

Detecting Images Generated by Deep Diffusion Models using their *Local Intrinsic Dimensionality*

Peter Lorenz¹, Ricard L. Durall¹, Janis Keuper^{1,2}
¹Fraunhofer ITWM, ²IMLA - Offenburg University
 Correspondence to peter.lorenz@itwm.fhg.de

Abstract

Diffusion models recently have been successfully applied for the visual synthesis of strikingly realistic appearing images. This raises strong concerns about their potential for malicious purposes. In this paper, we propose using the lightweight multi Local Intrinsic Dimensionality (multiLID), which has been originally developed in context of the detection of adversarial examples, for the automatic detection of synthetic images and the identification of the according generator networks. In contrast to many existing detection approaches, which often only work for GAN-generated images, the proposed method provides close to perfect detection results in many realistic use cases. Extensive experiments on known and newly created datasets demonstrate that the proposed multiLID approach exhibits superiority in diffusion detection and model identification.

Since the empirical evaluations of recent publications on the detection of generated images are often mainly focused on the “LSUN-Bedroom” dataset, we further establish a comprehensive benchmark for the detection of diffusion-generated images, including samples from several diffusion models with different image sizes.

The code for our experiments is provided at <https://github.com/deepfake-study/deepfake-multiLID>.

1. INTRODUCTION

Recently, denoising diffusion probabilistic models (DDPMs) [82, 37] have established a new paradigm in image generation thanks to their solid ability to synthesize high-quality images. As a result, a large number of studies have been exploring novel network architectures [83, 57, 20, 50, 69], alternative noise schedules to accelerate the sampling during inference [83, 57, 50, 53, 96, 76] and state-of-the-art text-to-image approaches [15, 67, 75, 69, 33, 73]. Furthermore, numerous image generation platforms, both

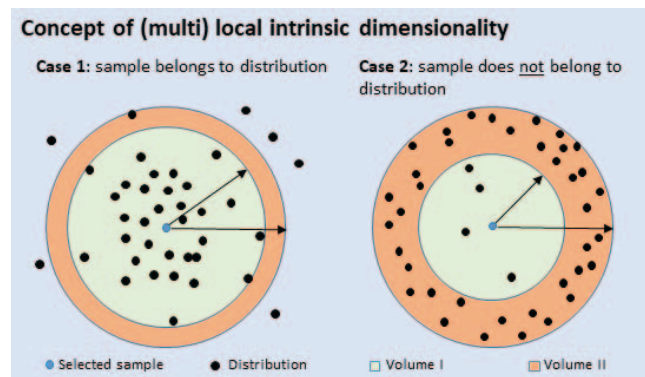


Figure 1: The underlying concept of the proposed method is to distinguish models by differences in the density of their internal feature distributions. LID estimates densities in the feature spaces of pre-trained CNNs, by computing fractions over the number of samples in given volumes: $|\text{volume I}|/|\text{volume II}| < 1$. The example above shows how this density measure indicates if the selected sample belong (left) or does not belong (right) to a reference distribution. Further details in section 3.1 and in the section B.

commercial and open-source, such as Midjourney [39], Dall-e 2 [38], Imagen [75], Dreambooth [73], and Stable Diffusion [69], have contributed to bringing this technology closer to people, boosting significantly its popularity. However, with the ease of generating content through diffusion models (DMs) at the click of a button, the presence of high-quality tampered content is growing leading to potential privacy issues [12, 101]. As the consumption of media expands to social media and deliberate modifications are made to spread false information [26], it becomes crucial to detect synthesized imagery. Although there are several detectors available for identifying non-natural images, most have not been designed for diffusion content due to fundamental differences in the generation process. For example, frequency-based approaches [99, 23, 68] have shown high detection scores when applied to images generated by generative adversarial networks (GANs), but they fail when

DMs are employed. The main reason for the phenomenon appears to be that GAN-generated images often exhibit distinct artifacts, characterized by a periodic, grid-like pattern, which is not present anymore in diffusion samples. In order to circumvent this problem, Wang et al. [94] introduced a novel representation for effectively detecting DM-generated images. Their approach involves analyzing the reconstruction error between synthetic and real images. Nonetheless, although the aforementioned methods exhibit promising results, they rely on a vast amount of data to be trained on. As a consequence, these systems might struggle when facing new scenarios with data scarcity. Additionally, none of them has proven to be able to distinguish different DM-generated images within the same context, i.e., dataset.

In this paper, our main objective is to identify synthetic content, in particular, DM-generated images. To that end, we introduce a novel pipeline consisting of i) forwarding the input images to an untrained ResNet [71] and extracting their feature-map representations; ii) computing multi local intrinsic dimensionality (multiLID) [52], a variant of the LID [54] on the resulting lower dimensional spaces; and iii) running a classifier to determine the nature of the input images given their multiLID. We show that this proposal can successfully distinguish between synthetic and natural images, as well as among different DM-generators, while requiring a relatively small training dataset, i.e., around 1,600 samples per class. To assess the effectiveness of our multiLID approach, we conduct an extended evaluation that encompasses images generated by various DMs, including unconditional and text-to-image generation setups, e.g., Glide [56], DDPM [37], Latent Diffusion [69], Palette [74], Stable Diffusion [69], VQ Diffusion [33] and Diffusion Transformer (DiT) [61]. We demonstrate that the multiLID representation has an effective identification capability through extensive experiments.

The three main contributions of our work can be summarized as follows:

- We introduce a lightweight method, i.e., multiLID, for diffusion-generated content identification, whose capabilities extend beyond real and synthetic image classification, as it can also determine the specific generative model.
- We evaluate the performance of our proposed method on numerous datasets from standardized ones, such as LSUN-Bedroom, to state-of-the-art such as CiFake and ArtiFact.
- We conduct a thorough study to assess and characterize the proposed methodology.

2. RELATED WORK

In this section, we provide a brief overview of recent diffusion models for image generation and discuss various DM-detection approaches.

2.1. Diffusion Models for Image Generation

Diffusion models have emerged as a powerful image generation paradigm, which was originally inspired by non-equilibrium thermodynamics [82]. Denoising diffusion probabilistic models (DDPMs), introduced by Ho et al. [37], have exhibited notable generative capabilities when compared to advanced GANs paradigm [45]. Song et al. [83] introduced the use of denoising diffusion implicit models (DDIMs) to speed up image generation while keeping a reasonable image quality trade-off. A later work, ablated diffusion model (ADM) [20] finds a much more effective architecture with classifier guidance. Finally, considering DDPMs as differential equations on manifolds, Liu et al. [50] proposed pseudo-numerical methods for diffusion models (PNDMs), which further enhance sampling efficiency and generation quality. In the quest for progress, the vector quantized diffusion model (VQD) [33] proposed a conditional variant of DDPM incorporating a variational quantized diffusion variational auto-encoder (VQ-VAE) [89] to model the latent space. Notably, the latent diffusion model (LDM) [69] has demonstrated superior robustness and efficiency compared to other diffusion models. LDMs employ a cross-attention mechanism inspired by transformers [24] to effectively combine text and image input sequences within the latent space. Building upon the foundation of LDM, the popular Stable Diffusion v2 has further enhanced generation performance while reducing computational requirements. Recently, Peebles and Xie [61] were able to replace the U-Net [70] backbone in LDMs with a vision transformer and establish a new paradigm called Diffusion Transformers (DiT). Built upon DiT, Gao et al. [28] proposed a Masked Diffusion Transformer (MDT) that consists of a mask latent modeling scheme to explicitly enhance the DMs' ability of contextual relation learning among object semantic parts in an image.

2.2. Detectors for Synthetic Images

The distinction between natural and synthetic images has captivated researchers since the advent of image generation. Durrall [23] discovered an approach to detect GAN-generated images based on classical frequency domain analysis. Later, Quian et al. introduced the Frequency in Face Forgery Network (F3-Net) [65]. The proposed framework is composed of two frequency-aware branches, one focused on mining subtle forgery patterns through frequency components partition, and the other aimed at extracting small-scale discrepancies of frequency statistics between real and synthetic images. For each branch, there is a pre-trained classifier to extract the features based on Xception architecture [27], and following both feature sources are combined to final deepfake detector. CNNDet [93] is another FFT-based detector employing a pre-trained Resnet50 and reutilize as a binary classifier. Their objective was to show

if there is a common pattern in the Fourier domain of different GAN models generated images to transfer to unknown synthetic data. Chai et al. [14] introduced the detector Patch-Forensics (Pa-Fo) and found that splitting images into patches to limit the receptive field of a classifier leveraged the ability to detect manipulated parts in an image. Self-Blended Images (SBI) [80] was built up on the idea of creating its own forgeries to learn generic and robust representation. To achieve that they used pre-trained EfficientNet-b4 [85] classifier and fine-tuned it on landmark mismatch, blending boundary, color mismatch, and frequency inconsistency features. With the emergence of DMs and their increasing dominance, traditional generative solutions like GANs have gradually been replaced. Studies by Dong et al. [21] and Ricker et al. [68] have shown that tailored GAN-generated image detectors have also become outdated, as they rely on extracting synthetic artifacts using frequency-aware features or trainable noise patterns within the amplitude and phase spectra domains [23, 5, 81], which are not that prominent in DM-generated images anymore. Wang et al. [94] discovered that DM-generated images exhibit features that are more easily reconstructed by pre-trained diffusion models compared to natural images. To identify such features, they presented Diffusion Reconstruction Error (DIRE). Guo et al. [35] and Guarnera et al. [34] proposed a hierarchical fine-grained labeling approach for forged or synthetic images, utilizing carefully designed training sets. The hierarchical formulation requires an extensive inclusion of forgery techniques in the training set, which can be challenging when having limited diversity in the training data. Amoroso et al. [2] explored the decoupling of semantic and style features in images, and demonstrated that synthetic images can display greater separability in the style domain. Nonetheless, the practicality of semantic-style disentangling is challenging, as it necessitates tailored training sets.

3. METHOD

In this paper, we conduct a thorough investigation of the multiLID method [52], originally developed for detecting adversarial examples, and validate its detection capability within the diffusion models context. Note that the direct application of multiLID on the images yields unsatisfactory results and therefore, we first employ an untrained ResNet18 [46] to extract low-dimensional features from the synthetic images. Then, we can apply multiLID on these extracted features and finally train a classifier, specifically a random forest model. The conceptual framework of our proposal is illustrated in fig. 2.

3.1. Preliminaries

In this section, we explain the background of the feature maps and the intrinsic dimensionality. Both are crucial for




Figure 2: Pipeline of our method. Generation: Synthetic images are generated or sampled from a dataset. Extraction: Image features from ResNet18 are extracted and then, their multiLID scores are calculated. Detection: A classifier (random forest) is trained on these multiLID scores to distinguish between synthetic and real (or synthetic).

understanding the multiLID. method.

The relevance of CNN Feature Maps cannot be underestimated in the framework of our method. Actually, the application of multiLID scores on the raw, high-dimensional data results in ineffective and uninformative outcomes. However, if we employ the extracted lower dimensional and structured feature maps, the performance dramatically boosts. The usage of neural networks to extract features is not new. In fact, extensive research has been conducted on the properties of CNN feature maps, primarily focused on natural images. In this regard, it is worth mentioning that the hypothesis suggesting that natural images lie on or near a low-dimensional manifold remains a topic of debate. However, as argued by Goodfellow et al. [31], there is at least some correctness in that assumption when it comes to images. This assertion is supported by two noteworthy observations. First, natural images exhibit local connectivity. In other words, each image is surrounded by other highly similar images that can be reached through image transformations such as contrast and brightness adjustments. Second, natural images appear to conform to a low-dimensional structure as the probability distribution of images is highly concentrated, i.e., randomly sampled pixels alone cannot assemble a meaningful image. The combination of natural scenes and sensor properties is widely believed to result in sparse and concentrated image distributions, as supported by several empirical studies on image patches [48, 22, 13]. In their seminal work, Olshausen et al. [58] demonstrated that natural images exhibit distinctive statistical regularities that differentiate them from random images. Understanding these regularities has practical implications, such as more efficient coding of natural images and serving as a valuable prior in the field of computer vision [62]. Furthermore, the low-dimensional manifold hypothesis has been extensively validated through rigorous experiments conducted on diverse image datasets [72, 78, 71, 86, 8]. In addition, Feferman et al. [25] proposed novel algorithms for systematically verifying the validity of this manifold hypothesis.

In the context of neural networks, Zhu et al. [102] presented a new neural network architecture that incorporates a low-dimensional manifold regularization term to improve the generalization performance of the model. The authors argued that the high-dimensional nature of neural networks can lead to overfitting and poor generalization. Moreover, neural networks heavily rely on low-dimensional textures and not on the shape information [29]. In the same vein, it has been suggested that natural images can be represented as mixtures of textures residing on a low-dimensional manifold [87, 88]. Gont et al. [30] discovered that neural network features possess low-dimensional characteristics, which are easy to learn. They also observed a decrease in the intrinsic dimension of features in the last layers of neural networks, with interesting dimensionality trends in the first layers. Shortly after, Pope et al. [63] found that common natural image datasets indeed have very low intrinsic dimensions relative to the high number of pixels in the images. In particular, they showed it with GAN-generated synthetic data.

Local Intrinsic Dimensionality (LID) is a method used to estimate the intrinsic dimensionality of a learned representation space. LID measures the average distance between a point and its neighboring points [3, 40] as illustrated in fig. 1. This is achieved through maximum likelihood estimation that can be calculated as follows: Consider a mini-batch \mathcal{B} of N examples, and let $r_i(x) = d(x, y)$ represent the Euclidean distance between the sample x and y its i -th nearest neighbor in \mathcal{B} . Then the LID can be approximated as:

$$\text{LID}(x) = - \left(\frac{1}{k} \sum_{i=1}^k \log \frac{d_i(x)}{d_k(x)} \right)^{-1}, \quad (1)$$

where k is a hyper-parameter that determines the number of nearest neighbors, and d is the distance metric employed.

Ma et al. [54] introduced LID to characterize adversarial examples. They argued that the average distance between samples and their neighbors in the learned latent space of a classifier exhibits distinct properties for adversarial and natural (not modified) samples. They assessed LID on the j -dimensional latent representations of a neural network $f(x)$, using the L_2 distance:

$$d_\ell(x, y) = \|f_\ell^{1..j}(x) - f_\ell^{1..j}(y)\|_2, \quad (2)$$

where $\ell \in L$ represents the feature maps, and computed a vector of LID values sample-wise:

$$\vec{\text{LID}}(x) = \{\text{LID}_{d_\ell}(x)\}_\ell^n. \quad (3)$$

They repeated this procedure for both natural and adversarial examples. Finally, a logistic regression classifier was trained to detect adversarial samples. The mathematical definition of the LID is in the section B.

3.2. Method - multiLID

The method *multiLID* [52] was designed to detect adversarial examples and is based on the LID. In this section, we explain which advantage multiLID has over the original LID method and its accompanying benefits. In practice, the statistical estimate of intrinsic dimensionality (ID) is not solely dependent on the chosen neighborhood size. Typically, the ID is evaluated on a mini-batch basis, where k nearest neighbors are determined from a random sample of points in the latent space. Although this approach might introduce some noise, it provides broader coverage of the space, while considering only a few neighbors for each ID evaluation. Consequently, the summation aggregates the relative growth rate over potentially large distances in the latent space, see eq. (1). Lorenz et al. [52] argue that this summation step combines locally discriminative information about the growth rate in close proximity, and with the growth rates computed from more distant points. To address this, they propose ‘‘unfolding’’ [52] the growth rate estimation. Instead of computing an aggregated (semi) local ID, they suggest calculating a feature vector, referred to as multiLID, for every sample x . The length of this feature vector is k , and it is defined as:

$$\vec{\text{multiLID}}_d(x)[i] = - \left(\log \frac{d_i(x)}{d_k(x)} \right), \quad (4)$$

where d represents the Euclidean distance.

By using the multiLID feature vector, multiLID aim to capture more fine-grained information about the relative growth rates at different distances for each sample. For example, let the number of nearest neighbors be $k = 10$ and we extract eight feature maps (from ReLU activation layers) per sample. Then, the multiLID feature vector has a length of $k \times 8 = 80$ for $k = 10$, while the LID algorithm would have a feature vector of 8 because it sums up the nearest neighbors. This approach allows us to consider the local growth rate information separately for each neighbor, without the need for aggregation.

4. EXPERIMENTS

In this section, we first introduce the used datasets, then the experimental setup, and finally we present and discuss an extensive collection of experiments.

4.1. Datasets

This subsection provides an overview of the datasets used in our study, including details on those that are publicly available and those that we created from pre-trained models. The datasets contain a range of image sizes, spanning from 32×32 to 768×768 pixels; and of heterogeneous

domains, such as faces animals, places, and even images with artistic style.

4.1.1 Public Datasets

The following datasets are *publicly available*:

CiFake dataset [7] offers a collection of real and synthetic images, comprising a total of 120,000 images. It combines 60,000 images sourced from the existing CIFAR-10 dataset [46] with an additional 60,000 DM-generated images. The generation of synthetics is carried out by a LDM model[69]. The dataset maintains the same classes as the original CIFAR-10 dataset.

ArtiFact is a large-scale image dataset [6], which includes a diverse collection of real and synthetic images from multiple categories: human/human faces, animal/animal faces, places, vehicles, art, and many other real-life objects. The real dataset comprises 8 subdatasets (ImageNet, AFHQ, CelebaHQ, COCO, FFHQ, Landscape, MetFaces, and LSUN (Bedroom, Car, Cat, Horse)) [19, 17, 42, 49, 45, 51, 43, 98] to ensure diversity. On the other hand, the synthetic dataset consists of DM-generated images from 25 distinct methods, including 13 GANs, 7 Diffusion, and 5 other miscellaneous generators. For our evaluation, we randomly select images from six diffusion models (Glide, DDPM, Latent Diffusion, Palette, Stable Diffusion, VQ Diffusion) [56, 37, 69, 74, 69, 33] and six GAN models (Big GAN, Gansformer, Gau GAN, Projected GAN, StyleGAN3, Taming Transformer) [10, 41, 60, 77, 44, 24] to conduct our evaluations. In total, we select 10,500 real and generated images with 5,250 images per category.

DiffusionDB is one of the first large-scale text-to-image dataset [95]. The images are generated by Stable Diffusion (SD) using prompts from users in a discord channel and the images exhibit an artistic style. In our study, we work with the subset “2m_random_5k” [91]. Since DiffusionDB does not provide a collection of real images, inspired by Xie et al. [97], we employ LAION-5B and SAC datasets (see below).

LAION-5B is a large-scale web-based dataset [79], which has over 5 billion images crawled from the Internet. The images are annotated by CLIP [66] in many different languages. Although this dataset provides different image sizes, we focus only on the high-resolution [47] subset. Note that the images are center cropped to fit the synthetic datasets. We use this dataset to compare synthetic images from DiffusionDB.

SAC (Simulacra Aesthetic Captions) dataset¹ [64] is created from various text-to-image diffusion models, such as CompVis latent GLIDE and Stable Diffusion. It comprises over 40,000 user-generated prompts, predominantly con-

¹The images in version 1.0 of SAC are provided as a subset in <https://s3.us-west-1.wasabisys.com/simulacrabot/sac.tar>. We only filter the images with size 512×512 pixels.

sisting of images with artistic styles. Xie et al. [97] observed that this dataset shares similarities with DiffusionDB and therefore, we use it as a real dataset to compare to DiffusionDB.

4.1.2 New Datasets

Additionally, we create new datasets from different models to further diversify and scale our evaluation. We extend these datasets referring in section E.

Stable Diffusion-v2.1 (SD-v2.1), we sample 2,000 images using the pre-trained model [69, 84]. In order to generate the samples, we collect and utilize prompts from LAION-5B. We employ the images from the LAION-5B dataset as a real dataset [79].

LSUN-Bedroom, we sample 2,000 images (for each method) using several pre-trained models from diffusers [90]. In particular, we leverage the following methods:

- {DDPM, DDIM, PNDM}-ema: The pre-trained model with the id “google/ddpm-ema-bedroom-256” includes DDPM, DDIM, and PNDM samplers.
- ADM: We download the pre-trained LSUN-Bedroom model of ADM [20] from the official repository [59].
- SD-v2.1: The pre-trained text-to-image model with the id “stabilityai/stable-diffusion-2-1” [69]. SD-v2.1 uses LDMs as a backend and additionally has integrated cross-attention to enable conditioning multi-modality [1].
- LDM: We use the pre-trained text-to-image model with the id “CompVis/ldmtext2im-large-256” [69].
- VQD: We use the pre-trained text-to-image model with the id “microsoft/vq-diffusion-ithq” [32].

As a real dataset, we employ the images from LSUN-Bedroom dataset [98] from huggingface [18]. We center-crop them to 256×256 pixels.

4.2. Experimental Setup

Data pre-processing: All experiments are conducted on the aforementioned datasets. First of all, we calculate the standard mean and standard deviation on the dataset and normalize the inputs. Once we have homogeneous data distribution, we feed the images into an untrained ResNet18 model² [36] to extract their features. Although the network is not trained [4, 11, 16], it already suffices to distill the main characteristics of the data. We have not observed a difference in the detector’s accuracy by using untrained or trained weights referring to section C. Then, we compute the multiLID scores from the extracted feature maps. Our

²As a ResNet18 implementation, we use the model provided by TIMM library <https://huggingface.co/docs/timm/index>. The selected layers are called: 1_conv2.1, 1_conv2.2, 2_conv2.1, 2_conv2.2, 3_conv2.1, 3_conv2.2, 4_conv2.1, 4_conv2.2., which has the advantage to manage all different image sizes.

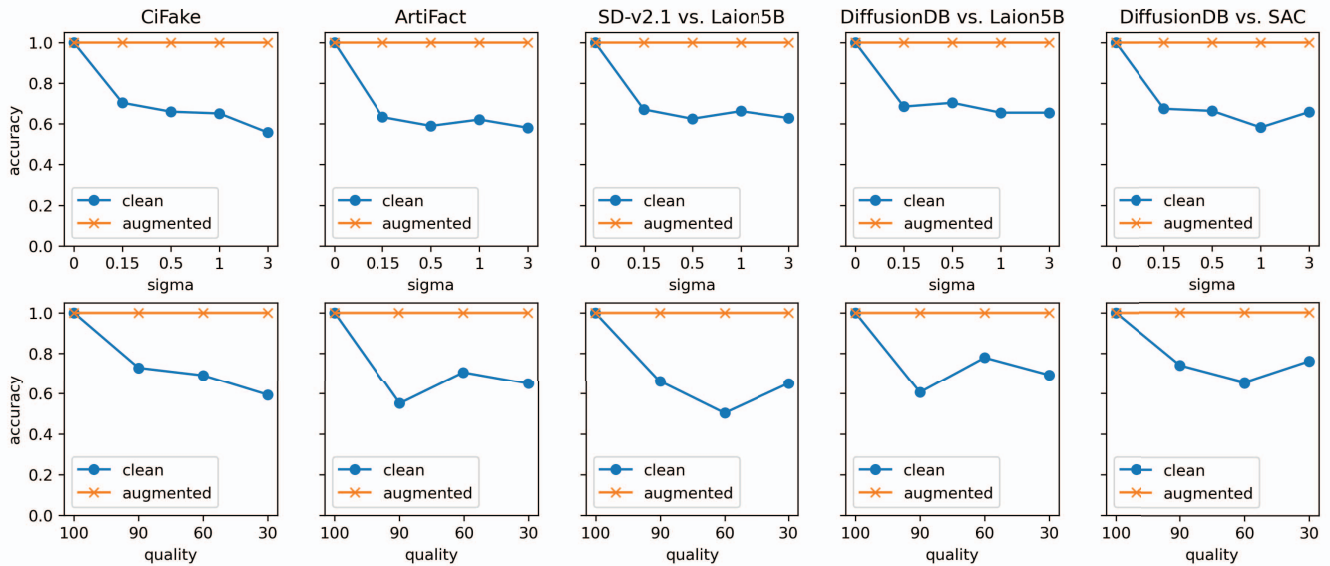


Figure 3: This figure contains the effects of data augmentation (top row: Gaussian blurring; bottom row: JPEG compression) on different datasets. To assess the multiLID performance, we calculate accuracy (ACC). In both cases the data augmentation is necessary to improve the detectors’s accuracy. We refer to the section E on the other datasets’ evaluations.

training data size has 1,600 samples per class unless otherwise specified in the experiments and finally we train a random forest classifier on the labeled multiLID scores.

Evaluation metrics: Following previous detection methods [100, 92, 93, 94], we also report the accuracy (ACC) in our experiments.

4.3. Classification

In this subsection, we present our findings across various datasets (refer to fig. 3). In real-world situations, images might undergo unidentifiable post-processing operations, such as compression and resizing. To assess the detectability of DM-generated images even after undergoing post-processing, we apply blurring and JPEG-compression techniques to both synthetic and authentic images, following the established procedure outlined in Wang et al. [92]. Following a similar approach to Wang et al. (2023) [94], we assess the resilience of multiLID in the context of two-class degradation: Gaussian blur and JPEG compression. The perturbations consist of five levels for Gaussian blur ($\sigma = 0, 0.15, 0.5, 1, 3$) and four levels for JPEG compression (quality = 100, 90, 60, 30). Furthermore, we augment the training data by incorporating these perturbations. In both scenarios, the multiLID algorithm demonstrates notably high accuracies when the training process involves data augmentation. Notice that accuracy results hold remain consistent regardless of the image size and dataset domain. In section E, we provide additional results from other datasets and conduct an ablation study on the effects of Gaussian blur and JPEG compression degradation.

4.4. Model Strength Assessment

In this subsection, we investigate the boundaries of our approach. In other words, we aim at gaining more insights about the strength of the algorithms depending on the number of samples and the entries (multiLID scores) of the feature vectors. Each extracted feature map of ResNet18’s selected layers ℓ results in 10 multiLID scores. This is indeed the case because we choose to compute the multiLID over the 10th nearest neighbors. Note that the whole length of the feature vector is $10 \times \ell = 80$. The first entries correspond to the first layers and the latter to the last layers of network. We evaluate the detection rates, in terms of accuracy, when using different numbers of samples and accumulating the entries over the feature vectors. In fig. 4, we benchmark our multiLID across two dimensions: i) the number of features; ii) the number of samples. We run this experiment five times to ensure reproducibility. We employ 2,000 samples per class, and our starting training-test split is 60-40%. This implies that the training split is equal to $4,000 \times 0.6 = 2,400$ and hence, 1,600 samples for the test set. Notice that while the training data will be decreased, the test set size keeps always the same (1,600 samples). We can observe how, independently of the dataset, our model only needs 800 synthetic images to learn to distinguish real and DM-generated images.

In addition, one can notice that the first eight entries of the feature vectors do not contribute to the detection, as the detection rate is always around 0.5 across the evaluations (see fig. 4). Similar results were observed by [30] as discussed in




Figure 4: These figures show the strength assessment. We conduct an ablation study of multiLID detection rates in accuracy (ACC). The x-axis represents the length of the feature vectors where the features from earlier layers to deeper layers are accumulated, more details in section 3.2. To that end, we gradually increase the number of samples and accumulate the feature vectors (from the first to the last layers). Each tile represents the mean over five independent runs. The variance can be found in fig. 7 in the section D.

section 3.1. Refer to the section G for an in-depth analysis. As the number of training samples increases, and the feature vector entries are larger than eight, then the detection accuracy becomes uniformly accurate. Moreover, we include the strength assessment over the variance in section D.

4.5. Identification and Transferability Capability Evaluation

In this section, we investigate the identification and transferability capabilities of the multiLID method. To address this objective, we raise the following questions: Can we achieve a dependable identification of each diffusion model through a multilabel classifier? If so, does the multiLID approach retain its transferability when applied to unfamiliar data originating from different models but belonging to the same domain?

To start answering the identification question, we explore the abilities of our approach to LSUN-Bedroom, as it has been widely used in previous literature [68, 94]. In fig. 5, we plot the confusion matrix from different DMs: {DDPM, DDIM, PNDM}-ema, LDM, SDv21, and VQD. The identification results indicate significantly high accuracy scores. Furthermore, we investigate other datasets, such as CelebaHQ (fig. 12a), LSUN-Cat (fig. 13a), LSUN-Church (fig. 14a), to examine the generalizability of the identification. Similarly to LSUN-Bedroom, the accuracy is perfect. Refer to the section F for the results.

Limitation of the Identification. We conduct a series of experiments on the ArtiFact dataset, which comprises 8 authentic datasets, 6 datasets generated from distinct GANs, and 6 datasets produced by different DMs. For these experiments, we utilize a total of 10,500 real images and an equivalent number of 10,500 images generated by GANs and DMs. We deviate from training a binary classifier solely for real and synthetic samples. Instead, we explore the classification of synthetic images originating from GANs and

DMs separately. While accurately distinguishing between real and synthetic images (i.e., GAN or Diffusion) poses no challenge for our approach, we encounter difficulty in reliably differentiating between GAN- and DM-generated images (refer to the left part of fig. 5).

Limitation of the Transferability. On the other hand, when it comes to transferability, we evaluate it in the form of a matrix. We conduct again our experiments on LSUN-Bedroom with different DM-generated images: {DDPM, DDIM, PNDM}-ema, LDM, SDv21, and VQD (see right fig. 5). Each classifier is trained on real and one of the diffusion-generated datasets. We transfer the datasets from other DM-generated datasets. As expected, the accuracy within the same dataset is accurate, however, the transferability is very low. As in the identification investigation, we validate our results on other datasets: CelebaHQ (fig. 12b), LSUN-Cat (fig. 13b), LSUN-Church (fig. 14b) datasets. We obtain the same pattern as for LSUN-Bedroom. Refer to the section F for these results.

Comparison to other Detectors. In table 1, it is evident that our method exhibits superior performance compared to other approaches, particularly when employing a limited quantity of training samples—specifically, 800 samples per class for both training and testing purposes. The reason is that multiLID does not necessitate fine-tuning neural networks, eliminating the need for extensive datasets. MultiLID operates without pre-trained networks, distinguishing it from other methodologies that rely on pre-trained weights from established classification architectures such as Xception, ResNet50, or EfficientNet-B4. Another existing method, as presented by Durall [23], operates without pre-trained weights and is based on simple Fourier analysis. However, as explained by [68], this type of frequency discrimination technique is unsuitable for working with DMs.

We extend our comparison by changing the classification from a binary to a multi-class scenario. From now on the




Figure 5: Limitation of the identification. Left: As described in section 4.1, our experiment is based on the ArtiFact and consists of 8 clean datasets, 6 GAN, and 6 DM-generated images. Center: Identification results on the dataset of LSUN-Bedroom. Right: Transferability results on the dataset of LSUN-Bedroom. The transferability is low, while the identification between clean and synthetic images is accurate. Referring to section F for more evaluations.

Table 1: This table compares the detection accuracy between synthetic and real data based on the LSUN-Bedroom dataset.

method	pre-trained	DDPM ema	DDIM ema	PNDM ema	ADM	LDM	SD-v2.1	VQD	total avg.
Durall [23]	✗	0.66	0.62	0.57	0.53	0.61	0.71	0.71	0.63
F3-Net [65]	✓	0.99	0.99	1.00	0.89	0.99	1.00	1.00	0.98
CNNDet [93]	✓	0.93	0.97	0.99	0.77	0.92	1.00	1.00	0.94
Pa-Fo [14]	✓	0.72	0.71	0.93	0.51	0.76	1.00	1.00	0.80
SBI [80]	✓	0.96	0.93	0.96	0.74	0.82	1.00	1.00	0.92
Ours (binary)	✗	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00

Table 2: This table compares the identification accuracy between synthetic and real (or other synthetic) data based on the LSUN-Bedroom dataset.

method	pre-trained	Real	DDPM ema	DDIM ema	PNDM ema	ADM	LDM	SD-v2.1	VQD	total avg.
Durall [23]	✗	0.13	0.25	0.08	0.16	0.09	0.21	0.65	0.52	0.26
F3-Net [65]	✓	0.96	0.93	0.83	0.99	0.81	0.99	1.00	1.00	0.94
CNNDet [93]	✓	0.73	0.88	0.93	0.99	0.76	0.97	1.00	1.00	0.91
Pa-Fo [14]	✓	0.67	0.88	1.00	1.00	1.00	1.00	1.00	1.00	0.94
SBI [80]	✓	0.46	0.83	0.80	0.97	0.72	0.91	1.00	1.00	0.84
Ours (Ident.)	✗	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00

detectors not only need to distinguish the real and synthetic samples but also among the different synthetics, i.e., different DMs. To that end, we keep the training parameters the same as for the binary case in table 1, but we modify the last layer in all the classifiers. table 2 shows the accuracy scores for all the methods. In general, we can see that the classification scores have decreased since having more classes usually poses a more complex problem for the classifiers. It is noteworthy that the "Real" class is the more troublesome. The multiLID solution exhibits solid detection results. To further assess this claim, we evaluate the experiment on more datasets in sections E and F in the appendix.

5. CONCLUSION

This paper focuses on the detection of diffusion-generated images. Driven by the observation that the grid-like pattern in the Fourier domain is not prominent anymore, we propose the usage of a local intrinsic method variant called multiLID for the examination of diffusion synthetic images. By leveraging multiLID, we seek to gain insights and improve the detection performance specifically in the context of diffusion model-generated images. Moreover, we aim to enhance the detection and identification of diffusion-generated images, addressing the shortcomings observed in previous detectors designed for GAN-generated images, such as FFT dependency or the need of large training datasets. To conduct an in-depth study, we train on publicly available as well as self-constructed datasets consisting of images from different types of diffusion models, such as unconditional, conditional, and text-to-image models. These datasets are specifically curated to enable the evaluation and analysis of DM-generated images. By including images from various DMs, we provide a more comprehensive and diverse set of data for studying and assessing the identification and transferability of diffusion-generated images.

Our extensive experimental results show that the multiLID image representation, significantly enhances the DM-generated identification of images, resulting in a highly effective approach for this particular task. On the other hand, our solution does not offer good transferability, which might reduce the detector’s applicability to unseen synthetics. Nonetheless, we believe that future work can mitigate this drawback.

Acknowledgement Thanks to Jay Wang, who suggested comparing his DiffusionDB with the artistic SAC dataset.

References

- [1] Jay Alammar. The Illustrated Stable Diffusion. <https://jalammar.github.io/illustrated-stable-diffusion>, 2022. [Online; accessed 10-July-2023].
- [2] Roberto Amoroso, Davide Morelli, Marcella Cornia, Lorenzo Baraldi, Alberto Del Bimbo, and Rita Cucchiara. Parents and children: Distinguishing multimodal deepfakes from natural images. *arXiv preprint arXiv:2304.00500*, 2023.
- [3] Laurent Amsaleg, Oussama Chelly, Teddy Furon, Stéphane Girard, Michael E Houle, Ken-ichi Kawarabayashi, and Michael Nett. Estimating local intrinsic dimensionality. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 29–38, 2015.
- [4] Alessio Ansuini, Alessandro Laio, Jakob H Macke, and Davide Zoccolan. Intrinsic dimension of data representations in deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [5] Vishal Asnani, Xi Yin, Tal Hassner, and Xiaoming Liu. Reverse engineering of generative models: Inferring model hyperparameters from generated images. *arXiv preprint arXiv:2106.07873*, 2021.
- [6] Md Awsafur Rahman, Bishmoy Paul, Najibul Haque Sarker, Zaber Ibn Abdul Hakim, and Shaikh Anowarul Fattah. Artifact: A large-scale dataset with artificial and factual images for generalizable and robust synthetic image detection. *arXiv e-prints*, pages arXiv–2302, 2023.
- [7] Jordan J Bird and Ahmad Lotfi. Cifake: Image classification and explainable identification of ai-generated synthetic images. *arXiv preprint arXiv:2303.14126*, 2023.
- [8] Matthew Brand. Charting a manifold. *Advances in neural information processing systems*, 15, 2002.
- [9] Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.
- [10] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [11] Shijie Cao, Lingxiao Ma, Wencong Xiao, Chen Zhang, Yunxin Liu, Lintao Zhang, Lanshun Nie, and Zhi Yang. Seernet: Predicting convolutional neural network feature-map sparsity through low-bit quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11216–11225, 2019.
- [12] Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramèr, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. *arXiv preprint arXiv:2301.13188*, 2023.
- [13] Gunnar Carlsson, Tigran Ishkhanov, Vin De Silva, and Afra Zomorodian. On the local behavior of spaces of natural images. *International journal of computer vision*, 76:1–12, 2008.
- [14] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. What makes fake images detectable? understanding proper-
ties that generalize. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, pages 103–120. Springer, 2020.
- [15] Wenhua Chen, Hexiang Hu, Chitwan Saharia, and William W Cohen. Re-imagen: Retrieval-augmented text-to-image generator. *arXiv preprint arXiv:2209.14491*, 2022.
- [16] Jeonghwan Cheon, Seungdae Baek, and Se-Bum Paik. Invariance of object detection in untrained deep neural networks. *bioRxiv*, pages 2022–09, 2022.
- [17] Yunje Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020.
- [18] Pedro Cuenca. LSUN Bedroom dataset. <https://huggingface.co/datasets/pcuenq/lsun-bedrooms>, 2023. [Online; accessed 10-July-2023].
- [19] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [20] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.
- [21] Chengdong Dong, Ajay Kumar, and Eryun Liu. Think twice before detecting gan-generated fake images from their spectral domain imprints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7865–7874, 2022.
- [22] David L Donoho and Carrie Grimes. Image manifolds which are isometric to euclidean space. *Journal of mathematical imaging and vision*, 23(1):5–24, 2005.
- [23] Ricard Durall, Margret Keuper, Franz-Josef Pfreundt, and Janis Keuper. Unmasking deepfakes with simple features. *arXiv preprint arXiv:1911.00686*, 2019.
- [24] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021.
- [25] Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, 2016.
- [26] German Federal Office for Information Security. Deep Fakes – Threats and Countermeasures. https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Informationen-und-Empfehlungen/Kuenstliche-Intelligenz/Deepfakes/deepfakes_node.html, 2023. [Online; accessed 14-June-2023].
- [27] C Fran et al. Deep learning with depth wise separable convolutions. In *IEEE conference on computer vision and pattern recognition (CVPR)*, 2017.
- [28] Shanghua Gao, Pan Zhou, Ming-Ming Cheng, and Shuicheng Yan. Masked diffusion transformer is a strong image synthesizer. *arXiv preprint arXiv:2303.14389*, 2023.

- [29] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*, 2018.
- [30] Sixue Gong, Vishnu Naresh Boddeti, and Anil K Jain. On the intrinsic dimensionality of image representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3987–3996, 2019.
- [31] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [32] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. *arXiv preprint arXiv:2111.14822*, 2021.
- [33] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10696–10706, 2022.
- [34] Luca Guarnera, Oliver Giudice, and Sebastiano Battiato. Level up the deepfake detection: a method to effectively discriminate images generated by gan architectures and diffusion models. *arXiv preprint arXiv:2303.00608*, 2023.
- [35] Xiao Guo, Xiaohong Liu, Zhiyuan Ren, Steven Grosz, Iacopo Masi, and Xiaoming Liu. Hierarchical fine-grained image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3155–3165, 2023.
- [36] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [37] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [38] David Holz. Dall-e 2. <https://labs.openai.com>, 2022. [Online; accessed 27-June-2023].
- [39] David Holz. Midjourney. <https://docs.midjourney.com/docs/model-versions>, 2022. [Online; accessed 26-June-2023].
- [40] Michael E Houle. Local intrinsic dimensionality i: an extreme-value-theoretic foundation for similarity applications. In *Similarity Search and Applications: 10th International Conference, SISAP 2017, Munich, Germany, October 4-6, 2017, Proceedings 10*, pages 64–79. Springer, 2017.
- [41] Drew A Hudson and Larry Zitnick. Generative adversarial transformers. In *International conference on machine learning*, pages 4487–4499. PMLR, 2021.
- [42] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [43] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33:12104–12114, 2020.
- [44] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Proc. NeurIPS*, 2021.
- [45] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- [46] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *arXiv*, 2009.
- [47] LAION. LAION High Resolution. <https://huggingface.co/datasets/laion/laion-high-resolution>, 2022. [Online; accessed 10-July-2023].
- [48] Ann B Lee, Kim S Pedersen, and David Mumford. The nonlinear statistics of high-contrast patches in natural images. *International Journal of Computer Vision*, 54:83–103, 2003.
- [49] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [50] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. *arXiv preprint arXiv:2202.09778*, 2022.
- [51] Elizaveta Logacheva, Roman Suvorov, Oleg Khomenko, Anton Mashikhin, and Victor Lempitsky. Deeplandscape: Adversarial modeling of landscape videos. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16*, pages 256–272. Springer, 2020.
- [52] Peter Lorenz, Margret Keuper, and Janis Keuper. Unfolding local growth rate estimates for (almost) perfect adversarial detection. *arXiv preprint arXiv:2212.06776*, 2022.
- [53] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *arXiv preprint arXiv:2206.00927*, 2022.
- [54] Xingjun Ma, Bo Li, Yisen Wang, Sarah M Erfani, Sudanthi Wijewickrema, Grant Schoenebeck, Dawn Song, Michael E Houle, and James Bailey. Characterizing adversarial subspaces using local intrinsic dimensionality. *arXiv preprint arXiv:1801.02613*, 2018.
- [55] Stefano Nembrini, Inke R König, and Marvin N Wright. The revival of the gini importance? *Bioinformatics*, 34(21):3711–3718, 2018.
- [56] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- [57] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.

- [58] Bruno A Olshausen and David J Field. Natural image statistics and efficient coding. *Network: computation in neural systems*, 7(2):333–339, 1996.
- [59] OpenAI. Guided Diffusion. <https://github.com/deepfake-study/guided-diffusion>, 2022. [Online; accessed 10-July-2023].
- [60] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Gaugan: semantic image synthesis with spatially adaptive normalization. In *ACM SIGGRAPH 2019 Real-Time Live!*, pages 1–1, 2019.
- [61] William Peebles and Saining Xie. Scalable diffusion models with transformers. *arXiv preprint arXiv:2212.09748*, 2022.
- [62] Gabriel Peyré. Manifold models for signals and images. *Computer vision and image understanding*, 113(2):249–260, 2009.
- [63] Phillip Pope, Chen Zhu, Ahmed Abdelkader, Micah Goldblum, and Tom Goldstein. The intrinsic dimension of images and its impact on learning. *arXiv preprint arXiv:2104.08894*, 2021.
- [64] John David Pressman, Katherine Crowson, and Simulacra Captions Contributors. Simulacra aesthetic captions. Technical Report Version 1.0, Stability AI, 2022. url <https://github.com/JD-P/simulacra-aesthetic-captions>.
- [65] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII*, pages 86–103. Springer, 2020.
- [66] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [67] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [68] Jonas Ricker, Simon Damm, Thorsten Holz, and Asja Fischer. Towards the detection of diffusion model deepfakes. *arXiv preprint arXiv:2210.14571*, 2022.
- [69] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [70] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [71] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.
- [72] Daniel L Ruderman. The statistics of natural images. *Network: computation in neural systems*, 5(4):517, 1994.
- [73] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023.
- [74] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022.
- [75] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- [76] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.
- [77] Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. Projected gans converge faster. *Advances in Neural Information Processing Systems*, 34:17480–17492, 2021.
- [78] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998.
- [79] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*, 2022.
- [80] Kaede Shiohara and Toshihiko Yamasaki. Detecting deepfakes with self-blended images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18720–18729, 2022.
- [81] Sergey Sinitisa and Ohad Fried. Deep image fingerprint: Accurate and low budget synthetic image detector. *arXiv preprint arXiv:2303.10762*, 2023.
- [82] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [83] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [84] StabilityAI. StableDiffusion version 2.1. <https://huggingface.co/stabilityai/stable-diffusion-2-1>, 2023. [Online; accessed 10-July-2023].
- [85] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.

- [86] Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.
- [87] Jonathan Vacher and Ruben Coen-Cagli. Combining mixture models with linear mixing updates: multilayer image segmentation and synthesis. *feedback*, 19:15, 2019.
- [88] Jonathan Vacher, Aida Davila, Adam Kohn, and Ruben Coen-Cagli. Texture interpolation for probing visual perception. *Advances in neural information processing systems*, 33:22146–22157, 2020.
- [89] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [90] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, and Thomas Wolf. Diffusers: State-of-the-art diffusion models, 2022.
- [91] Jay Wang. DiffusionDB: “2m_first_5ktrain”. https://huggingface.co/datasets/poloclub/diffusiondb/viewer/2m_first_5k/train, 2022. [Online; accessed 10-July-2023].
- [92] Sheng-Yu Wang, Oliver Wang, Andrew Owens, Richard Zhang, and Alexei A Efros. Detecting photoshopped faces by scripting photoshop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10072–10081, 2019.
- [93] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8695–8704, 2020.
- [94] Zhendong Wang, Jianmin Bao, Wengang Zhou, Weilun Wang, Hezhen Hu, Hong Chen, and Houqiang Li. Dire for diffusion-generated image detection. *arXiv preprint arXiv:2303.09295*, 2023.
- [95] Zijie J Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. Diffusiondb: A large-scale prompt gallery dataset for text-to-image generative models. *arXiv preprint arXiv:2210.14896*, 2022.
- [96] Daniel Watson, William Chan, Jonathan Ho, and Mohammad Norouzi. Learning fast samplers for diffusion models by differentiating through sample quality. In *International Conference on Learning Representations*, 2022.
- [97] Yutong Xie, Zhaoying Pan, Jinge Ma, Luo Jie, and Qiaozhu Mei. A prompt log analysis of text-to-image generation systems. In *Proceedings of the ACM Web Conference 2023*, pages 3892–3902, 2023.
- [98] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- [99] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *2019 IEEE international workshop on information forensics and security (WIFS)*, pages 1–6. IEEE, 2019.
- [100] Peng Zhou, Xintong Han, Vlad I Morariu, and Larry S Davis. Learning rich features for image manipulation detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1053–1061, 2018.
- [101] Derui Zhu, Dingfan Chen, Jens Grossklags, and Mario Fritz. Data forensics in diffusion models: A systematic analysis of membership privacy. *arXiv preprint arXiv:2302.07801*, 2023.
- [102] Wei Zhu, Qiang Qiu, Jiayi Huang, Robert Calderbank, Guillermo Sapiro, and Ingrid Daubechies. Ldmnet: Low dimensional manifold regularized neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2743–2751, 2018.