

FAKEPCD: Fake Point Cloud Detection via Source Attribution

Yiting Qu
CISPA Helmholtz Center for
Information Security
Saarbrücken, Germany
yiting.qu@cispa.de

Zhikun Zhang
CISPA Helmholtz Center for
Information Security
Saarbrücken, Germany
zhikun.zhang@cispa.de

Yun Shen
Netsapp
Bristol, United Kingdom
yun.shen@netapp.com

Michael Backes
CISPA Helmholtz Center for
Information Security
Saarbrücken, Germany
director@cispa.de

Yang Zhang
CISPA Helmholtz Center for
Information Security
Saarbrücken, Germany
zhang@cispa.de

ABSTRACT

To prevent the mischievous use of synthetic (fake) point clouds produced by generative models, we pioneer the study of detecting point cloud authenticity and attributing them to their sources. We propose an attribution framework FAKEPCD to attribute (fake) point clouds to their respective generative models (or real-world collections). The main idea of FAKEPCD is to train an attribution model that learns the point cloud features from different sources and further differentiates these sources using an attribution signal. Depending on the characteristics of the training point clouds, namely, sources and shapes, we formulate four attribution scenarios: close-world, open-world, single-shape, and multiple-shape, and evaluate FAKEPCD's performance in each scenario. Extensive experimental results demonstrate the effectiveness of FAKEPCD on source attribution across different scenarios. Take the open-world attribution as an example, FAKEPCD attributes point clouds to known sources with an accuracy of 0.82-0.98 and to unknown sources with an accuracy of 0.73-1.00. Additionally, we introduce an approach to visualize unique patterns (fingerprints) in point clouds associated with each source. This explains how FAKEPCD recognizes point clouds from various sources by focusing on distinct areas within them. Overall, we hope our study establishes a baseline for the source attribution of (fake) point clouds.¹

CCS CONCEPTS

• Security and privacy → Social aspects of security and privacy.

KEYWORDS

Fake Point Clouds, Source Attribution, Explainable Attribution

¹Our code is available at <https://github.com/YitingQu/FakePCD>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM AsiaCCS 2024, 1-5 July, 2024, Singapore

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

ACM Reference Format:

Yiting Qu, Zhikun Zhang, Yun Shen, Michael Backes, and Yang Zhang. 2024. FAKEPCD: Fake Point Cloud Detection via Source Attribution. In *Proceedings of the 2024 ACM Asia Conference on Computer and Communications Security (ACM AsiaCCS 2024)*, July 1–5, 2024, Singapore. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Point clouds are a set of data points in 3D space to represent 3D objects or scenes [15, 29]. They are commonly collected using 3D laser scanners [7, 16] and LiDAR (light detection and ranging) technology [8, 51]. Due to their ability in preserving geometric information in 3D space, point clouds have been widely used for 3D modeling scenarios, e.g., autonomous driving [15, 28, 32], robotics [9], and shape synthesis [10, 43].

Recently, deep learning techniques have accelerated the development of point cloud generative models [2, 17, 21, 31, 52]. These models can generate synthetic point clouds with different *shapes*, with each shape representing an object, such as an airplane, car, or chair. These synthetic (fake) point clouds are resembling the real-world ones, and it is difficult to distinguish them with human eyes [2, 17, 21, 31, 52](see Figure 1). This advancement eliminates the need for costly manual point cloud acquisition, thereby improving the training of various models in classification [35, 48], segmentation [8, 9], and object detection [51].

However, a coin has two sides. Despite their utility, synthetic point clouds might also be used maliciously. For instance, the data collector expects point clouds scanned by terrestrial laser scanning (TLS), the contractor, instead, provides synthetic point clouds to save cost, which may be inferior in quality and introduce the potential risk [25, 49]. In addition, an adversary could leverage LiDAR spoofing attacks [4] to inject synthetic point clouds (e.g., fake aircraft) into the UAV's LiDAR sensors. This can render the Intelligence Surveillance Reconnaissance (ISR) systems ineffective (i.e., collecting misleading military intelligence) during real-world operations. The above work exemplifies how synthetic point clouds can be generated and utilized for malicious purposes.

Research Problem. To reduce the potential security risks of synthetic point clouds, it is imperative to detect the authenticity of point clouds, i.e., the point clouds are collected from the real world or generated by different models. Identifying synthetic data from

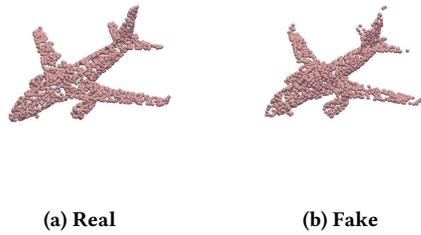


Figure 1: Examples of a real point cloud (airplane) and a fake point cloud. The fake point cloud is generated with Progressive Deconvolution Generation Network (PDGN) [17].

generative models allows developers to implement protective measures, such as watermarking, to prevent misuse. To this end, we pioneer study into the *source attribution* [54] of point clouds: *attributing point clouds to the corresponding sources* [57], i.e., collected from the real world or generated by a specific model. For instance, given two point clouds of the *shape* airplane displayed in Figure 1, we aim to identify whether the point cloud is real or generated with a particular model, i.e., PDGN [17].

Challenges. This attribution problem leads to non-trivial challenges. First, to attribute data to their sources, existing efforts [54] on synthetic images often assume that all sources at both training time and testing time are known, i.e., close-world attribution. However, in the real-world scenario, point clouds from unknown sources are often fused with those of known sources at the test time. We need to identify point clouds from known sources and recognize those from unknown sources. We term this more challenging scenario as open-world attribution (see formulation in Section 4). Secondly, different from synthetic images that can be sampled with one model, e.g., StyleGAN [18], for synthetic point clouds, training a new generative model is required for every shape. This suggests that with a single attribution model, we typically trace the source of point clouds if they share the same shape. However, enumerating all shapes can be resource-consuming. Similar to unknown sources, point clouds of new (unseen) shapes are often fused with seen shapes during source attribution. The above two points bring challenges to the source attribution of point clouds and are investigated in our study.

Our Work. We propose a (synthetic) point cloud detection and attribution framework, FAKEPCD, to verify the authenticity of point clouds and trace their sources. The idea is intuitive. We first learn the distribution of point clouds from different sources and extract point cloud features. Then, given a new point cloud, we calculate the feature distance and assign it to the specific source based on a threshold-based criterion. Specifically, FAKEPCD comprises three stages: close-world pre-training, open-world pre-training, and threshold-based assignment. We train the encoder to extract point cloud features in a fully supervised manner in the first stage to solve the close-world attribution problem. Based on the pre-trained encoder, we train the network in a supervised contrastive manner in the second stage to learn more generalizable features. Finally, in

the threshold-based assignment stage, we measure the distances of a testing example and each type of *anchors* (fixed point clouds from known sources) as the attribution signal. We classify a point cloud as an unknown source if the attribution signal is stronger than a threshold. Otherwise, we attribute it to the closest known source.

To evaluate FAKEPCD’s performance, we construct a dataset of 99K point clouds from six sources (five generative models + real-world collection), covering six common shapes, namely, *Airplane*, *Car*, *Chair*, *Bench*, *Cabinet*, and *Lamp*. We then categorize these point clouds along two dimensions: known and unknown sources, seen and unseen shapes. Point clouds of known sources and seen shapes serve as training data for FAKEPCD to learn point cloud features. In close-world attribution, we apply FAKEPCD to attribute point clouds to known sources. In open-world attribution, we evaluate FAKEPCD’s performance both in identifying point clouds from known sources and recognizing ones from unknown sources. In both attribution scenarios, we further study the FAKEPCD’s generalizability to attributing point clouds of unseen shapes.

Additionally, we explain how FAKEPCD attributes point clouds using *critical points* [35], *a set of points within a point cloud that can significantly affect a model’s prediction*. Using these critical points, we introduce a novel approach to visualize the unique patterns in point clouds from different sources. Specifically, we calculate critical points for each source and project them to the x-y plane as depth images. By stacking a certain number of depth images from the same source, we observe unique patterns associated with each generative model, namely, fingerprints.

Results. Experimental results show that FAKEPCD perfectly identifies point clouds from known sources in close-world attribution. It also achieves an accuracy ranging from 0.72 to 0.79 when attributing point clouds of unseen shapes, which are not included in its training data. For open-world attribution, FAKEPCD attributes point clouds to known sources with 0.82-0.98 accuracy and to unknown sources with 0.73-1.00 accuracy. In this scenario, FAKEPCD demonstrates generalizability when attributing point clouds of unseen shapes, especially from known sources, with 0.58-0.84 accuracy. Interestingly, we find that if an unseen shape closely resembles seen shapes in the training data, for example, Bench (unseen) being similar to Chair (seen), then the attribution performance for Bench tends to be higher compared to other unseen shapes. Furthermore, we visualize the unique patterns associated with each source with examples of Airplane, Car, and Chair. We find that FAKEPCD focuses on different areas within point clouds depending on their sources.

Contributions. Our contributions are summarized as follows:

- We provide the first work on point cloud authenticity and source attribution both in the close-world and open-world settings.
- We propose FAKEPCD to attribute a given point cloud to its source and recognize open-world examples. We also explore FAKEPCD’s generalizability in attributing point clouds of unseen shapes.
- We introduce an approach to visualize the unique patterns (fingerprints) within point clouds for different sources. This visualization illustrates how FAKEPCD attributes point clouds back to their respective sources.

2 PRELIMINARY

2.1 Point Clouds

A point cloud is a set of discrete points O . Each point $o \in O$ is a tuple that represents the x , y , and z coordinates. These points are recorded and connected to form the surface of a 3D object in Euclidean space (see Figure 1 for illustration of point clouds). They can be captured through various methods, such as laser scanning or photogrammetry (e.g., LiDAR and Kinect). Due to their ability to preserve geometric information in 3D space, they are commonly used in fields such as computer graphics [15], manufacturing [38], robotics [34], autonomous driving [28], remote sensing [24], and architecture [29]. The common research tasks on point clouds include classification (categorizing points in a point cloud into semantic classes, e.g., aircraft) [35], segmentation (dividing a point cloud into smaller semantic pieces, e.g., wings) [35], reconstruction (creating a 3D mesh or surface from a point cloud) [46], registration (aligning multiple point clouds to a single coherent representation) [34], etc.

2.2 Point Cloud Generative Models

Inspired by the success of generative models in the image domain [42], generative models for point clouds have emerged as a type of machine learning algorithms that can generate synthetic 3D point clouds [2, 17, 21, 31, 52]. These models are trained on large datasets of real-world point clouds to learn the underlying patterns and distributions that govern the generation of 3D shapes. The main challenge in this field is to synthesize high-quality point clouds that are diverse and faithful to the training data. To address the challenge, several types of generative models for point clouds have been developed, including Variational Autoencoders (VAEs) [21], Generative Adversarial Networks (GANs) [17], normalizing flow [52] and diffusion models [31]. These models tend to focus on improving the quality of generated shapes, enhancing efficiency, and extending the models to handle more complex data distributions. We refer the audience to Shi et al. [41] for a comprehensive survey of this research area.

3 THREAT MODEL

3.1 Adversary's Goal

In this paper, we consider adversaries that have the expertise to train and run point cloud generative models. Their goal is to produce fake point clouds with generative models and use them for malicious purposes. The application scenarios of these fake point clouds are introduced in Section 3.2). For instance, if their goal is to replace the original 3D design with generated point clouds to disrupt the manufacturing process, the adversary can launch (or outsource) cyber attacks to breach the targeted manufacturing system.

3.2 Application Scenarios

We consider three real-world application scenarios where point clouds might be used maliciously and source attribution can be an effective approach to identify the misuse. Their respective details are outlined below.

- **Inauthentic Point Cloud Scanning.** In this scenario, the data collector anticipates genuine point clouds acquired

through terrestrial laser scanning (TLS). The contractor, instead of scanning the real objects, provides artificially generated point clouds to save costs. These fake point clouds may be inferior in quality and can lead to models learning imprecise distributions [25, 49].

- **Manufacturing Disruption.** In this scenario, the adversary uses the generated point clouds to disrupt the manufacturing process. Take additive manufacturing (AM), commonly known as 3D printing in the consumer space, for instance. It relies on accurate 3D models for the printing process, allowing the machine to build the object layer by layer [12]. Due to the digital nature of AM (a network that connects computers and robots), the adversary may attack the AM supply chain via cyber-attack and replace the original design with the generated point clouds. Note that these generated point clouds are visually close to the real designs, which may not be noticed by operators due to automation. Nonetheless, AM processes are sensitive to small deviations from the design [36]. Consequently, the adversary can compromise the quality of the print parts and introduce undesired deficiencies. This practical scenario has been discussed in Ye et al. [53]
- **Counter Intelligence Surveillance and Reconnaissance.** In this scenario, the adversary leverages spoofing attacks on LiDAR sensors [4, 45] to inject generated point clouds (e.g., fake aircraft) in the LiDAR field of view of unmanned aerial vehicles (UAVs). This attack undermines Intelligence Surveillance Reconnaissance (ISR) systems, misleading real-world operations with falsified military intelligence. Here, we do not consider the case that the adversary may remove point regions from LiDAR readings [3] and, consequently, UAV may not identify the surveillance targets.

3.3 Attribution Capability

To reliably trace the generative models that originally generated the suspicious point clouds, we assume the following attribution capabilities.

- **Access to Authentic Point Clouds.** The attribution framework has access to adequate point clouds that are obtained from real-world objects. Note that *shape* information (e.g., Car, Airplane) naturally comes with the point clouds. We denote these point clouds as *real* point clouds for ease of presentation in the rest of the paper.
- **Access to Existing Generative Models.** The attribution framework has access to a list of known point cloud generative models. These models are our attribution targets. We denote these models as *sources* for ease of presentation in the rest of the paper. Note that the attribution framework can trivially generate point clouds from these sources to form its training data.

In Section 6.3, we show that these two capabilities are sufficient for the attribution framework to attribute point clouds to even unknown sources. This makes our attribution framework more practical in the real world.

Table 1: Summary of attribution scenarios.

	Close-World Attribution	Open-World Attribution
Single-Shape	a single shape; known sources	a single shape; known/unknown sources
Multi-Shape	seen/unseen shapes; known sources	seen/unseen shapes; known/unknown sources

4 ATTRIBUTION PROBLEM FORMULATION

To detect whether a point cloud is generated or collected from the real world, we solve this problem by attributing point clouds to the corresponding sources. As mentioned in Section 1, two major factors affect the attribution difficulty - settings (i.e., close-world attribution or open-world attribution) and shapes. We outline them below.

Settings. Formally, assume there are \mathcal{K} known sources labeled as $\{G_1, G_2, \dots, G_{\mathcal{K}}\}$ and one additional Real class labeled as R . Each source has N point clouds sampled from a given shape (e.g., Car). Given a point cloud $x_i \in X$ and its corresponding source y_i , we consider the following settings when conducting the attribution:

- **Close-World Attribution:** If x_i is a point cloud with y_i from the existing labels $\{G_1, G_2, \dots, G_{\mathcal{K}}, R\}$, the attribution is considered successful if the predicted label $\hat{y}_i = y_i$.
- **Open-World Attribution:** If x_i is a point cloud with $y_i \notin \{G_1, G_2, \dots, G_{\mathcal{K}}, R\}$, the attribution is considered successful if $\hat{y}_i = G_{\mathcal{K}+1}$, where $G_{\mathcal{K}+1}$ denotes the unknown source.

Shape. The shapes of point clouds $S = \{s_1, s_2, \dots, s_m\}$ in the training data also exert a great influence on the attribution outcome. Intuitively, the attribution performance is likely to be better on the point clouds of seen shapes (in the training data) compared to unseen shapes. To explore the model’s generalizability in attributing point clouds of unseen shapes, *multi-shape* modeling should also be considered. For each type of setting, we further consider the following two scenarios.

- **Single-Shape Scenario:** The attribution model is trained and tested on X^{s_i} , where X^{s_i} denotes the point clouds dataset of a single shape s_i . The attribution model is regarded as *single-shape model*, as it simply attributes point clouds of the same shape, e.g., Airplane.
- **Multiple-Shape Scenario:** The attribution model is trained on $S' = \{X^{s_1}, \dots, X^{s_l}\}$ and tested on $S'' = S' \cup \{X^{s_{l+1}}, \dots, X^{s_z}\}$ where $l < z \leq m$. We regard it as *multiple-shape model*, which attributes point clouds of multiple shapes simultaneously including unseen ones. For instance, our attribution model is trained on Airplane, Car, and Chair. At test time, we attribute an unseen Bench point cloud to its source.

Summary. We consider four attribution scenarios in the paper, which are summarized in Table 1. Specifically, we perform close-world attribution and open-world attribution in both the single-shape scenario and multiple-shape scenario respectively.

5 ATTRIBUTION FRAMEWORK

5.1 Overview

We present the attribution framework of (synthetic) point clouds, FAKEPCD. Figure 2 illustrates the workflow of FAKEPCD. The main idea of our framework is to train an encoder to capture the feature distance among different sources. The point cloud is then assigned to a source based on a threshold criterion. Concretely, the framework is composed of three stages and takes all the aforementioned four scenarios into account.

Close-World Pre-training (Section 5.2). Our attribution model consists of an encoder f and a classifier g , which is trained in a fully supervised manner. This stage alone can deal with the close-world attribution problem in both single-shape and multiple-shape scenarios in our evaluation. It also serves as the stepping-stone for open-world attribution. The process is illustrated in Part A of Figure 2.

Open-World Pre-training (Section 5.3). In this stage, we employ the supervised contrastive learning approach [19] to address the point cloud attribution problem. The attribution model in this stage consists of the same encoder f and a projection head h . Each point cloud is augmented via point cloud transformations to create its augmented version. We then feed both point clouds to the two identical models with sharing weights. Based on the pre-trained encoder from the first stage, we continue pre-training the encoder f with supervised contrastive loss. The process is illustrated in Part B of Figure 2.

Threshold-Based Assignment (Section 5.4). We first preserve the open-world pre-trained encoder f and projection head h in the attribution model. We then build an anchor set that contains a small set of point clouds from known sources. Given a testing sample, we obtain its feature vector from the attribution model and then measure the distance to each example in the anchor set. We then average these distances to obtain a mean distance value for each source. Using a threshold-based criterion (see Section 5.4), we assign the example to one of the existing sources or a new one (i.e., the unknown source). The process is illustrated in Part C of Figure 2.

Note. FAKEPCD is a framework. The users can use or fine-tune different encoders to instantiate their attribution models. For ease of presentation, we use an attribution model to describe technical details in different stages in this section, then use FAKEPCD to describe evaluation results in Section 6.

5.2 Close-World Pre-training

We train the attribution model that is composed of an encoder f and a classifier g in a fully-supervised manner. The parameters of the attribution model (encoder + classifier) are optimized based on the cross-entropy loss. Employing supervised learning with cross-entropy loss is a primary solution for addressing closed-world classification problems. We resolve the close-world attribution problem in both single-shape and multiple-shape scenarios in this stage.

5.3 Open-World Pre-training

Supervised learning generally does not yield good results in open-world attribution, especially under a threshold-based criterion, as it

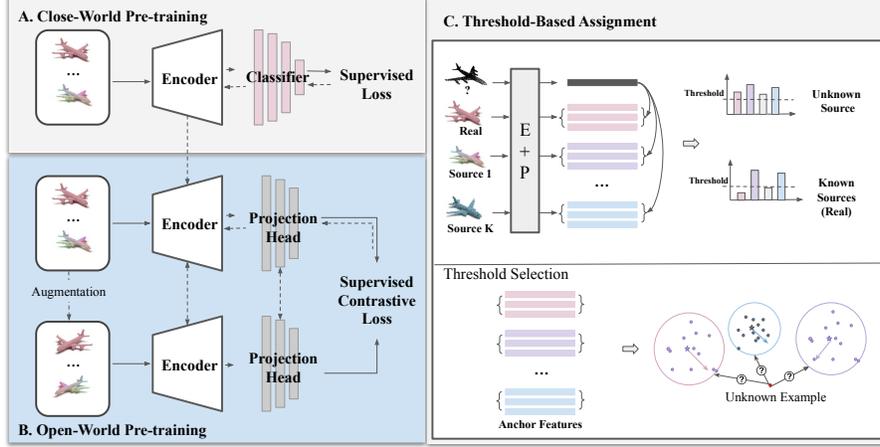


Figure 2: Workflow of FAKEPCD. “E+P” denotes the attribution model comprised of an encoder and a projection head. In the close-world pre-training stage, we train the attribution model in a fully supervised manner to solve the close-world attribution. In the open-world pre-training stage, models learn more differentiated point cloud features by supervised contrastive training. Threshold-based assignment enables us to attribute a point cloud to its source. By selecting point clouds of one-shape or multi-shape as the training dataset, this framework adapts to all the scenarios introduced in Section 4.

is not designed to handle open-world data. Our initial investigations show that the attribution model, when trained in a fully supervised manner, tends to misclassify point clouds from unknown sources to known ones with overwhelmingly high probabilities (see Section 6.3). Supervised contrastive learning [19], on the other hand, adapts the self-supervised contrastive approach to the supervised setting. With this learning paradigm, we first create an augmented version for each point cloud in the training set. We then feed both original and augmented point clouds to the encoder and projection head to obtain feature vectors. We optimize the encoder and projection head to maximize the similarity between feature vectors of positive pairs (same labels) while minimizing the similarity for negative pairs (different labels). We use the supervised contrastive loss (L) [19] to formulate the optimization process.

$$L = \min_{\theta_f, \theta_h} \left(\sum_{x_i \in X} \frac{-1}{|\mathbf{P}(x_i)|} \sum_{p \in \mathbf{P}(x_i)} \log \frac{\exp(z_i \cdot z_p / \tau)}{\sum_{a \in \mathbf{A}(x_i)} \exp(z_i \cdot z_a / \tau)} \right) \quad (1)$$

X represents the training set that contains original point clouds and their augmented versions. $\mathbf{A}(x_i) \equiv X \setminus x_i$, denotes the set of all examples excluding x_i . $\mathbf{P}(x_i) \equiv \{p \in \mathbf{A}(x_i) : \hat{y}_p = \hat{y}_i\}$ is the positive set containing examples with the same label as x_i . $z_i = h(f(x_i))$ is the output of the projection head, i.e., the feature vector of x_i . τ is the temperature scaling parameter.

We apply three types of augmentations commonly used with point clouds: *translation*, *jittering*, and *rotation*. Translation moves the origin of coordinates and builds a new coordinate frame. Jittering adds small Gaussian noise to the coordinates of points. Rotation turns point clouds with random angles along x , y , and/or z axis. Practically, when calculating feature vectors, we apply two identical encoders and projection heads, each processing either the original

point clouds or the augmented ones. During the optimization, we keep one pair of the encoder and projection head frozen and only update the other pair (see Figure 2). Note that the encoder can be randomly initialized or initialized with the checkpoint we obtained from the close-world pre-training stage. Finally, the encoders learn differentiated and generalized point cloud features. Based on these feature vectors, we further introduce a threshold-based criterion to assign point clouds to their respective sources.

5.4 Threshold-Based Assignment

We design the following threshold-based criterion to assign a testing example to either a specific known source or an unknown source. That is, *if its distances to all known sources are larger than a threshold, we recognize the example as a sample from an unknown source; otherwise, we assign it to the closest known source.* To this end, we first randomly sample point clouds from \mathcal{K} known sources, denoted as *Anchor Set*, with each source containing N point clouds. These anchor point clouds are then forwarded to the attribution model and obtain the feature vectors. Given a testing point cloud, x_i , we measure the Euclidean distance between its feature vector (z_i) and all feature vectors derived from the anchor set. We then group these distance values into \mathcal{K} clusters (C_j) based on their sources and calculate a mean distance value, $d_j(z_i)$, for each source (y_j). We compare these mean values to a selected threshold to attribute the point clouds.

$$d_j(z_i) = \frac{1}{|C_j|} \sum_{x_j \in C_j} \|z_i - z_j\|_2 \quad (2)$$

Threshold Selection. We select thresholds based on the intra-cluster distances of \mathcal{K} clusters in the anchor set. Intra-cluster distances are the distances between examples within the cluster to the cluster centroid. For example, for point clouds in the cluster (C_j), we calculate the Euclidean distance between the centroid and all point clouds in this cluster. We calculate the intra-cluster distances for \mathcal{K} clusters and obtain \mathcal{K} distance sequences, each representing the compactness of a specific source. An ideal threshold should be identified within these sequences. It should enable the assignment of a testing example from known sources to one of the \mathcal{K} clusters, while also assigning a testing example from unknown sources to a new source. However, the threshold selection is non-trivial as these \mathcal{K} distance sequences might have different distributions. To address this, we introduce another parameter, P -percentile, to control the distance threshold. The distance threshold is the smallest P -percentile value among all distance sequences. As such, threshold selection is transformed into a percentile selection problem. Note, multiple thresholds can be selected, with one corresponding to each cluster. Here, we use a unified threshold for simplicity, and this approach has already proven effective in Section 6.

6 EVALUATION

6.1 Experimental Setup

Real Point Cloud Dataset. For real-world collected point clouds, we use the dataset ShapeNetCore, a subset of ShapeNet [5]. It contains 55 object shapes such as Airplane, Car, Chair, Bench, Cabinet, and Lamp. We specifically choose these six shapes because they provide a sufficient number of samples in the dataset, ranging from 1,543 to 7,497 per shape. In contrast, other shapes are underrepresented with only hundreds of samples, making it hard to construct a sizable dataset. Each point cloud in the dataset is described by 15,000 points. To reduce computation complexity, we follow the previous studies [21, 35, 48, 52] and downsample each point cloud to 2,048 points.

Synthetic Point Cloud Dataset. We generate synthetic point clouds using five generative models: PointFlow [52], ShapeGF [2], Diffusion [31], PDGN [17], and SetVae [21]. These models cover a wide spectrum of generative techniques including Variational Autoencoders (VAEs) [2, 21], Generative Adversarial Networks (GANs) [17], normalizing flow [52] and diffusion [31]. We train each generative model using the real point clouds and generate 4,000 point clouds for Airplane, Car, and Chair, and 1,000 point clouds for Bench, Cabinet, and Lamp.

- **PointFlow.** We use PointFlow’s official implementation.² We set 128 as the feature dimension in the encoder and 512 as the hidden dimension in the CNF decoder. We train the generative models for 4,000 epochs with the Adam optimizer and set the initial learning rate to 0.002 which decays linearly after 2,000 epochs. When sampling generative examples for 6 shapes, we generate shapes and points conditioned on the shape from Gaussian priors.
- **ShapeGF.** We use ShapeGF’s official implementation.³ In detail, we first train the auto-encoder models for distribution

learning for 2,000 epochs and then train the corresponding generators for 5,000 epochs. Following the default configuration, the auto-encoder adopts 128 as the dimension of the latent code, and the generator takes a 256-dimensional vector from the Gaussian distribution and outputs the final 256-dimensional latent code. Adam optimizer is used with a learning rate of 0.0001. We acquire point clouds of six shapes individually by sampling the corresponding generators.

- **Diffusion.** We generate point clouds based on the officially released code.⁴ We select the flow-based model and set the latent dimension as 256. We then use the Adam optimizer to train the diffusion model for 80,000 iterations with an initial learning rate of 0.002. When generating new point clouds, we sample from the generative model with a prior drawn from a Normal distribution.
- **PDGN.** We use PDGN’s official implementation.⁵ PDGN adopts a progressive GAN model with multiple generators and discriminators. We use multiple resolutions of (256, 512, 1024, and 2048) point clouds from a 128-dimensional latent vector for the generator based on the original setting. Four PointNet-like modules serve as the discriminators. We train the model for 600 epochs with Adam optimizer and an initial learning rate of 0.0001. In the generating process, we sample point clouds by feeding noise vectors from the Gaussian distribution to the generator and record the results at the highest resolution.
- **SetVae.** We use the official implementation of SetVae.⁶ SetVae model has an encoder with *Induced Set Attention Blocks* and a generator with *Attentive Bottleneck Layers*. The parameters in these components are shared during both the generation and inference processes. We train the model for 8,000 epochs with the Adam optimizer and an initial learning rate of 10^{-3} that linearly decays to 0 after halfway through the training period.

Dataset Overview. Overall, we construct a dataset consisting of 99,280 point clouds from six sources (real-world collection + five generative models), covering six common shapes: Airplane, Car, Chair, Bench, Cabinet, and Lamp. We use the 3D visualization toolkit PyVista⁷ for point cloud rendering. Figure 12 in the Appendix displays generated point cloud examples of six shapes from each source.

Experimental Settings. We split the above dataset into training and test sets at a ratio of 0.6 for each source and shape. Furthermore, we leave out PDGN and SetVae as unknown sources and take the remaining four as known sources. Regarding the encoder, we adopt the feature extractor in PointNet [35] and DGCNN [48]. The classifier and projection head connected to the encoder is a multilayer perceptron (MLP). For the single-shape scenario, we train the attribution model on the Airplane, Car, and Chair independently and only test on the point clouds of the same shape. For the multiple-shape scenario, we train on three mixed shapes: Airplane,

²<https://github.com/stevenygd/PointFlow>

³<https://github.com/RuojinCai/ShapeGF>

⁴<https://github.com/luost26/diffusion-point-cloud>

⁵<https://github.com/fpthink/PDGN>

⁶<https://github.com/jw9730/setvae>

⁷<https://docs.pyvista.org/>

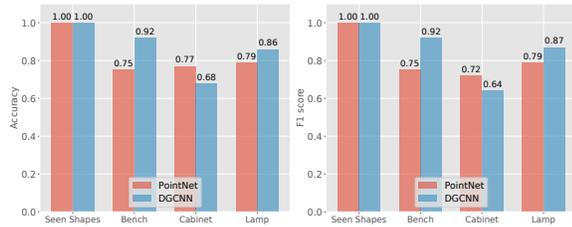


Figure 3: Attribution performance in multiple-shape scenarios when all the sources are included in the model’s training set. We report the accuracy (left figure) and F1 score (right figure).

Car, and Chair, and test on point clouds from all six shapes. Other experimental settings are outlined in their respective evaluations.

Evaluation Metrics. We measure the attribution performance for tracing both known and unknown sources, using the metrics of *accuracy* and *F1 score*.

6.2 Close-World Attribution

Setup. The attribution model in this scenario is composed of an encoder and a classifier. We employ a 3-layer MLP with (512, 256, K) neurons respectively as the classifier, where $K = 1 + 3$ with the Real source and three types of generative models (PointFlow, Diffusion, and ShapeGF). The above selection of known sources is consistent with the open-world attribution. The model is trained on 200 epochs at a batch size of 32. SGD is adopted with a learning rate of 0.1 and momentum of 0.9 to optimize the training procedure.

Single-Shape Results. We calculate the accuracy and F1 score when FAKEPCD attributes point clouds of a single shape, i.e., Airplane, Car, or Chair. We find that both FAKEPCD-PointNet and FAKEPCD-DGCNN perfectly attribute point clouds of all three shapes (accuracy=1, F1 score=1). This implies that the feature differences of point clouds from known sources can be well captured.

Multiple-Shape Results. We further test the performance of the above models on point clouds of unseen shapes, i.e., Bench, Cabinet, or Lamp. Figure 3 shows that FAKEPCD with both encoders still perfectly attribute point clouds of three seen shapes. FAKEPCD also achieves an accuracy between 0.68 and 0.92, and an F1 score ranging from 0.64 to 0.92 on unseen shapes. The results imply that the learned distinction among different sources can be transferred to unseen shapes. This also suggests that building a unified attribution model that is effective across various shapes is promising, especially when the training dataset contains a sufficient number of point clouds from diverse shapes. We also find that the DGCNN backbone outperforms the PointNet backbone in two out-of-shape shapes. We hypothesize that it is due to the robustness of the Graph Neural Network (GNN) employed by DGCNN, which better preserves the local spatial relationship among points.

Takeaways. FAKEPCD can reliably attribute seen point clouds to known sources. Moreover, the underlying model used by FAKEPCD learns sufficient knowledge during the close-world pre-training stage (see Part A in Figure 2) to attribute unseen shapes to known sources with decent accuracy.

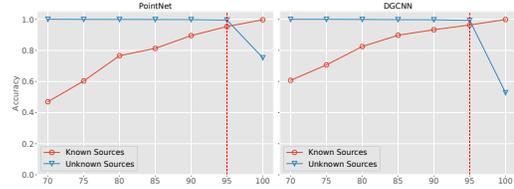


Figure 4: P -percentile selection when attributing point clouds (Car) using PointNet and DGCNN. A larger P -percentile leads to increased accuracy for known sources while decreased accuracy for unknown sources. The optimal trade-off is achieved at 95%.

6.3 Open-World Attribution

Setup. The attribution model in this scenario consists of an encoder and a projection head. For the encoder, we initialize the parameters with the checkpoint trained in close-world attribution. For the projection head, we adopt a 2-layer MLP with (512, 128) neurons. The model is trained for 300 epochs using an early stopping strategy and a batch size of 20. SGD is employed with a learning rate of 0.1 and a momentum of 0.9 to optimize the training procedure. In the assignment phase, we first construct the Anchors Set containing 100 random examples from the training set for each known source. For each unknown source, we use the test set of each shape introduced in Section 6.1. To determine the optimal threshold, we randomly sample 100 point clouds from the test set as a validation set and evaluate FAKEPCD’s performances on it across various P -percentiles. Take the shape Car as an example (see Figure 4), with a larger P -percentile, the accuracy for known sources increases, while it decreases for unknown sources. The optimal trade-off is achieved at 95%. Therefore, we adopt 95% as the optimal threshold for point clouds of Car and apply it to the remaining examples in the test set. We report the optimal thresholds for all six shapes in Table 3. We find that these optimal thresholds range from 70% to 95%, tending to be higher for seen shapes (75%-95%) and lower for unseen shapes (70%-75%). In practical scenarios, it is advisable to use thresholds within this range when a validation set of unknown sources is not available.

Baselines. Due to the absence of previous work in synthetic point cloud attribution (especially in open-world attribution), we compare our framework with several baselines that potentially work.

- **Fully-Supervised Training Approach** We train the PointNet encoder and DGCNN encoder in a supervised manner and leverage the logit probability for attribution. The models are trained for 200 epochs with a batch size of 32. We use Adam optimizer with a learning rate of 0.001. In the threshold assignment stage, we select the lowest logit probability of the training data in the ground-truth class, e.g., 0.95, and use it as a criterion to identify the sources of point clouds.
- **Projection-Based Approach** We project 3D point clouds onto 2D images and utilize a ResNet50 model to extract visual features. For each point cloud, we obtain three projections and concatenate their respective image features into one final feature, which is then fed into the final classification layer. The ResNet50 model is trained for 200 epochs with a

Table 2: Attribution performance (accuracy and F1 score) of FAKEPCD and baselines across three shapes. “P” and “D” denote the encoder (Pointnet and DGCNN). “Known” and “Unknown” indicate the attribution performance for known and unknown sources.

Method	Airplane		Car		Chair	
	Acc	F1 score	Acc	F1 score	Acc	F1 score
	Known	Unknown	Known	Unknown	Known	Unknown
Fully supervised	0.93 0.00	0.95 0.00	0.93 0.00	0.96 0.00	0.87 0.00	0.90 0.00
Projection-based	0.66 0.04	0.69 0.07	0.88 0.12	0.90 0.20	0.88 0.73	0.93 0.84
Siamese-P	0.14 0.25	0.19 0.40	0.83 0.83	0.90 0.90	0.88 0.80	0.92 0.90
Siamese-D	0.82 0.29	0.89 0.45	0.75 0.53	0.75 0.69	1.00 0.43	1.00 0.60
FAKEPCD-P	0.82 0.78	0.89 0.87	0.95 1.00	0.98 1.00	0.90 0.99	0.95 1.00
FAKEPCD-D	0.82 0.73	0.90 0.85	0.96 0.99	0.98 1.00	0.96 1.00	0.98 1.00

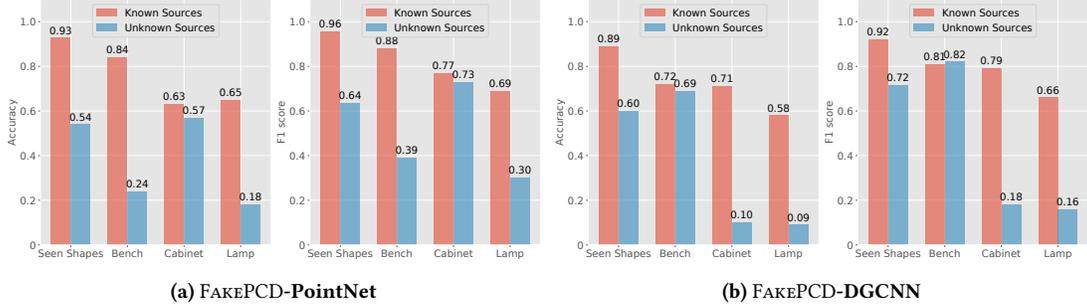


Figure 5: Attribution performance in the open-world scenario with unknown sources and unseen shapes. We report the accuracy and F1 score to evaluate FAKEPCD-PointNet and FAKEPCD-DGCNN respectively. “Seen Shapes” shows the average accuracy of Airplane, Car, and Chair.

Table 3: The Optimal P -percentile calculated with the validation set of each shape.

	Airplane	Car	Chair	Bench	Cabinet	Lamp
FAKEPCD-P	75%	95%	90%	70%	75%	70%
FAKEPCD-D	85%	95%	95%	70%	70%	70%

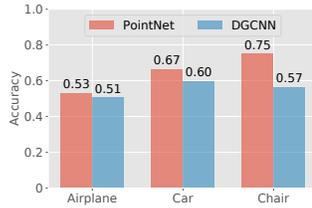


Figure 6: Accuracy of FAKEPCD in differentiating point clouds from two unknown sources: SetVae and PDGN.

batch size of 128, using Adam as the optimizer and a learning rate of 0.01. Regarding threshold assignment, we adopt the same selection strategy as that used in the fully-supervised training approach described above.

- **Siamese Training Approach** As one type of contrastive learning paradigm, a Siamese network learns to differentiate pairs of inputs [6]. A Siamese network consists of two identical encoders with the sharing parameters. It takes in a pair

of point clouds as input and outputs the logit probability that these originate from the same source. In every pair of point clouds, there is an equal probability of 50% of originating from the same source and 50% from different sources. We train the Siamese network for 200 epochs with a batch size of 20. We use the Adam optimizer, a learning rate of 0.001, and a weight decay of $1e-4$. In the attribution phase, we randomly select a pair of point clouds from a known source and a new source respectively, and the model asserts if they originate from the same source.

Single-Shape Results. Table 2 shows the comparison result of baselines and FAKEPCD across three shapes. FAKEPCD outperforms all baselines and consistently achieves high accuracy (0.73-1.00) and F1 score (0.85-1.00) when recognizing point clouds from unknown sources. For known sources, FAKEPCD achieves an accuracy between 0.82 and 0.96 and an F1 score ranging from 0.89 to 0.98. Take Car as an example, FAKEPCD-D achieves 0.98 and 1.00 F1 score on known and unknown sources respectively, while the best accuracy score from all baselines is 0.96 and 0.90. The best-performing baseline is the Siamese training approach, which shows a wide accuracy range from 0.14 to 1.00, occasionally matching FAKEPCD in certain shapes, such as Chair. For fully supervised and project-based baselines, we observe good attribution performance for known sources, however, they often fail to identify point clouds from unknown sources. This result suggests that contrastive learning (supervised contrastive learning in FAKEPCD and the Siamese baseline) has

a better performance in identifying open-world examples compared to the supervised training paradigm (fully supervised and projection-based baselines).

In the above evaluation, we use the optimal thresholds determined in Section 6.1. They are determined with the respective validation set, which has the same distribution as the testing examples. In the real world, this setting is often not realistic, as the validation set is not always available. To evaluate whether FAKEPCD can generalize to more unknown sources, we introduce two more generative models and test FAKEPCD’s performance on these new examples. The result reveals that FAKEPCD can still recognize point clouds from new unknown sources with an F1 score between 0.71-0.90. We present the details in Appendix A.

Multiple-Shape Results. Here, we evaluate the most challenging task: attributing unseen shapes to unknown sources. This scenario inevitably impacts attribution accuracy due to two-fold distractions: the unknown sources and the divergence in shape distribution caused by unseen shapes. The results are depicted in Figure 5. Although FAKEPCD fails to attribute unseen point clouds to unknown sources, it manages to attain decent performance for known sources (0.58-0.84 accuracy). For example, when attributing the unseen shape Bench, which resembles the shape of a Chair, the accuracy reaches 0.84 with FAKEPCD-PointNet and 0.72 with FAKEPCD-DGCNN for known sources.

Multiple Unknown Sources Differentiation. In the above evaluation, FAKEPCD classifies point clouds from all unknown sources into a single class, i.e., Unknown. We then investigate whether FAKEPCD can differentiate various unknown sources (PDGN and SetVae in our evaluation). To this end, we employ Gaussian clustering on the point cloud feature vectors from these two unknown sources. Specifically, we input these feature vectors into a Gaussian mixture model with two components to obtain the predicted clusters. From these clusters, we derive the classification results for the two sources. Figure 6 presents the differentiation results of three observed shapes (Airplane, Car, and Chair) from different unknown sources. The differentiation accuracy scores exceed the random guess baseline (0.5) for all three shapes. For example, FAKEPCD-PointNet achieves an accuracy of 0.75 in separating Chair features between PDGN and SetVae. This result demonstrates that FAKEPCD may further capture the differences among unknown sources.

Visual Analysis. We use t-SNE to project the point cloud features obtained from FAKEPCD to a 2-dimensional space. The visualizations are shown in Figure 13 and Figure 14 in the Appendix. Among all three shapes, Car and Chair demonstrate satisfactory clustering performances, where embedding projections from different sources are distinguishable from each other. We notice that it is difficult to separate Airplane point clouds between PointFlow and SetVae. Nevertheless, the projections from two unknown sources are separated for all three shapes, indicating that FAKEPCD can recognize more than one unknown source.

Takeaways. Based on the above analysis, we have several takeaways: 1) FAKEPCD can reliably attribute seen shapes to both known and unknown sources; 2) It can further distinguish point clouds from multiple unknown sources; 3) The attribution ability can generalize to point clouds of unseen shapes, in particular, those unseen shapes that are visually close to the seen ones.

7 EXPLAINABLE ATTRIBUTION

In this section, we discuss the nature of source attribution and introduce an approach to visualize the unique patterns in point clouds associated with each source. This explains how FAKEPCD recognizes point clouds from different sources.

7.1 Behind the Attribution

Generative models synthesize fake point clouds by sampling from a distribution that is learned using real-world data. The sampled distribution, in practice, is often biased toward the real-world data distribution. Take Generative Adversarial Neural Network (GAN) as an example. Due to the imperfect convergence during the training process, the generator and the discriminator often fall into the sub-optimal solution, such that GAN only partially describes the data distribution from the real world. The bias between the sampled distribution and real distribution has made the source attribution possible. The attribution model captures the unique bias by transforming point clouds into feature vectors and learning to differentiate them in a supervised or self-supervised fashion.

Critical Points. To explain how the attribution model differentiates point cloud feature vectors, we introduce *critical points* [35], which is defined as a set of points that significantly affect the feature vectors. Critical points contain the most geometric and semantic information about point clouds. We identify critical points of a point cloud by tracing the points that contribute the most to the maximum pooled feature obtained from PointNet. To ensure permutation invariance, which states that the order of points is invariant to the feature extraction, PointNet calculates the individual point feature or local geometric feature and then aggregates the local features into a global descriptor with a symmetric max-pooling function. For instance, PointNet transforms a point cloud of dimension (2048, 3) into a feature map of dimension (2048, 1024) before down-sampling into a global feature whose length is 1024 by max-pooling. We trace and record the points of these largest values in the feature map (input of the max-pooling layer) as critical points. Examples of critical points are demonstrated in Figure 15 in the Appendix. In this figure, we present three columns of point clouds, each representing a different type of airplane. Comparing the critical points within one column, we can see how FAKEPCD identified different points/areas in point clouds from different sources. More details on selecting these examples are presented in Appendix B.

7.2 Fingerprint Visualization

Fingerprints are persistent and identifiable patterns that could be distinguished from other sources/models. There are multiple ways to build fingerprints of image generative models [1, 54]. However, to our knowledge, none of the existing work explores the fingerprint associated with point cloud models. Here, we introduce how to visualize the fingerprints of the point cloud generative models and verify the uniqueness of the fingerprints qualitatively.

To build fingerprints, we first extract the critical points during the attribution process. Specifically, we randomly select 100 sets of critical points of a generative model (or collected from the real world) of the same shape, e.g., 100 airplanes. We then project the critical points to the same plane, e.g., x-y plane, and obtain 100 depth images. By stacking these depth images to form an average depth

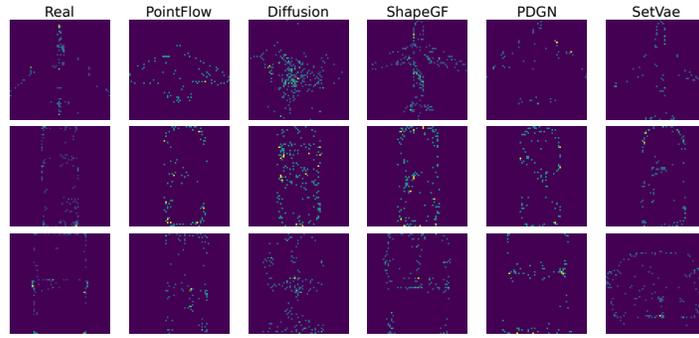


Figure 7: Visualization of six fingerprints on three shapes: Airplane, Car, and Chair, from top to bottom. These fingerprints are built by stacking 100 depth images projected from critical points from each source.

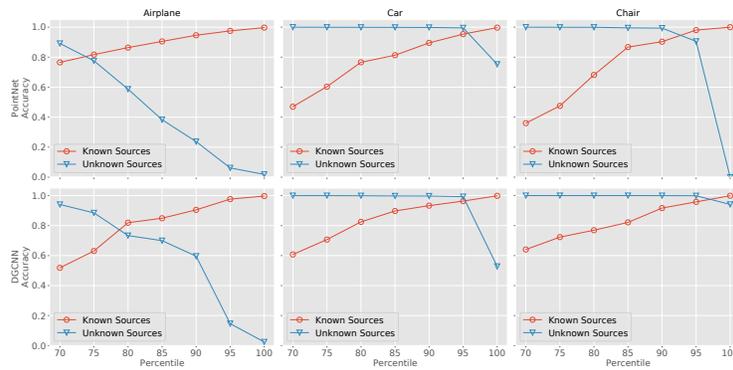


Figure 8: Attribution performances on known and unknown sources when the threshold is increasing. From left to right, we display the attribution results of Airplane, Car, and Chair.

image that shows the frequency spectra, we extract and visualize the unique patterns of each source.

Fingerprints of six different sources on three shapes are visualized in Figure 7. For the same shape shown in each row, every source presents a unique pattern that distinguishes it from other sources. Take Airplane from the Real source, as an example, points at the airplane’s head, end of the wings, and tail have higher frequencies than other areas. An airplane whose critical points set is similar to this pattern is likely to be attributed to the Real source. Similarly, for airplanes generated by ShapeGF, the critical points mainly describe the airplane’s skeleton. Overall, the above analysis enhances the understanding of how FAKEPCD conducts the source attribution task.

8 ABLATION STUDY

In this section, we conduct a series of ablation studies to investigate how attribution performance can be affected by several factors. We start with investigating how threshold selection in FAKEPCD influences the performances of both close-world attribution and open-world attribution. Second, we explore whether the feature vector dimension, i.e., the output dimension of encoding backbones, affects the attribution performance. In addition, we study the effect of the close-world pre-training stage in attributing point clouds compared to performing open-world pre-training from scratch (without the

close-world pre-training stage). Finally, we evaluate the robustness of FAKEPCD against common perturbations on the point clouds.

Effect of Threshold Selection. According to our threshold-based assignment principle, threshold selection can affect the attribution accuracy of both known sources and unknown sources. Figure 8 shows the trend of attribution accuracy in both settings when we enlarge the threshold by increasing the P -percentile gradually. As the threshold increases, the attribution accuracy of unknown sources demonstrates an evident rising trend, as the testing examples are more inclined to be assigned to known sources. Meanwhile, the accuracy of identifying known sources exhibits a reversed or steady trend depending on different shapes. For the attributing result of Airplane in the first column, we capture a trade-off relation with the intersection point at the range of 75%-80%. The accuracy on unknown sources maintains a high level for other shapes that can be differentiated more easily, and intersection appears after we set a higher percentile.

Effect of Feature Vector Dimension. We vary the dimension from 32 to 512 and evaluate the performance of FAKEPCD on Airplane, Car, and Chair. Figure 9 demonstrates the attribution performances on varying dimensions. For shapes such as Car and Chair, various dimensions do not significantly affect the attribution performance to a large extent. However, for point clouds that are relatively hard to differentiate, e.g., Airplane, we obtain the most

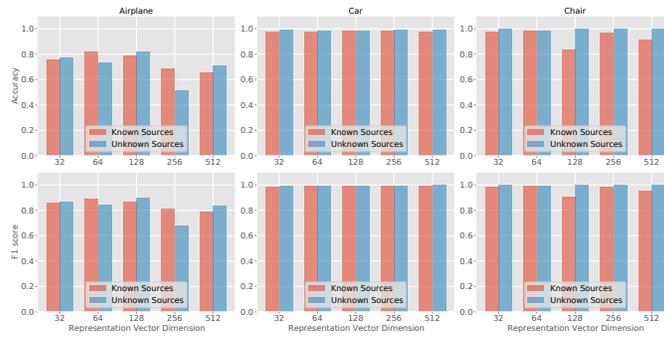


Figure 9: Attribution performances on both known and unknown sources when feature vector dimension is varied. Attribution performances of Airplane on both worlds are more balanced when the dimension is 128.

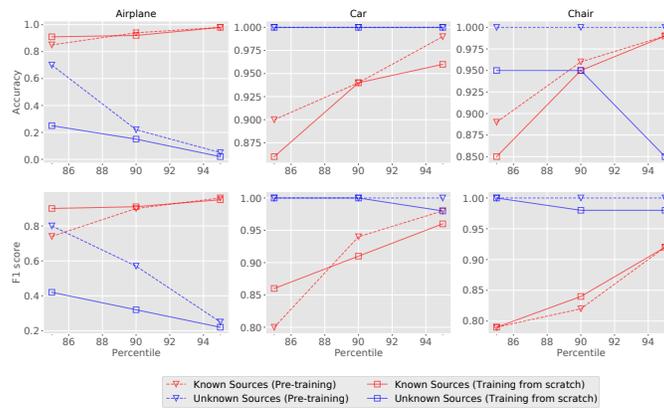


Figure 10: Attribution performances when the encoder is pre-trained and trained from scratch. Adopting a pre-trained encoder leads to better attribution accuracy when recognizing unknown sources.

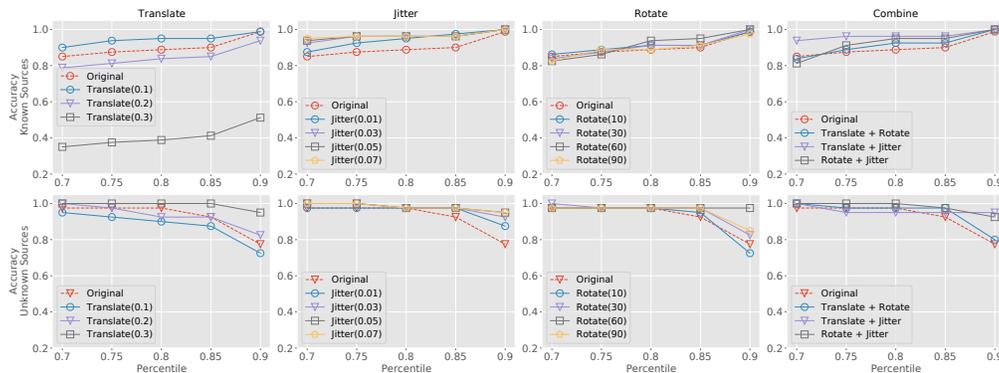


Figure 11: Attribution performances on Car against perturbations. Perturbations include *translate*, *jitter*, *rotate*, and the three combined. The attribution accuracy on both known and unknown sources is robust against most of the applied perturbations.

balanced attribution performance when the dimension size is 128, under the same threshold setting. Note that we adopt 128 as the feature vector dimension in the main experiments, following the work[19].

Effect of Close-World Pre-training Stage. To study how the encoder pre-training stage in FAKEPCD affects the open-world attribution, we compare the attribution performances when the encoder is randomly initiated or initiated with the parameters updated from the pre-training stage. Figure 10 displays the comparison results

of attributing three shapes, i.e., Airplane, Car, and Chair. We vary the P -percentile from 85 to 95 to gradually increase the threshold values. Figure 10 shows that FAKEPCD with the pre-training encoder performs almost equally as well as the training-from-scratch encoder when attributing data to known sources. However, the pre-training stage benefits the attribution accuracy of unknown sources to a certain extent.

Effect of Perturbations. We study the effect of common perturbations on point clouds to evaluate the robustness of FAKEPCD. Figure 11 shows the result after we exert different perturbations on the testing examples. We observe that *jitter*, *rotate*, and combined perturbations do not degrade attribution accuracy both on known and unknown sources, which demonstrates the resistance of FAKEPCD to these perturbations. However, we observe a decline when we translate point clouds in different coordinate frames, especially when the coordinate origin is distant from the old origin. The conjectured reason is that the point cloud features are not aligned in the feature space when point clouds are in different coordinate frames, resulting in a biased distance distribution for threshold-based attribution.

9 RELATED WORK

Point Cloud Generation. Generative models of point clouds can be broadly classified into three categories: Generative adversarial networks (GAN) based models, flow-based models, and variational auto-encoders (VAE) based models. To adapt GAN [14] in the point cloud domain, PC-GAN [26] modifies GAN model to learn a hierarchical sampling process for point cloud generation. As an end-to-end generation model, PDGN [17] proposes a progressive GAN network composed of multiple stacking deconvolution networks. Flow-based models [44] explicitly model the density function using continuous normalizing flow [37]. PointFlow [52] generates 3D point clouds by learning a distribution of distributions and models each distribution as an invertible parameterized transformation of 3D points from a prior distribution. Other novel works such as ShapeGF [2] model a shape by learning the gradient field of its log density, and then moving the points gradually from a generic prior distribution towards the surface. Diffusion models [31] generate point clouds in the perspective of thermodynamics.

Point Cloud Classification. Different from 2D images, 3D point cloud classification and segmentation tasks cannot be achieved with standard deep neural network models (i.e., MLP) due to irregular structure. PointNet [35] is the first work to directly forward raw point cloud data to the DNN model and obtain global features leveraging the symmetric max pooling function. Inspired by PointNet and convolution modules, follow-up works such as A-CNN [22] PointGrid [23], ShellNet [58], and PointCNN [27] learn the local features by adapting convolution modules to point clouds. Graph neural networks [39] are also applied in learning point cloud features. Wang et al. [48] propose DGCNN by constructing local graphs and learning the edge information that describes the relationships between a point and its neighbors.

Deepfake Image Detection and Attribution. Deepfake detection and attribution already emerged in 2D image [30, 54–56, 59] and text domain [11, 33]. For 2D images, existing work [47, 54, 59] has demonstrated that deepfake image detection could achieve a

favorable accuracy in the close-world attribution. Wang et al. [47] and Zhao et al. [59] detect fake images in the close world by training binary or multi-class classifiers. Ning et al. [54] attribute fake images to its GAN model by learning and analyzing image fingerprints. Previous studies [13, 55, 56] on tracing data in the open world also received much attention. Ning et al. [55, 56] actively embed fingerprints into generative models, such that open-world models can be verified by decoding and matching fingerprints. Recently, Sha et al. [40] carry out a systematic study on the detection and attribution of fake images generated by the latest text-to-image generation models, including stable diffusion, DALLE-2, GLIDE, etc. Inspired by the success of deepfake image detection and attribution, we explore the possibility of detecting fake point clouds in this work.

10 CONCLUSION

In this paper, we proposed an attribution framework FAKEPCD to pioneer the study of point cloud authentication and source attribution. This framework can attribute point clouds to both known and unknown sources. We also study the generalizability of attribution on different shapes by training single-shape models and multiple-shape models. Evaluation results verified the effectiveness of FAKEPCD when attributing both known and unknown sources, even shapes that are not included in the training set. The fingerprint visualization we introduced assists in understanding how our framework recognizes the critical points in point clouds and captures the uniqueness of each generative model. We hope that our study can also foster research in driving responsible and trustworthy artificial intelligence in point cloud generative models to prevent them from mischievous use.

Limitations and Future Work. This work has limitations. The first key area for improvement is to expand the number of generative models. In the main experiments, we currently only include four known sources and two unknown sources. In future work, we will explore incremental learning to incrementally involve more generative models on top of the trained attribution model. Second, when evaluating the performance of FAKEPCD, we rely on a small number of samples in the validation set to determine the optimal threshold, which is often not realistic in practice. We test two additional unknown sources in the Appendix and found that the threshold remains effective, although in a limited capacity, on more unknown sources. We will continue exploring the best approach to determine thresholds for recognizing closed-world and open-world examples.

ACKNOWLEDGMENTS

We thank all reviewers for their constructive suggestions. This work is partially funded by the European Health and Digital Executive Agency (HADEA) within the project “Understanding the individual host response against Hepatitis D Virus to develop a personalized approach for the management of hepatitis D” (D-Solve) (grant agreement number 101057917).

REFERENCES

- [1] <https://github.com/flairNLP>.
- [2] Ruojin Cai, Guandao Yang, Hadar Averbuch-Elor, Zekun Hao, Serge J. Belongie, Noah Snaveley, and Bharath Hariharan. Learning Gradient Fields for Shape

- Generation. In *European Conference on Computer Vision (ECCV)*, pages 364–381. Springer, 2020.
- [3] Yulong Cao, S Hrushikesh Bhupathiraju, Pirouz Naghavi, Takeshi Sugawara, Z Morley Mao, and Sara Rampazzi. You can't see me: physical removal attacks on lidar-based autonomous vehicles driving frameworks. In *USENIX Security Symposium (Usenix Security'22)*, 2022.
 - [4] Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Qi Alfred Chen, Kevin Fu, and Z. Morley Mao. Adversarial Sensor Attack on LiDAR-based Perception in Autonomous Driving. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 2267–2281. ACM, 2019.
 - [5] Angel X. Chang, Thomas A. Funkhouser, Leonidas J. Guibas, Pat Hanrahan, Qi-Xing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiang Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. *CoRR abs/1512.03012*, 2015.
 - [6] Xinlei Chen and Kaiming He. Exploring Simple Siamese Representation Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15750–15758. IEEE, 2021.
 - [7] Zhen Dong, Fuxun Liang, Bisheng Yang, Yusheng Xu, Yufu Zang, Jianping Li, Yuan Wang, Wenxia Dai, Hongchao Fan, Juha Hyypää, and Uwe Stilla. Registration of Large-scale Terrestrial Laser Scanner Point Clouds: A Review and Benchmark. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2020.
 - [8] Bertrand Douillard, James Patrick Underwood, Noah Kuntz, Vsevolod Vlaskine, Alastair James Quadros, Peter Morton, and Alon Frenkel. On the Segmentation of 3D LIDAR Point Clouds. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2798–2805. IEEE, 2011.
 - [9] Renaud Dubé, Daniel Dugas, Elena Stumm, Juan I. Nieto, Roland Siegwart, and Cesar Cadena. SegMatch: Segment based Place Recognition in 3D Point Clouds. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5266–5272. IEEE, 2017.
 - [10] Ge Gao, Mikko Lauri, Xiaolin Hu, Jianwei Zhang, and Simone Frintrop. CloudAAE: Learning 6D Object Pose Regression with On-line Data Synthesis on Point Clouds. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 11081–11087. IEEE, 2021.
 - [11] Anastasia Giachanou, Guobiao Zhang, and Paolo Rosso. Multimodal Multi-image Fake News Detection. In *International Conference on Data Science and Advanced Analytics (DSAA)*, pages 647–654. IEEE, 2020.
 - [12] Ian Gibson, David Rosen, Brent Stucker, Mahyar Khorasani, David Rosen, Brent Stucker, and Mahyar Khorasani. *Additive manufacturing technologies*, volume 17. Springer, 2021.
 - [13] Sharath Girish, Saksham Suri, Sai Saketh Rambhatla, and Abhinav Shrivastava. Towards Discovery and Attribution of Open-World GAN Generated Images. In *IEEE International Conference on Computer Vision (ICCV)*, pages 14094–14103. IEEE, 2021.
 - [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2672–2680. NIPS, 2014.
 - [15] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bannamoun. Deep Learning for 3D Point Clouds: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
 - [16] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11105–11114. IEEE, 2020.
 - [17] Le Hui, Rui Xu, Jin Xie, Jianjun Qian, and Jian Yang. Progressive Point Cloud Deconvolution Generation Network. In *European Conference on Computer Vision (ECCV)*, pages 397–413. Springer, 2020.
 - [18] Tero Karras, Samuli Laine, and Timo Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4401–4410. IEEE, 2019.
 - [19] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised Contrastive Learning. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
 - [20] Hyeonju Kim, Hyeonseung Lee, Woo Hyun Kang, Joun Yeop Lee, and Nam Soo Kim. Softflow: Probabilistic framework for normalizing flow on manifolds. *Advances in Neural Information Processing Systems*, 33:16388–16397, 2020.
 - [21] Jinwoo Kim, Jaehoon Yoo, Juho Lee, and Seunghoon Hong. SetVAE: Learning Hierarchical Composition for Generative Modeling of Set-Structured Data. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15059–15068. IEEE, 2021.
 - [22] Artem Komarichev, Zichun Zhong, and Jing Hua. A-CNN: Annularly Convolutional Neural Networks on Point Clouds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7421–7430. IEEE, 2019.
 - [23] Truc Le and Ye Duan. PointGrid: A Deep Network for 3D Shape Understanding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9204–9214. IEEE, 2018.
 - [24] Franz Leberl, Arnold Irschara, Thomas Pock, Philipp Meixner, Michael Gruber, Set Scholz, and Alexander Wiechert. Point clouds. *Photogrammetric Engineering & Remote Sensing*, 76(10):1123–1134, 2010.
 - [25] Ville V. Lehtola, Harri Kaartinen, Andreas Nüchter, Risto Kajaluoto, Antero Kukko, Paula Litkey, Eija Honkavaara, Tomi Rosnell, Matti T. Vaaja, Juho-Pekka Virtanen, Matti Kurkela, Aimad El Issaoui, Lingli Zhu, Anttoni Jaakkola, and Juha Hyypää. Comparison of the Selected State-Of-The-Art 3D Indoor Scanning and Point Cloud Generation Methods. *Remote. Sens.*, 2017.
 - [26] Chun Liang Li, Manzil Zaheer, Yang Zhang, Barnabás Póczos, and Ruslan Salakhutdinov. Point Cloud GAN. *CoRR abs/1810.05795*, 2018.
 - [27] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. PointCNN: Convolution On X-Transformed Points. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 828–838. NeurIPS, 2018.
 - [28] Ying Li, Lingfei Ma, Zilong Zhong, Fei Liu, Michael A. Chapman, Dongpu Cao, and Jonathan Li. Deep Learning for LiDAR Point Clouds in Autonomous Driving: A Review. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
 - [29] Shan Liu, Min Zhang, Pranav Kadam, and C-C Jay Kuo. *3D Point Cloud Analysis: Traditional, Deep Learning, and Explainable Machine Learning Methods*. Springer, 2021.
 - [30] Zhengzhe Liu, Xiaojuan Qi, and Philip H. S. Torr. Global Texture Enhancement for Fake Face Detection in the Wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8057–8066. IEEE, 2020.
 - [31] Shitong Luo and Wei Hu. Diffusion Probabilistic Models for 3D Point Cloud Generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2837–2845. IEEE, 2021.
 - [32] Anthony Ngo, Max Paul Bauer, and Michael M. Resch. Deep Evaluation Metric: Learning to Evaluate Simulated Radar Point Clouds for Virtual Testing of Autonomous Driving. *CoRR abs/2104.06772*, 2021.
 - [33] Aditya Pal, Chantat Eksombatchai, Yitong Zhou, Bo Zhao, Charles Rosenberg, and Jure Leskovec. PinnerSage: Multi-Modal User Embedding Framework for Recommendations at Pinterest. In *ACM Conference on Knowledge Discovery and Data Mining (KDD)*, pages 2311–2320. ACM, 2020.
 - [34] François Pomerleau, Francis Colas, Roland Siegwart, et al. A review of point cloud registration algorithms for mobile robotics. *Foundations and Trends® in Robotics*, 4(1):1–104, 2015.
 - [35] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85. IEEE, 2017.
 - [36] Prahalad K Rao, Jia Liu, David Roberson, Zhenyu Kong, and Christopher Williams. Online real-time quality monitoring in additive manufacturing processes using heterogeneous sensors. *Journal of Manufacturing Science and Engineering*, 137(6):061007, 2015.
 - [37] Danilo Jimenez Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. In *International Conference on Machine Learning (ICML)*, pages 1530–1538. JMLR, 2015.
 - [38] M Samie Tootooni, Ashley Dsouza, Ryan Donovan, Prahalad K Rao, Zhenyu Kong, and Peter Borgesen. Classifying the dimensional variation in additive manufactured parts from laser-scanned three-dimensional point cloud data using machine learning approaches. *Journal of Manufacturing Science and Engineering*, 139(9):091005, 2017.
 - [39] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The Graph Neural Network Model. *IEEE Transactions on Neural Networks*, 2009.
 - [40] Zeyang Sha, Zheng Li, Ning Yu, and Yang Zhang. De-fake: Detection and attribution of fake images generated by text-to-image diffusion models. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2023.
 - [41] Zifan Shi, Sida Peng, Yinghao Xu, Yiyi Liao, and Yujun Shen. Deep generative models on 3d representations: A survey. *arXiv preprint arXiv:2210.15663*, 2022.
 - [42] Connor Shorten and Taghi M Khoshgoufar. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019.
 - [43] Zhenbo Song, Wayne Chen, Dylan Campbell, and Hongdong Li. Deep Novel View Synthesis from Colored 3D Point Clouds. In *European Conference on Computer Vision (ECCV)*, pages 1–17. Springer, 2020.
 - [44] Michal Stypulkowski, Maciej Zamorski, Maciej Zieba, and Jan Chorowski. Conditional Invertible Flow for Point Cloud Generation. *CoRR abs/1910.07344*, 2019.
 - [45] Jiachen Sun, Yulong Cao, Qi Alfred Chen, and Z Morley Mao. Towards robust lidar-based perception in autonomous driving: General black-box adversarial sensor attack and countermeasures. In *USENIX Security Symposium (Usenix Security'20)*, 2020.
 - [46] George Vosselman, Sander Dijkman, et al. 3d building model reconstruction from point clouds and ground plans. *International archives of photogrammetry remote sensing and spatial information sciences*, 34(3/W4):37–44, 2001.
 - [47] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A. Efros. CNN-Generated Images Are Surprisingly Easy to Spot... for Now. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8692–8701. IEEE, 2020.

- [48] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic Graph CNN for Learning on Point Clouds. *ACM Transactions on Graphics*, 2019.
- [49] Steven Webster, Todd Du Bosq, Vinh Tran, Kirby Thomas, and Christopher May. Simulation of LIDAR Systems for Aerial Intelligence, Surveillance, and Reconnaissance. *STO-MP-MSG-149*, 2017.
- [50] Jianwen Xie, Yifei Xu, Zilong Zheng, Song-Chun Zhu, and Ying Nian Wu. Generative pointnet: Deep energy-based learning on unordered point sets for 3d generation, reconstruction and classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14976–14985, 2021.
- [51] Bin Yang, Wenjie Luo, and Raquel Urtasun. PIXOR: Real-time 3D Object Detection from Point Clouds. *CoRR abs/1902.06326*, 2019.
- [52] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge J. Belongie, and Bharath Hariharan. PointFlow: 3D Point Cloud Generation With Continuous Normalizing Flows. In *IEEE International Conference on Computer Vision (ICCV)*, pages 4540–4549. IEEE, 2019.
- [53] Zehao Ye, Chenang Liu, Wenmeng Tian, and Chen Kan. A deep learning approach for the identification of small process shifts in additive manufacturing using 3d point clouds. *Procedia Manufacturing*, 48:770–775, 2020.
- [54] Ning Yu, Larry Davis, and Mario Fritz. Attributing Fake Images to GANs: Learning and Analyzing GAN Fingerprints. In *IEEE International Conference on Computer Vision (ICCV)*, pages 7555–7565. IEEE, 2019.
- [55] Ning Yu, Vladislav Skripniuk, Sahar Abdelnabi, and Mario Fritz. Artificial Fingerprinting for Generative Models: Rooting Deepfake Attribution in Training Data. In *IEEE International Conference on Computer Vision (ICCV)*, pages 14448–14457. IEEE, 2021.
- [56] Ning Yu, Vladislav Skripniuk, Dingfan Chen, Larry Davis, and Mario Fritz. Responsible Disclosure of Generative Models Using Scalable Fingerprinting. *CoRR abs/2012.08726*, 2020.
- [57] Baiwu Zhang, Jin Peng Zhou, Iliia Shumailov, and Nicolas Papernot. On Attribution of Deepfakes. *CoRR abs/2008.09194*, 2021.
- [58] Zhiyuan Zhang, Binh-Son Hua, and Sai-Kit Yeung. ShellNet: Efficient Point Cloud Convolutional Neural Networks Using Concentric Shells Statistics. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1607–1616. IEEE, 2019.
- [59] Hanqing Zhao, Wenbo Zhou, Dongdong Chen, Tianyi Wei, Weiming Zhang, and Nenghai Yu. Multi-Attentional Deepfake Detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2185–2194. IEEE, 2021.

A GENERALIZABILITY TO MORE UNKNOWN SOURCES

In the open-world attribution, FAKEPCD presents satisfactory performance in identifying point clouds from two unknown sources: PDGN and SetVae. To evaluate whether FAKEPCD can generalize to more unknown sources, especially when their validation sets are not accessible, we introduce two additional generative models, GPointNet [50] and SoftFlow [20]. For each model, we generate 1,000 point clouds of the shape Airplane, Car, and Chair, respectively. We directly adopt the optimal thresholds obtained in Table 3 and measure FAKEPCD’s performances in attributing these samples. Table 4 displays the attribution performance. FAKEPCD identifies the point clouds from unknown sources with an accuracy ranging from 0.41 to 0.81 across three shapes. The best performance is observed in Car, where FAKEPCD-P achieves an accuracy of 0.81 and an F1 score of 0.9. This finding suggests that FAKEPCD may generalize to more unknown sources with previously determined thresholds in Table 3.

Table 4: Attribution performance of FAKEPCD in identifying point clouds from additional unknown sources: GPointNet [50] and Softflow [20].

	Airplane		Car		Chair	
	Acc	F1 score	Acc	F1 score	Acc	F1 score
FAKEPCD-P	0.41	0.58	0.55	0.71	0.81	0.90
FAKEPCD-D	0.57	0.73	0.51	0.68	0.43	0.60

B CRITICAL POINTS VISUALIZATION

We explain the intrinsic reason for the source attribution of point clouds in the main text. FAKEPCD attributes point clouds from different sources by identifying critical points. We also extract fingerprints of each source based on the critical points. Figure 15 illustrates the critical points visualization of the shape Airplane when we extract features with PointNet in the open-world attribution. For every shape, we select 3 representative point clouds from each source and also keep point clouds from one column as similar as possible to better compare the critical points distribution. We use *Chamfer Distance* to find out the similar point clouds in one column. Point clouds described only by red points are captured by the models, and the information provided by gray points is discarded in the max-pooling operation. Examples from the same source demonstrate a certain recognition pattern, e.g., in the Real source, critical points primarily describe the head and wings’ edge of the airplane; and in the ShapeGF source, these points are scattered at the airplane’s surface; in the Diffusion source, extra focus locates in the main body area. When we compare the examples from PointFlow and SetVae, their individual patterns are sometimes indistinguishable. This observation explains the relatively weaker performance when attributing Airplane. When we compare airplanes with similar outlooks (smaller Chamfer Distance) in one column, critical points draw slightly different outlines, which explains why FAKEPCD can recognize point clouds from different sources.

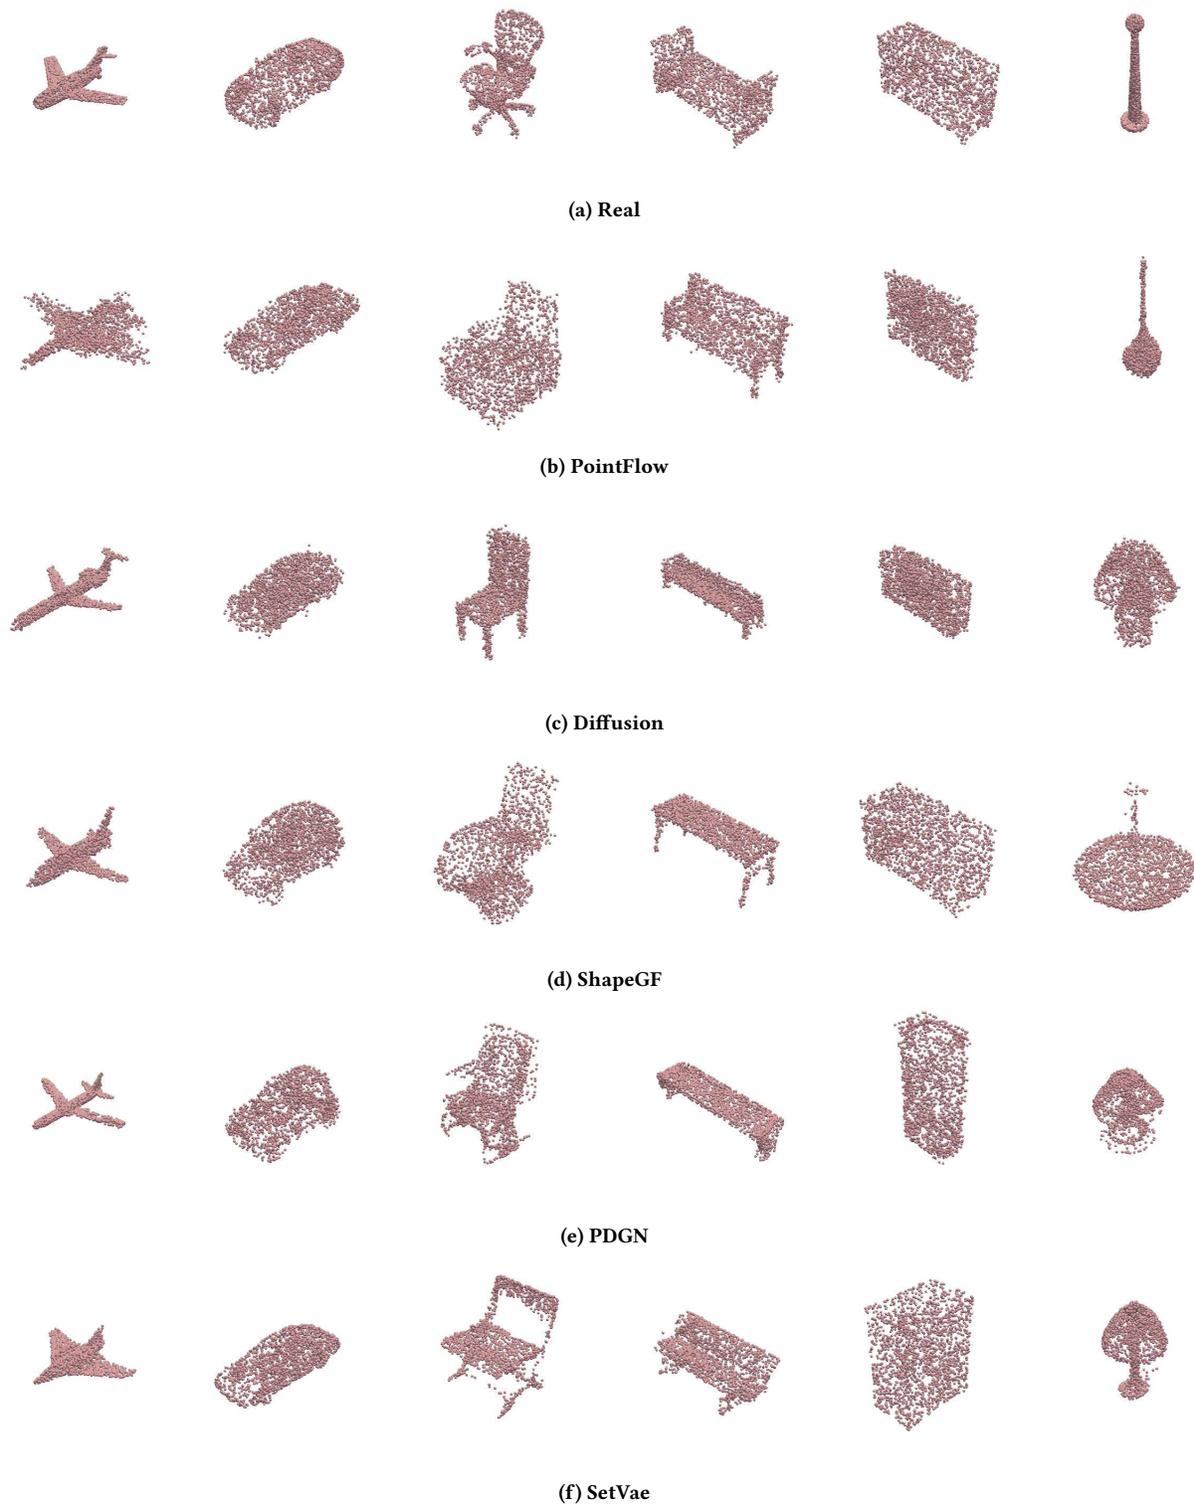


Figure 12: Generated examples from five generative models plus the real world. From left to right, we display the shapes: Airplane, Car, Chair, Bench, Cabinet, and Lamp

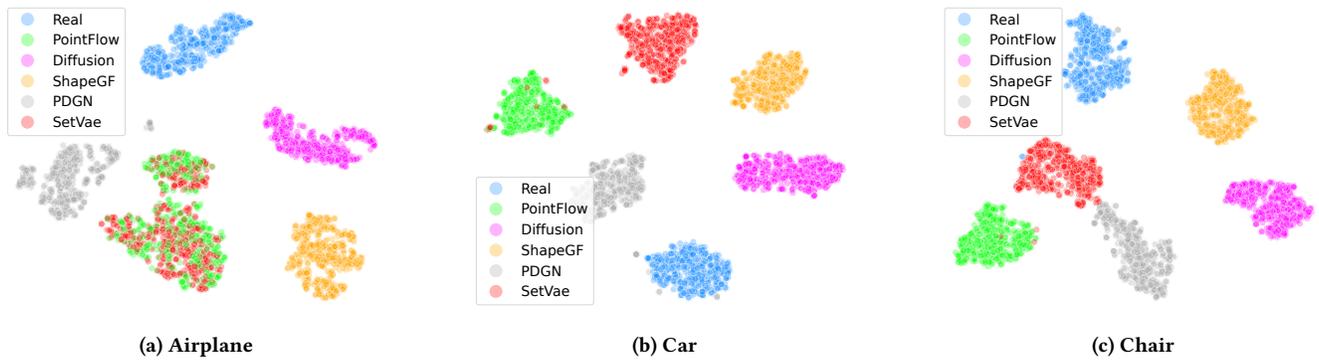


Figure 13: t-SNE visualization of clustered point cloud features obtained from FAKEPCD-PointNet (better viewed in color). PDGN and SetVae are unknown sources.

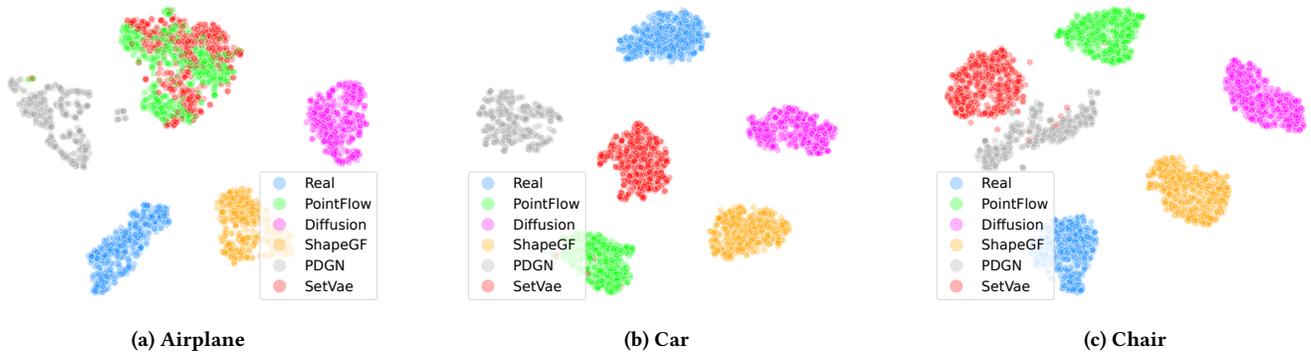


Figure 14: t-SNE visualization of clustered point cloud features obtained from FAKEPCD-DGCNN (better viewed in color). PDGN and SetVae are unknown sources.

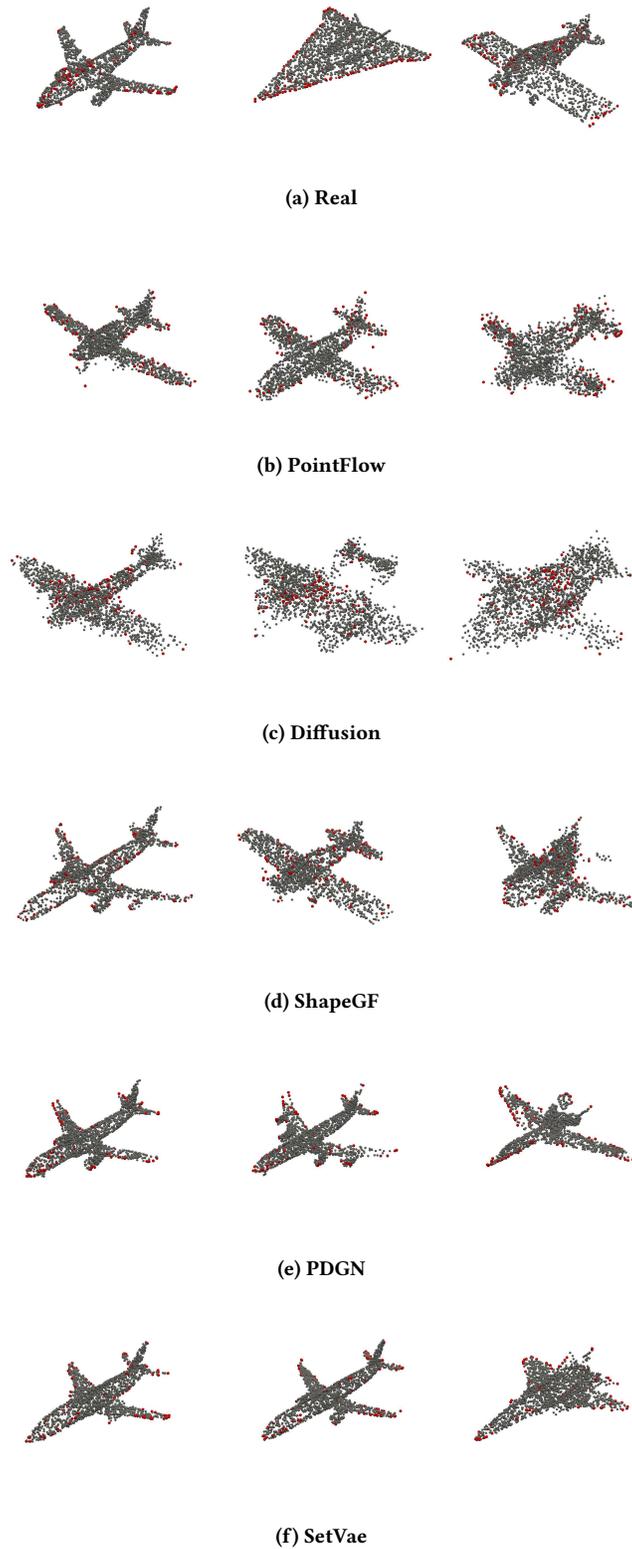


Figure 15: Critical points of Airplane. Grey points describe the original point cloud, and the red points are the identified critical points. Each row contains examples from one source and demonstrates certain commonalities. For example, the majority of critical points focus on the airplane head and wing area in the Real class, while the majority of those critical points are found on the main body in the Diffusion class.