

# CLExtract: Recovering Highly Corrupted DVB/GSE Satellite Stream with Contrastive Learning

Minghao Lin

University of Colorado Boulder  
minghao.lin@colorado.edu

Minghao Cheng

Independent Researcher  
minhalemail@gmail.com

Dongsheng Luo

Florida International University  
dluo@fiu.edu

Yueqi Chen

University of Colorado Boulder  
yueqi.chen@colorado.edu

**Abstract**—Since satellite systems are playing an increasingly important role in our civilization, their security and privacy weaknesses are more and more concerned. For example, prior work demonstrates that the communication channel between maritime VSAT and ground segment can be eavesdropped on using consumer-grade equipment. The stream decoder GSExtract developed in this prior work performs well for most packets but shows incapacity for corrupted streams. We discovered that such stream corruption commonly exists in not only Europe and North Atlantic areas but also Asian areas. In our experiment, using GSExtract, we are only able to decode 2.1% satellite streams we eavesdropped on in Asia.

Therefore, in this work, we propose to use a contrastive learning technique with data augmentation to decode and recover such highly corrupted streams. Rather than rely on critical information in corrupted streams to search for headers and perform decoding, contrastive learning directly learns the features of packet headers at different protocol layers and identifies them in a stream sequence. By filtering them out, we can extract the innermost data payload for further analysis. Our evaluation shows that this new approach can successfully recover 71-99% eavesdropped data hundreds of times faster speed than GSExtract. Besides, the effectiveness of our approach is not largely damaged when stream corruption becomes more severe.

## I. INTRODUCTION

Satellite systems are becoming infrastructures of modern civilization. It provides a wide range of services including media broadcasts that cover 100 million customers, Earth observation which contributes to environmental conservation efforts, and precise global positioning services. Especially in recent years, the New Space trend [18] significantly advances the development of organizations such as SpaceX and OneWeb that carry space missions like global broadband service. Nowadays, there are more than 2,000 operational satellites that orbit Earth and the market value exceeds \$150 billion a year [14].

A satellite system consists of three major components: ground segment, space segment, and ground-space communication. While all components are reported vulnerable in previous works (*e.g.*, [8][15][13][6]), for adversaries, ground-space communication is the most accessible attack surface.

Unlike firmware in ground-segment and satellite payload the reverse engineering of which requires physical access, ground-space communication relies on radio signals and covers a large area in the size of million square kilometers. As long as there is an antenna within the area and aligned to the satellite, attackers can eavesdrop on satellite streams and steal sensitive data.

A prior work [16] demonstrates this threat by eavesdropping on maritime VSAT communications in the North Atlantic. The stream decoder GSExtract developed in this work can extract between 40-60% of the GSE PDUs contained within the targeted streams. But for corrupted packets, GSExtract shows incapacity by recovering only 10-25% of them. We discovered that such stream corruption commonly exists for satellite communication in not only European areas but also Asian areas.

In our experiment, using GSExtract, we are only able to decode 2.1% satellite stream we eavesdropped on in Asia. This higher corruption is presumably because of surface reflection which is not an issue for maritime VSAT eavesdropping. However, in addition to surface reflection, we identified three more factors that universally influence satellite stream quality and are difficult to be eliminated.

The traditional Finite-State Machine (FSM) based decoding approach, like the one implemented in GSExtract, fundamentally cannot recover such highly-corrupted satellite streams. This is because it relies on critical information in networking packet headers (*e.g.*, length field, CRC-32) to perform decoding layer by layer. When this critical information is corrupted, the decoding can hardly carry on.

In this work, we propose to use a contrastive learning technique with data augmentation to recover satellite streams. Instead of relying on critical information for decoding, contrastive learning directly learns the features of packet headers at different layers and identifies them in a stream sequence. By filtering them out, we can extract the innermost payload that can be further analyzed by tools like Wireshark. Our approach further employs data augmentation to entitle the trained contrastive learning model with robustness against unseen corruptions. We implemented our approach and named it as CLExtract. Our experiment shows that CLExtract can successfully recover 71-99% eavesdropped data hundreds of times faster speed than GSExtract. Besides, the effective-

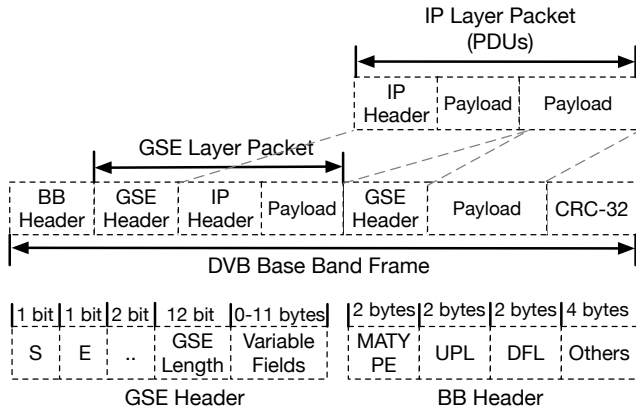


Fig. 1: Protocol layers of satellite stream and packet formats of each layer.

ness of CLExtract is not largely damaged when corruption becomes more severe.

Making eavesdropping more practical, we develop CLExtract to facilitate investigation in security and privacy issues in satellite communication. To foster future research, we will open-source CLExtract once this work is accepted. In summary, this paper makes the following contributions:

- Analysis into factors that cause satellite stream corruption and challenges of corrupted stream recovery.
- Proposed contrastive learning technique with data augmentation to recover corrupted satellite streams.
- Open-sourced implementation of the proposed technique and evaluation of its effectiveness and efficiency using eavesdropped data in the real world

In the following, we first describe the background in Section II, followed by the design overview in Section III. Then, we elaborate on the technical details in Section IV. The evaluation results and ethics are in Section V. Finally, we conclude our work in Section VII.

## II. BACKGROUND AND CHALLENGES

In this section, we first introduce the protocol layers of the satellite stream and packet formats of each layer. Then, we describe four major factors that can corrupt the integrity of DVB/GSE packets, especially in the scenario of eavesdropping. Finally, we discuss the challenges of recovering corrupted streams.

**DVB/GSE Protocol Format.** DVB (Digital Video Broadcasting) and its following DVB-S2 and S2X designed by ETSI (The European Telecommunications Standards Institute) are the de facto standard for the communication of most satellites. Data following this standard is encapsulated into continuous streams using GSE (Generic Stream Encapsulation) protocol. As shown in Figure 1, from innermost to outermost, the encapsulation consists of three protocol layers: IP layer, GSE layer, and DVB layer.

The IP layer supports common transport layer protocols (e.g., TCP, ICMP, and MPEG). The IP header has a field recording which transport layer protocol is used (e.g., 0x06

for TCP and 0x01 for ICMP). All transport layer data along with the IP header are encapsulated into PDUs (Protocol Data Units). PDUs are further divided into several slots and stored as data fields in continuous GSE packets. A GSE packet starts with a header of variable length. Figure 1 shows its layout. The fixed header is in two bytes, recording whether the current GSE packet is the starting packet (i.e.,  $S=1$ ) or the ending packet (i.e.,  $E=1$ ) for a complete PDU. Besides, it includes a length field that stores the size of the current GSE packet. Traditional FSM-based decoding approach relies on this length field to pinpoint the start position of the next GSE packet. More information like Fragment ID and Protocol Type is stored in variable fields. Due to the space limit, we omit details of this part. Readers can refer to [3] for more information. In the ending GSE packet of a complete PDU, there is a CRC-32 tail as the error detection code of the PDU.

With all GSE packets encapsulated, at the outermost DVB layer, several GSE packets are concatenated to fill in the data field of a Base Band (BB) Frame. The BB frame starts with a BB header the size of which are 10 bytes. The first two bytes of the header are the MATYPE field which describes the input stream format and the type of Mode Adaptation. The second two bytes are the UPL field storing user packet length in bits (up to 65535 bits). The following two bytes are the DFL field which holds the length of the data field (i.e., concatenated GSE packets) in bits (up to 58112 bits). Readers can refer to [2] for details.

**Four Factors Causing Stream Corruption.** Given a complete DVB/GSE packet, it is straightforward to decode it using a traditional FSM-based approach (e.g., GSEextract[16]) which decodes layer by layer according to information in headers (e.g., data field length). The innermost PDUs (i.e., IP packets) extracted in this approach can be further analyzed by tools like Wireshark. However, the quality of radio communication between the satellite and the ground segment, especially in the scenario of eavesdropping, is under the influence of many factors, which makes the FSM-based approach can hardly carry on.

The first factor is solar activities which result in perturbations of the ionosphere. This perturbation can change the density structure of the ionosphere by creating areas of enhanced density. This change reflects, refracts, and absorbs radio waves, leading to the loss of signal. According to NASA's report [1], when communication is interrupted, some satellites can be tumbled out of control for several hours, and weather images can be lost. The second factor causing signal attenuation and absorption are rainstorms, snow, and heavy winds near the ground, given that the signals from satellites are transmitted through air [4]. Higher-frequency bands tend to be more affected by rain because the wavelength itself is close in size to water molecules. The third factor is the surface reflection from distant reflectors such as mountains and large industrial infrastructures and local environmental effects including shadowing and blockage from objects and vegetation near the terminal [11]. The last factor is the quality of

the antenna. Attackers performing eavesdropping can seldom afford professional antennae and lack experience in antenna alignment. Therefore, the DVB/GSE packets received through consumer-grade antenna are more likely to be of low quality in comparison with the ground station.

**Challenges of Recovering Corrupted Stream.** Due to the nature of satellite radio communication, the four factors are not easy to eliminate. As we will show in the evaluation (Section V-A), oftentimes, the DVB/GSE packets received by eavesdroppers are highly corrupted. To make matters worse, eavesdroppers are passive attackers and cannot ask for re-transmission if discovering that the received packets are corrupted. How to decode corrupted DVB/GSE packets and extract PDUs is the key to the success of eavesdropping. In the following, we discuss technical challenges in detail.

The first challenge is broken stream. Eavesdropping may start from the middle of a BB frame or because of bad weather, packet headers are missed. With such a broken stream, it is difficult to determine where to start decoding. One brute-force solution is to treat every byte as a potential header. However, it is impractical in the real world because such decoding speed is too low to catch up with the transmission speed: with the best practice, we are only able to decode 609.71 MB in one hour, not identifying even one complete header in the highly corrupted stream we eavesdropped.

The second challenge is corrupted critical fields. One critical field is the length of packets. Even if we are able to identify the BB header or GSE header, for decoding, we need to determine the length of data fields at each layer so as to extract PDUs. Unfortunately, these length fields are corrupted and thus untrustworthy. Besides, signal loss can happen. Therefore, we cannot use the size of inner packets to correct the length field in outer packets. Without the right length, the decoding can hardly carry on. Another critical field is the protocol type in IP packets because once it is corrupted, Wireshark cannot determine which transport layer protocol is used and fails to analyze PDUs.

The third challenge is nonfunctional CRC-32. The error detection code is designed to examine if PDUs are corrupted. In a clear and high-quality communication channel, CRC-32 helps find corruption. However, our experiment results show that, in eavesdropping, corruption is so common that CRC-32 itself can also be problematic. Therefore, we cannot rely on it to correct our decoding results.

### III. DESIGN OVERVIEW

To resolve the technical challenges discussed above, in this work, we propose to employ state-of-the-art contrastive-learning techniques with data augmentations to decode and recover corrupted satellite streams. Instead of relying on critical fields or CRC-32 for decoding, contrastive learning directly learns the features of packet headers at different layers and identifies them in a stream sequence. By filtering them out, we can extract the innermost IP packets. Then we correct the transport layer protocol field in IP headers so that the extracted

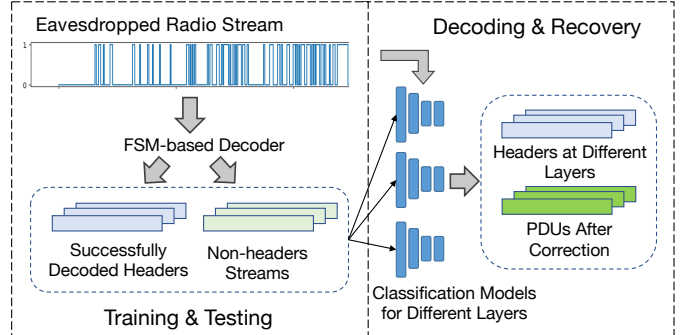


Fig. 2: The recovering workflow. We first train a classification model using the eavesdropped data and then apply this model to identify headers even if they are corrupted and extract PDUs after correction.

IP packets can be further analyzed by tools like Wireshark. Figure 2 shows the workflow of our approach.

To start, we first construct a training and testing dataset that consists of positive data (*i.e.*, successful decoded headers) and negative data (*i.e.*, non-headers streams). Applying FSM-based GSEextract on the eavesdropped data, we collect positive data - BB headers, GSE headers, and IP headers that can be successfully decoded. Although successful decoding does not mean these headers are 100% correct because the fields not used in FSM-based decoding can still be problematic, we deem these headers not corrupted and use them for training. This flaw in the training data, as we will describe in Section IV and show in the evaluation (Section V), won't influence the accuracy of our approach with the assistance of a pre-trained encoder network. Eliminating headers from successfully decoded streams, we obtain negative data - non-headers streams. With these two types of streams, we train our classification model. Given a sequence of bytes, this model can tell whether it is a header or not.

We trained three instances of this classification model to identify BB headers, GSE headers, and IP headers, respectively. We first apply the BB header model to divide the whole eavesdropped streams into BB frames. Within each BB frame, we apply the GSE header model to identify GSE layer packets. Then, we apply the IP header model to determine whether a GSE layer packet includes an IP header. With the extracted IP packets in hand, we use Hamming distance to correct the transport layer protocol field in the IP header if it is corrupted. After correction, the PDUs can be fed to tools like Wireshark for further analysis.

### IV. TECHNICAL DETAILS

In this section, we cover the technical details of our approach. First, we present the framework of our contrastive learning based classification model. Then, we take a scrutiny into the framework, elaborating on how it pre-trains an encoder network and how the encoder network is used to fine-tune a classification model.

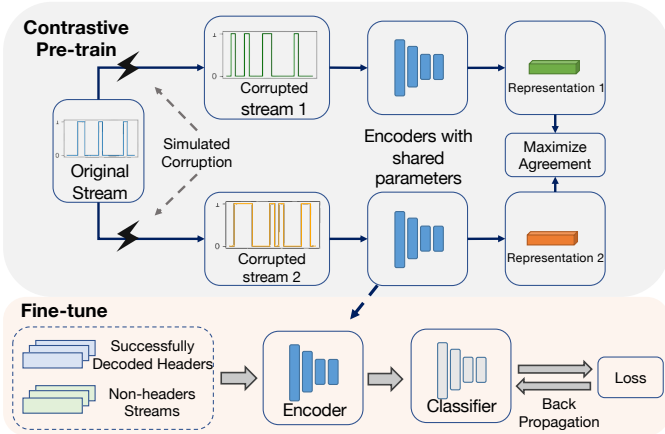


Fig. 3: The framework of proposed contrastive learning based classification model. The framework first pre-trains an encoder network and then uses it to fine-tune a supervised classification model.

#### A. Contrastive Learning Based Classification Framework

As a state-of-the-art machine learning paradigm, contrastive learning has achieved great success in the computer vision [7][10] and natural language processing domains [19][12]. Figure 3 illustrates the overall framework of our proposed contrastive learning based classification model. The first component of the framework is a self-supervised contrastive pre-train which generates an encoder network. This encoder network is used in the second component to fine-tune a classifier model which predicts whether the stream input is a header or not.

As the essential part of this framework, the encoder network learns a fixed dimensional vector representation for a stream input. Without technical details, Figure 4 illustrates the difference between representation with and without a pre-trained encoder network. A pre-trained encoder network can learn the features of non-corrupted headers in the meanwhile cluster corrupted headers with non-corrupted headers in the representation space. In comparison, an encoder network without pre-train will disperse corrupted and non-corrupted headers in the representation space, mixing them with non-header streams and thus influencing the effectiveness of fine-tuned classification model. Therefore, the encoder network with pre-train is robust to corruption noise, not only fixing the flaw of our training data mentioned in Section III, but also equipping the classification model with the capacity to identify corrupted headers.

#### B. Encoder Network and Contrastive Pre-train

**Architecture of Encoder network.** The encoder network maps the stream data to a fixed dimensional vector. Formally, we denote it by  $f_{\theta}(x) : \mathbb{R}^T \rightarrow \mathbb{R}^D$ . Here  $x$  is the input stream;  $T$  is the maximum length of possible headers;  $D$  is the dimensionality of the representation vector;  $\theta$  denotes the learnable parameters in the encoder network. Especially, for an input stream with a length smaller than  $T$ , we first pad it with 0x0 to make sure that all inputs are with the same length,

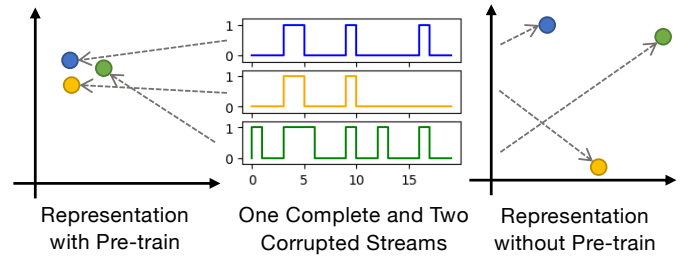


Fig. 4: The diagram illustrating the necessity of pre-train for a robust encoder network.

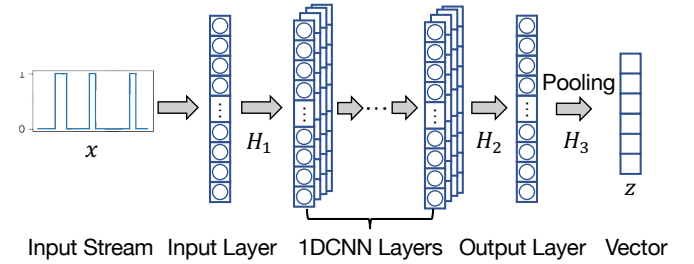


Fig. 5: The encoder network that maps the radio stream to a fixed dimensional vector.

*i.e.*,  $T$ . The architecture of the adopted encoder is shown in Figure 5. It consists of a fully connected layer (input layer), a 10-layer dilated 1-dimensional convolutional neural network (1DCNN) module [20], a fully connected layer (output layer), and a pooling layer. Compared to the vanilla 1DCNN, the dilated version has a large receptive field to capture long-range dependencies. Formally, we have

$$\begin{aligned}
 H_1 &= \text{InputLayer}(x) \\
 H_2 &= \text{1DCNNs}(H_1) \\
 H_3 &= \text{OutputLayer}(H_2) \\
 z &= \text{MeanPooling}(H_3),
 \end{aligned} \tag{1}$$

where  $H_1, H_2, H_3$  are hidden representations, and  $z \in \mathbb{R}^D$  is the representation vector.

**Data Augmentation by Simulating Corruption.** To make the encoder network robust to corruption noise, we introduce data augmentation techniques. More specifically, we simulate corruption based on successfully decoded headers and add corrupted headers to pre-train dataset. The communication corruption is usually modeled as white Gaussian noise. The noise added to each bit is independent of the others. In general, there are two types of corruption. One is bit flip and another is bit loss.

We simulate flip and loss with parameterized ratios  $\gamma_1, \gamma_2 \in [0, 1]$ . As shown in Algorithm 1, we first randomly sample an array  $\mathbf{d}$  of  $|x|$  real numbers from  $(0, 1)$  in Line 3, each corresponding to a bit in  $x$ . First, we determine if flip happens by comparing the sampled numbers with ratio  $\gamma_1$ . If  $r_i < \gamma_1$ , we flip the  $i$ -th bit in  $x$  (line 4-6). Then, we determine if loss happens by comparing  $r_i$  with  $\gamma_1$  and  $\gamma_2$ . If  $\gamma_1 < r_i < \gamma_2$ , we fill in the lost bit with 0 to simulate receivers (line 7-9). In

---

**Algorithm 1: Simulate Bit Flip & Bit Loss**

---

- 1: **Input:** a stream instance  $x$ , flip ratio  $\gamma_1$ , loss ratio  $\gamma_2$ ;
  - 2: **output:** Simulated corrupted stream instance  $x'$ ;
  - 3:  $\mathbf{r} \leftarrow$  random sample  $|x|$  numbers from Uniform( $(0, 1)$ );
  - 4:  $\mathbf{d}_1 \leftarrow \mathbf{r} - \gamma_1$ ;
  - 5:  $\mathbf{s}_1 \leftarrow \text{sign}(\mathbf{d}_1)$ ;
  - 6:  $x' \leftarrow x * \mathbf{s}_1 + (1 - x) * (1 - \mathbf{s}_1)$  ;
  - 7:  $\mathbf{d}_2 \leftarrow \mathbf{r} - (\gamma_1 + \gamma_2)$ ;
  - 8:  $\mathbf{s}_2 \leftarrow \text{sign}(\mathbf{d}_2)$ ;
  - 9:  $x'' \leftarrow x' * (1 + \mathbf{s}_2 - \mathbf{s}_1)$  ;
  - 10: **return**  $x''$ .
- 

other situations, we keep the bit. As such, the total corruption ratio is  $\gamma_1 + \gamma_2$  with  $\gamma_1$  for flip and  $\gamma_2$  for loss respectively.

**Contrastive Pre-train.** Drawing inspiration from recent self-supervised learning algorithms in computer vision and natural language processing, we adopt, a simple but effective contrastive learning framework, SimCLR framework [7] to learn representation vectors for input streams. SimCLR maximizes agreement between differently augmented streams of the same data example by using a contrastive loss in the latent space. Specifically, it contrasts the augmented streams generated from the same input (*i.e.*, positive pairs) by pulling them close in the representation space, while pushing apart the augmented streams generated from different inputs (*i.e.*, negative pairs).

Technically, for the contrastive pretrain objective, we follow SimCLR [7] to use the normalized temperature-scaled cross entropy loss (NT-XEnt) as the contrastive loss. Algorithm 2 summarizes our contrastive pretrain procedure. Specifically, we utilize the cosine similarity to define the similarity between two vectors  $\mathbf{z}_i$  and  $\mathbf{z}_j$ . Formally,

$$s(\mathbf{z}_i, \mathbf{z}_j) = \frac{\mathbf{z}_i \cdot \mathbf{z}_j}{\|\mathbf{z}_i\|_2 \cdot \|\mathbf{z}_j\|_2} \quad (2)$$

Given a batch of  $N$  stream instances, denote by  $\{x_i\}_{i=1}^N$ , we corrupt each instance twice to get  $2N$  augmented streams. The ones corrupted from the same  $x$  are considered positive pairs. For a positive pair  $(\mathbf{z}_i, \mathbf{z}_j)$ , we have its contrastive loss

$$\ell_{(i,j)} = -\log \frac{\exp(s(\mathbf{z}_i, \mathbf{z}_j)/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(s(\mathbf{z}_i, \mathbf{z}_k)/\tau)}, \quad (3)$$

where  $\tau$  is the temperature parameter and  $\mathbb{1}_{[k \neq i]}$  is the indicator function defined as follows.

$$\mathbb{1}_{[k \neq i]} = \begin{cases} 0 & \text{if } k = i \\ 1 & \text{if } k \neq i. \end{cases} \quad (4)$$

Intuitively, minimizing  $\ell_{(i,j)}$  encourages the model to identify the positive partner  $\mathbf{z}_j$  from  $2N$  vectors for a given  $\mathbf{z}_i$ . The batch loss is then computed by averaging all positive pairs in a mini-batch. Formally, we have

$$\mathcal{L} = \frac{1}{2N} \sum_{i=1}^N [\ell_{(2i-1, 2i)} + \ell_{(2i, 2i-1)}]. \quad (5)$$

### C. Supervised Fine-tune and Classification Model

After pre-train, the encoder network is robust to noise. We build our classifier on top of that to identify headers in the

---

**Algorithm 2: Contrastive Pre-train**

---

- 1: **Input:** a set of stream instances  $\{x_i\}$ , batch size  $N$ , temperature  $\tau$ , encoder networks  $f$ ;
  - 2: **output:** Pretrained encoder networks  $f$ ;
  - 3: **for** each minibatch  $\{x_i\}_{i=1}^N$  of stream instances **do**
  - 4:   **for** each instance  $x_i$  **do**
  - 5:      $\tilde{x}_{i1} \leftarrow$  apply simulated corruptions on  $x_i$ ;
  - 6:      $\tilde{x}_{i2} \leftarrow$  apply simulated corruptions on  $x_i$ ;
  - 7:      $\mathbf{z}_{2i-1} \leftarrow f(\tilde{x}_{i1})$ ;
  - 8:      $\mathbf{z}_{2i} \leftarrow f(\tilde{x}_{i2})$ ;
  - 9:   **end for**
  - 10:   compute batch loss  $\mathcal{L}$  with Eq. (5);
  - 11:   update parameters in  $f$  by minimizing  $\mathcal{L}$ .
  - 12: **end for**
  - 13: **return** encoder network  $f$ .
- 

---

**Algorithm 3: Supervised Fine-tune**

---

- 1: **Input:** a set of labeled stream instances  $\{(x_i, y_i)\}$ , batch size  $N$ , encoder networks  $f$ ;
  - 2: **Output:** classifier  $g$ ;
  - 3: Inputlayer, 1DCNNs  $\leftarrow$  extract from encoder network  $f$ ;
  - 4: classifier  $\leftarrow$  initiate 2 fully connected layers
  - 5: detector  $g \leftarrow$  stack Inputlayer, 1DCNNs, and classifier
  - 6: **for** each minibatch  $\{(x_i, y_i)\}_{i=1}^N$  **do**
  - 7:   **for** each instance  $x_i$  **do**
  - 8:      $\hat{y}_i \leftarrow g(x_i)$
  - 9:     compute classification loss  $\ell_i = \text{CrossEntropy}(\hat{y}_i, y_i)$ ;
  - 10:   **end for**
  - 11:   compute batch loss  $\mathcal{L}_{cls} = \frac{1}{N} \sum_{i=1}^N \ell_i$ ;
  - 12:   update parameters in classifier by minimizing  $\mathcal{L}_{cls}$ .
  - 13: **end for**
  - 14: **return** detector  $g$ .
- 

corrupted streams. For efficiency, we extract the Inputlayer and 1DCNNs from the pre-trained encoder network and stack them with 2 fully connected layers as the classifier. Then, we fine-tune the classification model with supervised successfully decoded headers, labeled with 1, and randomly sampled non-header streams, labeled with 0. During this step, we kept the Inputlayer and 1DCNNs frozen and don't update parameters in these layers. Since we have two labels, headers and non-headers, we adopt the binary cross entropy as the loss function. Finally, we fine-tune the model under supervision using Algorithm 3.

To recover the corrupted transport layer protocol field in IP headers, we adopt Hamming distance to calculate the distance between the potentially corrupted protocol value and the non-corrupted protocol value in the training dataset. The protocol value is automatically corrected to the target protocol with the smallest distance.

## V. EVALUATION AND ETHICS

### A. Experiment Setup and Ethics

In our experiment, we built a platform to eavesdrop on satellite communication data. This platform consists of a TBS-6903 Professional DVB-S2 Dual Tuner PCIe Card and a professionally customized Ku-band antenna that can automatically align the dish. The total cost of this platform is around \$15k. We

deployed the platform in a suburban area of a metropolis with more than 10 million of population in Asia, eavesdropping on the spectrum range from 11 GHz to 12.75 GHz which covers seven commercial satellites. The eavesdropping was conducted over 20 days from July to October. Finally, we received 23.6 GB DVB/GSE stream data.

Considering that sensitive information could be included in the stream, we followed ethical principles proposed by the prior work [16], not storing any data longer than necessary. Even for the learning model trained using the eavesdropped data, we deleted it immediately after completing the evaluation to prevent adversary data generation from the model (e.g., GAN [9]). In the experiment, we treated data units in IP packets as normal payloads and didn't make any attempt to decrypt them. For the sake of anonymity, we chose not to reveal the specific names of satellites and service providers that were eavesdropped on. Though we will open-source our implementation CLEExtract after this work is accepted, we withhold the publishing of training data. This is our best effort to bolster future research. Here, we advocate service providers or authorities to build a NVD-like database (National Vulnerability Database) [5] which stores anonymized and desensitized data. It will significantly advance research on space security.

From the eavesdropped data, we ran GSEExtract to extract BB frames that can be successfully decoded in an FSM-based approach. From these BB frames, we filter out BB headers, GSE headers, and IP headers - successfully decoded headers. Eliminating these headers, we obtain non-header streams. The two types of data construct our training and testing dataset. Using this approach, we collected in total 10471 BB frames (0.5 GB) from all streams (23.6 GB). This low success decoding rate (2.1%) indicates that corruption is very common in satellite communication. From such highly corrupted streams, as we will further show in our evaluation, tools like GSEExtract which are built upon the traditional FSM-based decoding approach can only recover a very small portion of data.

With the training and testing data in hand, we evaluate CLEExtract. We divided the whole data set into two parts: 2/3 of successfully decoded headers and non-headers stream to train our model, and the remaining 1/3 to measure its **effectiveness**. During data augmentation in contrastive pre-train, we set the flip ratio and loss ratio as 10% and 10% respectively. To **compare** CLEExtract with GSEExtract, we apply both to streams with different corruption degrees and corruption types. More specifically, we synthesize corruption using the same algorithm for data augmentation. Starting from 2%, we gradually increase the corruption degree ( $\gamma_1 + \gamma_2$ ) by a step of 2%, until 20%. In the meanwhile, we adjust the relative ratio flip( $\gamma_1$ ) : loss( $\gamma_2$ ) to 1:3, 1:1, and 3:1 to examine whether the types of corruption affect the robustness. Since the model is trained over successfully decoded streams with data augmentation, there is no label leakage to apply the model for corrupted streams which can be considered as a distinct dataset. To evaluate to which extent **pre-trained network**

with **data augmentation** improves the robustness, we compare the effectiveness of CLEExtract with and without data augmentation. Finally, we run CLEExtract and GSEExtract using a server with an Intel(R) Xeon(R) Gold 6258R CPU @ 2.70GHz, 1.5TB RAM, and an NVIDIA A100 80GB PCIe GPU, showing the efficiency of CLEExtract. Each experiment mentioned above is repeated for 10 rounds and we report the average results.

## B. Evaluation Results of CLEExtract

**Effectiveness.** We use four metrics to measure the effectiveness of CLEExtract. The first metric is ACC which stands for the ratio of the number of correct predictions to the total number of input samples. In our scenario, it is calculated as  $(TP+TN)/(TP+TN+FP+FN)$ . True positive (TP) means the headers identified by CLEExtract are indeed headers, and true negative (TN) means CLEExtract doesn't mistakenly identify non-header streams as headers. The second metric is Precision which is calculated as  $TP/(TP+FP)$ . The third metric is Recall which is  $TP/(TP+FN)$ . The last metric is F1 - the harmonic mean of Precision (P) and Recall (R). It is calculated as  $(2/F1 = 1/P + 1/R)$ . F1 is high if and only if both Precision and Recall are high.

In Table I, we show the effectiveness of CLEExtract when the relative corruption ratio flip( $\gamma_1$ ) : loss( $\gamma_2$ ) is 1:1. Overall, the four metrics for header identification at different layers are all very high no matter how corrupted the stream is - all above 0.71.3 while most are over 0.9. This indicates that CLEExtract can accurately identify headers in a corrupted satellite stream.

From the table, we can further observe that, in comparison with BB headers and IP headers, the effectiveness of CLEExtract in identifying GSE headers is relatively poor. The reasons are two-fold. On the one hand, unlike BB headers, the length of GSE headers is variable, ranging from 2 bytes to 12 bytes. In training, we have to pad the short GSE headers with 0x0 so that they can be uniformly represented in the encoder network. However, the padded 0x0s carry no information and mislead CLEExtract in classification if the non-header stream contains too many 0x0s. On the other hand, unlike IP headers the length of which is at least 20 bytes, the length of GSE headers is relatively small. Therefore, the classification model fails to learn enough features in training.

Enough if the effectiveness of CLEExtract in identifying GSE headers is not as good as identifying BB headers and IP headers, in comparison with GSEExtract, CLEExtract is much better. In Table II, we compare the numbers of identified GSE headers using CLEExtract with the number using GSEExtract. Due to the variable length and the small size nature mentioned above, the number of identified GSE headers by CLEExtract is slightly lower than GSEExtract when the corruption degree is 0.02. However, when corruption becomes more severe, regardless of the ratio of corruption types, CLEExtract can successfully identify much more GSE headers than GSEExtract (15707 vs. 516 when  $\gamma_1 + \gamma_2=20\%$  and  $\gamma_1 : \gamma_2=3:1$ ).

Corruption Degree ( $\gamma_1 + \gamma_2$ )	0.02	0.04	0.06	0.08	0.10	0.12	0.14	0.16	0.18	0.20
BB Headers										
ACC	0.984	0.981	0.976	0.971	0.964	0.955	0.946	0.935	0.921	0.914
Precision	0.977	0.976	0.974	0.973	0.973	0.969	0.966	0.962	0.959	0.957
Recall	0.992	0.986	0.979	0.969	0.957	0.941	0.927	0.908	0.883	0.870
F1	0.985	0.981	0.977	0.971	0.965	0.955	0.946	0.934	0.920	0.912
GSE Headers										
ACC	0.883	0.870	0.854	0.839	0.828	0.813	0.797	0.784	0.769	0.757
Precision	0.851	0.845	0.838	0.831	0.828	0.820	0.812	0.804	0.792	0.788
Recall	0.933	0.910	0.882	0.856	0.833	0.807	0.780	0.758	0.737	0.713
F1	0.890	0.877	0.859	0.843	0.830	0.813	0.796	0.780	0.764	0.748
IP Headers										
ACC	0.970	0.971	0.972	0.971	0.972	0.971	0.971	0.971	0.970	0.969
Precision	0.962	0.962	0.962	0.961	0.962	0.960	0.960	0.960	0.957	0.954
Recall	0.980	0.981	0.983	0.982	0.982	0.983	0.984	0.984	0.985	0.985
F1	0.971	0.972	0.972	0.972	0.972	0.971	0.972	0.972	0.971	0.969

TABLE I: Effectiveness and robustness of CLEExtract with different corruption degrees ( $\gamma_1 + \gamma_2$ ) while the ratio of corruption types is fixed to  $\gamma_1 : \gamma_2 = 1 : 1$ .

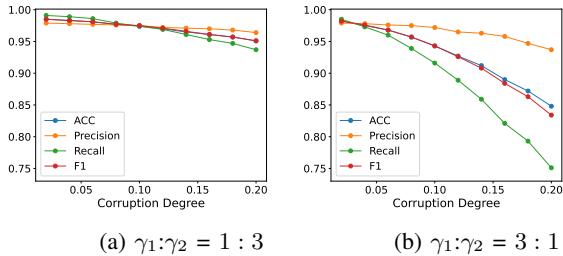


Fig. 6: Corresponds to Table I, effectiveness and robustness of BB header identification with different corruption degrees ( $\gamma_1 + \gamma_2$ ) and ratios ( $\gamma_1 : \gamma_2$ ). More corresponding results are in Figure 9 and 10.

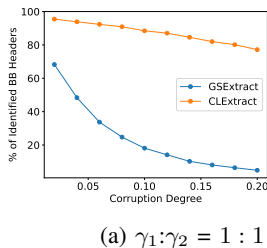


Fig. 7: Corresponds to Table II, comparison between CLEExtract and GSEExtract in BB header identification with different corruption degrees ( $\gamma_1 + \gamma_2$ ) and fixed ratio ( $\gamma_1 : \gamma_2$ ). More corresponding results are in Figure 11 and 12.

As Table I shows the results when the ratio of corruption types flip( $\gamma_1$ ) : loss( $\gamma_2$ ) is 1:1, we present the results for different ratios in Figure 6 for BB headers, Figure 9 for GSE headers, and Figure 10 for IP Headers (Figure 9 and 10 are in Appendix A). From the three figures, we can observe the effectiveness of CLEExtract in recovering corrupted streams, which aligns with the results in Table I and II.

**Robustness.** We measure the robustness of CLEExtract against different corruption degrees and ratios of corruption types. Still, from Table I, we can see that the four effectiveness metrics are not lowered too much when corruption becomes more severe. Taking BB header as an example, its ACC,

Precision, Recall, and F1 only drop 7%, 2%, 12%, and 7%, respectively, when the corruption degree increases from 2% to 20%. Similar trends are also presented in Figure 6, 9, and 10 that correspond to Table I.

From Table II, we can see that, in general, CLEExtract performs well when the ratio of corruption types changes. More specifically, CLEExtract identifies 16734 GSE headers when the corruption degree is 2% and flip( $\gamma_1$ ) : loss( $\gamma_2$ ) is 1:3. If we fix the corruption degree and switch the ratio to 3:1, the number is 16682 which is only 50 fewer. When the corruption degree increases to 20%, this gap becomes 1017 which takes over 4.2% of the total number of GSE headers. In comparison, GSEExtract is significantly influenced by the ratio of corruption types. From Table II, we can observe that the number of identified GSE headers drops 83.3% (2956 to 516) if the ratio switches from 1:3 to 3:1. It indicates that GSEExtract is more likely to be dis-functioned by bit flip than bit loss. This is because when a bit loss happens, there is still a 50% chance that the value is correct. However, when a bit flip happens, the value becomes completely wrong. For GSEExtract which heavily relies on critical information in headers, a wrong value can ruin the following decoding, leading to the missing of many headers. While for CLEExtract, it learns header features as a whole and is less influenced by the information loss of bit flip.

In Figure 7, we present the comparison in BB header identification with fixed  $\gamma_1 : \gamma_2$ . We can observe that when the corruption degree is 0.2, GSEExtract almost malfunctions while CLEExtract can identify around 80% BB headers. Similar results are shown in BB header and IP header identification with different ratios (Figure 11 and 12 in Appendix A). As such, we conservatively conclude that CLEExtract shows strong robustness and can recover satellite streams even if it is corrupted up to at least 20%.

**Data Augmentation.** Recall that we employ data augmentation to improve the robustness of CLEExtract, here, we evaluate to which extent data augmentation achieves this goal.

Corruption Degree ( $\gamma_1 + \gamma_2$ )	0.02	0.04	0.06	0.08	0.10	0.12	0.14	0.16	0.18	0.20
Flip( $\gamma_1$ ) : Loss( $\gamma_2$ ) = 1 : 1										
GSEextract	17905	13639	9849	7231	5070	3838	2738	1997	1535	1114
CLEextract	16684	16663	16637	16601	16541	16444	16461	16340	16253	16275
Flip( $\gamma_1$ ) : Loss( $\gamma_2$ ) = 1 : 3										
GSEextract	18911	15617	12918	10562	8851	7040	5427	4509	3626	2956
CLEextract	16734	16758	16755	16776	16697	16723	16704	16705	16672	16724
Flip( $\gamma_1$ ) : Loss( $\gamma_2$ ) = 3 : 1										
GSEextract	16942	11087	7214	4598	2826	1981	1411	990	639	516
CLEextract	16682	16615	16511	16448	16336	16170	15995	15998	15697	15707

TABLE II: Comparison with GSEextract in terms of the number of identified GSE headers, with different corruption degrees ( $\gamma_1 + \gamma_2$ ) and different ratios of corruption types ( $\gamma_1 : \gamma_2$ ). The total number of GSE headers is 24287.

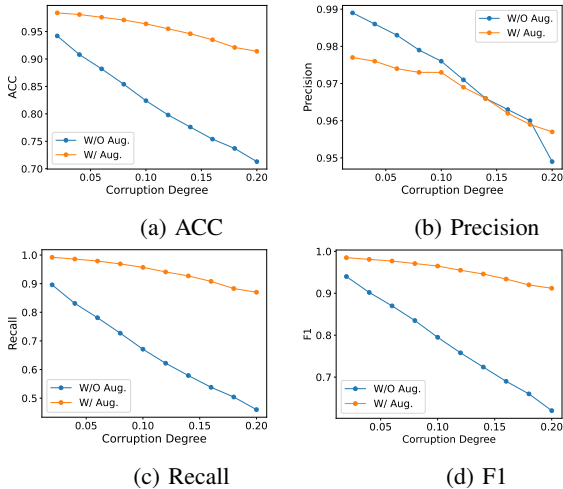


Fig. 8: Effectiveness and robustness of BB header identification w/ data augmentation ( $\gamma_1 + \gamma_2 = 20\%$ )

In Figure 8 (and Figure 13 14 in Appendix A), we present the effectiveness and robustness of identifying BB headers, GSE headers, and IP headers, with and without data augmentation. We can clearly see that data augmentation significantly improves the effectiveness and robustness of CLEextract with the orange line lying over the green line. The only exception is the Precision metric. Actually, this is within expectation. Intuitively, data augmentation blurs the boundary between corrupted headers and non-header streams. Therefore, the model with data augmentation will mistakenly identify non-header streams as corrupted headers, increasing FP when calculating the Precision metric ( $TP / (TP + FP)$ ). However, it does not mean that data augmentation makes CLEextract perform poorer. On the one hand, the difference between Precision with and without data augmentation is less than 5%. On the other hand, CLEextract aims to balance between Precision and Recall (*i.e.*, not missing TP and FN at the same time). Therefore, in terms of F1 which is a comprehension of Precision and Recall, data augmentation improves the effectiveness and robustness.

**Efficiency.** In our evaluation, we compare CLEextract with GSEextract in terms of efficiency. With all eavesdropped data,

it takes GSEextract 10,391.52 seconds to process and identify in total 10,471 BB headers, 24,287 GSE headers, and 10,479 IP headers - our training and testing set. In other words, the identification rate is 0.1 BB headers, 0.23 GSE headers, and 0.1 IP headers per second.

Using this dataset, it takes near 14.5 hours to train all models in CLEextract. Applying our models, from the whole eavesdropped data, CLEextract spends 19,870.47 seconds processing and identifying in total 498,819 BB headers, 1,156,412 GSE headers, and 497,566 IP headers. That is to say, CLEextract can identify 25.10 BB headers, 58.20 GSE headers, and 25.04 IP headers per second. Admitted that the headers pinpointed by CLEextract can be false positives, from the statistics above, we can roughly estimate that CLEextract is hundreds of times faster than GSEextract. In addition to more headers identified by CLEextract, another reason for the efficiency improvement is that CLEextract can be paralleled while GSEextract must perform decoding in a sequential fashion.

## VI. LIMITATION AND FUTURE WORK

This work has the following limitations. First, the effectiveness of GSE header identification is not as good as that of BB header identification and IP header identification. The reasons, as discussed in Section V-B, are the variable length and the small size nature of GSE header. To address this problem, in the future, we plan to include bounding box regression and intersection over union techniques [17] to further improve CLEextract. Second, we evaluate the robustness of GSEextract and CLEextract up to 20% corruption degree. We don't measure the performance when the corruption becomes more severe because 20% is high enough to render the data payload meaningless. However, as future work, we plan to further raise the corruption degree to examine how CLEextract performs in an extreme case. Third, to compare the efficiency of GSEextract and CLEextract, we count the numbers of identified headers in one second. These numbers are accurate for GSEextract but include false positives for CLEextract. We cannot eliminate false positives because we don't have the ground-truth of eavesdropped satellite streams. Though we can obtain a rough estimation from these numbers, to do compare more precisely, in the future, we plan to simulate the radio



communication with the ground-truth. We will continue to open source our simulation platform to foster future research.

## VII. CONCLUSION

In this work, we design CLExtract to decode and recover highly corrupted satellite streams. CLExtract uses a contrastive learning technique with data augmentation to learn the features of packet headers at different protocol layers and identify them in a stream sequence. By filtering them out, CLExtract extracts the innermost data payload that can be further analyzed by tools like Wireshark. Compared with the state-of-the-art GSExtract, CLExtract can successfully recover 71-99% more eavesdropped data hundreds of times faster than GSExtract. Moreover, the effectiveness of CLExtract is not largely damaged when corruption becomes more severe.

## REFERENCES

- [1] "Safeguarding Our Satellites From the Sun," 2006, [https://www.nasa.gov/mission\\_pages/stereo/news/stereo\\_1859.html](https://www.nasa.gov/mission_pages/stereo/news/stereo_1859.html).
- [2] "Digital Video Broadcasting (DVB); Frame structure channel coding and modulation for a second generation digital transmission system for cable systems (DVB-C2)," 2015, [https://dvb.org/wp-content/uploads/2019/12/al38\\_dvb-c2\\_spec.pdf](https://dvb.org/wp-content/uploads/2019/12/al38_dvb-c2_spec.pdf).
- [3] "Digital Video Broadcasting (DVB); Generic Stream Encapsulation (GSE); Part 1: Protocol," 2015, [https://www.etsi.org/deliver/etsi\\_ts/102600\\_102699/10260601/01.02.01\\_60/ts\\_10260601v010201p.pdf](https://www.etsi.org/deliver/etsi_ts/102600_102699/10260601/01.02.01_60/ts_10260601v010201p.pdf).
- [4] "Does Weather Affect Satellite Internet?" 2022, <https://www.bestsatelliteoptions.com/resources/does-weather-affect-satellite-internet>.
- [5] "National Vulnerability Database," 2023, <https://nvd.nist.gov/>.
- [6] N. Boschetti, N. G. Gordon, and G. Falco, "Space Cybersecurity Lessons Learned from the ViaSat Cyberattack," in *AIAA Accelerating Space Commerce, Exploration, and New Discovery (Ascend)*, 2022.
- [7] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*, 2020, pp. 1597–1607.
- [8] G. Giuliani, T. Ciussani, A. Perrig, and A. Singla, "ICARUS: Attacking low Earth orbit satellite networks," in *The Proceedings of the 2021 USENIX Annual Technical Conference (USENIX ATC)*, 2021.
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communication of ACM*, p. 139–144, oct 2020.
- [10] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *CVPR*, 2020, pp. 9729–9738.
- [11] S. D. Ilcev, "Surface Reflection and Local Environmental Effects in Maritime and other Mobile Satellite Communications," *the International Journal on Marine Navigation and Safety of Sea Transportation*, 2011.
- [12] L. Logeswaran and H. Lee, "An efficient framework for learning sentence representations," *arXiv preprint arXiv:1803.02893*, 2018.
- [13] M. Manulis, C. P. Bridges, R. Harrison, V. Sekar, and A. Davis, "Cyber security in New Space," in *International Journal of Information Security (IJIS)*, 2020.
- [14] U. of Concerned Scientists, "UCS Satellite Database," <https://www.ucsusa.org/nuclear-weapons/space-weapons/satellite-database>, 2019.
- [15] J. Pavur, "On Detecting Deception in Space Situational Awareness," in *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security (AsiaCCS)*, 2021.
- [16] J. Pavur, D. Moser, M. Strohmeier, V. Lenders, and I. Martinovic, "A tale of sea and sky on the security of maritime vsat communications," in *2020 IEEE Symposium on Security and Privacy (SP)*, 2020.
- [17] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [18] M. Weinzierl, "Space, the Final Economic Frontier," *Journal of Economic Perspectives*, vol. 32, no. 2, pp. 173–192, 2018.

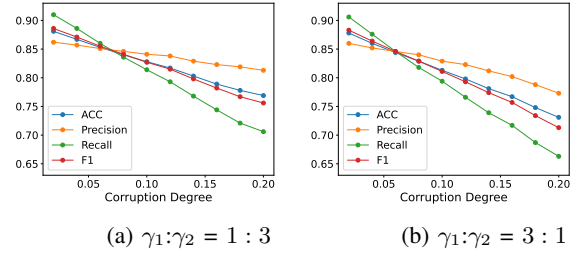


Fig. 9: Corresponds to Table I, effectiveness and robustness of GSE header identification with different corruption degrees ( $\gamma_1 + \gamma_2$ ) and ratios ( $\gamma_1 : \gamma_2$ ).

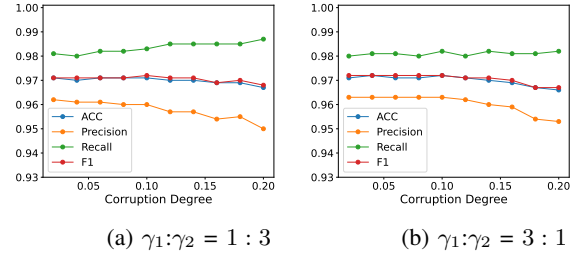


Fig. 10: Corresponds to Table I, effectiveness and robustness of IP header identification with different corruption degrees ( $\gamma_1 + \gamma_2$ ) and ratios ( $\gamma_1 : \gamma_2$ ).

- [19] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "Xlnet: Generalized autoregressive pretraining for language understanding," *Advances in neural information processing systems*, vol. 32, 2019.
- [20] Z. Yue, Y. Wang, J. Duan, T. Yang, C. Huang, and B. Xu, "Ts2vec: Towards universal representation of time series," *arXiv preprint arXiv:2106.10466*, 2021.

## APPENDIX

Due to space limit, we move some evaluation results to Appendix. For effectiveness and robustness evaluated over testing data, refer to Figure 9 and 10. For comparison between CLExtract and GSExtract, refer to Figure 11 and 12. For data augmentation, refer to Figure 13 and 14.

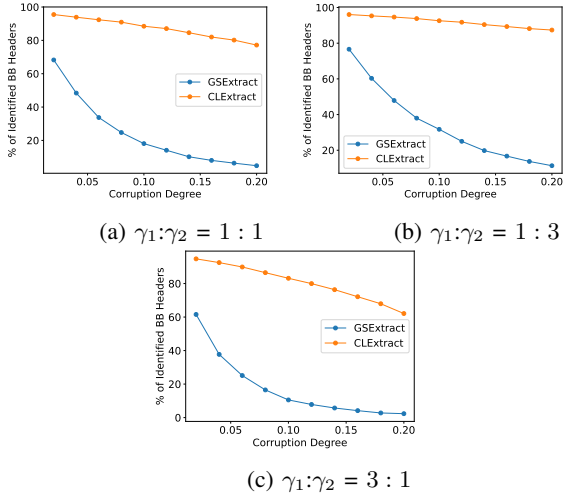


Fig. 11: Corresponds to Table II, comparison between CLExtract and GSEExtract in BB header identification with different corruption degrees ( $\gamma_1 + \gamma_2$ ) and ratios ( $\gamma_1 : \gamma_2$ ).

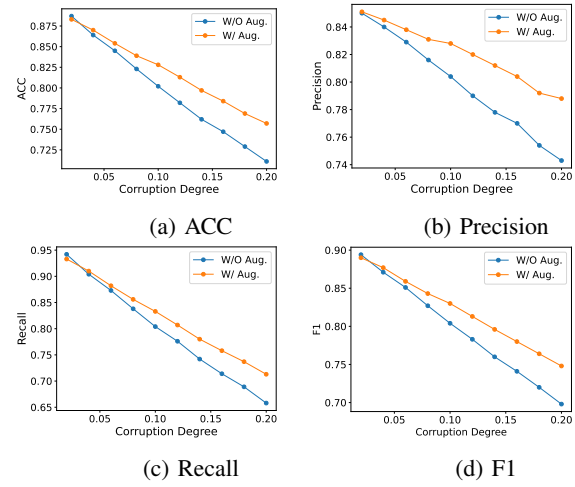


Fig. 13: Corresponds to Figure 8, effectiveness and robustness of GSE header identification w/ data augmentation ( $\gamma_1 + \gamma_2 = 20\%$ )

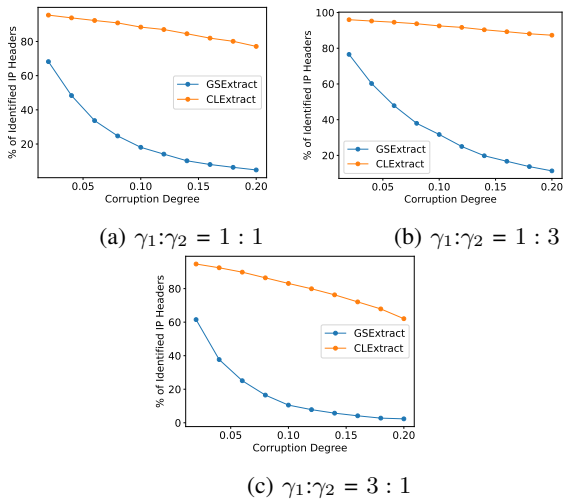


Fig. 12: Corresponds to Table II, comparison between CLExtract and GSEExtract in IP header identification with different corruption degrees ( $\gamma_1 + \gamma_2$ ) and ratios ( $\gamma_1 : \gamma_2$ ).

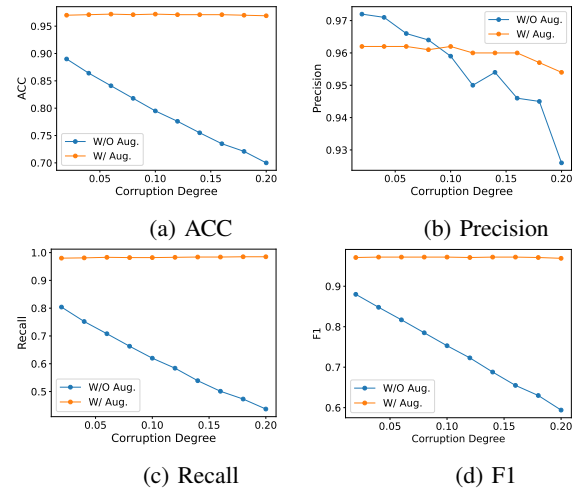


Fig. 14: Corresponds to Figure 8, effectiveness and robustness of IP header identification w/ data augmentation ( $\gamma_1 + \gamma_2 = 20\%$ )