22. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: AISTATS (2010)

23. Hahn, D., Aphorpe, N., Feamster, N.: Detecting compressed cleartext traffic from consumer internet of things devices (2018)
24. Hahn, D., Aphorpe, N., Feamster, N.: Detecting Compressed Cleartext Traffic from Consumer Internet of Things Devices. [arXiv:1805.02722](https://arxiv.org/abs/1805.02722) [cs], May 2018. <http://arxiv.org/abs/1805.02722>
25. Kharraz, A., Kirda, E.: Redemption: real-time protection against ransomware at end-hosts. In: RAID (2017)
26. Kirda, E.: Unveil: a large-scale, automated approach to detecting ransomware (keynote). In: SANER (2017)
27. Klambauer, G., Unterthiner, T., Mayr, A., Hochreiter, S.: Self-normalizing neural networks. CoRR abs/1706.02515 (2017). <http://arxiv.org/abs/1706.02515>
28. LeCun, Y., Bottou, L., Orr, G.B., Müller, K.R.: Efficient backprop. In: Neural Networks: Tricks of the Trade (1998)
29. Lee, H., Ge, R., Ma, T., Risteski, A., Arora, S.: On the ability of neural nets to express distributions. In: Kale, S., Shamir, O. (eds.) Proceedings of the 30th Conference on Learning Theory, COLT 2017, Amsterdam, The Netherlands, 7–10 July 2017. Proceedings of Machine Learning Research, vol. 65, pp. 1271–1296. PMLR (2017). <http://proceedings.mlr.press/v65/lee17a.html>
30. Malhotra, P.: Detection of encrypted streams for egress monitoring. Master of Science, Iowa State University, Ames (2007). <https://lib.dr.iastate.edu/rtd/14632/>
31. Mamun, M.S.I., Ghorbani, A.A., Stakhanova, N.: An entropy based encrypted traffic classifier. In: Qing, S., Okamoto, E., Kim, K., Liu, D. (eds.) ICICS 2015. LNCS, vol. 9543, pp. 282–294. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-29814-6\\_23](https://doi.org/10.1007/978-3-319-29814-6_23)
32. Mbol, F., Robert, J.-M., Sadighian, A.: An efficient approach to detect Torrent-Locker ransomware in computer systems. In: Foresti, S., Persiano, G. (eds.) CANS 2016. LNCS, vol. 10052, pp. 532–541. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-48965-0\\_32](https://doi.org/10.1007/978-3-319-48965-0_32)
33. Mehnaz, S., Mudgerikar, A., Bertino, E.: Rwgard: a real-time detection system against cryptographic ransomware. In: Research in Attacks, Intrusions, and Defenses. RAID 2018 (2018)
34. Palisse, A., Durand, A., Le Bouder, H., Le Guernic, C., Lanet, J.-L.: Data aware defense (DaD): towards a generic and practical ransomware countermeasure. In: Lipmaa, H., Mitrokotsa, A., Matulevičius, R. (eds.) NordSec 2017. LNCS, vol. 10674, pp. 192–208. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-70290-2\\_12](https://doi.org/10.1007/978-3-319-70290-2_12)
35. Park, B., Savoldi, A., Gubian, P., Park, J., Lee, S.H., Lee, S.: Data extraction from damage compressed file for computer forensic purposes. *Int. J. Hybrid Inf. Technol.* **1**(4), 14 (2008)
36. Rukhin, A., et al.: A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications. Special Publication 800-22r1a, NIST, April 2010
37. Trottier, L., Giguere, P., Chaib-draa, B.: Parametric exponential linear unit for deep convolutional neural networks. In: 2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA) (2017)
38. Wallace, G.K.: The jpeg still picture compression standard. *IEEE Trans. Consum. Electron.* **38**(1), xviii–xxxiv (1992)
39. Walls, R.J., Learned-Miller, E., Levine, B.N.: Forensic triage for mobile phones with DECODE. In: USENIX Security Symposium (2011)
40. Wang, R., Shoshitaishvili, Y., Kruegel, C., Vigna, G.: Steal this movie - automatically bypassing DRM protection in streaming media services. In: USENIX (2013)

41. Wang, Y., Zhang, Z., Guo, L., Li, S.: Using entropy to classify traffic more deeply. In: 2011 IEEE Sixth International Conference on Networking, Architecture, and Storage, pp. 45–52, July 2011
42. Zhang, H., Papadopoulos, C., Massey, D.: Detecting encrypted botnet traffic. In: 2013 Proceedings IEEE INFOCOM, pp. 3453–3458, April 2013