Zero-Shot Detection of AI-Generated Images

Davide Cozzolino¹, Giovanni Poggi¹, Matthias Nießner², and Luisa Verdoliva^{1,2}

¹ University Federico II of Naples, 80125 Naples, Italy
 ² Technical University of Munich, 85748 Garching, Germany
 {davide.cozzolino,poggi,verdoliv}@unina.it, niessner@tum.de

Abstract. Detecting AI-generated images has become an extraordinarily difficult challenge as new generative architectures emerge on a daily basis with more and more capabilities and unprecedented realism. New versions of many commercial tools, such as DALLE, Midjourney, and Stable Diffusion, have been released recently, and it is impractical to continually update and retrain supervised forensic detectors to handle such a large variety of models. To address this challenge, we propose a zero-shot entropy-based detector (ZED) that neither needs AI-generated training data nor relies on knowledge of generative architectures to artificially synthesize their artifacts. Inspired by recent works on machinegenerated text detection, our idea is to measure how surprising the image under analysis is compared to a model of real images. To this end, we rely on a lossless image encoder that estimates the probability distribution of each pixel given its context. To ensure computational efficiency, the encoder has a multi-resolution architecture and contexts comprise mostly pixels of the lower-resolution version of the image. Since only real images are needed to learn the model, the detector is independent of generator architectures and synthetic training data. Using a single discriminative feature, the proposed detector achieves state-of-the-art performance. On a wide variety of generative models it achieves an average improvement of more than 3% over the SoTA in terms of accuracy. Code is available at https://grip-unina.github.io/ZED/.

1 Introduction

The quality of AI-generated images has improved tremendously in recent years, to the point where they are virtually indistinguishable from real images upon visual inspection. In addition, the latest generators are widely available online and allow easy creation and retouching of images based on simple textual prompts. All this opens the way to endless application opportunities in a variety of fields, from the creative arts to industries of all kinds. However, on the flip side, such tools can be also used for malicious purposes, thus posing serious threats to our society. For example, pre-trained generators can be easily optimized to generate fake works by a specific artist [31], or used to orchestrate effective, large-scale disinformation campaigns to influence public opinion in advanced democracies [20]. These immediate risks create an urgent need for reliable and automated detection of AI-generated images [41].



Fig. 1: ZED leverages the intrinsic model of real images learned by a state-of-the-art lossless image coder. For real images, the model is correct and the actual coding cost is close its expected value. Synthetic images have different statistics than real images, so they "surprise" the encoder, and the actual coding cost differs significantly from its expected vale. This is evident from the graphic on the right that shows how the coding cost gap increases for synthetic images much more than for real ones when predicting high resolution details from low resolution data.

Until very recently, supervised learning paradigms dominated the image forensics community, with deep models trained on large datasets of real and fake images [64]. These approaches, however, are tailored to specific domains and are difficult to generalize to unseen deepfake samples. In the seminal paper by Wang et al. [74], it is shown that a simple detector trained only on ProGAN images from 20 different categories generalizes well to other images created by different generative adversarial networks (GAN) thanks to suitable augmentation. However, performance still suffers on images generated by prompt-driven diffusion models (DM). Similarly, a detector suitably trained on Latent DM images performs well on all other DM images but fails to generalize properly on GAN images [10]. To reduce the dependence on training data, recent works [2, 11, 51, 67] rely on general-purpose features extracted by pre-trained visual-language models, such as CLIP (Contrastive Language-Image Pre-Training) [56]. Despite the good performance, these methods still depend on the choice of the training dataset. A recent trend to improve generalization is based on few-shot methods [12, 17, 33] which can partially solve the problem, but still require some prior knowledge of the target models, even if limited to a few images. With this work we make a step further and develop an approach that is not influenced at all by newer and previously unseen generative models.

To this end, we propose a zero-shot detection method that only requires real images for learning their underlying distribution. Our key idea is to use lossless coding and a multi-resolution prediction strategy for computing conditional distributions of all image pixels at three different levels of resolution. Given such distributions, we compute statistics related to the actual and expected coding cost. If the image is coherent with the predicted distribution (*no surprise*), then there is no mismatch and the image under analysis is labelled as real. We expect synthetic images to be characterized by a higher coding cost under the distribution of real images (see Fig. 1). Based on this intuition, we design discriminative features that measure how well the image under test fits the model of real images embedded in the encoder. Even by using a single feature, we can obtain significant performance above 95% in terms of AUC for several recent models, such as DALL·E, Midjourney, and SDXL.

In summary, the main contributions of this paper are the following:

- we propose a zero-shot detector of artificially generated images: no fake images are necessary for training which guarantees independence from any specific generation method;
- this is the first work that exploits an implicit model of real images, learnt for lossless encoding to address image forensics task;
- our experiments show on a wide variety of generative models that even using a single feature the proposed detector provides state-of-the-art results (+3.4%) in terms of accuracy).

2 Related work

Supervised learning. The problem of distinguishing synthetic images from real ones is commonly formulated as a binary classification task. State-of-the-art methods explicitly or implicitly exploit forensic artifacts by leveraging a large amount of real and generated images. Some of them rely on semantic flaws, such as face asymmetries [4] or incorrect perspective, lighting, shadows [21, 22, 65]. However, technology advances very quickly and such errors will very likely disappear in next-generation tools. Therefore, most methods focus on low-level and inconspicuous artifacts [9, 18]. Major efforts have been made to prevent conventional supervised detectors from overfitting the training data. Popular recipes include using datasets as varied as possible with intense augmentation [74], pre-training models on large general-purpose datasets [46], preserving fine-grain details of images [7,27], exploiting high-frequency artifacts in the spatial [43, 68, 72] or Fourier domain [18, 24, 78], leveraging inter-pixel correlation discrepancies [71, 79], adopting inversion techniques [1, 75].

With the advent of diffusion models that presents significant architectural differences with GANs, the importance to design methods that work equally well on known and unknown sources became even more evident [10]. An important finding was the increased generalization that could be achieved using pre-trained large vision-language models, such as CLIP-ViT [51]. In this case only a lightweight linear classifier is trained on top of these features to adapt to the forensic task. Very good performance is obtained on DMs even if the network was trained only on GANs. Other methods also show the potential of such approach [2, 11, 59], sometimes including multimodal features [44, 67].

Some supervised methods assume to have only real images available and create the synthetic images needed for training by simulating the artifacts introduced by a generator, for example by passing real images through an autoencoder [24,34,78]. The more generative architectures are simulated, the more effective is the detector. Of course, the performance degrades on images generated by an architecture not considered in the simulation phase. Differently from all these methods our approach does not require collecting or generating synthetic images thus avoiding any type of dependence on this class.

Few-shot/incremental learning. A significant step towards improved generalization is the use of few-shot or incremental learning strategies [12, 17, 33, 47]. Along this path, a recent work [19] proposes to regularly re-train a detector on new synthetic generators in the very same temporal order of their release, as in a real-world scenario. Results show a good generalization to unseen models, but only as long as the architecture of new generators is similar to that of old ones. Although few-shot methods represent an important progress in reducing the dependence on training data, the ultimate goal is to remove this dependence entirely to ensure maximum generalization. In pursuit of this goal, in this work we propose a truly zero-shot detector.

Zero-shot learning. Only a few very recent papers avoid training on synthetic data altogether. A solution was proposed in [60] based on the observation that synthetic images are reconstructed more accurately than real images by a latent DM autoencoder. The main limitation is that the method only reliably detects images generated by latent diffusion models. The method in [30], instead, exploits the fact that small perturbations of [real/synthetic] images correspond to [small/large] variations in the embedding space of a pre-trained large model. Differently from these strategies our work takes inspiration from some interesting proposals that have recently appeared for synthetic text detection [25, 29, 49, 69]. They exploit the fact that LLMs (Large Language Models) work by generating the probability distribution of the next token given the previous ones. In the generation phase, new tokens are sequentially added to a sentence based on these distributions. In the analysis phase, one can replicate the process for a given sentence under test and measure how well the actual tokens match the predicted ones. A good match suggests that the sentence was indeed generated by an LLM. Although inspired by these methods, our zero-shot synthetic image detector differs from them because it leverages a model of *real* images and does not depend in any way on synthetic data or generators. Moreover, to build the model we take advantage of the remarkable field-proved ability of lossless encoders to accurately describe pixels based on their context.

3 Method

3.1 Background

Here we provide some background on zero-shot methods that leverage large pre-trained language models for machine-generated text detection. They exploit the native functionality of these models to provide next-token predictions [29]. Before a string of characters s can be processed by a language model, it must be parsed into a sequence of tokens (mostly words). The tokenizer T outputs a list of indices

$$T: s \to \{x_0, x_1, \dots, x_L\},\tag{1}$$

where $x_i \in \{1, ..., n\}$ is the index of the *i*-th token of the sequence, addressing a size-*n* vocabulary of tokens. The language model operates by predicting the next

index/token given the list of previous ones, thereby allowing for the generation of a full sentence given just a short prompt. Actually, language models output more information than just the index of the most likely token. Given the list of previous indices $X_i = \{x_0, \ldots, x_{i-1}\}$, they provide the probability of all possible values of the current one, that is, $P(x_i = k | X_i)$, for $k = 1, \ldots, n$.

The idea is to exploit this functionality to measure the conformity of the string under analysis to the LLM intrinsic model of language. That is, these methods try to answer the question "How likely is it that this sentence was generated by my LLM?" Hence they compute (for free) the likelihood of the given list of indices under the probability distribution learned by the LLM

$$P(x_0, \dots, x_L) = P(x_0) \cdot P(x_1 | x_0) \cdot \dots \cdot P(x_L | x_0, \dots, x_{L-1}) = P(x_0) \prod_{i=1}^L P(x_i | X_i)$$
(2)

In practice, the negative log-likelihood (also called log-perplexity) is computed instead, that is (neglecting x_0)

$$\text{NLL} = -\sum_{i=1}^{L} \log P(x_i | X_i) \tag{3}$$

If the *i*-th observed index x_i was very likely to come after the previous ones, namely, it is not surprising, its contribution to the NLL is close to 0. On the contrary, if it was unlikely to appear, given the previous ones (an anomaly) it impacts significantly on the NLL. Overall, a sequence with low NLL is likely to have been generated by the LLM, and will be therefore detected as synthetic. Of course, this basic description is only meant to convey the general concepts, the reader is referred to the literature [26] for more details.

3.2 From Text to Images

When we try to translate the above concepts into the realm of images, we run into a big problem: the most effective and popular image generation engines *do not* provide anything similar to the next token distribution observed in the case of LLMs. Indeed, there exist some autoregressive synthesis methods [45,58] that could be adapted to this task, but their generation approach is very different from those of the most popular GAN- and DM-based methods. Therefore in this work we change perspective or, better said, we now assume the correct one-class perspective, and look for a model of real images, rather than synthetic ones. Armed with such a model, we will be able to decide whether a given image is unsurprising, therefore real, or somewhat anomalous, therefore synthetic, regardless of the specific generation model used to create it.

Now, the concepts of prediction, surprise, perplexity, along with information measure and entropy, are pervasive in the literature on image coding, part of information theory. Lossless image encoders typically include a predictor that, given a suitable context, estimates the value of the target pixel, and an entropy encoder that efficiently represents prediction errors. Indeed, by analyzing the recent literature in the field we managed to single out a tool that perfectly suits our needs, the Super-Resolution based lossless Compressor (SReC) proposed by Cao *et al.* [6], which provides a computationally lightweight tool for predicting the distribution of image pixels at multiple resolution.

3.3 Super-resolution based Lossless Compressor

Here we present a high-level description of SReC, focusing only on the aspects more relevant for our purposes. The interested reader is referred to the original paper for details [6]. The general idea is to train a neural network to predict the current pixel, $x_{i,j}$, given a set of previously coded pixels, and encode the difference between the true pixel value and its prediction. However, this purely autoregressive formulation is highly impractical, as it implies long encoding/decoding times. Therefore, SReC uses a multi-resolution prediction strategy. A low-resolution version $y^{(1)}$ of the original image $x^{(0)}$ is built through 2×2 average pooling, that is

$$y_{i,j}^{(1)} = \frac{x_{2i,2j}^{(0)} + x_{2i+1,2j}^{(0)} + x_{2i,2j+1}^{(0)} + x_{2i+1,2j+1}^{(0)}}{4}$$
(4)

Then, each four-pixel group of the high-resolution image is predicted based only on the low-resolution image, independent of other groups at the same resolution level, allowing for parallel processing and high-speed encoding. Since the fourth pixel of a group is known, given the other three and the low resolution image, the conditional joint distribution of the group reads

$$P(x_{2i,2j}^{(0)}, x_{2i+1,2j}^{(0)}, x_{2i,2j+1}^{(0)} | Y_{i,j}^{(1)}) = P(x_{2i,2j}^{(0)} | Y_{i,j}^{(1)}) \cdot P(x_{2i+1,2j}^{(0)} | x_{2i,2j}^{(0)}, Y_{i,j}^{(1)}) \\ \cdot P(x_{2i,2j+1}^{(0)} | x_{2i,2j}^{(0)}, x_{2i+1,2j}^{(0)}, Y_{i,j}^{(1)})$$
(5)

where $Y_{i,j}^{(1)}$ is the relevant context in the lower resolution image, that is a receptive field centered on $y_{i,j}^{(1)}$. Each term in this factorization is estimated by a dedicated convolutional neural network (CNN). In particular, a parametric distribution is assumed, given by the mixture of K discrete logistic distributions,

$$P(x|X) = \sum_{k=1}^{K} w_k \text{logistic}(x|\mu_k, s_k)$$
(6)

where $\operatorname{logistic}(x|\mu, s) = \sigma(\frac{x-\mu+0.5}{s}) - \sigma(\frac{x+\mu+0.5}{s})$ is the difference of two sigmoid functions, with position parameter μ and scale parameter s, and K = 10 is always assumed. The CNN takes the context X of the pixel of interest as input and outputs the weights of the mixture together with the position and scale parameters of all logistics. In turn, these parameters allow one to compute the desired distribution. This whole process is replicated on two more lower-resolution scales, for a total of four levels, the lowest resolution, an 8×8 subsampled "prompt"



Fig. 2: NLL and Entropy. We compute the spatial distribution of NLL and Entropy at three resolutions. For real images (top) the paired maps are very similar at all scales: when the uncertainty on a pixel (entropy) grows, also the coding cost (NLL) does. Therefore, the NLL-Entropy difference maps are all very dark. For synthetic images (bottom) NLL and Entropy maps are not always similar, because the model is not correct, and hence the difference maps are brighter, especially the high-resolution map.

image, coded in clear, and three higher resolution images, each one predicted from its lower resolution version. All networks are trained to minimize the cross entropy between the predicted model probability $P_{\theta}(x)$ and the empirical data distribution P(x) given by the training image dataset. We mention in passing that this loss is closely related to the log-perplexity considered for text synthesis.

To summarize, SReC provides us with a lightweight tool for computing conditional distributions of all image pixels at three different levels of resolution, and therefore to compute all kinds of statistics that can expose the mismatch between a test image and the learned model. Considering that SReC achieves state-of-the-art performance in lossless image compression, one can also argue that the learned model of real images is very accurate. Given this tool, we can now design a zero-shot detector of synthetic images.

3.4 Features and Decision Statistics

Let $x \in \{0, \ldots, 255\}^{N \times M \times 3}$ be the image under test. In our multi-resolution framework, this will be the highest-resolution version, $x^{(0)} = x$. Through 2×2 average pooling, we generate a lower resolution version $y^{(1)} = \operatorname{avpool}(x^{(0)})$, and then, through rounding, its integer-valued version $x^{(1)} = \operatorname{round}(y^{(1)})$. The process is repeated, and eventually we have four integer versions of the image $\{x^{(0)}, x^{(1)}, x^{(2)}, x^{(3)}\}$, together with three non-integer versions $\{y^{(1)}, y^{(2)}, y^{(3)}\}$. In the context of lossless coding, the lowest resolution version, $x^{(3)}$, must be sent in clear together with the rounding bits at levels 3, 2, and 1, but we mention this only for completeness and for a more compelling interpretation of results. The CNNs trained on real images provide the predicted probability distribution



Fig. 3: Extracting decision statistics. The full resolution image $x^{(0)}$ is downsampled three times. The lowest-resolution version, $x^{(3)}$, feeds the level-2 CNN, which outputs the probability distributions of level-2 pixels. These distributions, together with the actual level-2 pixels, are used to compute the level-2 coding cost NLL⁽²⁾ and its expected value $H^{(2)}$. All these steps are then repeated for levels 1 and 0. Eventually, NLLs and entropies are combined to compute the decision statistics.

for all pixels³ of levels 0, 1, and 2

$$P(x_{i,j}^{(l)} = k | X_{i,j}^{(l)}) \tag{7}$$

where $k \in \{0, \ldots, 255\}$ and $X_{i,j}^{(l)}$ is the context for pixel $x_{i,j}^{(l)}$, including a portion of the lower-resolution image $y^{(l+1)}$ and possibly some same-resolution neighbors of the current pixel. Given the above distribution, we compute the negative log likelihood and the entropy at each pixel

$$NLL_{i,j}^{(l)} = -\log P(x_{i,j}^{(l)} | X_{i,j}^{(l)}) H_{i,j}^{(l)} = -\sum_{k} P(k | X_{i,j}^{(l)}) \log P(k | X_{i,j}^{(l)})$$
(8)

These quantities are shown in Fig.2 for two sample images, real and synthetic. Then, through spatial averaging, we obtain the corresponding quantities for the images at all resolution levels $\text{NLL}^{(l)} = \langle \text{NLL}^{(l)}_{i,j} \rangle$ and $H^{(l)} = \langle H^{(l)}_{i,j} \rangle$, for l = 0, 1, 2. These are the features associated by the system to input image and our decision statistics will be suitable combinations of them.

Before going on, it is convenient to give a physical interpretation of these quantities. Each NLL can be interpreted as the actual coding cost for the corresponding image. While each entropy can be interpreted as the expected value of the coding cost given the context, when the image is coherent with the predicted distribution. In the presence of a mismatch, NLL – H > 0, on the average, with a gap that increases with increasing distribution mismatch. Our fundamental assumption is that the trained CNNs provide a good model of real images, and synthetic images tend not to follow the same model. Therefore, we expect that synthetic images are characterized by higher coding cost, hence higher NLL, under this distribution. This observation would lead us to use the NLLs as decision

³ More precisely, all color components of all pixels, but to simplify notations, in the following we will neglect color and treat the image as if grayscale.

statistics. However, the coding cost does not depend only on the distribution mismatch but also (predominantly) on the intrinsic information content of the image, measured by the entropy. A complex image, say a photo of a crowd, is more difficult to encode/describe than a smooth image, say a blue sky, no matter what model we use. Therefore, to get rid of this bias, we consider the coding cost gap, defined as the difference $D^{(l)} = \text{NLL}^{(l)} - H^{(l)}$, as decision statistic. Hence, for each image, we have three basic decision statistics, one for each resolution level. It is worth observing that some forms of normalization are adopted for machine generated text detection as well [29, 49, 70]. A block diagram of our method is shown in Fig.3.

A sample graph of the coding cost gap is shown in Fig.1, on the right. For real images and three families of synthetic images we report the average gap (solid line) plus/minus its standard deviation (colored band) for the various resolutions levels. Two important observations can be made. First of all, the level-0 coding cost gap, concerning the full resolution image, seems to be much more discriminant than the others. Moreover, the gap grows much faster for synthetic images than for real images when going from level 1 to level 0. Therefore, as decision statistics we will consider both $D^{(0)}$ (the level-0 coding cost gap) and $\Delta^{01} = D^{(0)} - D^{(1)}$ (its slope). In addition, in preliminary experiments we observed that synthetic images are sometimes characterized by a coding cost much *lower* rather than much higher than expected, that is the NLL is much lower than the entropy. This is also an anomaly, which signals the likely synthetic nature of the image. Therefore, besides the above statistics we also consider their absolute values $|D^{(0)}|$ and $|\Delta^{(01)}|$. These observations are supported by the sample graphical analysis shown in Fig.5 in the ablation study.

4 Results

4.1 Datasets and Metrics

We benchmarked our model on a large variety of synthetic generators both GANs and DMs: GauGAN [53], BigGAN [5], StarGAN [8], StyleGAN2 [38], Diffusion-GAN [76], GigaGAN [35], GALIP [73], DDPM [32], ADM [16], GLIDE [50], Stable Diffusion [62, 63], DiT [54], DeepFloyd-IF [39], Stable Diffusion XL [55], DALL E [14], DALL E 2 [57], DALL E 3 [52], Midjourney V5 [48], and Adobe Firefly [23]. We collected images from publicly available datasets [3,10,51,74] and generated additional images as needed when they were not publicly available. We ensured that all datasets included pristine and synthetic images with similar semantic content, both compressed and uncompressed, to avoid any kind of bias (see Fig 4). For some synthetic generators we have multiple datasets, built on the basis of different real image datasets LSUN [77], FFHQ [37], ImageNet [15], COCO [42], LAION [66] and RAISE [13]. This is a fortunate circumstance: we kept them carefully separate as this allows us to analyze how the performance of a detector depends on the class of real images used in the synthesis phase. Overall we used a total of 29k synthetic images and 6k real images. More details on the generated and actual images are provided in the supplementary material.



Fig. 4: Examples of real and AI-generated images of different categories used in our experiments. Top: real images from LSUN, FFHQ, ImageNET and COCO. Bottom: generated images from DiffusionGAN, StyleGAN2, DiT and SDXL.

Following other papers [11, 43, 51] we measure performance using the area under the ROC curve (AUC) and the balanced accuracy. We also show the influence of the threshold selection on the performance.

4.2 Ablation Study

Features analysis. First, we want to provide a better insight into the role and importance of the features described in Section 3.4: $D^{(0)}$ (the 0-level coding cost gap), its slope $\Delta^{01} = D^{(0)} - D^{(1)}$ and their absolute values. To this end, we consider the set of real and synthetic (DALL E 2, GLIDE, Midjourney, SDXL) images of the Synthbuster dataset [3]. We note, in passing, that this dataset includes only uncompressed images, which dispels any possible doubt that our method exploits some JPEG compression bias between real and fake images [28]. Some selected scatter plots and graphs are shown in Fig.5. The rightmost box shows that encoding cost (NLL) and entropy (H) alone are not very discriminating, even if computed at the more informative level 0 (high resolution). In contrast, their difference, the 0-level coding cost gap $D^{(0)}$, seems to separate the different classes quite well (central box), in particular the real class (violet) from the others. Note that the level-1 gap (not shown) is not equally discriminating, and the level-2 gap, plotted on the y axis, turns out to be essentially useless. In the third box we plot the empirical distributions of $D^{(0)}$ for the various classes. This representation makes the good separability of the classes further clear but also highlights an unexpected phenomenon: GLIDE images group mostly to the left of the real class, that is, they have a lower-than-expected coding cost. Although not in line with our initial hypotheses, this fact nevertheless represents an anomaly, which can be detected by thresholding the absolute value of the statistic rather than the statistic itself.



Fig. 5: Decision statistics. NLL and entropy by themselves are not discriminant (left). Their difference (center) is much more useful for detection, but only at high resolution, $D^{(0)}$, while $D^{(1)}$ is less discriminant and $D^{(2)}$ basically useless. Right box shows histograms of $D^{(0)}$ for real and synthetic images. Note that for GLIDE, $D^{(0)}$ is negative, on the average. Good discrimination is still possible based on the absolute value.



Fig. 6: AUC of proposed method as a function of decision statistic (see Section 3.4) and dataset of real images used to train the lossless encoder: Open Images, LAION, COCO, and their augmented versions (*). Synthetic test images are selected to match the corresponding real test images: ImageNet (top), and LAION (bottom).

Influence of the real class. To better understand the role of the real dataset used to train the lossless encoder, we perform an experiment in which we vary it. Along with the original encoder pre-trained on the Open Images dataset [40] (about 338k high-resolution images), we consider two other versions, trained from scratch on the LAION dataset [66] (\simeq 117k), and the COCO dataset [42] (\simeq 106k), respectively, using the same hyperparameters as [6]. Additionally, we consider versions (marked with *) trained on the same datasets, augmented with JPEG compressed images with quality between 80 and 100. We compute the performance in terms of AUC on two different datasets of synthetic and

Table 1: Reference methods. For each one we indicate the key idea, the datasets of real and synthetic images used for training with their sizes, whether or not augmentation is used, the test strategy.

Acronym [ref]	Idea/Approach	Training Real/Fake	$\operatorname{Size}(\mathbf{K})$	Augment.	Test Strategy
Wang2020 [74]	High diversity	LSUN/ProGAN	360/360	\checkmark	global pooling
PatchFor. [7]	Patch-based	${ m CelebA}, { m FF}/{ m various}$	84/272		resizing
Liu2022 [43]	Noise-based	LSUN/ProGAN	360/360	\checkmark	global pooling
Corvi2023 [10]	No-downsampling	$\rm COCO, LSUN/Latent$	180/180	\checkmark	global pooling
LGrad [72]	Gradient-based	LSUN/ProGAN	72/72	\checkmark	resizing
DIRE [75]	Inversion	LSUN-Bed/ADM	40/40		resizing
DE-FAKE [67]	Prompt-based	LSUN/Stable Diff.	20/20		resizing
Ojha2023 [51]	CLIP	LSUN/ProGAN	360/360	\checkmark	cropping
NPR [71]	Residual	LSUN/ProGAN	72/72		resizing
AEROBLADE $[60]$	AE rec. error	- / -	- / -		global distance

real images, where this latter class comes from ImageNet [15] (Fig.6, top) or LAION [66] (Fig.6, bottom). We can observe that the best and more uniform results across the four decision statistics are obtained using COCO^{*}, while training on Open Images guarantees good performance if the real class is LAION, but bad performance if it is ImageNet. Additional results are included in the supplementary material.

4.3 SoTA Comparison

In our analysis we include only methods with code and/or pre-trained models publicly available on-line. Eventually, we included 7 CNN-based methods [7,10, 43,71,72,74,75], 2 CLIP-based methods [51,67] and a training-free method [60]. A brief summary of these techniques is provided in Tab.1, while a more detailed description is given in the supplementary material. For a fair comparison we avoid testing on ProGAN [36] and Latent Diffusion [61], because a good number of these supervised methods were trained on datasets that include images from these generators. Even so, we have a total of 30 datasets for testing. Results are reported in Tab.2 in terms of AUC, with the best figure for each dataset highlighted in bold. Note that each row is characterized by the name of the generator (e.g., GauGAN) and by a single letter that recalls the set of real images used to train it: S for LSUN, F for FFHQ, I for ImageNet, C for COCO, L for LAION, R for RAISE. This detail allows us to study how the performance depends on the real dataset (but with synthetic images from the same generator and with semantic content aligned with real images).

First of all, we observe that for most reference methods the average AUC does not exceed 80%. Notable exceptions are the CLIP-based Ojha2023 (88.4%) and the CNN-based Corvi2023 (89.4%). Interestingly, some methods show very different performance when the real class changes. This may be due to JPEG bias as already suggested in [28,60]. A deeper analysis on this point is presented

13

	Real data	Wang2020	PatchFor.	Liu2022	Corvi2023	LGrad	DIRE	DEFAKE	Ojha2023	NPR	AEROBLADE	$Ours \ D^{(0)}$	$\operatorname{Ours}\left D^{(0)} \right $	Ours Δ^{01}	$Ours \left \Delta^{01} \right $
GauGAN	С	98.9	80.8	99.7	83.8	81.6	99.9	43.8	100.	89.1	55.1	99.8	99.8	99.9	99.7
$\operatorname{Big}GAN$	Ι	92.7	85.5	94.7	83.4	77.2	99.8	59.0	99.6	86.8	51.9	92.3	88.6	95.9	92.6
$\operatorname{Star} \operatorname{GAN}$	F	94.7	100.	99.9	95.9	73.9	40.4	45.9	99.7	81.5	84.0	100.	100.	100.	100.
$\operatorname{StyleGAN2}$	\mathbf{S}	98.1	83.8	99.7	89.1	99.8	58.3	39.1	96.7	100.	30.0	96.6	96.1	96.7	96.5
	\mathbf{F}	94.9	85.1	99.9	58.4	82.7	55.5	47.6	91.0	71.3	60.1	43.1	87.7	41.1	88.7
$\operatorname{GigaGAN}$	Ι	73.7	61.0	97.3	50.5	76.4	99.9	64.3	94.6	82.4	47.5	72.4	68.1	72.4	68.1
	\mathbf{C}	79.5	84.0	99.6	90.9	76.7	99.9	87.9	97.6	95.5	80.6	96.5	94.0	96.7	93.4
Diff.GAN	S	89.8	92.6	99.5	96.6	99.5	49.8	44.8	97.4	100.	43.9	99.4	99.4	99.5	99.5
GALIP	\mathbf{C}	89.7	98.2	94.3	87.7	56.7	100.	75.6	98.6	90.7	65.0	98.4	96.3	99.7	99.4
DALL E	L	66.4	71.7	95.0	98.3	95.2	99.8	55.9	97.3	99.5	24.1	99.2	95.8	98.2	95.4
DDPM	F	31.6	98.4	22.8	100.	9.8	23.1	50.5	77.7	92.4	81.7	76.6	25.2	93.8	79.6
ADM S	S	67.6	67.6	70.6	80.3	81.1	52.0	37.4	88.2	94.1	53.1	49.5	53.5	69.4	71.0
	I	61.0	81.9	94.4	81.1	72.7	99.5	69.1	85.3	78.5	80.3	87.8	90.5	95.3	92.1
	С	64.8	97.4	96.3	97.2	81.5	99.9	92.4	88.8	95.4	98.0	47.8	88.5	91.1	91.1
GLIDE	R	32.2	95.0	56.6	86.5	50.6	42.9	92.2	72.8	63.3	87.7	23.2	89.4	51.1	65.1
	L	72.6	74.1	90.8	86.9	90.3	100.	60.2	95.3	99.8	68.7	54.5	84.2	93.8	88.5
DiT	Ι	58.6	83.1	88.0	100.	56.2	99.6	87.4	77.8	78.4	99.8	89.4	84.3	94.9	91.0
Stable D. 1.4	С	68.2	86.1	95.3	100.	54.7	99.9	93.3	97.9	76.5	99.8	48.4	74.8	54.6	71.4
	R	37.9	61.8	73.4	100.	50.0	37.6	88.0	87.7	43.0	96.9	99.4	98.7	97.0	97.2
Stable D. 2	С	56.5	78.6	94.2	100.	62.8	99.3	97.9	82.3	89.3	99.9	83.0	90.3	84.5	89.1
	R	50.2	38.7	34.8	100.	41.4	35.5	80.7	89.5	44.0	97.4	98.5	96.8	95.8	95.9
SDXL R	C	83.8	60.8	89.3	100.	89.3	99.5	94.0	80.0	99.3	87.9	99.9	99.9	99.9	99.8
	R	54.3	68.4	31.1	100.	57.2	47.1	84.4	85.1	76.7	69.7	100.	100.	99.1	99.2
DeepIF	С	78.0	62.7	72.2	99.9	68.8	98.9	96.9	92.9	91.6	81.9	91.7	82.3	88.4	79.4
DALLE 2	C	88.5	52.4	98.9	88.2	78.6	99.9	80.6	97.1	90.0	59.3	100.	100.	100.	99.9
	R	64.8	41.9	70.4	69.4	58.6	44.7	70.9	95.2	39.5	32.8	100.	100.	100.	100.
DALLE 3	C	65.0	47.3	99.5	100.	88.4	99.9	96.2	86.4	97.7	99.7	99.7	99.5	98.3	98.2
	R	10.9	52.7	0.2	60.8	37.9	47.6	92.4	36.4	48.7	48.3	79.1	66.7	78.0	78.1
Midjourney	R	40.2	57.8	40.7	100.	56.3	51.0	78.1	66.2	77.0	99.0	99.7	99.3	98.5	98.5
Adobe Firefly	ĸ	84.8	49.4	11.8	98.0	40.6	57.4	81.4	97.5	32.1	52.8	73.6	41.2	80.8	80.4
AVG		68.3	73.3	77.0	89.4	68.2	74.6	72.9	88.4	80.1	71.2	83.3	86.4	88.8	90.0

Table 2: AUC for reference and proposed methods. Best score in bold with a 0.5% margin. S = LSUN, F = FFHQ, I = ImageNet, C = COCO, L = LAION, R = RAISE.

in the supplementary material. The proposed zero-shot approach goes above 80% with all decision statistics, reaching the top value of 90.0% when $|\Delta^{01}|$ is used. Obviously, this is a very good result, but what makes it especially valuable is the absence of any dependence on the generators' models. This point is further stressed by the fact that the AUC remains extremely stable across all test sets, with a minimum of 65.1% on GLIDE-R. On the contrary, the best competitor, Corvi2023, has a long score of top results but also some very poor ones. suggesting a certain instability, likely due to the presence/absence of specific artifacts in the test images, and eventually the risk of not adapting to models of new conception. We want also to draw the reader's attention on the already mentioned case of GLIDE and on the fact that the proposed method exhibits wildly different results with different decision statistics. In particular, with $|D^{(0)}|$ the AUC is 89.4% as opposed to the already mentioned 65.1% with $|\Delta^{01}|$. This suggests there may be better ways to exploit the basic NLL⁽¹⁾ and $H^{(1)}$, possibly jointly at all levels, to synthesize a better and more stable decision statistics.

Finally, in Fig.7, we report the accuracy as a function of the decision threshold for the best methods. A separate curve is shown for each real image dataset by



Fig. 7: Balanced accuracy as a function of the detection threshold. For each dataset of real images, we average accuracy over all associated synthetic generators. The dotted vertical line indicates the global optimal threshold and the \times symbol the corresponding accuracy. Note that only for the proposed method all peaks are very close, indicating the presence of a single threshold. Charts for other methods are reported in the Suppl.

averaging over the associated synthetic generators. Unlike AUC, the accuracy critically depends on the selection of a good threshold and some calibration data may be needed for this purpose. Note that only for the proposed method there is a single good threshold that ensures near-optimal accuracy for all datasets.

4.4 Limitations

Our work was developed to detect whether an image has been fully generated and not to detect local manipulations. However, it could be easily extended to accomplish this task since we already compute a map of local pixel-wise statistics. Furthermore, our approach relies on a model of the real class learned by the encoder. If real images do not satisfy this model, the approach may not perform correctly. For example, if images are highly compressed or resized (as is the case on the web), statistical analysis may not be reliable.

5 Conclusion

We introduced a novel zero-shot forensic detector to distinguish AI-generated images from real ones. Unlike most current methods, our approach does not require fake images during training, which ensures generalization to yet unknown generative models. The idea is to exploit an implicit model of real images and classify off-model images as synthetic. To this end, we leverage an appropriate lossless encoder, trained only on real images, that can predict the probability distribution of each pixel given its context. Synthetic images are expected to not respect this distribution, thus revealing their artificial nature. Our experiments show that the proposed detector is consistently competitive with detectors trained in supervised modality, and outperforms them in terms of generalization ability. We believe that our approach is an important stepping stone towards effective forensic tools that can operate without relying on domain- or methodspecific training data. Future work will focus on making the method robust to the most common forms of image impairment, so as to make it suitable for in the wild application. Acknowledgments. We gratefully acknowledge the support of this research by a TUM-IAS Hans Fischer Senior Fellowship, the ERC Starting Grant Scan2CAD (804724), and a Google Gift. This material is also based on research sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under agreement number FA8750-20-2-1004. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government. In addition, this work has received funding by the European Union under the Horizon Europe vera.ai project, Grant Agreement number 101070093.

References

- Albright, M., McCloskey, S.: Source Generator Attribution via Inversion. In: CVPR Workshop. pp. 96-103 (2019)
- Amoroso, R., Morelli, D., Cornia, M., Baraldi, L., Del Bimbo, A., Cucchiara, R.: Parents and Children: Distinguishing Multimodal DeepFakes from Natural Images. ACM Trans. Multimedia Comput. Commun. Appl. (2024)
- 3. Bammey, Q.: Synthbuster: Towards Detection of Diffusion Model Generated Images. IEEE Open Journal of Signal Processing (2023)
- Boháček, M., Farid, H.: A geometric and photometric exploration of GAN and Diffusion synthesized faces. In: CVPR Workshop. pp. 874—883 (2023)
- 5. Brock, A., Donahue, J., Simonyan, K.: Large Scale GAN Training for High Fidelity Natural Image Synthesis. In: ICLR (2018)
- Cao, S., Wu, C.Y., Krähenbühl, P.: Lossless Image Compression through Super-Resolution. arXiv preprint arXiv:2004.02872v1 (2020)
- Chai, L., Bau, D., Lim, S.N., Isola, P.: What Makes Fake Images Detectable? Understanding Properties that Generalize. In: ECCV. pp. 103-120 (2020)
- Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In: CVPR. pp. 8789–8797 (2018)
- Corvi, R., Cozzolino, D., Poggi, G., Nagano, K., Verdoliva, L.: Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In: CVPR Workshop. pp. 973–982 (2023)
- Corvi, R., Cozzolino, D., Zingarini, G., Poggi, G., Nagano, K., Verdoliva, L.: On the detection of synthetic images generated by diffusion models. In: ICASSP. pp. 1–5 (2023)
- Cozzolino, D., Poggi, G., Corvi, R., Nießner, M., Verdoliva, L.: Raising the Bar of AI-generated Image Detection with CLIP. In: CVPR Workshop. pp. 4356–4366 (2024)
- Cozzolino, D., Thies, J., Rössler, A., Riess, C., Nießner, M., Verdoliva, L.: Forensictransfer: Weakly-supervised domain adaptation for forgery detection. arXiv preprint arXiv:1812.02510 (2018)
- Dang-Nguyen, D.T., Pasquini, C., Conotter, V., Boato, G.: RAISE: A Raw Images Dataset for Digital Image Forensics. In: ACM MMSys. p. 219-224 (2015)

- 16 Cozzolino et al.
- Dayma, B., Patil, S., Cuenca, P., Saifullah, K., Abraham, T., Lê Khàc, P., Melas, L., Ghosh, R.: DALL-E Mini (2021). https://doi.org/10.5281/zenodo.5146400, https://github.com/borisdayma/dalle-mini
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. In: CVPR. pp. 248-255 (2009)
- Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. NeurIPS 34, 8780-8794 (2021)
- 17. Du, M., Pentyala, S., Li, Y., Hu, X.: Towards Generalizable Deepfake Detection with Locality-Aware AutoEncoder. In: CIKM. pp. 325-334 (2020)
- Durall, R., Keuper, M., Keuper, J.: Watch Your Up-Convolution: CNN Based Generative Deep Neural Networks Are Failing to Reproduce Spectral Distributions. In: CVPR. pp. 7890-7899 (2020)
- Epstein, D.C., Jain, I., Wang, O., Zhang, R.: Online Detection of AI-Generated Images. In: ICCV Workshop. pp. 382-392 (2023)
- Epstein, Z., Hertzmann, A., Herman, L., Mahari, R., Frank, M.R., Groh, M., Schroeder, H., Akten, A.S.M., Fjeld, J., Farid, H., Leach, N., Pentland, A.S., Russakovsky, O.: Art and the science of generative AI: A deeper dive. arXiv preprint arXiv:2306.04141 (2023)
- 21. Farid, H.: Lighting (in) consistency of paint by text. arXiv preprint arXiv:2207.13744 (2022)
- 22. Farid, H.: Perspective (in) consistency of paint by text. arXiv preprint arXiv:2206.14617 (2022)
- 23. Firefly, A.: https://www.adobe.com/sensei/generative-ai/firefly.html (2023)
- Frank, J., Eisenhofer, T., Schönherr, L., Fischer, A., Kolossa, D., Holz, T.: Leveraging Frequency Analysis for Deep Fake Image Recognition. In: ICML. pp. 3247–3258 (2020)
- Gehrmann, S., Strobelt, H., Rush, A.M.: GLTR: Statistical detection and visualization of generated text. In: 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations. pp. 111-116 (2019)
- Ghosal, S.S., Chakraborty, S., Geiping, J., Huang, F., Manocha, D., Bedi, A.S.: Towards possibilities & impossibilities of AI-generated text detection: A survey. arXiv preprint arXiv:2310.15264 (2023)
- Gragnaniello, D., Cozzolino, D., Marra, F., Poggi, G., Verdoliva, L.: Are GAN generated images easy to detect? A critical analysis of the state-of-the-art. In: ICME. pp. 1-6 (2021)
- Grommelt, P., Weiss, L., Pfreundt, F.J., Keuper, J.: Fake or JPEG? Revealing Common Biases in Generated Image Detection Datasets. arXiv preprint arXiv:2403.17608 (2024)
- Hans, A., Schwarzschild, A., Cherepanova, V., Kazemi, H., Saha, A., Goldblum, M., Geiping, J., Goldstein, T.: Spotting LLMs With Binoculars: Zero-Shot Detection of Machine-Generated Text. In: ICML (2024)
- He, Z., Chen, P.Y., Ho, T.Y.: RIGID: A Training-free and Model-Agnostic Framework for Robust AI-Generated Image Detection. arXiv preprint arXiv:2405.20112 (2024)
- Heikkilä, M.: This artist is dominating AI-generated art. and he's not happy about it. MIT Technology Review (2022)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS 33, 6840–6851 (2020)
- Jeon, H., Bang, Y.O., Kim, J., Woo, S.: T-GD: Transferable GAN-generated Images Detection Framework. In: ICML. vol. 119, pp. 4746-4761 (2020)

- Jeong, Y., Kim, D., Ro, Y., Kim, P., Choi, J.: FingerprintNet: Synthesized Fingerprints for Generated Image Detection. In: ECCV. pp. 76-94 (2022)
- Kang, M., Zhu, J.Y., Zhang, R., Park, J., Shechtman, E., Paris, S., Park, T.: Scaling up gans for text-to-image synthesis. In: CVPR. pp. 10124-10134 (2023)
- Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive Growing of GANs for Improved Quality, Stability, and Variation. In: ICLR (2018)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR. pp. 4401-4410 (2019)
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of StyleGAN. In: CVPR. pp. 8110-8119 (2020)
- Konstantinov, M., Shonenkov, A., Bakshandaeva, D., Schuhmann, C., Ivanova, K., Klokova, N.: https://www.deepfloyd.ai/deepfloyd-if (2023)
- Krasin, I., Duerig, T., Alldrin, N., Ferrari, V., Abu-El-Haija, S., Kuznetsova, A., Rom, H., Uijlings, J., Popov, S., Veit, A., et al.: OpenImages: A public dataset for large-scale multi-label and multi-class image classification. Dataset available from https://github.com/openimages (2017)
- 41. Lin, L., Gupta, N., Zhang, Y., Ren, H., Liu, C.H., Ding, F., Wang, X., Li, X., Verdoliva, L., Hu, S.: Detecting multimedia generated by large ai models: A survey. arXiv preprint arXiv:2204.06125 (2024)
- 42. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: ECCV. pp. 740-755 (2014)
- 43. Liu, B., Yang, F., Bi, X., Xiao, B., Li, W., Gao, X.: Detecting generated images by real images. In: ECCV. pp. 95–110 (2022)
- Liu, H., Tan, Z., Tan, C., Wei, Y., Wang, J., Zhao, Y.: Forgery-aware Adaptive Transformer for Generalizable Synthetic Image Detection. In: CVPR. pp. 10770– 10780 (2024)
- Mahajan, S., Roth, S.: PixelPyramids: Exact Inference Models from Lossless Image Pyramids. In: ICCV. pp. 6639–6648 (2021)
- 46. Mandelli, S., Bonettini, N., Bestagini, P., Tubaro, S.: Detecting GAN-generated Images by Orthogonal Training of Multiple CNNs. In: ICIP. pp. 3091-3095 (2022)
- 47. Marra, F., Saltori, C., Boato, G., Verdoliva, L.: Incremental learning for the detection and classification of GAN-generated images. In: WIFS. pp. 1–6 (2019)
- 48. Midjourney: https://www.midjourney.com/home (2023)
- Mitchell, E., Lee, Y., Khazatsky, A., Manning, C.D., Finn, C.: DetectGPT: Zero-Shot Machine-Generated Text Detection using Probability Curvature. In: ICML. pp. 24950-24962 (2023)
- Nichol, A.Q., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., Mcgrew, B., Sutskever, I., Chen, M.: GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diff. Models. In: ICML. pp. 16784–16804 (2022)
- 51. Ojha, U., Li, Y., Lee, Y.J.: Towards universal fake image detectors that generalize across generative models. In: CVPR. pp. 24480-24489 (2023)
- 52. OpenAI: https://openai.com/dall-e-3 (2023)
- 53. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: CVPR. pp. 2337-2346 (2019)
- 54. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: ICCV. pp. 4195-4205 (2023)
- Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: SDXL: Improving latent diffusion models for high-resolution image synthesis. In: ICLR (2024)

- 18 Cozzolino et al.
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748-8763 (2021)
- 57. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical Text-Conditional Image Generation with CLIP Latents. arXiv preprint arXiv:2204.06125 (2022)
- Reed, S.E., van den Oord, A., Kalchbrenner, N., Colmenarejo, S.G., Wang, Z., Chen, Y., Belov, D., de Freitas, N.: Parallel multiscale autoregressive density estimation. In: ICML. pp. 2912–2921 (2017)
- Ricker, J., Damm, S., Holz, T., Fischer, A.: Towards the detection of diffusion model deepfakes. In: VISAPP. pp. 446-457 (2024)
- Ricker, J., Lukovnikov, D., Fischer, A.: AEROBLADE: Training-Free Detection of Latent Diffusion Images Using Autoencoder Reconstruction Error. In: CVPR. pp. 9130-9140 (2024)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
- 62. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: https://github. com/CompVis/stable-diffusion (2022)
- 63. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: https://github. com/Stability-AI/stablediffusion (2022)
- Rössler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Nießner, M.: Faceforensics++: Learning to detect manipulated facial images. In: ICCV. pp. 1-11 (2019)
- 65. Sarkar, A., Mai, H., Mahapatra, A., Lazebnik, S., Forsyth, D.A., Bhattad, A.: Shadows Don't Lie and Lines Can't Bend! Generative Models don't know Projective Geometry... for now. In: CVPR. pp. 28140-28149 (2024)
- Schuhmann, C., Kaczmarczyk, R., Komatsuzaki, A., Katta, A., Vencu, R., Beaumont, R., Jitsev, J., Coombes, T., Mullis, C.: LAION-400M: Open Dataset of CLIP-Filtered 400 Million Image-Text Pairs. In: NeurIPS (2021)
- Sha, Z., Li, Z., Yu, N., Zhang, Y.: DE-FAKE: Detection and Attribution of Fake Images Generated by Text-to-Image Generation Models. In: ACM SIGSAC. pp. 3418-3432 (2023)
- Sinitsa, S., Fried, O.: Deep Image Fingerprint: Towards Low Budget Synthetic Image Detection and Model Lineage Analysis. In: WACV. pp. 4067–4076 (2024)
- 69. Solaiman, I., Brundage, M., Clark, J., Askell, A., Herbert-Voss, A., Wu, J., Radford, A., Krueger, G., Kim, J.W., Kreps, S., et al.: Release Strategies and the Social Impacts of Language Models. arXiv preprint arXiv:1908.09203 (2019)
- 70. Su, J., Zhuo, T.Y., Wang, D., Nakov, P.: DetectLLM: Leveraging Log Rank Information for Zero-Shot Detection of Machine-Generated Text. In: Conference on Empirical Methods in Natural Language Processing (2023)
- Tan, C., Zhao, Y., Wei, S., Gu, G., Liu, P., Wei, Y.: Rethinking the Up-Sampling Operations in CNN-based Generative Network for Generalizable Deepfake Detection. In: CVPR. pp. 28130-28139 (2024)
- 72. Tan, C., Zhao, Y., Wei, S., Gu, G., Wei, Y.: Learning on Gradients: Generalized Artifacts Representation for GAN-Generated Images Detection. In: CVPR. pp. 12105-12114 (2023)
- 73. Tao, M., Bao, B.K., Tang, H., Xu, C.: Galip: Generative adversarial clips for textto-image synthesis. In: CVPR. pp. 14214–14223 (2023)
- 74. Wang, S.Y., Wang, O., Zhang, R., Owens, A., Efros, A.A.: CNN-generated images are surprisingly easy to spot... for now. In: CVPR. pp. 8692-8701 (2020)

19

- 75. Wang, Z., Bao, J., Zhou, W., Wang, W., Hu, H., Chen, H., Li, H.: DIRE for Diffusion-Generated Image Detection. ICCV pp. 22445-22455 (2023)
- 76. Wang, Z., Zheng, H., He, P., Chen, W., Zhou, M.: Diffusion-GAN: Training GANs with Diffusion. In: ICLR (2023)
- 77. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. arXiv preprint arXiv:1506.03365 (2015)
- Zhang, X., Karaman, S., Chang, S.F.: Detecting and Simulating Artifacts in GAN Fake Images. In: WIFS. pp. 1-6 (2019)
- 79. Zhong, N., Xu, Y., Qian, Z., Zhang, X.: Rich and Poor Texture Contrast: A Simple yet Effective Approach for AI-generated Image Detection. arXiv preprint arXiv:2311.12397v1 (2023)