Parents and Children: Distinguishing Multimodal DeepFakes from Natural Images

ROBERTO AMOROSO*, University of Modena and Reggio Emilia, Italy
DAVIDE MORELLI*, University of Modena and Reggio Emilia, Italy and University of Pisa, Italy
MARCELLA CORNIA, University of Modena and Reggio Emilia, Italy
LORENZO BARALDI, University of Modena and Reggio Emilia, Italy
ALBERTO DEL BIMBO, University of Florence, Italy
RITA CUCCHIARA, University of Modena and Reggio Emilia, Italy and IIT-CNR, Italy

Recent advancements in diffusion models have enabled the generation of realistic deepfakes from textual prompts in natural language. While these models have numerous benefits across various sectors, they have also raised concerns about the potential misuse of fake images and cast new pressures on fake image detection. In this work, we pioneer a systematic study on deepfake detection generated by state-of-the-art diffusion models. Firstly, we conduct a comprehensive analysis of the performance of contrastive and classification-based visual features, respectively extracted from CLIP-based models and ResNet or ViT-based architectures trained on image classification datasets. Our results demonstrate that fake images share common low-level cues, which render them easily recognizable. Further, we devise a multimodal setting wherein fake images are synthesized by different textual captions, which are used as seeds for a generator. Under this setting, we quantify the performance of fake detection strategies and introduce a contrastive-based disentangling method that lets us analyze the role of the semantics of textual descriptions and low-level perceptual cues. Finally, we release a new dataset, called COCOFake, containing about 1.2M images generated from the original COCO image-caption pairs using two recent text-to-image diffusion models, namely Stable Diffusion v1.4 and v2.0.

CCS Concepts: • **Computing methodologies** → **Image representations**; **Matching**; *Computer vision tasks.*

Additional Key Words and Phrases: multimodal deepfakes, vision-and-language, generative models

ACM Reference Format:
Roberto Amoroso, Davide Morelli, Marcella Cornia, Lorenzo Baraldi, Alberto Del Bimbo, and Rita Cucchiara. 2024. Parents and Children: Distinguishing Multimodal DeepFakes from Natural Images. *ACM Trans. Multimedia Comput. Commun. Appl.* (2024)

1 INTRODUCTION

Machine-generated images have gained extensive popularity in the digital world due to the popularity of GANs [25, 35, 36, 48] and diffusion models [15, 52, 55, 58]. While image generation tools can be employed for lawful goals, such as assisting content creators, generating simulated datasets, or enabling multimodal interactive applications, they have raised concerns regarding their potential for illegal and malicious purposes [2, 8, 14, 30]. These include the forgery of natural images, the generation of images in support of fake news, and the generation of NSFW contents [50, 60]. In

*Both authors contributed equally to this research.

Authors' addresses: Roberto Amoroso, University of Modena and Reggio Emilia, Modena, Italy, roberto.amoroso@unimore.it; Davide Morelli, University of Modena and Reggio Emilia, Modena, Italy and University of Pisa, Pisa, Italy, davide.morelli@unimore.it; Marcella Cornia, University of Modena and Reggio Emilia, Modena, Italy, marcella.cornia@unimore.it; Lorenzo Baraldi, University of Modena and Reggio Emilia, Modena, Italy, lorenzo.baraldi@unimore.it; Alberto Del Bimbo, University of Florence, Florence, Italy, alberto.delbimbo@unifi.it; Rita Cucchiara, University of Modena and Reggio Emilia, Modena, Italy and IIT-CNR, Pisa, Italy, rita.cucchiara@unimore.it.

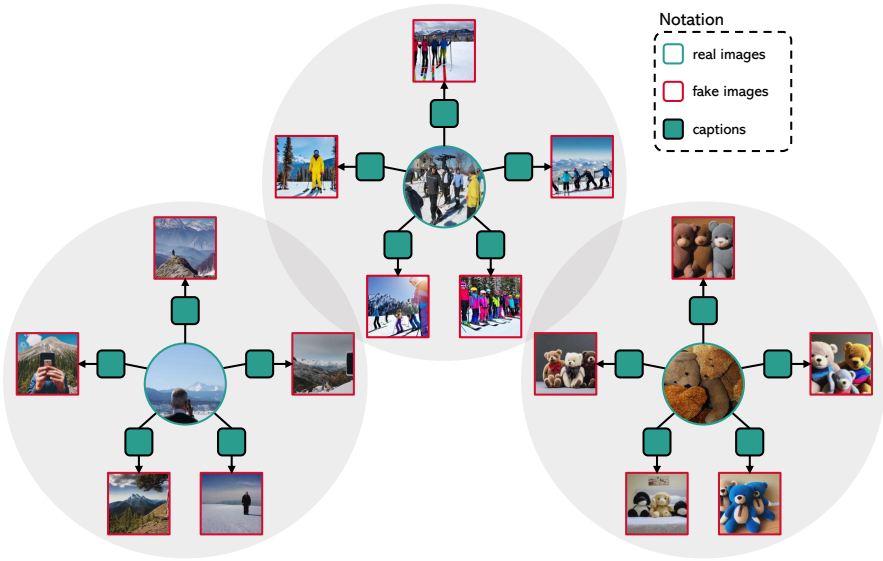

Fig. 1. Overview of our multimodal deepfakes detection setting, in which five subsets of the semantics contained in a given image are employed to generate as many fake images.

this context, assessing the authenticity of images becomes a fundamental goal for security and for guaranteeing the trustworthiness of AI algorithms.

Most of the past approaches for deepfake detection have employed perceptual cues [19, 22, 73], including frequency analysis, the detection of artifacts, or pixel discontinuities. Furthermore, a significant portion of the early studies has focused exclusively on fake faces [40, 41, 56]. Today's generators [16, 23, 52, 53, 55, 58] are general-purpose, text-driven, and exhibit higher generation quality. If we look at images generated by Stable Diffusion [55] (a few examples are reported in Sec. 3.5), we might notice that some of them appear hyper-realistic and, thus, easily recognizable, while others contain semantic anomalies. However, most of them are realistically plausible.

In this paper, we aim at developing a systematic study on deepfake detection, in an era when generated content is becoming increasingly realistic and text-driven. We do this in a multimodal setting that enables us to examine deepfake detection from both a perceptual and a semantic perspective. Specifically, given an image, we consider different textual descriptions and fake images generated by using each of the descriptions as a prompt (Fig. 1). In this manner, we build clusters sharing similar semantics, containing one real image and multiple fake images. Under this setting, we first train a classifier to recognize deepfakes and investigate the effectiveness of different visual features extracted from both contrastive-based backbones like CLIP [51] and classification-based ones such as ResNet [31] and ViT-based networks [18] trained on ImageNet. Surprisingly, we find out that high-level contrastive-based features learned on image and text pairs are very effective in discriminating between real and generated images. We hypothesize that low-level perceptual features also percolate into such descriptors, even though they are trained at a semantic level.

While these findings might be effective in defending us from current generators, we can expect that tomorrow's generators will increase their quality and become less detectable via low-level features. Thus, we devise a contrastive-based disentanglement strategy that enables to remove the contribution of low-level features. This approach establishes a more complex setting in which generated images cannot be distinguished at a perceptual level. Under this setting, we propose and

discuss a general procedure for discriminating between fake and real images based on semantic information. To evaluate the effectiveness of the proposed method, we introduce a new dataset, namely COCOFake, which comprises approximately 1.2M images generated from the original COCO image-caption pairs using both Stable Diffusion v1.4 and v2.0 as text-to-image generative models.

Contributions In summary, the main contributions of this work are as follows:

- We develop a framework that utilizes machine-generated variants of natural images to investigate the detectability of diffusion model-generated images at the semantic and perceptual levels. By filtering the semantic content of natural images through natural language descriptions, we create a dataset of machine-generated images that can be used to investigate the performance of fake detection against modern diffusion models.
- We demonstrate that contrastive-based features can be effectively employed for fake detection against modern diffusion models, with high recognition rates.
- We propose a contrastive-based disentanglement approach to distinguish between low-level and semantic features in modern visual extractors. This allows us to distinguish between natural images and the generated ones using only semantic cues while neglecting the perceptual ones. This is important for the future development of more realistic generators.
- We generate and release the COCOFake dataset[1], which contains over 1.2M fake images linked to natural images through captions. This dataset can be used to test and evaluate the performance of fake detection algorithms against diffusion model-generated images and assess their robustness in detecting fake images generated by different text-to-image generative models.

2 RELATED WORK

General Deepfake Detection. In recent years, with the growth and diffusion of generative models, several research efforts [13, 67] have been made to effectively detect synthetic images generated by GANs [25, 35, 36, 48, 75] and other deep learning-based architectures [39, 64]. While initial works did not concentrate on the generalization capabilities of deepfake detectors [47, 56], subsequent approaches [5, 11, 24, 26, 46, 68] focused instead on the development of generic detectors that can be applied to different generators, thus avoiding the need to have a specific detector for each generative model. On the same line, different solutions [19, 22, 73] proposed to detect deepfakes based on the spectrum of GAN-generated images. In fact, CNN-based generative models usually leave a distinguishable fingerprint over generated images, due to transposed convolutions [19, 73], up-sampling operations [6, 22], and the spectral bias of convolution layers [20, 37]. Some works in similar directions also focused on associating fake images to the corresponding generator among several known GANs [33, 72] or extending deepfake detection to the video domain [12, 27–29, 71]. In the latter case, deepfakes are usually generated by partially manipulating original videos with existing tools for face swapping and other sophisticated algorithms for audio manipulation. Research efforts in this domain have mainly been dedicated to improving deepfake detection performance with the integration of multiple modalities, such as spatial rich model filters [28, 45] and audio traces [7, 71] in both cases combined with RGB features.

Detection of Deepfakes Generated with Diffusion Models. While all aforementioned methods are tailored for detecting deepfakes generated by GANs or other visual forgery tools, a few works extended the analysis to deepfake images coming from diffusion models [15, 49, 52, 55, 58]. Among them, Wolter et al. [69] proposed to detect fake images based on their wavelet-packet representations taking into account features from the pixel and frequency space. Ricker et al. [54] evaluate the

[1]The dataset can be downloaded at this link: https://github.com/aimagelab/COCOFake.

performance of state-of-the-art detectors and also tackle the frequency domain, analyzing different factors that influence the spectral properties of these images, discovering that GANs and diffusion models produce images with different characteristics that require adaptation of existing classifiers to ensure reliable detection. Similarly, Corvi et al. [10] introduced an analysis of the forensics traces left by common diffusion models and investigated whether deepfake detectors tailored for GANs can also distinguish images generated by diffusion models. Finally, Sha et al. [63] analyzed and compared deepfakes generated by different text-to-image diffusion models, investigating the possibility of correctly attributing deepfake images to the diffusion model that generated them. Overall, these studies highlight the need for developing detection methods that can effectively detect deepfakes generated by various types of generative models, including diffusion models.

Datasets for Deepfake Detection. The availability of large datasets has played a crucial role in the development of deepfake detection techniques. One of the most widely used datasets is FaceForensics++ [56], which contains videos of real and fake faces generated using several generative models. The dataset provides both raw and manipulated videos with different compression rates and resolutions, allowing the evaluation of deepfake detection methods under different scenarios. Another popular dataset is Celeb-DF [41], which contains videos of celebrities manipulated using different techniques including GANs and face swapping. Celeb-DF also provides several levels of difficulty, ranging from low-quality to high-quality forgeries, making it suitable for evaluating both traditional and advanced deepfake detection methods. Other datasets have been proposed, such as DeeperForensics-1.0 [32], which contains manipulated videos generated using multiple GAN-based models, and DFDC [17], composed of thousands of videos of real and fake faces.

Despite the availability of these datasets, there is still a need for more diverse and challenging datasets that reflect the increasing sophistication of deepfake generation methods. In particular, while current datasets mainly focus on faces, there is a lack of datasets for detecting deepfakes in other types of images, such as natural scenes. The proposed COCOFake dataset aims to address this limitation by providing a large-scale dataset of natural images and their corresponding synthetic images generated by diffusion models, along with natural language captions linking them. This allows for the evaluation of deepfake detection methods in a more complex and diverse context and also enables the development of methods that can identify semantic inconsistencies between natural and synthetic images.

3 PROPOSED METHOD

3.1 Notation and Preliminaries

We propose a framework for studying and detecting multimodal generated fake images, which encompasses the identification and separation of their perceptual and semantic components. In the rest of the paper, we will employ the following notation: I_R will indicate a natural (real) image, C a textual description (*i.e.*, a caption), and I_F will indicate a fake image produced by a generator. Under this setting, a *parent* real image I_R can be the seed for N different *children* fake images $I_{F,i}$ given a set of textual descriptions $\{C_i\}$ of I_R, with $i = 1, ..., N$, by using each of the descriptions as prompt for the generator.

Semantic and Style Components of an Image. The information content of an image can be credited to many factors. For simplicity, we assume that an image I, regardless of its authenticity, embodies two information contributions, namely a *semantic component* $\mathcal{H}_{sem}(I)$ and a perceptual or *style component* $\mathcal{H}_{sty}(I)$. The former represents the content that could be expressed in a textual sentence, while the latter describes the image appearance, encompassing elements such as colors, textures, brightness, and low-level visual cues. Given a real image I_R, we can therefore express its

total information $\mathcal{H}$ as a function of its semantic and style components, as follows:

$$\mathcal{H}(I_R) = f(\mathcal{H}_{sem}(I_R), \mathcal{H}_{sty}(I_R)). \tag{1}$$

However, when an image is described through a natural language sentence, only a portion of its semantics is actually conveyed inside the caption. In other words, natural language descriptions act as a filter for the semantic content of the image. Hence, we introduce $\Delta\mathcal{H}_{sem}(I, C)$ to represent the portion of semantic information described by a caption C. By analogy, we could say that the textual descriptions of an image act as DNA fragments that can be utilized to generate an offspring of images.

Generating Offspring with Natural Language Utterances. From an input image I_R we can, therefore, extract N semantic information subsets $\Delta\mathcal{H}^i_{sem}(I_R, \cdot)$ and feed them to a generator obtaining N different fake images $I_{F,i}$, with $i = 1, ..., N$. We define *semantic cluster* the ensemble of the starting real image I_R and the offspring of N fake images $I_{F,i}$ generated from it. For instance, given a real image dataset such as COCO [42], containing K images, each represented by $N = 5$ captions, we could create K clusters of $N + 1$ images with one parent and N children.

3.2 Learning to Discriminate Real and Fake images

Once a dataset in the aforementioned form has been built, we first measure to what extent real and generated images can be discriminated independently from their membership to a semantic cluster. Instead of doing this by learning ad-hoc visual features, we investigate the usage of state-of-the-art pre-trained visual models. In other words, given a dataset containing both real and generated images, we develop a model that identifies real images by using visual features extracted with a pre-trained backbone. Regarding the generation of the images, in the following, we will employ Stable Diffusion [55], which is freely available and represents a state-of-the-art approach. Nevertheless, the approach could be easily extended to other generators.

To evaluate the discriminative power of current pre-trained visual features, we model the discriminator as a two-class linear classifier, so that input visual features are only linearly projected before taking the final decision on their realism. Formally, given a real image $I_R \in \mathbb{R}^{3\times H\times W}$ and an image encoder $E_I : \mathbb{R}^{3\times H\times W} \rightarrow \mathbb{R}^D$, we extract a vectorial image feature F_I as

$$F_I = E_I(I_R). \tag{2}$$

The features F_I are then fed into a linear layer $L : \mathbb{R}^D \rightarrow \mathbb{R}$, whose output is thresholded to classify between *real* (*i.e.*, 0) and *fake* (*i.e.*, 1) images. As it will be discussed in the experimental section, our findings indicate that this is (still) a relatively simple task even when employing a state-of-the-art generator. This is, most likely, due to the fact that fake images are slightly different in terms of low-level cues with respect to real images.

3.3 Semantic Preservation Analysis

As a second analysis, we investigate the preservation of semantic information across both real and generated fake images. To do so, we consider a multimodal embedding space, in which both images and texts can be projected. Specifically, we verify if, starting from a generated image, we can retrieve the particular caption used as prompt during its generation. In other words, we test if the subset of the real semantic information $\Delta\mathcal{H}_{sem}(I_F, C)$ associated with a caption C is still recognizable in the visual features extracted from the generated image.

Formally, given a caption C describing a real image I_R, and a textual encoder E_T, we tokenize and extract the textual features F_T as:

$$F_T = E_T(C). \tag{3}$$

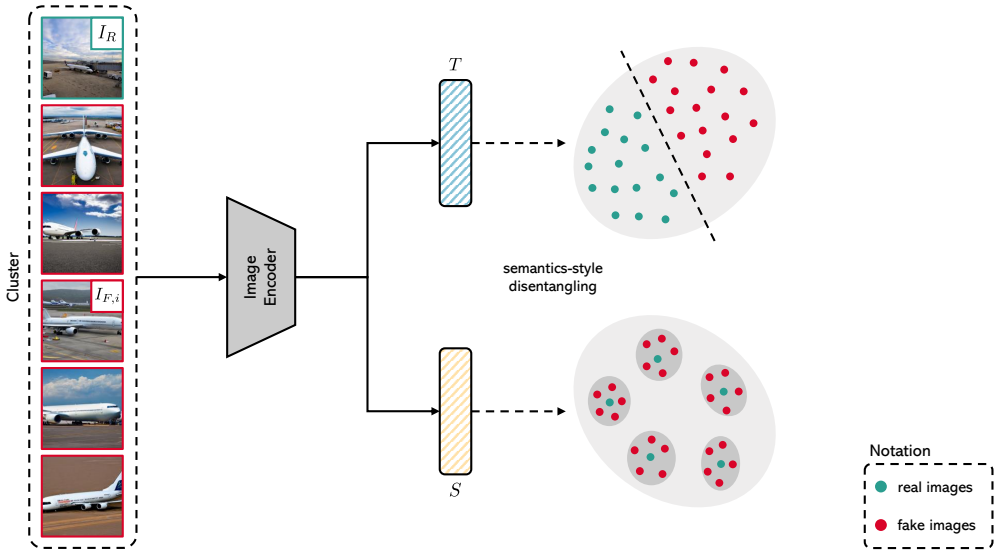

Fig. 2. Schema of our approach for disentangling semantics and style for deepfake detection.

For each visual feature of a given fake image I_F, we verify the ability to retrieve the corresponding textual feature used to create I_F through the generator model.

As it will be shown in the experimental section, we find out that (a) the alteration of low-level cues induced by the generator does not affect the semantic contribution coming from the original image, and (b) the semantic contribution of the generator does not obfuscate the original semantic content.

3.4 Disentangling Semantics and Style

As the detection of fake images is likely promoted by the difference in low-level cues between generated and real images, we finally investigate a more challenging setting in which the style component induced by the generator is disentangled and removed. To do so, we learn a model which identifies the style component of the generator which is common to all generated images. We then measure whether, after eliminating such a component, the remaining semantic information is sufficient to discriminate between real and fake images. Noticeably, this corresponds to a more challenging setting where all the common low-level traits left by the generator are removed and not employed to perform deepfake classification. In other words, this also corresponds to recognizing fakes generated by an "ideal" generator that does not leave common low-level traits.

To perform this analysis, we propose a new contrastive-based learning model that can project images in a semantic space and in a style space (Fig. 2). For a good style-semantic disentanglement we expect that, in the style embedding latent space, the feature vectors of real images should be separated from features of fake images in a cluster-agnostic way, while in the semantic embedding latent space the cluster compactness should be preserved. Specifically, we train two separate linear projections T and S, where T focuses on style while S on semantics. For the T layer we aim at increasing the distance between fake and real elements, regardless of their membership in a specific cluster. For the S layer, instead, we want to create compact clusters of elements sharing the same semantic content, while increasing the distance among two fake elements or two real elements.

We express these requirements through two loss components $\mathcal{L}_c$ and $\mathcal{L}_{fr}$. The former attracts elements of the same cluster, while the latter attracts elements having the same label (*i.e.*, real and fake). From here, we can define the losses needed to train T and S, respectively, as follows:

$$\begin{aligned} \mathcal{L}_T &= \mathcal{L}_{fr} - \mathcal{L}_c, \\ \mathcal{L}_S &= \mathcal{L}_c - \mathcal{L}_{fr}. \end{aligned} \tag{4}$$

To implement both $\mathcal{L}_c$ and $\mathcal{L}_{fr}$, we leverage a Supervised Contrastive Loss [38], defined as follows:

$$\mathcal{L}_{SupCon} = \sum_{i \in I} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(\mathcal{F}_i \mathcal{F}_p^\top / \tau)}{\sum\limits_{a \in A(i)} \exp(\mathcal{F}_i \mathcal{F}_a^\top / \tau)}, \tag{5}$$

where $i \in I \equiv \{1, ..., N+1\}$ represents the index of an arbitrary sample, $\mathcal{F}$ are ℓ_2-normalized input features of a given image, τ is a temperature parameter, $A(i) \equiv I / \{i\}$. $P(i)$ is the set of indices of all items sharing the same label of i, and $|P(i)|$ is its cardinality.

Depending on the nature of the labels used in the training of the supervised contrastive loss, we can implement repulsive and attractive forces in the form of the loss components $\mathcal{L}_c$ and $\mathcal{L}_{fr}$. In $\mathcal{L}_c$, in particular, we assign the same label to elements belonging to the same cluster, while in $\mathcal{L}_{fr}$ we assign the same label to all real samples, and the same label to all fake images. The objective of $\mathcal{L}_c$ is to attract elements of the same cluster, while $\mathcal{L}_{fr}$ pushes real and fake images.

3.5 The COCOFake Dataset for Multimodal Deepfake Recognition

In literature, to the best of our knowledge, there are no multimodal datasets containing texts, real and fake images that are compatible with our multimodal setting. Thus, we generate and release the COCOFake dataset, an extension of COCO [42]. Each real image in COCOFake is paired with five fake images that are conditionally generated based on each of the captions associated with the same image. We employ the Stable Diffusion model [55] as our generator. Specifically, we create two different versions of our dataset, one based on Stable Diffusion v1.4[2] and the other based on Stable Diffusion v2.0[3]. Both text-to-image generators have been pre-trained on the English image-text pairs of the LAION-5B dataset [61] and finetuned on the LAION-Aesthetics subset[4]. While Stable Diffusion v1.4 is based on the CLIP ViT-L/14 text encoder [51], the 2.0 version exploits the OpenCLIP ViT-H/14 one [51]. During image generation, we employ the safety checker module to reduce the probability of explicit images and disable the invisible watermarking of the outputs to prevent easy identification of the images as machine-generated.

Overall, referring to the splits defined in [34] and typically employed in image captioning literature [1, 4, 59], the COCO dataset comprises 113,287 training images, 5,000 validation, and 5,000 test images. Preserving the same splits, COCOFake is composed of 679,722 training images, 30,000 validation, and 30,000 test images for each version of Stable Diffusion, thus comprising more than 1.2M generated images (*i.e.*, around 600k for each version of Stable Diffusion). Sample real-generated image clusters from the COCOFake dataset are shown in Fig. 3. For each example, we present the real image alongside the five fake images generated from each of the five captions from the original COCO dataset. As it can be seen, the generated images are generally coherent with the corresponding caption. However, in some cases, the generated images are overly realistic with brighter colors and a more professional photographic style than the real counterpart. This can be attributed to the dataset employed in the finetuning phase (*i.e.* the LAION-Aesthetics subset) of the Stable Diffusion model [55], used to generate fake images. In Fig. 4 we report less realistic

[2]https://huggingface.co/CompVis/stable-diffusion-v1-4
[3]https://huggingface.co/stabilityai/stable-diffusion-2-base
[4]https://laion.ai/blog/laion-aesthetics/

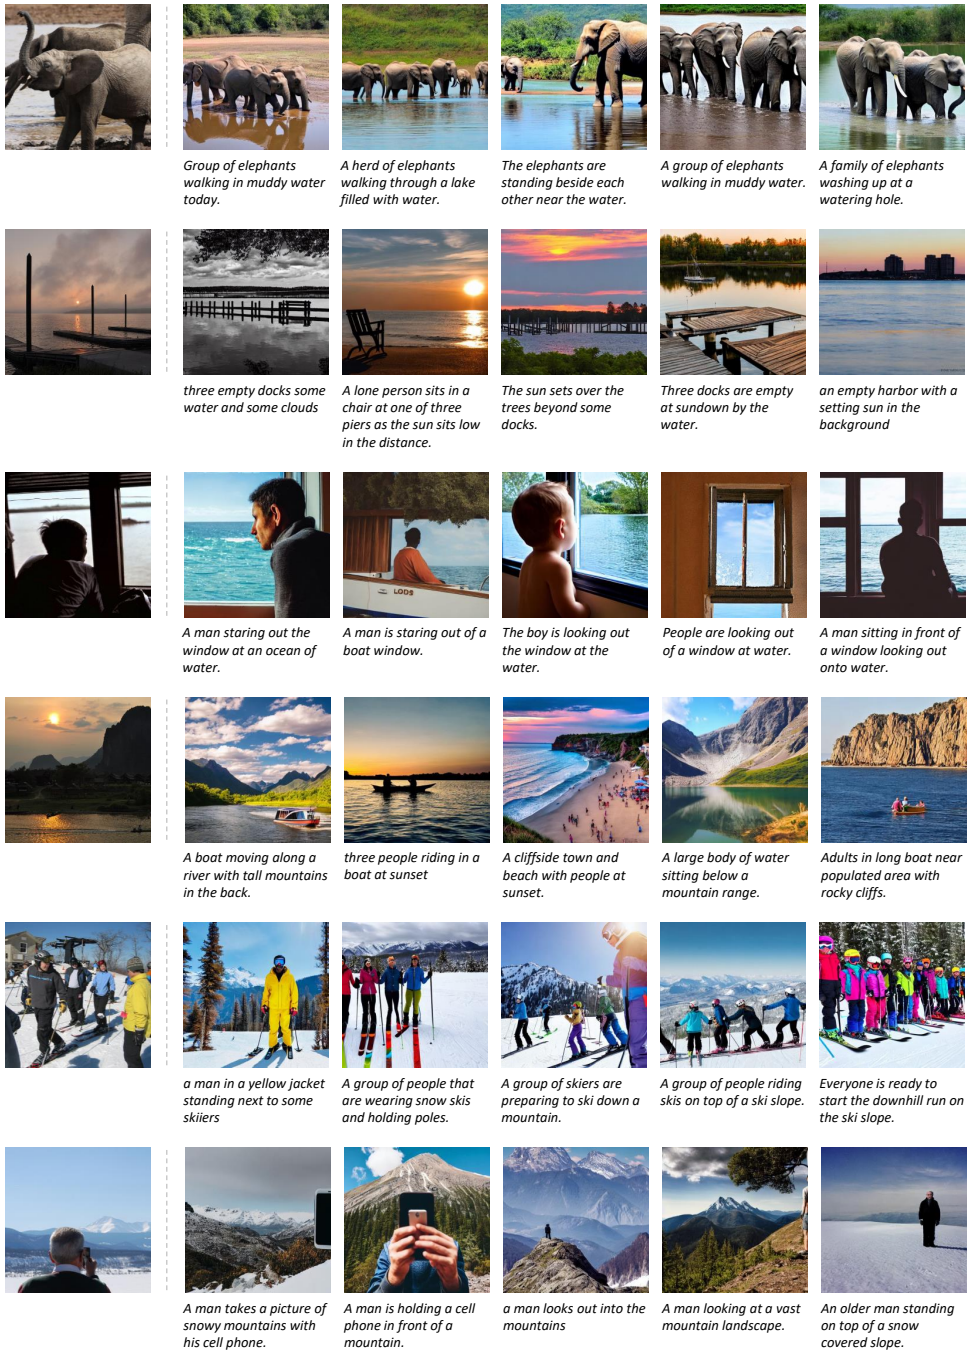

Fig. 3. Sample images from COCOFake. The leftmost column shows the original (real) image, while the remaining ones show fake images generated by Stable Diffusion v1.4 from each of the five COCO captions.

examples from the COCOFake dataset, again showing the original image and the five fake images with the corresponding captions. Failure cases include hallucinating the semantic content of the

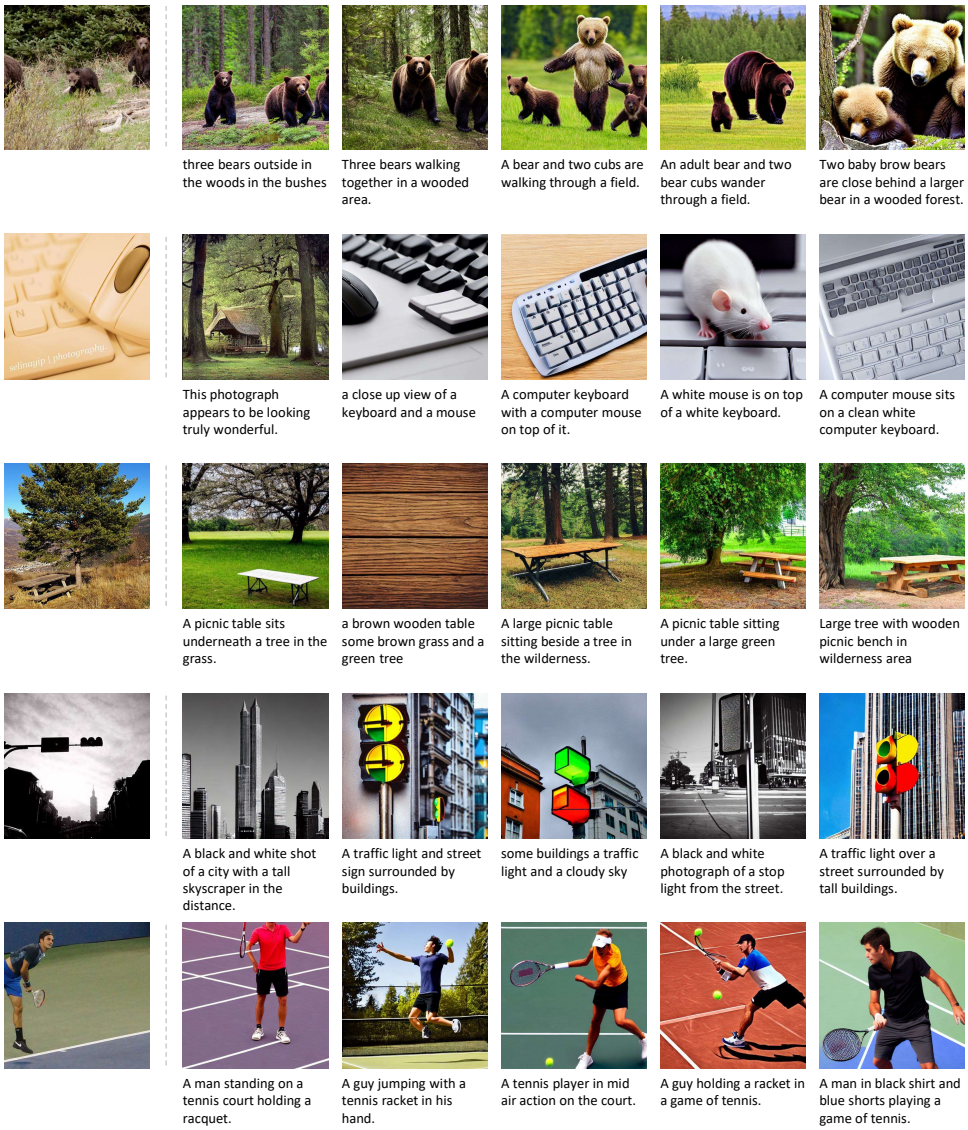

Fig. 4. Less realistic images from COCOFake. The leftmost column shows the original (real) image, while the remaining ones show fake images generated by Stable Diffusion v1.4 from each of the five COCO captions.

caption (first two rows), incorrect understanding of the caption (third row), abstract rendering of objects (traffic lights in the third row), and unrealistic rendering of human poses (last row).

In our experiments, we evaluate deepfake detection performance under a standard setting in which we train the model on images generated by one Stable Diffusion version and test on images generated by the same model. Furthermore, to assess the robustness of our analysis, we also consider the generalization capabilities to images generated by a text-to-image diffusion model different from the one used during training. Under this setting, we compare the performance of our method on images generated by different versions of Stable Diffusion, providing insights into the impact of the generative model on the deepfake detection performance.

4 EXPERIMENTAL EVALUATION

4.1 Implementation Details

Image Encoders. We test two families of backbones: the first are trained for classification on ImageNet [57], while the second are trained on a cross-modal setting on large-scale datasets using contrastive-based loss functions. Due to the nature of the task these networks were trained for, only the latter family provides also text encoders E_T. Specifically, we employ a ResNet [31] model with 48 convolutional layers and a Vision Transformer (ViT) [18] architecture in its B/32 configuration. The ViT encoder takes as input squared patches extracted from the input image and consists of a sequence of multi-head self-attention layers [66]. Both these architectures are trained on the ImageNet dataset [57] that contains around 1.3M images.

As cross-modal architectures, we use two models coming from CLIP [51]. In particular, we employ CLIP RN50 and CLIP ViT-B/32 models, both pre-trained on the OpenAI WebImageText (WIT) dataset, composed of 400 million image-text pairs collected from the web. Moreover, we employ the open source implementation of CLIP (*i.e.*, OpenCLIP [70]), trained with a post-ensemble method for improving robustness to out-of-distribution samples. In our experiments, we consider two versions of the OpenCLIP ViT-B/32 model: one trained on the LAION-400M dataset [62] that contains 400 million CLIP-filtered image-text pairs crawled from the web and the other trained on the larger LAION-2B composed of 2 billion image-text pairs [61].

Linear Probing Details. In our experiments, we also conduct linear probes. In this case, we follow the approach of [51] and employ the features extracted from the backbones to train a logistic regression model with ℓ_2 penalty and LBFGS solver [3, 74]. To balance the training samples, we employ one randomly extracted fake image for each cluster.

Disentanglement Architecture and Training Details. When disentangling semantics and styles, we train the two linear layers S and T, which perform a linear projection to the same dimensionality of the backbone visual features. To train these layers, we employ AdamW [43] as optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We use a batch size of 1,024 and a learning rate of 0.001, training all models for 25 epochs.

4.2 Metrics

To assess the performance of our proposed methodology and evaluate spatial relationships between elements in the embedding spaces, we employ seven different metrics. These aim to quantify the capability to discriminate between real and fake images and to quantify disentanglement.

Min and Max Intra-Cluster Distance Accuracy. These two metrics are employed to evaluate the relative spatial positions of the elements inside a cluster. In particular, for each cluster, we measure the distances between the real image and each of the fake images belonging to the cluster. We then check how many times the real image is the item having the minimum or maximum distance with respect to all the others in the cluster. In other words, for each cluster, the min distance accuracy scores if the real image feature is on average the nearest to all the fake image features, while the max distance accuracy scores if it is the most distant one.

Overall and Full Cluster Accuracy. These two metrics measure the real/fake classification accuracy both over the entire dataset and inside each cluster. The former metric is cluster-independent and is computed using all the elements of a dataset split (*i.e.*, validation, test). The latter, instead, is a cluster-based metric that scores if all elements of a cluster are correctly classified as real or fake, and the metric is then averaged across all clusters.

Overall AUC. As reported in previous deepfake detection literature [10, 46], this metric is used along with accuracy to evaluate how well a deepfake detection model can distinguish between real

Table 1. Minimum and maximum distance accuracy on validation and test sets of COCOFake, using different visual backbones. Results are reported using images generated by both Stable Diffusion v1.4 and v2.0.

		Validation Set (SD v1.4)		Test Set (SD v1.4)		Validation Set (SD v2.0)		Test Set (SD v2.0)	
Backbone	**Dataset**	Min Dist. Accuracy	Max Dist. Accuracy	Min Dist. Accuracy	Max Dist. Accuracy	Min Dist. Accuracy	Max Dist. Accuracy	Min Dist. Accuracy	Max Dist. Accuracy
RN50	ImageNet	8.50	23.58	8.82	24.82	5.98	29.62	6.62	30.16
ViT-B/32	ImageNet	6.84	23.12	6.88	23.88	5.12	29.18	4.92	30.00
CLIP RN50	OpenAI WIT	3.72	38.48	3.60	41.24	2.40	46.72	2.20	48.28
CLIP ViT-B/32	OpenAI WIT	3.30	38.88	3.24	40.10	2.92	42.08	2.98	44.18
OpenCLIP ViT-B/32	LAION-400M	5.28	31.94	5.00	32.02	4.58	34.06	4.62	36.02
OpenCLIP ViT-B/32	LAION-2B	1.40	42.80	1.72	44.00	1.88	42.64	1.78	43.80

and fake images. In our setting, it is computed using all the elements of the validation or test set of our dataset.

Exact Pair and Intra-Cluster Retrieval. These metrics are used to evaluate the goodness of the retrieval task (see Sec. 3.3), in which given a generated image we seek to retrieve its parent caption. The former metric is a recall@k computed considering as ground-truth, for each fake image, the caption used for generating it. The latter, instead, is a recall@k that measures for a given fake image if the retrieved caption matches one of the five captions of the cluster the image belongs to.

4.3 Performance of Visual Features

Unsupervised Classification. We start by assessing the capabilities of existing image features to discriminate between real and generated images, in an unsupervised setting. We employ the min and max distance accuracy metrics defined above and check the presence of spatial relationships between real and generated images inside each cluster.

Results are reported in Table 1 on the test and validation sets of both Stable Diffusion v1.4 and v2.0. We employ six different visual backbones, namely two ResNet-50 pre-trained on ImageNet and OpenAI WIT and four ViT-B/32 pre-trained on ImageNet, OpenAI WIT, LAION-400M, and LAION-2B. As it can be seen, according to the features extracted from the aforementioned backbones, the real image of each cluster tends to be the one with maximum distance with respect to all the other elements. This suggests that these features are discriminative for the task of deepfake classification and that they percolate low-level features that allow for distinction between real and generated items inside of each semantic cluster. Noticeably, the maximum distance accuracy increases when considering backbones trained on multimodal datasets compared to backbones trained on classification, suggesting that image-text matching promotes the percolation of perceptual features.

Comparing the results when using fake images generated by the two considered Stable Diffusion versions, it can be noticed that Stable Diffusion v2.0 exhibits an improvement over v1.4 as evidenced by an increase in the maximum distance metric and a decrease in the minimum distance metric. This suggests that the features extracted from v2.0 are better separable and hence the generated images are more easily detected.

Linear Probing. Following the approach popularized by [51], we train a linear projection through logistic regression on top of the features extracted from the aforementioned backbones. We perform this experiment by training on both Stable Diffusion v1.4 and v2.0 images, and testing either on the validation and test sets containing images generated by the same Stable Diffusion version used during training or on the validation and test sets containing images generated by the Stable Diffusion model not used to train the linear projection.

Table 2. Overall and full cluster accuracy results on the validation and test sets, using linear probing and features of different backbones trained on the COCOFake training set. Results are reported using images generated by both Stable Diffusion v1.4 and v2.0.

		Validation Set (SD v1.4 → SD v1.4)		**Test Set (SD v1.4 → SD v1.4)**		**Validation Set (SD v1.4 → SD v2.0)**		**Test Set (SD v1.4 → SD v2.0)**	
Backbone	**Dataset**	Overall Accuracy	Full Cluster Accuracy	Overall Accuracy	Full Cluster Accuracy	Overall Accuracy	Full Cluster Accuracy	Overall Accuracy	Full Cluster Accuracy
RN50	ImageNet	90.31	57.56	90.62	57.94	81.71	34.94	82.31	35.84
ViT-B/32	ImageNet	87.64	47.62	87.16	47.32	76.71	24.68	77.31	26.92
CLIP RN50	OpenAI WIT	99.07	94.60	99.17	95.30	93.54	69.08	93.74	69.64
CLIP ViT-B/32	OpenAI WIT	99.11	94.84	98.97	94.24	94.41	72.30	94.72	73.62
OpenCLIP ViT-B/32	LAION-400M	97.88	88.18	97.83	87.80	83.30	38.48	84.32	40.74
OpenCLIP ViT-B/32	LAION-2B	99.68	98.01	99.64	97.84	98.88	93.68	98.96	94.08
		Validation Set (SD v2.0 → SD v2.0)		**Test Set (SD v2.0 → SD v2.0)**		**Validation Set (SD v2.0 → SD v1.4)**		**Test Set (SD v2.0 → SD v1.4)**	
Backbone	**Dataset**	Overall Accuracy	Full Cluster Accuracy	Overall Accuracy	Full Cluster Accuracy	Overall Accuracy	Full Cluster Accuracy	Overall Accuracy	Full Cluster Accuracy
RN50	ImageNet	91.07	59.84	91.45	61.44	91.08	60.44	91.33	60.76
ViT-B/32	ImageNet	85.55	42.92	86.12	44.90	84.89	41.50	84.49	39.60
CLIP RN50	OpenAI WIT	98.67	92.56	98.68	92.60	98.57	91.94	98.66	92.48
CLIP ViT-B/32	OpenAI WIT	98.56	92.04	98.48	91.48	98.58	92.02	98.48	91.76
OpenCLIP ViT-B/32	LAION-400M	95.03	74.70	95.57	77.42	97.40	85.62	97.29	84.88
OpenCLIP ViT-B/32	LAION-2B	99.52	97.16	99.59	97.54	99.47	96.80	99.41	96.56

Results are reported in Table 2 in terms of overall accuracy and full cluster accuracy. As it can be seen, all the selected visual features exhibit a significant capability in linearly discriminating real and fake images, on the validation and test sets of the COCOFake dataset when considering both Stable Diffusion v1.4 and v2.0. In continuity with the previous experiment, we observe that contrastive-based visual backbones showcase significantly higher accuracy levels, up to 98.01% and 97.16% of full cluster accuracy respectively on the validation set with Stable Diffusion v1.4 and v2.0 images, and up to 99.68% and 99.52% overall accuracy on the same split. This further confirms the observation that contrastive-based backbones extract and project into their embedding space, low-level and perceptual features that allow discriminating current deepfakes. To assess the robustness of the method, we further test the trained classifiers on the data generated by the Stable Diffusion model not used during training (*i.e.*, Stable Diffusion v2.0 for the linear projection trained on the 1.4 version, and Stable Diffusion v1.4 for the linear projection trained on the 2.0 version). As it can be observed in the right part of Table 2, the trained classifier performs comparably also in this setting with an overall accuracy close to or greater than 99% in all cases. In particular, training on Stable Diffusion v2.0 images generalizes slightly better on images generated by Stable Diffusion v1.4 than the opposite direction with 99.47% and 96.80% of overall and full cluster validation accuracy compared to 98.88% and 93.68% obtained when testing the linear projection trained on Stable Diffusion v1.4 images on the validation set with images generated by the 2.0 version. Overall, these experiments show that the pre-trained visual backbones exhibit high discrimination power when identifying deepfakes.

In light of the high accuracy levels of the aforementioned experiment, in Fig. 5 we report sample misclassified images. It can be noted, in particular, that fake images incorrectly classified as authentic (right side of the figure) depict close-ups and artistic drawings, whose authenticity is visually harder to guarantee.

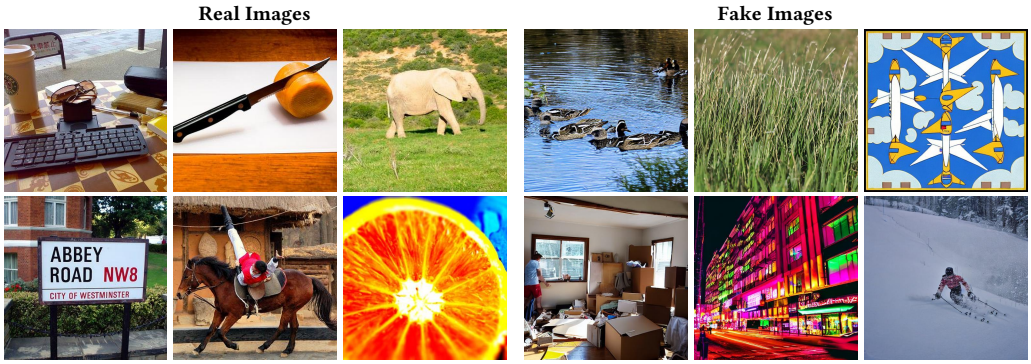

Fig. 5. Sample misclassification errors on both real (left) and fake (right) images, using OpenCLIP ViT-B/32 trained on LAION-2B as the visual encoder.

Table 3. Exact pair and intra-cluster retrieval results. Results are reported using images generated by both Stable Diffusion v1.4 and v2.0.

		Validation Set (SD v1.4)						Test Set (SD v1.4)					
		Exact Pair			Intra-Cluster			Exact Pair			Intra-Cluster		
Backbone	**Dataset**	R@1	R@3	R@5	R@1	R@3	R@5	R@1	R@3	R@5	R@1	R@3	R@5
CLIP RN50	OpenAI WIT	31.33	49.05	56.93	41.91	58.46	66.01	30.98	48.38	56.42	42.09	58.35	65.93
CLIP ViT-B/32	OpenAI WIT	32.12	50.43	58.36	43.34	60.15	67.42	31.96	49.67	57.51	43.24	59.3	66.78
OpenCLIP ViT-B/32	LAION-400M	36.48	55.36	63.28	47.17	63.62	70.73	35.53	54.49	62.56	46.72	62.92	70.22
OpenCLIP ViT-B/32	LAION-2B	40.34	59.44	67.18	50.78	66.64	73.58	39.57	58.78	66.18	50.46	66.34	73.03
		Validation Set (SD v2.0)						**Test Set (SD v2.0)**					
		Exact Pair			Intra-Cluster			Exact Pair			Intra-Cluster		
Backbone	**Dataset**	R@1	R@3	R@5	R@1	R@3	R@5	R@1	R@3	R@5	R@1	R@3	R@5
CLIP RN50	OpenAI WIT	33.05	51.17	59.21	44.73	61.32	69.05	32.53	59.96	58.89	44.67	61.43	68.65
CLIP ViT-B/32	OpenAI WIT	34.70	53.48	61.31	46.73	.63.26	70.49	34.20	52.73	60.94	46.30	62.62	69.99
OpenCLIP ViT-B/32	LAION-400M	42.62	62.31	69.67	53.66	69.71	76.24	42.07	61.74	69.26	53.04	69.06	75.88
OpenCLIP ViT-B/32	LAION-2B	48.67	67.68	74.77	58.39	73.76	80.07	47.83	67.25	74.22	58.24	73.60	79.53

Semantic Preservation. We then conduct the retrieval-based analysis anticipated in Sec. 3.3, in which we look for the original caption used to generate a particular image inside of a multimodal embedding space. The objective of this experiment is to assess whether the semantic information contained in the caption is preserved after the generation and to what extent the generation process alters semantic features.

Results are reported in Table 3, using the exact pair and intra-cluster retrieval metrics and considering validation and test sets containing Stable Diffusion v1.4 and v2.0 images. Surprisingly, retrieving the exact caption used to generate an image is not always easy, and the process is successful only in 40% of the cases when selecting a proper backbone. Even when considering all captions of the same clusters as positives, moreover, we observe a recall@1 of around 50%, again highlighting the difficulty of the task. The results are slightly higher when performing the experiment on the COCOFake version with Stable Diffusion v2.0 images, achieving 48.67% and 58.39% in terms of exact pair and intra-cluster retrieval on the validation set of the dataset. This suggests that the 2.0 version of Stable Diffusion can generate images more semantically aligned with the corresponding captions than the 1.4 version, probably due to the more powerful text

Table 4. AUC and accuracy results on the semantic space S and on the style space T. These results are obtained by training on the COCOFake training set with Stable Diffusion v1.4 images under the disentanglement setting and evaluating on test set of the COCOFake dataset, using data extracted from both Stable Diffusion v1.4 and v2.0.

		Test Set (SD v1.4 → SD v1.4)					
Backbone	**Dataset**	Overall AUC S	Overall Accuracy S	Full Cluster Accuracy S	Overall AUC T	Min Dist. Accuracy T	Max Dist. Accuracy T
RN50	ImageNet	74.93	62.96	8.64	98.45	0.42	89.08
ViT-B/32	ImageNet	68.19	64.04	8.46	96.60	1.30	76.26
CLIP RN50	OpenAI WIT	80.73	74.76	21.40	99.87	0.00	98.46
CLIP ViT-B/32	OpenAI WIT	71.29	67.48	12.90	99.74	0.20	98.14
OpenCLIP ViT-B/32	LAION-400M	70.27	66.84	10.98	99.45	0.10	94.48
OpenCLIP ViT-B/32	LAION-2B	78.00	72.62	17.32	99.93	0.06	99.39
		Test Set (SD v1.4 → SD v2.0)					
Backbone	**Dataset**	Overall AUC S	Overall Accuracy S	Full Cluster Accuracy S	Overall AUC T	Min Dist. Accuracy T	Max Dist. Accuracy T
RN50	ImageNet	74.05	58.53	6.62	98.15	0.52	89.48
ViT-B/32	ImageNet	68.46	63.00	8.86	94.92	1.78	72.84
CLIP RN50	OpenAI WIT	77.58	64.77	12.74	99.71	0.12	96.42
CLIP ViT-B/32	OpenAI WIT	70.98	62.66	10.06	99.30	0.26	94.60
OpenCLIP ViT-B/32	LAION-400M	70.87	68.32	12.02	98.25	0.52	83.98
OpenCLIP ViT-B/32	LAION-2B	76.49	71.92	16.98	99.86	0.04	98.70

encoder used in Stable Diffusion v2.0 (*i.e.*, OpenCLIP ViT-H/14). Nonetheless, these results point out that current generators produce images with partially altered semantic features, and are also in line with the previous observation that contrastive-based extractors percolate low-level features.

4.4 Semantic-Style Disentangling Results

We then turn our attention to evaluating the semantic-style disentanglement approach, in which we aim at training two separate embedding spaces, one storing semantic information and the second focusing on style information. We evaluate the semantic projection in terms of overall AUC and full cluster and overall classification accuracy, and the style projection in terms of overall AUC and minimum and maximum distance accuracy. Specifically, this is done by performing linear probing on top of the two disentangled projections S and T, following the approach described in Sec. 4.3, and computing AUC and overall and full cluster accuracy scores. Instead, minimum and maximum distance accuracy are directly computed on the T projection, to evaluate the relative spatial positions of the elements inside each cluster after disentangling semantics and style.

Results are reported in Table 4 and Table 5 on the COCOFake test set for all the aforementioned backbones, training the semantic-style disentanglement on the training set respectively with Stable Diffusion v1.4 and v2.0 images. In both cases, we observe that, in the T space which focuses on style, real and fake images can be properly distinguished, as the real image is always far apart from the generated ones. On the contrary, this does not happen in the S space, which focuses on semantics, and in which all elements belonging to the same cluster are pulled together, independently of their authenticity. Still, the identification of deepfakes is feasible even in this more challenging space, although with lower AUC and accuracy scores (*i.e.*, with an AUC up to 86% and an accuracy of up to 80%.) As this corresponds to testing a more challenging generator that leaves fewer lower-level traces, we believe this result might offer interesting insights for future works. Similar but slightly

Table 5. AUC and accuracy results on the semantic space S and on the style space T. These results are obtained by training on the COCOFake training set with Stable Diffusion v2.0 images under the disentanglement setting and evaluating on test set of the COCOFake dataset, using data extracted from both Stable Diffusion v1.4 and v2.0.

		Test Set (SD v2.0 → SD v2.0)					
Backbone	**Dataset**	Overall AUC S	Overall Accuracy S	Full Cluster Accuracy S	Overall AUC T	Min Dist. Accuracy T	Max Dist. Accuracy T
RN50	ImageNet	79.30	68.01	13.04	98.43	0.54	89.58
ViT-B/32	ImageNet	69.20	66.31	11.40	95.80	1.94	72.94
CLIP RN50	OpenAI WIT	85.54	80.71	31.92	99.79	0.04	97.92
CLIP ViT-B/32	OpenAI WIT	74.51	68.98	14.20	99.76	0.08	97.60
OpenCLIP ViT-B/32	LAION-400M	72.64	68.51	12.80	99.02	0.38	90.52
OpenCLIP ViT-B/32	LAION-2B	82.69	76.60	23.94	99.87	0.04	99.20
		Test Set (SD v2.0 → SD v1.4)					
Backbone	**Dataset**	Overall AUC S	Overall Accuracy S	Full Cluster Accuracy S	Overall AUC T	Min Dist. Accuracy T	Max Dist. Accuracy T
RN50	ImageNet	76.87	67.91	12.62	97.54	0.60	83.70
ViT-B/32	ImageNet	67.36	65.77	10.16	94.45	2.34	69.00
CLIP RN50	OpenAI WIT	83.00	78.67	27.10	99.76	0.06	97.66
CLIP ViT-B/32	OpenAI WIT	72.48	68.79	45.36	99.73	0.05	97.88
OpenCLIP ViT-B/32	LAION-400M	69.85	65.39	9.60	99.32	0.10	94.14
OpenCLIP ViT-B/32	LAION-2B	82.58	78.76	26.44	99.86	0.08	99.34

lower results can also be observed when testing on images generated by the Stable Diffusion version not used during training, with an overall AUC up to 83% and an overall accuracy up to 79%. When instead considering the overall AUC computed over the T projection, we can notice that the best results are above 99% across almost all settings, thus confirming the proper distinction between real and fake images in the T space.

The structure of the two spaces can be further visualized in Fig. 6, in which we report 2D t-SNE visualizations [65] of the feature space of the OpenCLIP ViT-B/32 LAION-2B backbone, before and after disentanglement and for both Stable Diffusion v1.4 and Stable Diffusion v2.0. In the original embedding space, as provided by the backbone, real and generated samples appear to be mostly overlapped, even if we do not observe a complete overlap – which is in line with the results presented in Table 1 and Table 2. After the disentanglement, instead, the geometry of the T and S spaces appears completely different: the T space clearly separates real and fake data (with the exception of a few outliers), while in the S space we can observe a complete overlap between real and generated samples and a tendency to group into semantic clusters.

A closer visualization of the original feature space and the embedding spaces produced by the two projections is reported in Fig. 7. In this case, we report, on each row, the relative positioning of eight sample clusters from the COCOFake test set with Stable Diffusion v1.4 images. As it can be seen, the two proposed projections are again effective both in separating real and fake images and in promoting the clustering of images sharing similar semantics regardless of their authenticity.

4.5 Robustness Analysis to Image Transformations

As shown in recent literature [9, 21, 44], in addition to evaluating deepfake detection methods in a standard setting, it is also important to assess their robustness to image transformations, which may cause a severe performance degradation in some cases. To this aim, we replicate the experiment

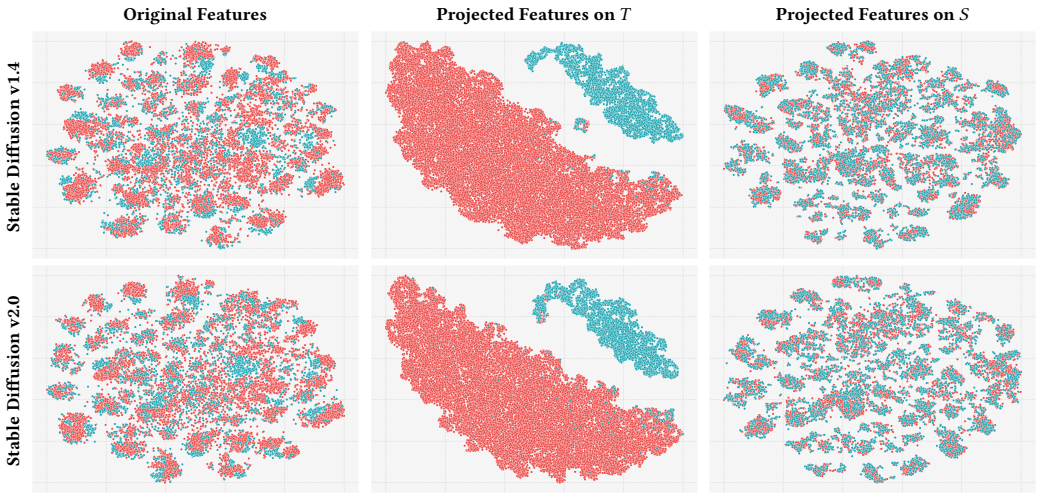

Fig. 6. t-SNE visualizations over the validation set using the original visual features from the OpenCLIP ViT-B/32 LAION-2B backbone (left), the features projected on the T space (style) after disentanglement (middle), and the features projected on the S space (semantics) after disentanglement (right), using Stable Diffusion v1.4 (top) and v2.0 (bottom). Red dots indicate fake images, blue dots indicate real images.

described in Sec. 4.4 by testing on real and fake images that have undergone one of three considered image transformation techniques (*i.e.*, Gaussian blur, JPEG compression, and resize). Specifically, we consider the disentangled spaces trained on non-transformed Stable Diffusion v2.0 images and evaluate on the corresponding test set where one image transformation is applied to all real and fake images, using a kernel size of 3 for Gaussian blurring, an image compression rate of 60 for JPEG compression, and an image edge size equal to 64 pixels for resizing.

Results are shown in Table 6 in terms of the previously described AUC and accuracy evaluation metrics. Notably, while all image transformations cause a slight deterioration in performance, applying JPEG compression or scaling images to a lower resolution leads to the most drastic degradation of the final results, especially considering the results on the T space with an overall AUC of 89% to 97% for JPEG compression and 87% to 94% for resizing. Conversely, Gaussian blur does not significantly impact deepfake detection performance with an overall AUC above 98%.

4.6 Comparison with Other Methods

Finally, we compare our results with existing deepfake detection methods specifically tailored to recognize fake images from GAN-based generators. Specifically, we include in the comparison the models proposed by Wang *et al.* [68] which are based on a ResNet-50 model trained with different image transformations (*i.e.*, Gaussian blur and JPEG compression) and DetectGAN [46] based on an ensemble of different CNNs. For both competitors, we use the pre-trained weights downloaded from the official repositories provided by the authors.

Table 7 reports the results in terms of overall AUC, overall accuracy, and full cluster accuracy on the validation and test sets of COCOFake, using images generated by both versions of Stable Diffusion. Our results are obtained after disentangling semantics and style and by performing linear probing on the style space T which is in charge of distinguishing real and fake images. As it can be seen, both competitors fail to effectively discriminate fake images from real ones with an overall AUC around 40% for the model proposed in [68] and 55% for the DetectGAN approach [46],

Fig. 7. t-SNE visualizations on sampled clusters from the Stable Diffusion v1.4 test set using features extracted from the OpenCLIP ViT-B/32 architecture pre-trained on LAION-2B. We report the original features from the visual backbone (left), the features projected on the T space (style) after disentanglement (middle), and the features projected on the S space (semantics) after disentanglement (right). Dots indicate fake images, triangles indicate real images. Images from the same cluster are shown with the same color.

when tested on Stable Diffusion v1.4 images. On the contrary, all versions of our model achieve AUC scores greater than 99% confirming the effectiveness of the T space in correctly detecting deepfakes.

5 CONCLUSION

This paper proposes a multimodal setting for deepfake detection and analysis, in which real and generated images sharing the same semantics are paired into semantic clusters. In our setting, different semantic projections of a given image, expressed through captions, are employed to generate fake images. Employing the popular Stable Diffusion model as generator, we investigated the performance of contrastive and classification-based visual features, highlighting that diffusion-based deepfakes share common low-level features that make them easily identifiable.

Table 6. AUC and accuracy results on the semantic space S and on the style space T. These results are obtained by training on the COCOFake training set with Stable Diffusion v2.0 images under the disentanglement setting and evaluating on test set of the COCOFake dataset, using different image transformations.

		Gaussian Blur (SD v2.0 → SD v2.0)					
Backbone	**Dataset**	Overall AUC S	Overall Accuracy S	Full Cluster Accuracy S	Overall AUC T	Min Dist. Accuracy T	Max Dist. Accuracy T
CLIP RN50	OpenAI WIT	77.44	74.28	19.42	99.26	0.12	90.96
CLIP ViT-B/32	OpenAI WIT	70.41	59.65	7.72	99.48	0.16	94.70
OpenCLIP ViT-B/32	LAION-400M	71.20	68.75	12.52	98.27	0.56	86.28
OpenCLIP ViT-B/32	LAION-2B	79.31	75.16	21.38	99.80	0.12	98.50
		JPEG Compression (SD v2.0 → SD v2.0)					
Backbone	**Dataset**	Overall AUC S	Overall Accuracy S	Full Cluster Accuracy S	Overall AUC T	Min Dist. Accuracy T	Max Dist. Accuracy T
CLIP RN50	OpenAI WIT	61.64	55.05	5.10	88.60	3.62	57.62
CLIP ViT-B/32	OpenAI WIT	64.77	57.14	6.00	89.38	4.04	54.30
OpenCLIP ViT-B/32	LAION-400M	69.22	62.01	8.26	96.97	0.82	82.56
OpenCLIP ViT-B/32	LAION-2B	69.06	69.75	13.62	93.32	2.28	65.66
		Resize (SD v2.0 → SD v2.0)					
Backbone	**Dataset**	Overall AUC S	Overall Accuracy S	Full Cluster Accuracy S	Overall AUC T	Min Dist. Accuracy T	Max Dist. Accuracy T
CLIP RN50	OpenAI WIT	62.75	70.05	10.50	87.70	3.20	56.82
CLIP ViT-B/32	OpenAI WIT	67.78	31.70	0.36	90.85	2.70	62.28
OpenCLIP ViT-B/32	LAION-400M	71.61	41.34	1.74	94.32	2.16	72.50
OpenCLIP ViT-B/32	LAION-2B	75.12	30.71	0.70	86.72	6.00	40.54

Further, we proposed an approach to disentangle semantic and perceptual information, based on supervised contrastive learning. Under this setting, we investigated the classification of authenticity in a semantic space in which low-level cues left by the generator are removed, thus tackling a more challenging scenario. As a complementary contribution, we also collected and released the COCOFake dataset, containing about 1.2M images generated from COCO using both Stable Diffusion 1.4 and 2.0. We believe that our work can shed further light on the development of deepfake detection strategies, also in consideration of the constant evolution of generator models.

ACKNOWLEDGMENTS

We acknowledge the CINECA award under the ISCRA initiative, for the availability of high-performance computing resources and support. This work has been supported by the Horizon Europe project "European Lighthouse on Safe and Secure AI (ELSA)" (HORIZON-CL4-2021-HUMAN-01-03), co-funded by the European Union, and by the PNRR project "Future Artificial Intelligence Research (FAIR)", co-funded by the Italian Ministry of University and Research.

REFERENCES

[1] Manuele Barraco, Sara Sarto, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2023. With a Little Help from your own Past: Prototypical Memory Networks for Image Captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision.*

[2] Miles Brundage, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, Paul Scharre, Thomas Zeitzoff, Bobby Filar, et al. 2018. The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation. *arXiv preprint arXiv:1802.07228* (2018).

[3] Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. 1995. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16, 5 (1995), 1190–1208.

Table 7. Comparison with existing deepfake detection methods in terms of overall AUC, overall accuracy, and full cluster accuracy. Our results are obtained by performing linear probing on the style space T. Results are reported on the validation and test sets of COCOFake, using images extracted from both Stable Diffusion v1.4 and v2.0.

	Validation Set (SD v1.4)			**Test Set (SD v1.4)**		
Model	Overall AUC	Overall Accuracy	Full Cluster Accuracy	Overall AUC	Overall Accuracy	Full Cluster Accuracy
Wang *et al.* (RN50 Blur+JPEG 0.5) [68]	40.61	83.26	0.34	41.29	83.26	0.48
Mandelli *et al.* (DetectGAN) [46]	54.55	83.12	4.78	54.84	83.09	5.06
Ours (CLIP RN50)	99.85	98.79	93.32	99.87	98.89	93.90
Ours (CLIP ViT-B/32)	99.79	98.63	92.26	99.74	98.47	91.58
Ours (OpenCLIP ViT-B/32 - LAION-400M)	99.44	97.08	84.34	99.45	97.21	85.00
Ours (OpenCLIP ViT-B/32 - LAION-2B)	**99.93**	**99.44**	**96.68**	**99.93**	**99.39**	**96.44**
	Validation Set (SD v2.0)			**Test Set (SD v2.0)**		
Model	Overall AUC	Overall Accuracy	Full Cluster Accuracy	Overall AUC	Overall Accuracy	Full Cluster Accuracy
Wang *et al.* (RN50 Blur+JPEG 0.5) [68]	53.05	83.32	0.40	53.53	83.35	0.48
Mandelli *et al.* (DetectGAN) [46]	64.26	83.55	4.98	64.79	83.55	5.28
Ours (CLIP RN50)	99.79	98.38	91.06	99.82	98.29	90.80
Ours (CLIP ViT-B/32)	99.76	98.29	90.50	99.69	98.14	90.10
Ours (OpenCLIP ViT-B/32 - LAION-400M)	99.02	97.31	85.30	99.14	97.28	84.98
Ours (OpenCLIP ViT-B/32 - LAION-2B)	**99.87**	**99.26**	**95.68**	**99.87**	**99.30**	**96.02**

[4] Davide Caffagni, Manuele Barraco, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2023. SynthCap: Augmenting Transformers with Synthetic Data for Image Captioning. In *Proceedings of the International Conference on Image Analysis and Processing*.

[5] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. 2020. What Makes Fake Images Detectable? Understanding Properties that Generalize. In *Proceedings of the European Conference on Computer Vision*.

[6] Keshigeyan Chandrasegaran, Ngoc-Trung Tran, and Ngai-Man Cheung. 2021. A Closer Look at Fourier Spectrum Discrepancies for CNN-Generated Images Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[7] Harry Cheng, Yangyang Guo, Tianyi Wang, Qi Li, Xiaojun Chang, and Liqiang Nie. 2023. Voice-Face Homogeneity Tells Deepfake. *ACM Transactions on Multimedia Computing, Communications, and Applications* 20, 3 (2023), 1–22.

[8] Robert Chesney and Danielle Citron. 2019. Deepfakes and the new disinformation war: The coming age of post-truth geopolitics. *Foreign Affairs* 98 (2019), 147.

[9] Federico Cocchi, Lorenzo Baraldi, Samuele Poppi, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2023. Unveiling the Impact of Image Transformations on Deepfake Detection: An Experimental Analysis. In *Proceedings of the International Conference on Image Analysis and Processing*.

[10] Riccardo Corvi, Davide Cozzolino, Giada Zingarini, Giovanni Poggi, Koki Nagano, and Luisa Verdoliva. 2023. On The Detection of Synthetic Images Generated by Diffusion Models. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*.

[11] Davide Cozzolino, Diego Gragnaniello, Giovanni Poggi, and Luisa Verdoliva. 2021. Towards Universal GAN Image Detection. In *Proceedings of the International Conference on Visual Communications and Image Processing*.

[12] Davide Cozzolino, Andreas Rössler, Justus Thies, Matthias Nießner, and Luisa Verdoliva. 2021. ID-Reveal: Identity-Aware DeepFake Video Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

[13] Davide Cozzolino, Justus Thies, Andreas Rössler, Christian Riess, Matthias Nießner, and Luisa Verdoliva. 2018. ForensicTransfer: Weakly-supervised Domain Adaptation for Forgery Detection. *arXiv preprint arXiv:1812.02510* (2018).

[14] Rita Cucchiara, Lorenzo Baraldi, Marcella Cornia, and Sara Sarto. 2024. Video Surveillance and Privacy: A Solvable Paradox? *Computer* 57, 3 (2024), 91–100.

[15] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems*.

[16] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, and Jie Tang. 2021. CogView: Mastering Text-to-Image Generation via Transformers. In *Advances in Neural Information Processing Systems*.

[17] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. 2020. The DeepFake Detection Challenge (DFDC) Dataset. *arXiv preprint arXiv:2006.07397* (2020).

[18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Proceedings of the International Conference on Learning Representations*.

[19] Ricard Durall, Margret Keuper, and Janis Keuper. 2020. Watch Your Up-Convolution: CNN Based Generative Deep Neural Networks Are Failing to Reproduce Spectral Distributions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[20] Tarik Dzanic, Karan Shah, and Freddie Witherden. 2020. Fourier Spectrum Discrepancies in Deep Network Generated Images. In *Advances in Neural Information Processing Systems*.

[21] Pierre Fernandez, Alexandre Sablayrolles, Teddy Furon, Hervé Jégou, and Matthijs Douze. 2022. Watermarking Images in Self-Supervised Latent Spaces. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*.

[22] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. 2020. Leveraging Frequency Analysis for Deep Fake Image Recognition. In *Proceedings of the International Conference on Machine Learning*.

[23] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. 2022. Make-A-Scene: Scene-Based Text-to-Image Generation with Human Priors. In *Proceedings of the European Conference on Computer Vision*.

[24] Sharath Girish, Saksham Suri, Sai Saketh Rambhatla, and Abhinav Shrivastava. 2021. Towards Discovery and Attribution of Open-World GAN Generated Images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

[25] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*.

[26] Diego Gragnaniello, Davide Cozzolino, Francesco Marra, Giovanni Poggi, and Luisa Verdoliva. 2021. Are GAN generated images easy to detect? A critical analysis of the state-of-the-art. In *Proceedings of the IEEE International Conference on Multimedia and Expo*.

[27] Zhihao Gu, Yang Chen, Taiping Yao, Shouhong Ding, Jilin Li, Feiyue Huang, and Lizhuang Ma. 2021. Spatiotemporal Inconsistency Learning for DeepFake Video Detection. In *Proceedings of the ACM International Conference on Multimedia*.

[28] Bing Han, Xiaoguang Han, Hua Zhang, Jingzhi Li, and Xiaochun Cao. 2021. Fighting Fake News: Two Stream Network for Deepfake Detection via Learnable SRM. *IEEE Transactions on Biometrics, Behavior, and Identity Science* 3, 3 (2021), 320–331.

[29] Bing Han, Jianshu Li, Wenqi Ren, Man Luo, Jian Liu, and Xiaochun Cao. 2023. SIGMA-DF: Single-Side Guided Meta-Learning for Deepfake Detection. In *Proceedings of the ACM International Conference on Multimedia Retrieval*.

[30] Douglas Harris. 2018. Deepfakes: False Pornography Is Here and the Law Cannot Protect You. *Duke Law & Technology Review* 17 (2018), 99.

[31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[32] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. 2020. DeeperForensics-1.0: A Large-Scale Dataset for Real-World Face Forgery Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[33] Matthew Joslin and Shuang Hao. 2020. Attributing and Detecting Fake Images Generated by Known GANs. In *Proceedings of the IEEE Security and Privacy Workshops*.

[34] Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[35] Tero Karras, Samuli Laine, and Timo Aila. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[36] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[37] Mahyar Khayatkhoei and Ahmed Elgammal. 2022. Spatial Frequency Bias in Convolutional Generative Adversarial Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

[38] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. In *Advances in Neural Information Processing Systems*.

[39] Durk P Kingma and Prafulla Dhariwal. 2018. Glow: Generative Flow with Invertible 1x1 Convolutions. In *Advances in Neural Information Processing Systems*.

[40] Yuezun Li and Siwei Lyu. 2018. Exposing DeepFake Videos By Detecting Face Warping Artifacts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*.

[41] Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, and Siwei Lyu. 2020. [Celeb-DF: A Large-Scale Challenging Dataset for DeepFake Forensics]. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[42] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *Proceedings of the European Conference on Computer Vision*.

[43] Ilya Loshchilov and Frank Hutter. 2018. Decoupled Weight Decay Regularization. In *Proceedings of the International Conference on Learning Representations*.

[44] Yuhang Lu and Touradj Ebrahimi. 2024. Assessment Framework for Deepfake Detection in Real-World Situations. *EURASIP Journal on Image and Video Processing* 2024, 1 (2024), 6.

[45] Yuchen Luo, Yong Zhang, Junchi Yan, and Wei Liu. 2021. Generalizing Face Forgery Detection With High-Frequency Features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[46] Sara Mandelli, Nicolò Bonettini, Paolo Bestagini, and Stefano Tubaro. 2022. Detecting GAN-generated Images by Orthogonal Training of Multiple CNNs. In *Proceedings of the International Conference on Image Processing*.

[47] Francesco Marra, Diego Gragnaniello, Davide Cozzolino, and Luisa Verdoliva. 2018. Detection of GAN-Generated Fake Images over Social Networks. In *Proceedings of the IEEE Conference on Multimedia Information Processing and Retrieval*.

[48] Mehdi Mirza and Simon Osindero. 2014. Conditional Generative Adversarial Nets. *arXiv preprint arXiv:1411.1784* (2014).

[49] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2021. GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models. *arXiv preprint arXiv:2112.10741* (2021).

[50] Samuele Poppi, Tobia Poppi, Federico Cocchi, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2024. Safe-CLIP: Removing NSFW Concepts from Vision-and-Language Models. *arXiv preprint arXiv:2311.16254* (2024).

[51] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*.

[52] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. *arXiv preprint arXiv:2204.06125* (2022).

[53] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-Shot Text-to-Image Generation. In *Proceedings of the International Conference on Machine Learning*.

[54] Jonas Ricker, Simon Damm, Thorsten Holz, and Asja Fischer. 2024. Towards the Detection of Diffusion Model Deepfakes. In *Proceedings of the International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*.

[55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[56] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. 2019. FaceForensics++: Learning to Detect Manipulated Facial Images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

[57] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* (2015).

[58] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *arXiv preprint arXiv:2205.11487* (2022).

[59] Sara Sarto, Manuele Barraco, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2023. Positive-Augmented Contrastive Learning for Image and Video Captioning Evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[60] Patrick Schramowski, Manuel Brack, Björn Deiseroth, and Kristian Kersting. 2023. Safe Latent Diffusion: Mitigating Inappropriate Degeneration in Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[61] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson,

Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. LAION-5B: An open large-scale dataset for training next generation image-text models. In *Advances in Neural Information Processing Systems*.

[62] Christoph Schuhmann, Robert Kaczmarczyk, Aran Komatsuzaki, Aarush Katta, Richard Vencu, Romain Beaumont, Jenia Jitsev, Theo Coombes, and Clayton Mullis. 2021. LAION-400M: Open Dataset of CLIP-Filtered 400 Million Image-Text Pairs. In *Advances in Neural Information Processing Systems Workshops*.

[63] Zeyang Sha, Zheng Li, Ning Yu, and Yang Zhang. 2023. DE-FAKE: Detection and Attribution of Fake Images Generated by Text-to-Image Diffusion Models. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*.

[64] Arash Vahdat and Jan Kautz. 2020. NVAE: A deep hierarchical variational autoencoder. In *Advances in Neural Information Processing Systems*.

[65] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9, 11 (2008), 2579–2605.

[66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*.

[67] Luisa Verdoliva. 2020. Media forensics and deepfakes: an overview. *IEEE Journal of Selected Topics in Signal Processing* 14, 5 (2020), 910–932.

[68] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. 2020. CNN-Generated Images Are Surprisingly Easy to Spot... for Now. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[69] Moritz Wolter, Felix Blanke, Raoul Heese, and Jochen Garcke. 2022. Wavelet-packets for deepfake image analysis and detection. *Machine Learning* (2022), 1–33.

[70] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. 2022. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

[71] Wenyuan Yang, Xiaoyu Zhou, Zhikai Chen, Bofei Guo, Zhongjie Ba, Zhihua Xia, Xiaochun Cao, and Kui Ren. 2023. AVoiD-DF: Audio-Visual Joint Learning for Detecting Deepfake. *IEEE Transactions on Information Forensics and Security* 18 (2023), 2015–2029.

[72] Ning Yu, Larry S Davis, and Mario Fritz. 2019. Attributing Fake Images to GANs: Learning and Analyzing GAN Fingerprints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

[73] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. 2019. Detecting and Simulating Artifacts in GAN Fake Images. In *Proceedings of the International Workshop on Information Forensics and Security*.

[74] Ciyou Zhu, Richard H Byrd, Peihuang Lu, and Jorge Nocedal. 1997. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Software* 23, 4 (1997), 550–560.

[75] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.