

Leveraging Representations from Intermediate Encoder-blocks for Synthetic Image Detection

Christos Koutlis[✉] and Symeon Papadopoulos[✉]

Information Technologies Institute @ CERTH, Thessaloniki, Greece

Abstract. The recently developed and publicly available synthetic image generation methods and services make it possible to create extremely realistic imagery on demand, raising great risks for the integrity and safety of online information. State-of-the-art Synthetic Image Detection (SID) research has led to strong evidence on the advantages of feature extraction from foundation models. However, such extracted features mostly encapsulate high-level visual semantics instead of fine-grained details, which are more important for the SID task. On the contrary, shallow layers encode low-level visual information. In this work, we leverage the image representations extracted by intermediate Transformer blocks of CLIP’s image-encoder via a lightweight network that maps them to a learnable forgery-aware vector space capable of generalizing exceptionally well. We also employ a trainable module to incorporate the importance of each Transformer block to the final prediction. Our method is compared against the state-of-the-art by evaluating it on 20 test datasets and exhibits an average +10.6% absolute performance improvement. Notably, the best performing models require just a single epoch for training (~8 minutes). Code available at <https://github.com/mever-team/rine>.

Keywords: synthetic image detection · AI-generated image detection · intermediate representations

1 Introduction

The advancements in the field of synthetic content generation disrupt digital media communications, and pose new societal and economic risks [28]. Generative Adversarial Networks (GAN) [21], the variety of successor generative models [24] and the latest breed of Diffusion models [14] are capable of producing highly realistic images that often deceive humans [34], and lead to grave risks ranging from fake pornography to hoaxes, identity theft, and financial fraud [50]. Therefore, it becomes extremely challenging to develop reliable synthetic content detection methods that can keep up with the latest generative models.

The scientific community has recently put a lot of effort on the development of automatic solutions to counter this problem [38, 50]. More precisely, the discrimination of GAN-generated from real images using deep learning models has gained a lot of interest, although it has been found that such approaches struggle to generalize to other types of generative models [13]. Synthetic Image Detection




Fig. 1: Feature space visualization for unseen data with t-SNE.

(SID) is typically performed based either on image-level [6, 52] or frequency-level features [19, 25, 43]. Also, there are methods [11, 12] that analyse the traces left by generative models to provide useful insights on how to address the SID task.

Recently, the utilization of features extracted by foundation models, such as CLIP [44], has been found to yield surprisingly high performance in the SID task, with minimal training requirements [41]. However, these features are extracted from the last model layer, known to capture high-level semantics. Instead, low-level image features deriving from intermediate layers, which are known to be highly relevant for synthetic image detection [4, 12], have not been explored to date, despite their potential to further boost performance.

To address the SID task more effectively, we propose the RINE model by leveraging Representations from Intermediate Encoder-blocks of CLIP. Specifically, we collect the image representations provided by intermediate Transformer blocks that carry low-level visual information and project them with learnable linear mappings to a forgery-aware vector space. Additionally, a Trainable Importance Estimator (TIE) module is used to incorporate the impact of each intermediate Transformer block in the final prediction. We train our models only on images generated by ProGAN (our experiments consider 1, 2, and 4 object classes for training) and evaluate it on 20 test datasets including images generated from GAN, Diffusion as well as other (deepfake, low-level vision, perceptual loss, DALL-E) generative models, surpassing by an average +10.6% the state-of-the-art performance. Fig. 1 provides a comparison between the feature space of RINE and that of state-of-the-art methods, for GAN, Diffusion and real images, showcasing the discriminating ability of the proposed representation. Also, it is noteworthy that (i) we reach this level of performance with only 6.3M learnable parameters, (ii) trained only for 1 epoch, which requires ~ 8 minutes, and (iii) we surpass the state-of-the-art by +9.4%, when training only on a single ProGAN object class, while the second-to-best model uses 20.

2 Related work

Synthetic Image Detection (SID) has recently emerged as a special but increasingly interesting field of DeepFake¹ or synthetic media detection [38, 50]. Early studies investigated fingerprints left by GAN generators, similarly to the fingerprints left by cameras [35], letting them not only to distinguish between real

¹ Often used as an umbrella term for all kinds of synthetic or manipulated content

and synthetic images but also to attribute synthetic images to specific GAN generators [37, 54]. Other studies focused on GAN-based image-to-image translation detection [36] and face manipulation detection [47], but still the considered evaluation sets contained images generated from the same generators as the ones used for training, neglecting to assess generalizability. Such approaches have been shown to struggle with detecting images coming from unseen generators [13].

Subsequent works focused on alleviating this issue through patch-based predictions [6], auto-encoder based architectures [13] and co-occurrence matrices computed on the RGB image channels [39]. Next, Wang et al. [52] proposed a simple yet effective approach, based on appropriate pre-processing and intense data augmentation, which is capable of generalizing well to many GAN generators by simply training on ProGAN images [29]. Additionally, traces of GAN generators in the frequency domain were exploited for effectively addressing the SID task [19, 55], while the use of frequency-level perturbation maps was proposed in [26] to force the detector to ignore domain-specific frequency-level artifacts. Image gradients estimated by a pre-trained CNN model have been shown to produce artifacts that can generalize well to unseen GAN generated images [49]. Removing the down-sampling operation from the first ResNet’s layer, that potentially eliminates synthetic traces, together with intense augmentation has shown promising results [12]. Another approach to improve generalization has been to select training images based on their perceptual quality; this was shown to be effective especially in cross-domain generalization settings [17].

Representations extracted by foundation models are surprisingly effective on SID, generalizing equally well on GAN and Diffusion model generated data [41]. Motivated by this success, we hypothesize that further performance gains are possible by leveraging representations from intermediate layers, which carry low-level visual information, in addition to representations from the final layer that primarily carry high-level semantic information. Our approach is different from deep layer aggregation that was explored in several computer vision tasks [53], as well as for deepfake detection [27], as we process the frozen layers’ outputs of a foundation model, while deep layer aggregation is a training technique that concatenates layer outputs to produce predictions during the training of a CNN.

Finally, the performance drop of synthetic image detectors in real-world conditions typically performed when uploading content to social media (e.g., cropping and compression) has been analysed in [12, 36]. Our method is also effective in such conditions as demonstrated by the experimental results in Sec. 5.6.

3 Methodology

3.1 Representations from Intermediate Encoder-blocks

Fig. 2 illustrates the RINE architecture. Let us consider a batch of b input images $\mathbf{X} \in \mathbb{R}^{b \times 3 \times w \times h}$, with w width, and h height, and targets $\mathbf{Y} = \{y_i\}_{i=1}^b \in \mathbb{R}^b$, with $y_i \in \{0, 1\}$. The images are reshaped into a sequence of p flattened patches $\mathbf{X}_p \in \mathbb{R}^{b \times p \times (P^2 \cdot 3)}$, where P denotes patch side length and $p = w \cdot h / P^2$. Then, \mathbf{X}_p




Fig. 2: The RINE architecture. A batch of b images is processed by CLIP’s image encoder. The concatenation of the n d -dimensional CLS tokens (one from each Transformer block) is first projected and then multiplied with the blocks’ scores, estimated by the Trainable Importance Estimator (TIE) module. Summation across the second dimension results in one feature vector per image. Finally, after the second projection and the consequent classification head modules, two loss functions are computed. Binary cross-entropy \mathcal{L}_{CE} directly optimizes SID, while the contrastive loss $\mathcal{L}_{Cont.}$ assists the training by forming a dense feature vector cluster per class.

is linearly projected to d dimensions, the learnable d -dimensional CLS token is concatenated to the projected sequence, and positional embeddings are added to all $p+1$ tokens in order to construct the image encoder’s input, $\mathbf{Z}_0 \in \mathbb{R}^{b \times (p+1) \times d}$. CLIP’s image encoder [44], being a Vision Transformer (ViT) [18], consists of n successive Transformer encoder blocks [51], each processing the output of the previous block as shown in Eq. (1):

$$\begin{aligned}\tilde{\mathbf{Z}}_l &= \text{MSA}(\text{LN}(\mathbf{Z}_{l-1})) + \mathbf{Z}_{l-1} \\ \mathbf{Z}_l &= \text{MLP}(\text{LN}(\tilde{\mathbf{Z}}_l)) + \tilde{\mathbf{Z}}_l\end{aligned}\quad (1)$$

where $l=1, \dots, n$ denotes the block’s index, $\mathbf{Z}_l \in \mathbb{R}^{b \times (p+1) \times d}$ denotes the l th block’s output, MSA denotes Multi-head Softmax Attention [51], LN denotes Layer Normalization [2], and MLP denotes a Multi-Layer Perceptron with two GELU [23] activated linear layers of $4 \cdot d$ and d number of units, respectively.

Then, we define the Representations from Intermediate Encoder-blocks (RINE) \mathbf{K} , as the concatenation of CLS tokens from the corresponding n blocks:

$$\mathbf{K} = \oplus \{ \mathbf{Z}_l^{[0]} \}_{l=1}^n \in \mathbb{R}^{b \times n \times d} \quad (2)$$

where \oplus denotes concatenation, and $\mathbf{Z}_l^{[0]} \in \mathbb{R}^{b \times 1 \times d}$ denotes the CLS token from the output of block l . We keep CLIP, which produces these features, frozen during training and use \mathbf{K} to construct discriminative features for the SID task.

3.2 Trainable modules

The representations \mathbf{K} are processed by a projection network (\mathcal{Q}_1 in Fig. 2):

$$\mathbf{K}_m = \text{ReLU}(\mathbf{K}_{m-1} \mathbf{W}_m + \mathbf{b}_m) \in \mathbb{R}^{b \times n \times d'} \quad (3)$$

where $m=1,\dots,q$ denotes the index of the network’s layer (with $\mathbf{K}_0 = \mathbf{K}$), $\mathbf{W}_m \in \mathbb{R}^{d' \times d'}$ (except $\mathbf{W}_1 \in \mathbb{R}^{d \times d'}$) and $\mathbf{b}_m \in \mathbb{R}^{d'}$ define the linear mapping, and ReLU denotes the Rectified Linear Unit [1] activation function. After each layer, dropout [48] with rate 0.5 is applied.

Additionally, each of the d' learned features can be more relevant for the SID task at various processing stages, thus we employ a Trainable Importance Estimator (TIE) module to adjust their impact to the final decision. More precisely, we consider a randomly-initialized learnable variable $\mathbf{A} = \{\alpha_{lk}\} \in \mathbb{R}^{n \times d'}$ the elements of which will estimate the importance of feature k at processing stage (i.e., Transformer block) l . This is used to construct one feature vector per image i , as a weighted average of feature k across processing stages l :

$$\tilde{\mathbf{K}}^{(ik)} = \sum_l^n \mathcal{S}(\mathbf{A})^{(lk)} \cdot \mathbf{K}_q^{(ilk)} \quad (4)$$

where (\cdot) denotes tensor indices, $\mathbf{K}_q = \{\mathbf{K}_q^{(ilk)}\} \in \mathbb{R}^{b \times n \times d'}$ is the output of projection \mathcal{Q}_1 , and \mathcal{S} denotes the Softmax activation function acting across the first dimension of \mathbf{A} .

Finally, a second projection network (\mathcal{Q}_2 in Fig. 2) with the same architecture as the first one (cf. Eq. (3)) takes as input $\tilde{\mathbf{K}} \in \mathbb{R}^{b \times d'}$ and outputs $\tilde{\mathbf{K}}_q \in \mathbb{R}^{b \times d'}$, which is consequently processed by the classification head that predicts the final output \hat{y}_i (probability to be fake). The classification head consists of two $d' \times d'$ ReLU-activated dense layers and one $d' \times 1$ dense layer which produces the logits.

3.3 Objective function

We consider the combination of two objective functions to optimize the parameters of the proposed model. The first is the binary cross-entropy loss [20], which measures the classification error and directly optimizes the SID objective:

$$\mathcal{L}_{CE} = - \sum_{i=1}^b y_i \log \hat{y}_i + (1 - y_i) \cdot \log(1 - \hat{y}_i)$$

The second is the Supervised Contrastive Learning loss (SupContrast) [32], which is considered in order to assist the training process by bringing closer the feature vectors inside $\tilde{\mathbf{K}}_q$ that share targets and move apart the rest, defined as:

$$\mathcal{L}_{Cont.} = - \sum_{i=1}^b \frac{1}{G(i)} \sum_{g \in G(i)} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_g / \tau)}{\sum_{a \in A(i)} \exp(\mathbf{z}_i \cdot \mathbf{z}_a / \tau)}$$

where $A(i) = \{1, \dots, i-1, i+1, \dots, b\}$, $G(i) = \{g \in A(i) : y_g = y_i\}$, $\mathbf{z}_j = \{\tilde{\mathbf{K}}_q^{(jk)}\}_{k=1}^{d'} \in \mathbb{R}^{d'}$, the \cdot denotes dot product, and τ is a temperature parameter. We combine the two loss functions by a tunable factor ξ , as shown in Eq. (5):

$$\mathcal{L} = \mathcal{L}_{CE} + \xi \cdot \mathcal{L}_{Cont.} \quad (5)$$

4 Experimental setup

4.1 Datasets

Following the training protocol of related work [41, 49, 52], we use for training ProGAN [29] generated images and the corresponding real images of the provided dataset. Similarly with previous works [26, 49], we consider three settings with ProGAN-generated training data from 1 (horse), 2 (chair, horse) and 4 (car, cat, chair, horse) object classes. For testing, this is the first study to consider 20 datasets, with generated and real images, combining the evaluation datasets used in [41, 49, 52]. Specifically, the evaluation sets are from ProGAN [29], StyleGAN [30], StyleGAN2 [31], BigGAN [5], CycleGAN [56], StarGAN [9], GauGAN [42], DeepFake [47], SITD [7], SAN [15], CRN [8], IMLE [33], Guided [16], LDM [46] (3 variants), Glide [40] (3 variants), and DALL-E [45].

4.2 Implementation details

The training of RINE is conducted with batch size 128 and learning rate 1e-3 for only 1 epoch using the Adam optimizer. To assess the existence of potential benefit from further training we also conduct a set of experiments with 3, 5, 10, and 15 epochs (cf. Sec. 5.4). For each of the 3 training settings (1-class, 2-class, 4-class), we consider a hyperparameter grid to obtain the best performance, namely $\xi \in \{0.1, 0.2, 0.4, 0.8\}$, $q \in \{1, 2, 4\}$, and $d' \in \{128, 256, 512, 1024\}$. Additionally, we consider two CLIP variants, namely ViT-B/32 and L/14, for the extraction of representations. The models' hyperparameters that result in the best performance (presented in Sec. 5) are illustrated in Tab. 1, along with the corresponding number of trainable parameters, and required training duration. For the training images, following the best practice in related work, we apply Gaussian blurring and JPEG compression with probability 0.5, then random cropping to 224×224 , and finally random horizontal flip with probability 0.5. The validation and test images are only center-cropped at 224×224 . Resizing is omitted both during training and testing as it is known to eliminate synthetic traces [12]. All experiments are conducted using one NVIDIA GeForce RTX 3090 Ti GPU.

# classes	backbone	ξ	q	d	d'	# params.	training time
1	ViT-L/14	0.1	4	1024	1024	10.52M	~2 min.
2	ViT-L/14	0.2	4	1024	128	0.28M	~4 min.
4	ViT-L/14	0.2	2	1024	1024	6.32M	~8 min.

Table 1: Hyperparameter configuration of the best 1-, 2-, and 4-class models. The number of learnable parameters per model (in millions), and the training time (in minutes) is also reported.

4.3 State-of-the-art methods

The state-of-the-art SID methods we consider in our comparison include the following:

1. Wang [52]: A standard ResNet-50 [22] architecture pre-trained on ImageNet is fine-tuned on SID with appropriate selection of pre- and post-processing, as well as data augmentations.
2. Patch-Forensics [6]: A truncated (at block 2) Xception [10] pre-trained on ImageNet is fine-tuned on SID for each image patch, finally aggregating the individual patch-level predictions.
3. FrePGAN [26]: Frequency-level perturbation maps are generated by an adversarially trained network making fake images hard to distinguish from real ones, then a standard ResNet-50 [22] is trained on the SID task.
4. LGrad [49]: Image gradients, computed using a pre-trained deep network, are processed by a standard ResNet-50 [22] pre-trained on ImageNet, which is fine-tuned on the SID task.
5. DMID [12]: A ResNet-50 [22] without down-sampling in the first layer is trained on the SID task using intense augmentation.
6. Universal Fake Detector (UFD) [41]: Simple linear probing on top of (last layer) features extracted from CLIP’s ViT-L/14 [44] image encoder.

For all methods except FrePGAN, we compute performance metrics on the evaluation datasets presented in Sec. 4.1, using the publicly available checkpoints provided by their official repositories. For FrePGAN, that does not provide publicly available code and models, we present the scores that are reported in their paper [26].

4.4 Evaluation protocol

We evaluate the proposed architecture with accuracy (ACC) and average precision (AP) metrics on each test dataset, following previous works for comparability purposes. For the calculation of accuracy no calibration is conducted, we consider 0.5 as threshold to all methods. Best models are identified by the maximum sum ACC+AP. We also report the average (AVG) metric values across the test datasets to obtain summary evaluations.

5 Results

5.1 Comparative analysis

In Tab. 2 and Tab. 3, we present the performance scores (ACC & AP respectively) of our method versus the competing ones. Our 1-class model outperforms all state-of-the-art methods irrespective of training class number. On average, we surpass the state-of-the-art by +9.4% ACC & +4.3% AP with the 1-class model, by +6.8% ACC & +4.4% AP with the 2-class model, and by +10.6%

method	# cl.	Generative Adversarial Networks							Low level vision				Perceptual loss				Latent Diffusion			Glide			AVG
		Pro-GAN	Style-GAN	Style-GAN2	Big-GAN	Cycle-GAN	Star-GAN	Gau-GAN	Deep-fake	SITD	SAN	CRN	IMLE	Guided steps	200 CFG steps	100 CFG steps	100 CFG steps	27	27	10	DALL-E		
Wang [52] (prob. 0.5)	20	100.0	66.8	64.4	59.0	80.7	80.9	79.2	51.3	55.8	50.0	85.6	92.3	52.1	51.1	51.4	51.3	53.3	55.6	54.2	52.5	64.4	
Wang [52] (prob. 0.1)	20	100.0	84.3	82.8	70.2	85.2	91.7	78.9	53.0	63.1	50.0	90.4	90.3	60.4	53.8	55.2	55.1	60.3	62.7	61.0	56.0	70.2	
Patch-Forensics [6]	†	66.2	58.8	52.7	52.1	50.2	96.9	50.1	58.0	54.4	50.0	52.9	52.3	50.5	51.9	53.8	52.0	51.8	52.1	51.4	57.2	55.8	
FrePGAN [26]	1	95.5	80.6	77.4	63.5	59.4	99.6	53.0	70.4	-*	-	-	-	-	-	-	-	-	-	-	-		
FrePGAN [26]	2	99.0	80.8	72.2	66.0	69.1	98.5	53.1	62.2	-	-	-	-	-	-	-	-	-	-	-	-		
FrePGAN [26]	4	99.0	80.7	84.1	69.2	71.1	99.9	60.3	70.9	-	-	-	-	-	-	-	-	-	-	-	-		
LGrad [49]	1	99.4	96.1	94.0	79.6	84.6	99.5	71.1	63.4	50.0	44.5	52.0	52.0	67.4	90.5	93.2	90.6	80.2	85.2	83.5	89.5	78.3	
LGrad [49]	2	99.8	94.5	92.1	82.5	85.5	99.8	73.7	61.5	46.9	45.7	52.0	52.1	72.1	91.1	93.0	91.2	87.1	90.5	89.4	88.7	79.4	
LGrad [49]	4	99.9	94.8	96.1	83.0	85.1	99.6	72.5	56.4	47.8	41.1	50.6	50.7	74.2	94.2	95.9	95.0	<u>87.2</u>	<u>90.8</u>	<u>89.8</u>	88.4	79.7	
DMID [12]	20	100.0	99.4	92.9	96.9	92.0	99.5	94.8	54.1	<u>90.6</u> **	55.5	100.0	100.0	53.9	58.0	61.1	57.5	56.9	59.6	58.8	71.7	77.6	
UFD [41]	20	99.8	79.9	70.9	95.1	98.3	95.7	99.5	71.7	71.4	51.4	57.5	70.0	70.2	94.4	74.0	95.0	78.5	79.0	77.9	87.3	80.9	
RINE (Ours)	1	99.8	88.7	86.9	<u>99.1</u>	99.4	98.8	99.7	82.7	84.7	72.4	93.4	<u>96.9</u>	77.9	<u>96.9</u>	83.5	97.0	83.8	87.4	85.4	91.9	<u>90.3</u>	
	2	99.8	84.9	76.7	98.3	99.4	99.6	99.9	66.7	91.9	67.8	83.5	96.8	69.6	96.8	80.0	<u>92.3</u>	83.6	86.0	84.1	<u>92.3</u>	87.7	
	4	100.0	88.9	<u>94.5</u>	99.6	<u>99.3</u>	99.5	<u>99.8</u>	<u>80.6</u>	<u>90.6</u>	<u>68.3</u>	89.2	90.6	<u>76.1</u>	98.3	88.2	98.6	88.9	92.6	90.7	95.0	91.5	

* Hyphens denote scores that are neither reported in the corresponding paper nor the code and models are available in order to compute them.

** We applied cropping at 2000x1000 on SITD [7] for DMID [12] due to GPU memory limitations.

† Patch-Forensics has been trained on ProGAN data but not on the same dataset as the rest models. For more details please refer to [6].

Table 2: Accuracy (ACC) scores of baselines and our model across 20 test datasets. The second column (# cl.) presents the number of used training classes. Best performance is denoted with **bold** and second to best with underline. Our method yields +10.6% average accuracy compared to the state-of-the-art.

ACC & +4.5% AP with the 4-class model. In terms of ACC, we obtain the best score in 14 out of 20 test datasets, and simultaneously the first and second best performance in 10 of them. In terms of AP, we obtain the best score in 15 out of 20 test datasets, and simultaneously the first and second best performance in 14 of them. The biggest performance gain is on the SAN dataset [15] (+16.9% ACC). Considering more training classes does not reduce generalization (cf. supplementary material).

method	# cl.	Generative Adversarial Networks							Low level vision				Perceptual loss				Latent Diffusion			Glide			AVG
		Pro-GAN	Style-GAN	Style-GAN2	Big-GAN	Cycle-GAN	Star-GAN	Gau-GAN	Deep-fake	SITD	SAN	CRN	IMLE	Guided steps	200 CFG steps	100 CFG steps	100 CFG steps	27	27	10	DALL-E		
Wang [52] (prob. 0.5)	20	100.0	98.0	97.8	88.2	96.8	95.4	98.1	64.8	82.2	56.0	99.4	99.7	69.9	65.9	66.7	66.0	72.0	76.5	73.2	66.3	81.7	
Wang [52] (prob. 0.1)	20	100.0	99.5	99.0	84.5	93.5	98.2	89.5	87.0	68.1	53.0	99.5	99.5	73.2	71.2	73.0	72.5	80.5	84.6	82.1	71.3	84.0	
Patch-Forensics [6]	†	94.6	79.3	77.6	83.3	74.7	99.5	83.2	71.3	91.6	39.7	<u>99.9</u>	98.9	58.7	68.9	73.7	68.7	50.6	52.8	48.4	66.9	74.1	
FrePGAN [26]	1	99.4	90.6	93.0	60.5	59.9	100.0	49.1	81.5	-*	-	-	-	-	-	-	-	-	-	-	-		
FrePGAN [26]	2	<u>99.2</u>	92.0	94.0	61.8	70.3	100.0	51.0	80.6	-	-	-	-	-	-	-	-	-	-	-	-		
FrePGAN [26]	4	<u>99.9</u>	89.6	98.6	71.1	74.4	100.0	71.7	91.9	-	-	-	-	-	-	-	-	-	-	-	-		
LGrad [49]	1	<u>99.9</u>	99.6	99.5	88.9	94.4	100.0	82.0	79.7	44.1	45.7	82.3	82.5	71.1	97.3	98.0	97.2	90.1	93.6	92.0	96.9	86.7	
LGrad [49]	2	100.0	99.6	99.6	92.6	94.7	<u>99.9</u>	83.2	71.6	42.4	45.3	66.1	80.9	75.6	97.2	98.1	97.2	94.6	96.5	95.8	96.5	86.4	
LGrad [49]	4	100.0	<u>99.8</u>	<u>99.9</u>	90.8	94.0	100.0	79.5	72.4	39.4	42.2	63.9	69.7	79.5	99.1	99.1	99.2	93.3	95.2	95.0	97.3	85.5	
DMID [12]	20	100.0	100.0	<u>99.8</u>	98.6	100.0	<u>99.8</u>	94.7	99.8 **	87.7	100.0	100.0	73.0	86.8	89.4	87.3	86.5	89.9	89.0	96.1	93.9		
UFD [41]	20	100.0	97.3	97.5	99.3	<u>99.8</u>	99.4	100.0	84.4	89.9	62.6	94.5	98.3	89.5	99.3	92.5	<u>99.3</u>	95.3	95.6	95.0	97.5	94.3	
RINE (Ours)	1	100.0	99.1	99.7	99.9	100.0	100.0	100.0	<u>97.4</u>	95.8	91.9	98.5	<u>99.9</u>	<u>95.7</u>	<u>99.8</u>	98.0	99.9	98.9	99.3	99.1	<u>99.3</u>	98.6	
	2	100.0	99.5	99.6	99.9	100.0	100.0	100.0	96.4	<u>97.5</u>	<u>93.1</u>	98.2	99.8	<u>95.7</u>	99.9	98.0	99.9	98.9	99.0	98.8	99.6	98.7	
	4	100.0	99.4	100.0	99.9	100.0	100.0	100.0	97.9	<u>97.2</u>	94.9	97.3	99.7	96.4	<u>99.8</u>	<u>98.3</u>	99.9	<u>98.8</u>	<u>99.3</u>	<u>98.9</u>	<u>99.3</u>	98.8	

* Hyphens denote scores that are neither reported in the corresponding paper nor the code and models are available in order to compute them.

** We applied cropping at 2000x1000 on SITD [7] for DMID [12] due to GPU memory limitations.

† Patch-Forensics has been trained on ProGAN data but not on the same dataset as the rest models. For more details please refer to [6].

Table 3: Average precision (AP) scores of baselines and our model across 20 test datasets. The second column (# cl.) presents the number of used training classes. Best performance is denoted with **bold** and second to best with underline. Our method yields +4.5% mean average precision (mAP) compared to the state-of-the-art.

	1-class		2-class		4-class		AVG	
	ACC	AP	ACC	AP	ACC	AP	ACC	AP
all generators								
w/o contr. loss	87.3	98.3	87.9	98.5	90.0	98.8	88.4	98.5
w/o TIE	85.0	97.8	86.4	98.5	90.5	98.8	87.3	98.3
w/o intermediate	78.9	93.1	81.1	94.7	82.5	94.8	80.8	94.2
full	90.3	98.6	87.7	98.7	91.5	98.8	89.8	98.7
non-GAN generators								
w/o contr. loss	83.4	97.4	84.6	97.7	86.3	98.2	84.7	97.8
w/o TIE	80.0	96.6	82.0	97.8	87.0	98.2	83.0	97.5
w/o intermediate	73.4	90.0	76.3	92.4	77.0	92.4	75.5	91.6
full	87.2	98.0	84.3	98.1	88.3	98.3	86.6	98.1

Table 4: Ablation analysis compares the full architecture with it after removing the contrastive loss, the TIE module, and the intermediate representations.

5.2 Ablations

In Tab. 4, we present an ablation study, where we remove three of RINE’s main components, namely the intermediate representations, the TIE, and the contrastive loss, one-by-one. To be more precise, “w/o intermediate” means that we use only the last layer’s features (equivalent to the SotA method [41]). We measure the RINE’s performance on 20 test datasets, after removing each of the three components, and present the average ACC and AP for the 1-, 2-, and 4-class models, as well as the average across the three models. In addition, we present generalizability results by averaging across the non-GAN generators. The results demonstrate the positive impact of all proposed components. Also, ablating intermediate representations yields the biggest performance loss reducing the metrics to the previous SotA levels (i.e., [41]). In supplementary material the interested reader can find ablation of the fusion mechanism and backbone.

5.3 Hyperparameter analysis

Fig. 3(a)-(c) illustrates ACC & AP boxplots for ViT-B/32 vs. ViT-L/14 in the 1-, 2-, and 4-class settings. Each boxplot is built from 48 scores (ACC/AP) obtained from the experiments of all combinations of ξ , q , and d' . It is clear that CLIP ViT-L/14 outperforms ViT-B/32, thus from now on we report results of ViT-L/14 only. Fig. 3(d)-(f) illustrates the impact of the contrastive loss factor ξ , showing little variability across the different values, in all training settings. Fig. 4 presents ACC boxplots for each (q, d') combination of the trainable part of the model. For $q=1$, little variability is observed for different d' choices, while for $q=2$ and $q=4$, there is a slight performance increase with d' on average. The hyperparameter choices we have made for the proposed models are based on the maximum ACC & AP and not their average or distribution.

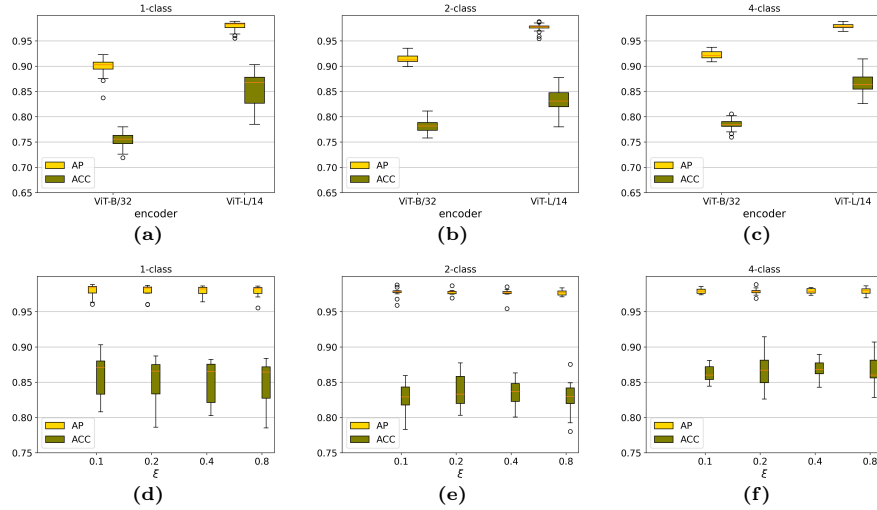


Fig. 3: Boxplots of ACC and AP expressing the impact of image encoder ViT-B/32 vs. ViT-L/14 for the (a) 1-class, (b) 2-class, and (c) 4-class settings, and impact of the contrastive loss factor ξ for the (d) 1-class, (e) 2-class, and (f) 4-class settings.

5.4 The effect of training duration

As stated in Sec. 4.2, the proposed models are trained for only one epoch. Here, we experimentally assess the potential benefit from further training. We identify the 3 best-performing configurations in the 1-epoch case for each of the 1-, 2-, and 4-class settings. Then, we train the corresponding models for 15 epochs and evaluate them in the 3rd, 5th, 10th, and 15th epoch (reducing learning rate by a factor of 10 at 6th and 11th epoch). Fig. 5 presents their average and maximum performance, across the 20 test datasets, at different training stages. Neither the average nor the maximum performance improve if we continue training; instead, in the 1- and 4-class training settings performance drops.

5.5 The effect of training set size

Here, we present a set of experiments in order to assess the effect of the training set size to our models' performance. Specifically, we re-train (again for only one epoch) the best models (1-class, 2-class, 4-class) using 20%, 50%, 80%, and 100% of the training data. For the 1-class, these percentages correspond to 7K, 18K, 28K, and 36K images, for the 2-class to 14K, 36K, 57K, and 72K images, and for the 4-class to 28K, 72K, 115K, and 144K images. In Fig. 6, we present the obtained performance (ACC & AP) on the 20 test datasets. We observe that the results are very close to each other, especially in terms of AP, showcasing the effectiveness of the proposed method even in limited training data settings.




Fig. 4: Boxplots of ACC for all different architectures of the trainable part of the model. X-axis shows the number of trainable parameters as well as the tuple (q, d') .




Fig. 5: Maximum and average performance (ACC & AP) across the 20 evaluation datasets, of the 3 best-performing configurations at 3, 5, 10, and 15 epochs, for the (a) 1-class, (b) 2-class, and (c) 4-class training settings.

5.6 Robustness to common image perturbations

Following common practice [41, 49, 52, 54], we assess our method’s robustness to perturbations typically applied in social media by applying blurring, cropping, compression, noise addition, and their combination to each evaluation sample with 0.5 probability (experimental setting directly taken from [19]). The results in Tab. 5 indicate high robustness against cropping and compression, with minimal performance losses, and a good level of robustness against blurring and noise addition. Their combination considerably reduce performance but at similar levels to that of current SotA without applying perturbations (cf. Tab. 2).

5.7 Importance of intermediate stages

In Sec. 5.2, we show that the intermediate representations is the most important component of the proposed method. Here, we estimate the importance of each block by analysing the importances $\mathbf{A} = \{\alpha_{lk}\}$ estimated by TIE (cf. Eq. (4)):

$$l_k^* = \underset{l}{\operatorname{argmax}}(\alpha_{lk}), \text{ for } k=1, \dots, d'$$

$$f_l = \frac{|\{k \mid l_k^* = l\}|}{d'}, \text{ for } l=1, \dots, n$$

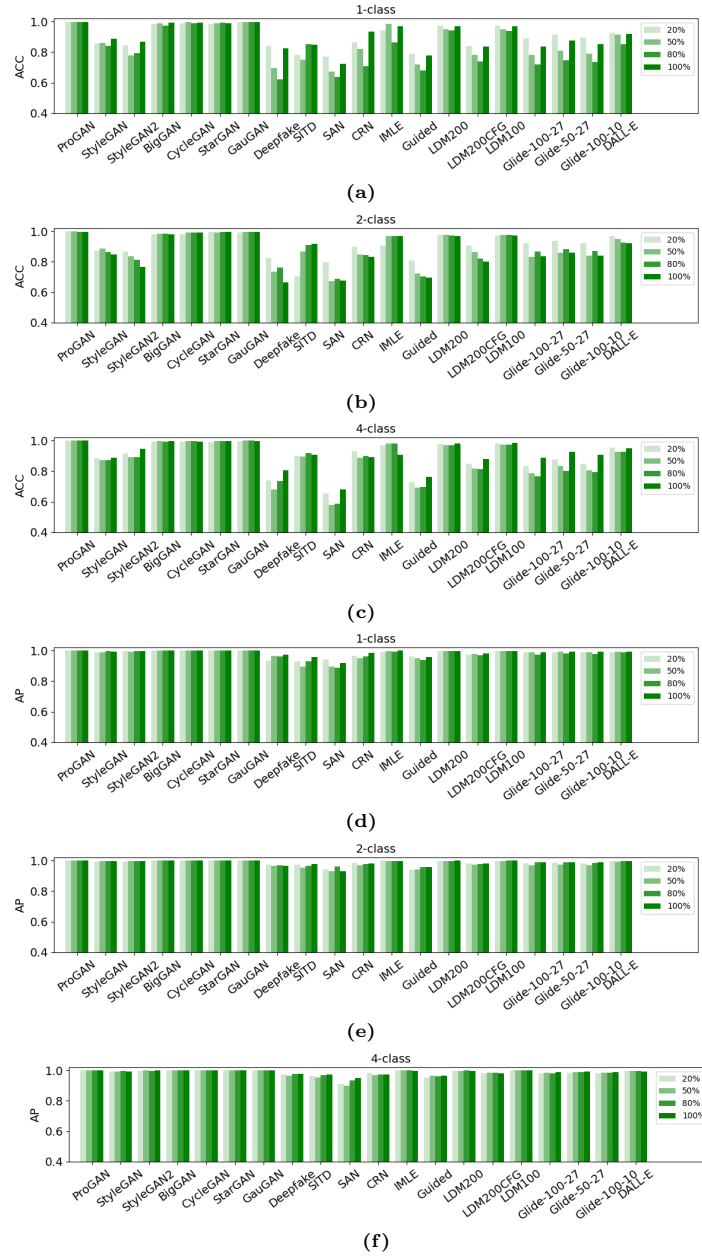


Fig. 6: ACC on the 20 evaluation datasets when the proposed model is trained with 20%, 50%, 80%, and 100% of the training data, for the (a) 1-class, (b) 2-class, and (c) 4-class settings. In (d)-(f) the same for AP.

	Generative Adversarial Networks								Low level vision				Perceptual loss				Latent Diffusion			Glide				AVG
	Pro-	Style-	Style-	Big-	Cycle-	Star-	Gau-	Deep-	SITD	SAN	CRN	IMLE	Guided	200 steps	200 CFG	100 steps	100 27	50 27	100 10	DALL-E				
	perturb.	GAN	GAN	GAN2	GAN	GAN	GAN	GAN													fake			
ACC																								
1-class																								
blur	96.3	78.3	76.2	91.3	90.3	94.1	93.9	72.3	88.9	61.4	83.8	86.5	77.1	93.3	78.5	93.2	77.5	79.0	77.0	88.9	83.9			
crop	99.8	85.7	81.0	98.9	99.5	98.9	99.8	79.8	89.2	63.9	91.8	94.1	74.9	96.5	81.3	96.5	82.2	86.3	84.7	91.9	88.8			
compress	99.1	81.2	77.3	96.9	99.0	98.4	99.7	78.7	89.2	62.8	88.6	96.1	72.1	94.4	76.3	94.2	79.3	81.0	81.6	89.0	86.7			
noise	97.5	77.0	73.8	91.8	95.7	94.0	98.2	80.1	88.6	63.5	80.1	89.8	78.1	93.7	75.4	94.2	83.1	85.2	83.2	88.4	85.6			
combined	88.8	72.3	71.7	80.7	87.7	90.1	89.3	75.3	88.3	62.8	76.3	80.0	80.2	87.4	72.2	88.4	78.4	81.1	79.2	82.6	80.6			
no	99.8	88.7	86.9	99.1	99.4	98.8	99.7	82.7	84.7	72.4	93.4	96.9	77.9	96.9	83.5	97.0	83.8	87.4	85.4	91.9	90.3			
2-class																								
blur	97.2	75.4	68.1	90.3	94.3	95.4	97.0	60.7	88.9	59.6	80.9	91.3	72.0	92.7	74.0	91.8	77.2	78.2	76.8	87.4	82.5			
crop	99.8	81.3	72.5	98.0	99.4	99.6	99.8	64.5	89.7	61.4	85.3	96.9	67.2	97.0	79.2	97.2	83.1	85.3	83.4	92.5	86.7			
compress	99.4	78.0	70.1	94.6	99.4	98.5	99.7	64.0	88.9	61.6	82.2	96.7	66.2	94.5	74.4	94.0	79.2	81.4	80.3	88.7	84.6			
noise	97.7	74.2	67.5	89.5	96.2	93.5	98.8	69.0	88.6	62.6	73.3	88.4	71.7	92.8	71.0	91.6	79.5	80.3	79.5	87.1	82.6			
combined	90.6	69.6	65.4	81.4	88.2	91.9	92.5	66.7	89.2	61.6	75.5	88.3	74.7	87.6	68.2	86.3	76.1	78.5	76.0	81.3	79.5			
no	99.8	84.9	76.7	98.3	99.4	99.6	99.9	66.7	91.9	67.8	83.5	96.8	69.6	96.8	80.0	97.3	83.6	86.0	84.1	92.3	87.7			
4-class																								
blur	97.7	86.9	90.4	93.9	91.5	98.5	95.1	72.3	91.7	57.8	81.0	81.6	77.6	92.8	81.2	92.3	83.2	87.2	85.2	92.2	86.5			
crop	100.0	88.5	93.4	99.5	99.4	99.5	99.8	79.2	92.2	62.1	87.7	88.3	74.9	98.0	88.3	98.4	88.6	92.4	90.1	95.0	90.8			
compress	99.9	87.5	89.2	98.1	99.2	99.6	99.7	72.8	91.4	61.9	88.0	93.0	70.6	95.2	79.8	95.9	85.8	88.4	87.2	91.2	88.7			
noise	97.7	80.0	78.4	91.3	93.9	94.8	98.0	78.2	91.9	59.6	71.9	79.8	74.5	92.5	74.8	92.7	83.6	86.0	84.2	88.1	84.6			
combined	91.9	75.8	78.9	82.6	88.7	91.2	92.9	72.5	92.5	58.7	77.5	85.0	79.1	89.1	73.4	89.2	82.8	85.2	83.5	82.5	82.7			
no	100.0	88.9	94.5	99.6	99.3	99.5	99.8	80.6	90.6	68.3	89.2	90.6	76.1	98.3	88.2	98.6	88.9	92.6	90.7	95.0	91.5			
AP																								
1-class																								
blur	99.7	94.6	93.6	97.6	98.6	99.5	99.8	93.2	96.1	76.1	94.1	98.9	91.0	98.6	91.2	98.6	93.1	94.0	92.9	96.8	94.9			
crop	100.0	98.8	99.5	99.9	100.0	99.9	100.0	96.9	96.5	83.5	98.2	99.9	95.3	99.8	98.0	99.8	99.1	99.4	99.4	99.3	98.2			
compress	100.0	98.5	99.0	99.7	99.9	99.9	100.0	96.1	96.0	82.3	95.9	99.6	94.0	99.7	96.2	99.7	98.6	98.8	98.8	99.1	97.6			
noise	99.8	92.5	93.0	98.2	99.2	98.7	99.9	89.1	95.8	82.9	90.6	96.9	95.0	99.0	92.6	99.1	96.8	97.3	96.5	97.7	95.5			
combined	97.8	84.9	84.3	92.1	97.0	97.6	98.2	84.0	95.9	77.8	87.2	94.6	92.2	95.7	84.0	96.0	88.7	89.9	87.7	92.1	90.9			
no	100.0	99.1	99.7	99.9	100.0	100.0	100.0	97.4	95.8	91.9	98.5	99.9	95.7	99.8	98.0	99.9	98.9	99.3	99.1	99.3	98.6			
2-class																								
blur	99.8	94.7	92.2	97.7	99.6	99.8	99.9	90.6	96.2	77.3	90.7	98.2	91.2	98.8	91.3	98.5	93.8	93.9	93.4	97.4	94.7			
crop	100.0	99.4	99.4	99.9	100.0	100.0	100.0	95.8	96.8	86.4	97.4	99.7	94.7	99.9	98.1	99.9	98.8	99.0	98.9	99.6	98.2			
compress	100.0	99.1	98.8	99.7	100.0	100.0	100.0	94.9	96.3	85.2	96.5	99.7	94.1	99.7	96.0	99.7	98.3	98.5	98.4	99.2	97.7			
noise	99.8	93.1	91.7	98.2	99.5	99.3	99.9	84.6	96.4	84.3	91.7	98.0	93.7	99.2	91.7	99.2	97.2	96.9	96.7	98.3	95.5			
combined	97.8	84.8	81.0	92.8	98.0	98.2	98.5	78.3	96.6	77.2	87.3	96.2	90.4	96.5	84.1	95.8	90.4	91.4	90.7	93.4	91.0			
no	100.0	99.5	99.6	99.9	100.0	100.0	100.0	96.4	97.5	93.1	98.2	99.8	95.7	99.9	98.0	99.9	98.9	99.0	98.8	99.6	98.7			
4-class																								
blur	99.9	97.1	98.4	98.7	99.6	99.9	99.8	93.5	97.3	74.2	95.1	99.5	90.5	98.4	93.1	98.1	94.4	96.0	94.9	97.7	95.8			
crop	100.0	99.5	100.0	99.9	100.0	100.0	100.0	97.8	96.6	86.9	97.5	99.7	95.6	99.8	98.3	99.9	99.0	99.4	99.0	99.5	98.4			
compress	100.0	99.4	99.9	99.9	100.0	100.0	100.0	96.1	97.2	86.0	95.4	99.5	94.3	99.6	96.0	99.6	98.6	98.9	98.6	98.8	97.9			
noise	99.9	96.5	97.6	98.7	98.9	99.1	99.9	88.0	97.4	86.5	82.7	91.5	94.5	99.0	93.0	98.8	97.8	98.1	97.8	98.0	95.7			
combined	98.2	89.4	91.7	93.3	97.9	98.2	98.4	82.1	97.3	77.9	88.1	94.6	93.3	96.6	87.6	96.3	92.4	94.9	93.7	93.5	92.8			
no	100.0	99.4	100.0	99.9	100.0	100.0	100.0	97.9	97.2	94.9	97.3	99.7	96.4	99.8	98.3	99.9	98.8	99.3	98.9	99.3	98.8			

Table 5: Performance (ACC & AP) after applying common image perturbations.




Fig. 7: Frequency (f_l) of each Transformer block obtaining the maximum importance over feature elements.

Fig. 7 illustrates the frequency f_l of Transformer block l obtaining maximum importance, for $l=1, \dots, n$. For the 1-class and 4-class models $n=24$ and $d'=1024$, while for the 2-class model $n=24$ and $d'=128$. We observe that there are many Transformer blocks obtaining higher maximum importance frequency than the last one in all training settings, supporting the proposed method’s motivation.

5.8 The effect of training data origin on detecting Diffusion images

We additionally assess our method’s performance on Diffusion data produced by commercial tools, and more precisely on the Synthbuster dataset [3]. Due to the images’ size we consider the ten-crop validation technique and average the predictions. Tab. 6 illustrates RINE’s performance on Synthbuster when trained on ProGAN data from [52], and on Latent Diffusion Model (LDM) data from [12]. As expected training on Diffusion data provides big performance gains.

training set	DALL-E 2	DALL-E 3	SD1.3	SD1.4	SD2	SD XL	Glide	Firefly	Midjourney
ProGAN (from [52])	64.6/86.7	21.1/30.7	66.3/91.3	65.8/91.4	49.0/66.8	53.3/73.4	42.6/55.7	70.8/99.3	34.2/39.5
LDM (from [12])	89.8/96.2	47.2/32.5	96.4/100.0	96.4/100.0	93.5/98.5	96.3/99.8	90.0/96.3	85.3/93.4	92.4/97.4

Table 6: Performance (ACC/AP) of our method on Synthbuster.

6 Conclusions

In this work, inspired by the recent advances on the use of large-scale pre-trained visual features, we propose RINE that leverages representations from intermediate encoder-blocks of CLIP, as more effective to SID. Our experiments demonstrate the superiority of our approach compared to the state-of-the-art. As an added benefit, the proposed approach requires only one epoch, translating to ~ 8 minutes of training time, and limited training data to achieve maximum performance, while it is robust to common image transformations.

Acknowledgments

This work is partially funded by the Horizon Europe project vera.ai (GA No. 101070093) and H2020 project AI4Media (GA No. 951911).

References

1. Agarap, A.F.: Deep learning using rectified linear units (relu). arXiv preprint arXiv:1803.08375 (2018)
2. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
3. Bammey, Q.: Synthbuster: Towards detection of diffusion model generated images. *IEEE Open Journal of Signal Processing* (2023)
4. Bayar, B., Stamm, M.C.: Constrained convolutional neural networks: A new approach towards general purpose image manipulation detection. *IEEE Transactions on Information Forensics and Security* **13**(11), 2691–2706 (2018)
5. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019 (2019)
6. Chai, L., Bau, D., Lim, S.N., Isola, P.: What makes fake images detectable? understanding properties that generalize. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI*. pp. 103–120. Springer (2020)
7. Chen, C., Chen, Q., Xu, J., Koltun, V.: Learning to see in the dark. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3291–3300 (2018)
8. Chen, Q., Koltun, V.: Photographic image synthesis with cascaded refinement networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1511–1520 (2017)
9. 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: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8789–8797 (2018)
10. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1251–1258 (2017)
11. Corvi, R., Cozzolino, D., Poggi, G., Nagano, K., Verdoliva, L.: Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 973–982 (2023)
12. Corvi, R., Cozzolino, D., Zingarini, G., Poggi, G., Nagano, K., Verdoliva, L.: On the detection of synthetic images generated by diffusion models. In: *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1–5. IEEE (2023)
13. 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)
14. Croitoru, F.A., Hondru, V., Ionescu, R.T., Shah, M.: Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
15. Dai, T., Cai, J., Zhang, Y., Xia, S.T., Zhang, L.: Second-order attention network for single image super-resolution. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11065–11074 (2019)
16. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)

17. Dogoulis, P., Kordopatis-Zilos, G., Kompatsiaris, I., Papadopoulos, S.: Improving synthetically generated image detection in cross-concept settings. In: Proceedings of the 2nd ACM International Workshop on Multimedia AI against Disinformation. pp. 28–35 (2023)
18. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021 (2021)
19. Frank, J., Eisenhofer, T., Schönherr, L., Fischer, A., Kolossa, D., Holz, T.: Leveraging frequency analysis for deep fake image recognition. In: International conference on machine learning. pp. 3247–3258. PMLR (2020)
20. Good, I.J.: Rational decisions. *Journal of the Royal Statistical Society: Series B (Methodological)* **14**(1), 107–114 (1952)
21. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* **27** (2014)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
23. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)
24. Iglesias, G., Talavera, E., Díaz-Álvarez, A.: A survey on gans for computer vision: Recent research, analysis and taxonomy. *Computer Science Review* **48**, 100553 (2023)
25. Jeong, Y., Kim, D., Min, S., Joe, S., Gwon, Y., Choi, J.: Bihpf: Bilateral high-pass filters for robust deepfake detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 48–57 (2022)
26. Jeong, Y., Kim, D., Ro, Y., Choi, J.: Freggan: robust deepfake detection using frequency-level perturbations. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 1060–1068 (2022)
27. Jevnisek, A., Avidan, S.: Aggregating layers for deepfake detection. In: 2022 26th International Conference on Pattern Recognition (ICPR). pp. 2027–2033. IEEE (2022)
28. Karnouskos, S.: Artificial intelligence in digital media: The era of deepfakes. *IEEE Transactions on Technology and Society* **1**(3), 138–147 (2020)
29. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. In: International Conference on Learning Representations (2018)
30. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019)
31. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8110–8119 (2020)
32. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. *Advances in neural information processing systems* **33**, 18661–18673 (2020)
33. Li, K., Zhang, T., Malik, J.: Diverse image synthesis from semantic layouts via conditional imle. 2019 iee. In: CVF International Conference on Computer Vision (ICCV). pp. 4219–4228 (2019)

34. Lu, Z., Huang, D., Bai, L., Qu, J., Wu, C., Liu, X., Ouyang, W.: Seeing is not always believing: Benchmarking human and model perception of ai-generated images. In: Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track (2023)
35. Lukas, J., Fridrich, J., Goljan, M.: Digital camera identification from sensor pattern noise. *IEEE Transactions on Information Forensics and Security* **1**(2), 205–214 (2006)
36. Marra, F., Gragnaniello, D., Cozzolino, D., Verdoliva, L.: Detection of gan-generated fake images over social networks. In: 2018 IEEE conference on multimedia information processing and retrieval (MIPR). pp. 384–389. IEEE (2018)
37. Marra, F., Gragnaniello, D., Verdoliva, L., Poggi, G.: Do gans leave artificial fingerprints? In: 2019 IEEE conference on multimedia information processing and retrieval (MIPR). pp. 506–511. IEEE (2019)
38. Mirsky, Y., Lee, W.: The creation and detection of deepfakes: A survey. *ACM Computing Surveys (CSUR)* **54**(1), 1–41 (2021)
39. Nataraj, L., Mohammed, T.M., Chandrasekaran, S., Flenner, A., Bappy, J.H., Roy-Chowdhury, A.K., Manjunath, B.: Detecting GAN generated fake images using co-occurrence matrices. In: Alattar, A.M., Memon, N.D., Sharma, G. (eds.) *Media Watermarking, Security, and Forensics 2019*, Burlingame, CA, USA, 13-17 January 2019. Society for Imaging Science and Technology (2019)
40. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., Sabato, S. (eds.) *International Conference on Machine Learning, ICML 2022*, 17-23 July 2022, Baltimore, Maryland, USA. *Proceedings of Machine Learning Research*, vol. 162, pp. 16784–16804. PMLR (2022)
41. Ojha, U., Li, Y., Lee, Y.J.: Towards universal fake image detectors that generalize across generative models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 24480–24489 (2023)
42. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2337–2346 (2019)
43. Qian, Y., Yin, G., Sheng, L., Chen, Z., Shao, J.: Thinking in frequency: Face forgery detection by mining frequency-aware clues. In: *European conference on computer vision*. pp. 86–103. Springer (2020)
44. 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: *International conference on machine learning*. pp. 8748–8763. PMLR (2021)
45. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: *International Conference on Machine Learning*. pp. 8821–8831. PMLR (2021)
46. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
47. Rossler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Nießner, M.: Face-forensics++: Learning to detect manipulated facial images. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 1–11 (2019)
48. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**(1), 1929–1958 (2014)

49. Tan, C., Zhao, Y., Wei, S., Gu, G., Wei, Y.: Learning on gradients: Generalized artifacts representation for gan-generated images detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12105–12114 (2023)
50. Tolosana, R., Vera-Rodriguez, R., Fierrez, J., Morales, A., Ortega-Garcia, J.: Deep-fakes and beyond: A survey of face manipulation and fake detection. *Information Fusion* **64**, 131–148 (2020)
51. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
52. Wang, S.Y., Wang, O., Zhang, R., Owens, A., Efros, A.A.: Cnn-generated images are surprisingly easy to spot... for now. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8695–8704 (2020)
53. Yu, F., Wang, D., Shelhamer, E., Darrell, T.: Deep layer aggregation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2403–2412 (2018)
54. Yu, N., Davis, L.S., Fritz, M.: Attributing fake images to gans: Learning and analyzing gan fingerprints. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7556–7566 (2019)
55. Zhang, X., Karaman, S., Chang, S.F.: Detecting and simulating artifacts in gan fake images. In: 2019 IEEE international workshop on information forensics and security (WIFS). pp. 1–6. IEEE (2019)
56. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)