

# Improving Synthetically Generated Image Detection in Cross-Concept Settings

Pantelis Dogoulis

Information Technologies Institute, CERTH  
Thermi, Thessaloniki, Greece  
dogoulis@iti.gr

Ioannis Kompatsiaris

Information Technologies Institute, CERTH  
Thermi, Thessaloniki, Greece  
ikom@iti.gr

Giorgos Kordopatis-Zilos

Czech Technical University in Prague  
Information Technologies Institute, CERTH  
kordogeo@fel.cvut.cz

Symeon Papadopoulos

Information Technologies Institute, CERTH  
Thermi, Thessaloniki, Greece  
papadop@iti.gr

## ABSTRACT

New advancements for the detection of synthetic images are critical for fighting disinformation, as the capabilities of generative AI models continuously evolve and can lead to hyper-realistic synthetic imagery at unprecedented scale and speed. In this paper, we focus on the challenge of generalizing across different concept classes, e.g., when training a detector on human faces and testing on synthetic animal images – highlighting the ineffectiveness of existing approaches that randomly sample generated images to train their models. By contrast, we propose an approach based on the premise that the robustness of the detector can be enhanced by training it on realistic synthetic images that are selected based on their quality scores according to a probabilistic quality estimation model. We demonstrate the effectiveness of the proposed approach by conducting experiments with generated images from two seminal architectures, StyleGAN2 and Latent Diffusion, and using three different concepts for each, so as to measure the cross-concept generalization ability. Our results show that our quality-based sampling method leads to higher detection performance for nearly all concepts, improving the overall effectiveness of the synthetic image detectors.

## CCS CONCEPTS

• **Computing methodologies** → **Computer vision; Image manipulation; Machine learning.**

## KEYWORDS

synthetically generated images, deepfake detection, generalization

### ACM Reference Format:

Pantelis Dogoulis, Giorgos Kordopatis-Zilos, Ioannis Kompatsiaris, and Symeon Papadopoulos. 2023. Improving Synthetically Generated Image Detection in Cross-Concept Settings. In *2nd ACM International Workshop on Multimedia AI against Disinformation (MAD '23)*, June 12, 2023, Thessaloniki, Greece. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3592572.3592846>

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

MAD '23, June 12, 2023, Thessaloniki, Greece

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0187-0/23/06...\$15.00

<https://doi.org/10.1145/3592572.3592846>

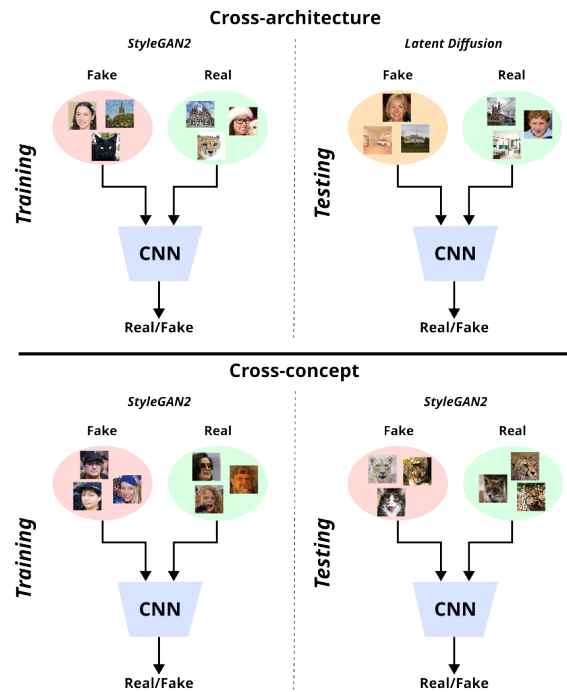


Figure 1: Overview of two generalization evaluation settings: *Cross-architecture* (top), where training and testing are performed on images generated by different generative models, *Cross-concept* (bottom), where training and testing are performed on images of different concepts generated by the same generative model.

## 1 INTRODUCTION

Generative AI models constitute a rapidly growing area, and have already found their place in a wide range of applications [28]. Advances have been driven in part by the development of new algorithms and techniques, such as Generative Adversarial Networks (GANs) [11, 19–22] and recently Diffusion Models [7, 16, 29, 32], which have made it possible to generate increasingly realistic and high-quality synthetic data, reaching unprecedented levels of photorealism. At the same time, generative models can potentially be used by malicious users that spread disinformation across several social platforms. Also, a major concern is that recognizing whether an image is real or fake becomes increasingly challenging [9].

Distinguishing between generated and authentic images has attracted the interest of numerous researchers in the field of multimedia forensics [12, 36]. The most widely used approach for detecting generated images involves training a neural network on a binary classification task (real vs fake) using a large corpus of labelled images. A key component in this development process is the application of a set of carefully selected image augmentations during the training phase, as demonstrated in prior work [12, 36]. Our proposed approach relies on this supervised learning setup.

Additionally, the generalization of the detectors are mostly studied with respect to the different generative architectures [12, 35]. The main goal is to create a detector that is effective in detecting fake images from different architectures. However, herein we take a different avenue and define generalization as the ability to detect synthesized images that depict different concept classes, not necessarily generated from a different architecture (Figure 1). We refer to this capability as generalization in *cross-concept settings*, which is often referred to as *Domain Generalization* in literature [34]. We study the behavior of the detector in such a way that it could be trained in one concept class of real and fake images, e.g., *human faces*, and can then be used to distinguish between real and fake images also in other concepts, e.g., *animal faces*. We empirically find that using standard practices from the literature to train our detectors is not effective enough to generalize on cross-concept scenarios.

In this paper, we tackle the cross-concept generalization challenge of synthetic image detectors by proposing a sampling strategy for the selection of generated images used for training. Prior work [12, 35] relies on generating a large number of images that are used for training. This is equivalent to random sampling, as no criterion for the selection of the generated images is applied. Instead, we assess the quality of the images based on a probabilistic method that provides a Quality Calculation (QC) score. Hence, we rank a large pool of generated images according to their scores and then select the top-k images in terms of quality to train our detectors. This is under the assumption that high quality images will lead the network to focus less on the artifacts of the generative process and more on the characteristics that are invariant to the image content. In that way, we should be able to build more robust detectors in the cross-concept scenario. We evaluate our method using fake images generated by StyleGAN2 [20] and the unconditional module of the recently introduced Latent Diffusion [29]. We also evaluate on three concept classes for each generative architecture to evaluate cross-concept generalization. When training with the proposed sampling strategy, the performance is considerably improved compared to random selection.

Our contribution can be summarized in the following:

- We demonstrate the lack of generalization of state-of-the-art detectors in the cross-concept scenario.
- We propose a sampling strategy that considers image quality scoring for sampling training data.
- We demonstrate improved performance using the proposed approach in the cross-concept settings of three concept classes for two generative architectures.
- We provide our code publicly available to facilitate future research on the field<sup>1</sup>.

<sup>1</sup><https://github.com/dogoulis/qc-sgid>

## 2 RELATED WORK

Detecting synthetically generated images has become an increasingly important task due to the widespread use of numerous generative models to create fake images. In recent years, the research community has proposed various methods for detecting such images. In this section, we discuss the two main categories of methods that have been proposed in the literature: feature-based and frequency-based. Moreover, we analyze two recent works that also consider generated images based on diffusion models.

### 2.1 Feature-based methods

Feature-based methods for detecting synthetic images focus on extracting visual features from the images and using them to train a classifier. One of the most well-known methods in this category is the use of Convolutional Neural Networks (CNNs) trained on a large dataset of real and fake images [35]. In this work, Wang et al. show that a simple ResNet-50 classifier trained with a strong augmentation scheme can generalize to different generative architectures with high scores. Furthermore, Gagnaniello et al. [12] propose some modifications to the network architecture and pre-processing pipeline to improve the performance of the ResNet-50 classifier. The relation between augmentations during training and detection performance is also analyzed in [25]. Moreover, Chai et al. [2] propose using a patch-based CNN that focuses on local regions for improved performance. Du et al. [8] follow a similar local-based approach, training an autoencoder architecture that produces more generalizable results. Similarly, Ju et al. [18] achieve good generalization by analyzing both local (patch-based feature selection) and global features (spatial information). An alternative approach is introduced in [40], where the authors construct a novel loss function under an unsupervised learning framework to boost the generalization of the detector. Specifically, they propose a contrastive loss function that encourages the network to learn discriminative features for real and fake images using a small number of unlabeled images from the target domain.

### 2.2 Frequency-based methods

Frequency-based methods for detecting synthetic images focus on analyzing the color and frequency characteristics of the images. Initial approaches establish the presence of artifacts that are both related to the generative model as well as its specific parameters [26, 38]. One common approach is to extract co-occurrence matrices in both the spatial and frequency domains and feed them as input to CNN models [1, 27]. Both of these works report good results with respect to the generalization of the detector across different generative architectures. Zhang et al. [41] propose detecting artifacts in the spectrum space and use them as input to a ResNet34 network. Similarly, in [10], the authors show that upsampling operations create artifacts in the frequency domain and optimize a simple classifier to detect synthetically generated images. A different approach is followed in [23], where the authors observe the existence of geometric grids in the magnitude and phase spectrograms of the generated images and use them for the detection of fake images.




Figure 2: Overview of the proposed approach. The synthetically generated images (in the pink ellipsis) are evaluated based on a quality assessment model ( $QC(\cdot)$ ) trained on real images (in the green ellipsis). The top- $k$  generated images are then selected and provided, along with real images, as input to a ResNet-50 model, which is trained to discriminate between real and fake images.

### 2.3 Detection of diffusion generated images

Recently, the advent of diffusion models has stimulated the need for new methods to detect the corresponding synthetic images. Corvi et al. [4] conduct a comparative study of state-of-the-art detectors for GAN-generated images, evaluating their performance on the recently introduced diffusion-based generative methods. They show that current detectors cannot generalize to diffusion-based generative architectures. Similarly, Sha et al. [31] perform a comparative analysis of current state-of-the-art detectors and conclude that they are not effective in detecting images generated by recent GAN models. They also propose a multimodal detector that uses images and text prompts as input and show that this technique can improve detection performance.

Overall, while there have been many proposed methods for detecting synthetic images, the task remains challenging, and new techniques are needed to keep up with the rapidly evolving landscape of generative models. However, none of these works address the problem of generalization across different concepts.

## 3 PROPOSED APPROACH

Our method can be split into two main parts. The first part is the sampling strategy, where the quality of the generated images is quantified and the best of them are selected to compose the training set. The second part is the training phase, where we use a CNN network and a pre-processing pipeline following the one proposed by two seminal works in the literature [12, 35]. The use of quality quantification for composing the training set is arguably the most important part of our contribution. This builds on the intuitive assumption that if a network is trained on fake images of high perceptual quality, then it will focus on less obvious details and artifacts of the generative process. Hence, it will learn characteristics that are invariant to the content of the image and, as a result, boost the ability to detect fake images from different concepts.

### 3.1 Quantifying quality in generated images

There are different ways to assess the quality of AI-generated images automatically. The most well-known such metrics are Fréchet Inception Distance (FID) [15] and the Inception Score (IS) [30]. The first measures the distance between the distributions of real and generated images in the feature space. The second measures the quality and diversity of generated images by looking at how well they are classified by a pre-trained Inception network [33]. This metric is generally more sensitive to changes in image quality than the FID. However, both methods measure the similarity between distributions of real and fake images, so they do not fit our needs. In this work, we use a metric that has been proposed in [13], and we refer to it as Quality Calculation (QC). The intuition behind this approach is similar to the one behind the FID and IS. The quality of generated images is quantified based on their similarity to a pre-defined real image distribution. First, a feature extractor  $f(\cdot)$  is used to extract features from the real images. Then, the real image distribution is approximated by a Gaussian Mixture Model (GMM), which is a weighted sum of Gaussian distributions with parameters  $\mu^i$  and  $\Sigma^i$ . Hence, the quality of the input image is calculated using the trained GMM. More formally, for an input image  $I$ , the quality score of the image is calculated based on:

$$p(\mathbf{x}|\lambda) = \sum_{i=1}^M \mathbf{w}^i g(\mathbf{x}|\mu^i, \Sigma^i) \quad (1)$$

where  $\mathbf{x} = f(I)$  and  $\mathbf{w}^i$  are the weights of the mixture model that satisfy the constraint  $\sum_{i=1}^M \mathbf{w}^i = 1$ . The function  $g(\cdot)$  represents the Gaussian density components of the model parameterized by the mean vector  $\mu^i$  and the covariance matrix  $\Sigma^i$ . In a more general form, the parameters of this model can be denoted as  $\lambda = \{\mathbf{w}^i, \mu^i, \Sigma^i\}$ . The estimation of  $\lambda$  is given by the Expectation-Maximization (EM) algorithm [5]. Then, for each generated

image  $I_g$ , its quality score  $QC(I_g)$  is given by:

$$QC(I_g) = p(f(I_g)|\lambda) \quad (2)$$

The advantage of using the QC method when calculating the quality of the generated images is that it can produce a final ranking of images. Contrary to the FID score, which produces a single scalar that denotes the quality of a set of generated images, QC ranks the generated images from highest to lowest quality.

### 3.2 Synthetically generated image detection

Figure 2 displays an overview of the proposed approach. First, we train the QC models [13] for each of the different concepts based on the corresponding datasets of real images. Then, we select the top- $k$  generated images with the highest quality and the same amount of real images in order to form the training dataset. The next step is to train a classifier network to distinguish between real and fake images. The classifier network and the training process are based on the methodology proposed in [35]. A ResNet-50 [14] model is initialized with the weights of ImageNet [6] and trained using the dataset compiled in the previous step. Additionally, during training, we apply a set of augmentations including several geometric augmentations (i.e., random crop and resize), Gaussian blurring, and JPEG compression. The network minimizes the binary cross-entropy loss.

## 4 EVALUATION SETUP

In this section, we discuss the evaluation pipeline for our experiments. Specifically, we present the datasets, evaluation methodology, and implementation details used.

### 4.1 Datasets

We obtain real images from publicly available datasets, including FFHQ [21], AFHQ [3], and LSUN [37]. The FFHQ dataset contains images of human faces, the AFHQ dataset contains images of dogs, cats, and a general class of wildlife, while the LSUN dataset contains nearly one million images of 10 scene categories and twenty object classes. Then, we employ pretrained diffusion and GAN models to generate artificially generated images for different classes to evaluate the cross-concept scenario. More specifically, we use the StyleGAN2 [22] model to generate images from pretrained networks in FFHQ, AFHQ and LSUN-churches, while we also use the Latent Diffusion [29] model to generate images from pretrained networks in FFHQ, LSUN-bedrooms, and LSUN-churches. We denote each dataset as  $\mathcal{X} + \mathcal{X}^A$ , where  $\mathcal{X}$  is the set of real images of a specific concept class and  $\mathcal{X}^A$  is the corresponding set of fake images generated with architecture  $A$ , which takes values either  $G$  for StyleGAN2 or  $D$  for Latent Diffusion. Regarding the different concepts, we use  $\mathcal{H}$  for human,  $\mathcal{A}$  for animal,  $\mathcal{C}$  for church, and  $\mathcal{B}$  for bedroom datasets.

### 4.2 Evaluation methodology and metric

We assess the detection performance of our model on each concept class by utilizing the AUC metric. This particular scoring method does not rely on any predetermined threshold, making it suitable for assessing the robustness and generalization of the detector models. For each concept, we perform three training sessions and report the mean AUC and its standard deviation. For comparison, we train

the ResNet-50 classifier following the same training process but with random sampling for fake images instead, i.e., without using the proposed sampling strategy.

### 4.3 Implementation details

For our sampling strategy, we initially generate 20K total images for each concept and architecture, and then select the top 10K images based on their QC score. The test set consists of 2K images for each concept and it is randomly selected before the QC selection. We sampled the same number of real images in order to form the training and test datasets. The GMM model that was implemented, used the Inception model [33] as a feature extractor. Moreover, we used 50 Gaussian components and a batch size of 50 instances.

For preprocessing, during training, we implement the augmentations from Section 3.2 with an output size of  $224 \times 224$  pixels. While during testing, we only resize images to  $224 \times 224$ .

For our detector, we use a ResNet-50 [14] pretrained on ImageNet [6]. The model is trained using the AdamW optimizer [24]. The learning rate was equal to  $10^{-3}$ , and a step scheduler with 5 epochs was used. Weight decay is also applied with a factor of  $5 \cdot 10^{-5}$ . Additionally, a drop path [17] rate of 0.1 is employed to prevent overfitting, which randomly drops entire paths (i.e., sequences of layers) in the model during training. All training and evaluation processes were carried out on a server with one NVIDIA GeForce RTX 3060 GPU.

## 5 EXPERIMENTS

In this section, we present the experimental results of the proposed method and demonstrate the generalization improvements of the models when trained with it. Our primary interest lies in the cross-concept scenario, where a network is trained on instances of a specific concept class and evaluated on unseen instances from a different concept class. This evaluation serves as a measure of the model's ability to generalize beyond its training domain and is a critical aspect of assessing the effectiveness of the proposed method in terms of cross-concept generalization.

### 5.1 Cross-concept evaluation

Table 1 presents the AUC scores of the detectors trained with random sampling and the proposed method, for generated images based on StyleGAN2 and Latent Diffusion. To ensure the validity and reliability of our findings, we run three training sessions for each concept and report the mean and standard deviation of AUC.

We first analyse the results of our intra-class experiments, where we train and evaluate a detector in the same concept class. The corresponding runs are coloured in gray in Table 1. Our findings indicate that there is no significant difference in the performance of the detector when using images of higher quality, due to the perfect performance in almost all cases (with the exception of the Bedroom concept in Latent Diffusion). This result is consistent with previous studies that have shown that a robust augmentation scheme is sufficient to train very accurate detectors within the same domain [12, 35]. In our study, we used the same augmentation scheme for both randomly selected and quality-based selected subsets, which explains the lack of notable difference in the detector's performance between the two subsets.

training	sampling	test		
		$\mathcal{H}+\mathcal{H}^G$	$\mathcal{A}+\mathcal{A}^G$	$C+C^G$
$\mathcal{H}+\mathcal{H}^G$	random	<b>100.0±0.0</b>	89.1±1.00	77.5±3.42
	QC (ours)	<b>100.0±0.0</b>	<b>94.8±1.22</b>	<b>83.5±0.94</b>
$\mathcal{A}+\mathcal{A}^G$	random	61.8±1.36	<b>100.0±0.0</b>	<b>89.1±0.71</b>
	QC (ours)	<b>66.4±1.08</b>	<b>100.0±0.0</b>	86.5±0.15
$C+C^G$	random	53.7±0.22	58.6±3.51	<b>100.0±0.0</b>
	QC (ours)	<b>59.4±0.19</b>	<b>67.1±0.16</b>	<b>100.0±0.0</b>

(a) StyleGAN2

training	sampling	test		
		$\mathcal{H}+\mathcal{H}^D$	$\mathcal{B}+\mathcal{B}^D$	$C+C^D$
$\mathcal{H}+\mathcal{H}^D$	random	<b>100.0±0.0</b>	64.4±2.31	66.7±0.88
	QC (ours)	<b>100.0±0.0</b>	<b>74.1±0.68</b>	<b>77.5±0.03</b>
$\mathcal{B}+\mathcal{B}^D$	random	52.1±1.47	96.3±0.79	<b>99.7±0.03</b>
	QC (ours)	<b>56.2±5.13</b>	<b>99.6±0.02</b>	99.4±0.02
$C+C^D$	random	54.2±4.08	96.3±1.19	<b>100.0±0.0</b>
	QC (ours)	<b>58.5±1.52</b>	<b>98.9±0.58</b>	<b>100.0±0.0</b>

(b) Latent Diffusion

**Table 1: AUC of detection model trained on randomly selected samples or based on the proposed QC score for each concept and generative model. Mean and standard deviation of three training sessions with different seeds are reported. Bold indicates the best performance between the proposed and the random baseline. Gray colour indicates intra-concept evaluation.**

training	test set	quality quartile ↓			
		0-25%	25-50%	50-75%	75-100%
$\mathcal{H}+\mathcal{H}^G$	$\mathcal{A}+\mathcal{A}^G$	93.4±0.98	92.8±1.32	93.2±1.17	89.8±0.89
	$C+C^G$	83.0±1.07	79.2±1.03	80.9±0.83	77.8±0.92
$\mathcal{A}+\mathcal{A}^G$	$\mathcal{H}+\mathcal{H}^G$	66.8±1.02	66.2±1.15	63.9±0.97	62.7±0.98
	$C+C^G$	84.1±0.32	84.9±0.17	90.3±0.11	94.0±0.13
$C+C^G$	$\mathcal{H}+\mathcal{H}^G$	56.1±0.19	59.9±0.18	59.5±0.21	58.8±0.17
	$\mathcal{A}+\mathcal{A}^G$	62.2±0.13	64.1±0.17	67.1 ±0.14	69.8±0.24

(A) StyleGAN2

training	test set	quality quartile ↓			
		0-25%	25-50%	50-75%	75-100%
$\mathcal{H}+\mathcal{H}^D$	$\mathcal{B}+\mathcal{B}^D$	70.1±0.72	73.6±0.67	74.8±0.68	72.5±0.69
	$C+C^D$	74.1±0.05	73.2±0.04	78.6±0.07	77.4±0.02
$\mathcal{B}+\mathcal{B}^D$	$\mathcal{H}+\mathcal{H}^D$	52.1±5.42	56.9±5.34	55.3±5.22	57.7±4.88
	$C+C^D$	99.2±0.02	98.3±0.07	99.4±0.04	99.4±0.02
$C+C^D$	$\mathcal{H}+\mathcal{H}^D$	55.7±1.55	56.4±1.52	57.8±1.48	57.3±1.42
	$\mathcal{B}+\mathcal{B}^D$	98.2±0.64	98.4±0.72	98.4±0.58	99.3±0.71

(b) Latent Diffusion

**Table 2: AUC of detection model for test subsets grouped based on their QC score for the two generative models. Mean and standard deviation of three training sessions with different seeds are reported.**

Next, we discuss the results of our proposed methodology in the case where a detector is trained on generated and real images from a concept class and is evaluated on images of a different class. Regarding the experiments on StyleGAN2 sets, it is evident that our proposed method clearly improves the robustness of the detector in almost all cases except in the case when training on Animals ( $\mathcal{A}+\mathcal{A}^G$ ) and testing on Churches ( $C+C^G$ ). In all other cases, our proposed approach outperforms the baseline by more than 5% in terms of mean AUC, reaching up to 8.5%. Furthermore, the standard deviation of the runs with the proposed QC sampling strategy is relatively low, being less than 1.5% in all cases, and often significantly lower compared with the runs where random sampling was used. This implies that using our strategy leads to more consistent and reliable models. Similar conclusions can be drawn from the experiments on Latent Diffusion sets. Specifically, our method surpasses the baseline in all but one case, i.e., training on Bedrooms ( $\mathcal{B}+\mathcal{B}^D$ ) and testing on Churches ( $C+C^D$ ), where the difference is marginal. The AUC score is improved by almost 10% in some cases. Also, the standard deviation is generally low, except for one case.

In summary, our experimental comparison provides compelling evidence of the effectiveness of the proposed method in improving the generalization performance across different concepts.


## 5.2 Results when test image quality varies

We also evaluate our model in the case of selecting different test set composition in terms of image quality. Specifically, the aim is to evaluate whether images of better quality are more difficult to detect or not. Hence, we split the test sets into four quartiles denoted as 0-25% (lowest quality), 25-50%, 50-75%, and 75-100% (highest quality) based on their QC score. We observe that in several cases of StyleGAN2 generated images, the detector achieves better AUC scores for low- or medium-quality images. This means that a test set would benefit from an automatic selection step where higher quality images are retained, since it would be more challenging. Instead, in the case of images generated using the Latent Diffusion model, there is no apparent association between the test set composition in terms of image quality and the detection performance.

## 5.3 Results when the detector is trained with many classes and architectures

In this section, we evaluate the performance of the model when trained based on QC sampling in three different datasets:

- $\mathcal{G}$ : Dataset consists of images from all concepts, generated by StyleGAN2, and the corresponding sets of real images.
- $\mathcal{D}$ : Dataset consists of images from all concepts, generated by Latent Diffusion, and the corresponding sets of real images.
- $\mathcal{G} + \mathcal{D}$ : Combination of the other two datasets.



**Figure 3: Visualization of the magnitude spectrograms for the different classes of the two generative architectures using the denoising network proposed in [39].**

training	test		
	$\mathcal{G}$	$\mathcal{D}$	$\mathcal{G} + \mathcal{D}$
$\mathcal{G}$	99.9±0.00	60.5±1.53	77.3±0.88
$\mathcal{D}$	59.2±0.29	99.9±0.00	74.6±1.07
$\mathcal{G} + \mathcal{D}$	99.9±0.00	99.9±0.01	99.9±0.00

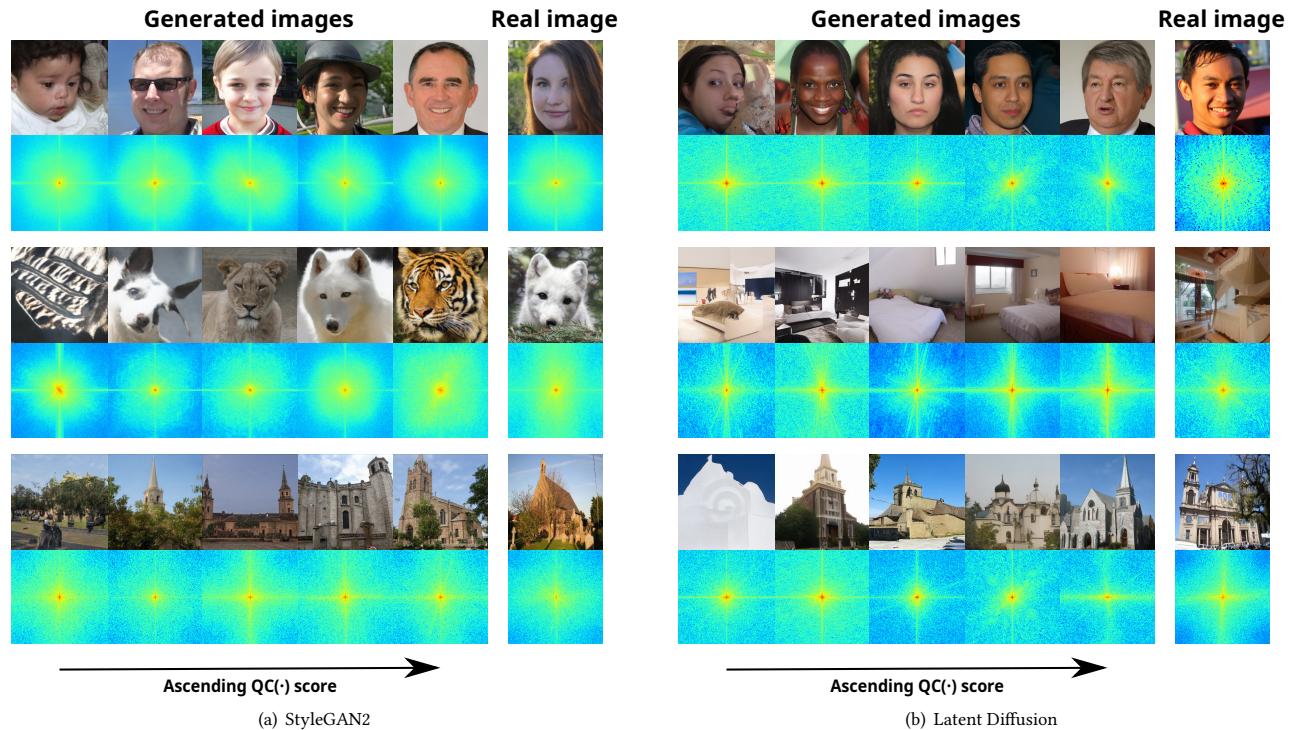
**Table 3: AUC of detection model when trained on all the different classes and architectures. Mean and standard deviation of three sessions with different seeds are reported.**

The findings are presented in Table 3. Our analysis reveals that the performance of the detector significantly improves when it is trained on a diverse set of generated images, including those produced by StyleGAN2 and Latent Diffusion models, and evaluated on a range of test sets. Conversely, we observe that detectors trained solely on images generated by a single architecture demonstrate limited generalization ability when tested on images generated by other architectures [4, 31]. Furthermore, our investigation indicates that the semantic content of each dataset is a crucial factor affecting the detection of generated images, and the detector’s efficacy is significantly enhanced with the exploitation of a greater variety of classes during training.

#### 5.4 Qualitative artifact analysis

In this section, we visualize the artifacts that are calculated using the Fast Fourier Transform for different generated classes of the same architecture. Following [4], we first randomly select 1000 generated images per concept and model architecture. Then using the denoising function proposed in [39], we transform each image  $X$  to its denoised version  $f(X)$  and then compute the residual  $R(X) = X - f(X)$ . Next, we average the residuals and apply a 2D Fast Fourier Transform in order to obtain the magnitude and phase spectrograms. The magnitude spectrograms are a good indicator for the analysis of the artifacts introduced by a generative model.

As seen in Figure 3, differences exist between the generated images of each concept. It is evident that generated images from models pretrained on datasets with similar structures present more common characteristics. For instance, FFHQ and AFHQ, which are used for Humans  $\mathcal{H}^G$  and Animals  $\mathcal{A}^G$ , respectively, are high-quality and high-resolution datasets consisting of images with centered and aligned faces. This produces similar cloudy magnitude spectrograms for these two cases. Similar spectrograms also appear in the case of the Bedrooms  $\mathcal{B}^D$  and Churches  $\mathcal{C}^D$  generated from the models that were trained on the LSUN-Bedroom and LSUN-Church datasets, which contain images of significantly lower resolution and quality compared to FFHQ and AFHQ. Nevertheless, all images are preprocessed in the same way in our experiments; hence, we expect



**Figure 4: Sampled image and magnitude spectrogram of the top 100%, 75%, 50%, 25%, 1% generated images ranked based on their QC score and a real image from the same concept with its magnitude spectrogram.**

such differences to be mitigated during our evaluations. This observation reinforces the main assumption of this work, which is that images produced by generative architectures are class-dependent. This means that the classification of real and fake images is affected by the semantic image content [34, 35].

Moreover, Figure 4 illustrates five generated images for each generation model and concept as well as one real image from the same concept. Specifically, we rank the generated images from each concept, and then we select the images ranked on 1%, 25%, 50%, 75%, and 100% quartiles of the total image quality ranking. This means that the image ranked in the 1% is the image with the best quality, the image ranked in the 100% quartile has the worst quality, while the images located within the 25%, 50%, and 75% quartiles denote images of progressively lower quality as their respective quartile rankings increase. It is noticeable that moving from the worst to the best image ranked by the QC score, the perceptual quality is significantly better. Furthermore, we provide the corresponding magnitude spectrograms of each image in the second row of the figure. Interestingly, we observe that the image with the best quality has a very similar spectrogram to the spectrogram of the real image, while as we move to the image with the worst quality, the spectrogram appears to be noisier. This finding supports our main intuition that training with higher-quality images can facilitate the learning of subtle and less distinguishable artifacts by the network.

## 6 CONCLUSION

In this work, we address the challenge of generalization for synthetic image detection in the cross-concept scenario. We observe

that even state-of-the-art detectors that use strong augmentation schemes to improve their robustness to different generative architectures still lack the ability to generalize to unseen concept classes. To address this, we propose a novel approach that uses a probabilistic method to quantify the quality of generated images and shows that training detectors with higher-quality images can significantly improve their generalization ability. We evaluate our method using fake images generated by two state-of-the-art generative models, StyleGAN2 and Latent Diffusion, and demonstrate that our approach improves the detection performance when training with randomly selected generated images. Also, we enhance the basic intuition by visualising the magnitude spectrograms of quality-based selected generated images. Although our results are promising, there is still considerable room for improvement in the development of detectors that can generalize better across different concept classes and generative models. To this end, we plan to carry out more extensive experiments involving more concepts and generative models, and investigating the potential of frequency-based methods and ensembles of models to further increase the robustness and accuracy of the detectors.

## ACKNOWLEDGMENTS

This work was partially funded by the projects “vera.ai: VERification Assisted by Artificial Intelligence” (GA no. 101070093) and “AI4Media: A European Excellence Centre for Media, Society and Democracy” (GA no. 951911).

## REFERENCES

- [1] Mauro Barni, Kassem Kallas, Ehsan Nowroozi, and Benedetta Tondi. 2020. CNN detection of GAN-generated face images based on cross-band co-occurrences analysis. In *Proceedings of the IEEE international Workshop on Information Forensics and Security*.
- [2] 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*.
- [3] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [4] Riccardo Corvi, Davide Cozzolino, Giada Zingarini, Giovanni Poggi, Koki Nagano, and Luisa Verdoliva. 2022. On the detection of synthetic images generated by diffusion models. *arXiv preprint arXiv:2211.00680* (2022).
- [5] Arthur P Dempster, Nan M Laird, and Donald B Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* (1977).
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [7] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. In *Proceedings of the Advances in Neural Information Processing Systems*.
- [8] Mengnan Du, Shiva Pentylala, Yuening Li, and Xia Hu. 2019. Towards generalizable forgery detection with locality-aware autoencoder. In *Proceedings of the ACM International Conference on Information and Knowledge Management*.
- [9] Hany Farid. 2016. *Photo forensics*.
- [10] 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*.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020).
- [12] 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*.
- [13] Shuyang Gu, Jianmin Bao, Dong Chen, and Fang Wen. 2020. Giqa: Generated image quality assessment. In *Proceedings of the European Conference on Computer Vision*.
- [14] 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*.
- [15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the Advances in Neural Information Processing Systems*.
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. In *Proceedings of the Advances in Neural Information Processing Systems*.
- [17] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. 2016. Deep networks with stochastic depth. In *Proceedings of the European Conference on Computer Vision*.
- [18] Yan Ju, Shan Jia, Lipeng Ke, Hongfei Xue, Koki Nagano, and Siwei Lyu. 2022. Fusing Global and Local Features for Generalized AI-Synthesized Image Detection. In *Proceedings of the IEEE International Conference on Image Processing*.
- [19] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *Proceedings of the International Conference on Learning Representations*.
- [20] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2021. Alias-free generative adversarial networks. In *Proceedings of the Advances in Neural Information Processing Systems*.
- [21] 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*.
- [22] 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*.
- [23] Bo Liu, Fan Yang, Xiuli Bi, Bin Xiao, Weisheng Li, and Xinbo Gao. 2022. Detecting Generated Images by Real Images. In *Proceedings of the European Conference on Computer Vision*.
- [24] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *Proceedings of the International Conference on Learning Representations*.
- [25] Sara Mandelli, Nicolò Bonettini, Paolo Bestagini, and Stefano Tubaro. 2022. Detecting Gan-Generated Images by Orthogonal Training of Multiple CNNs. In *Proceedings of the IEEE International Conference on Image Processing*.
- [26] Francesco Marra, Diego Gragnaniello, Luisa Verdoliva, and Giovanni Poggi. 2019. Do gans leave artificial fingerprints?. In *Proceedings of the IEEE Conference on Multimedia Information Processing and Retrieval*.
- [27] Lakshmanan Nataraj, Tajuddin Manhar Mohammed, BS Manjunath, Shivkumar Chandrasekaran, Arjuna Flenner, Jawadul H Bappy, and Amit K Roy-Chowdhury. 2019. Detecting GAN generated fake images using co-occurrence matrices. *Electronic Imaging* (2019).
- [28] 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).
- [29] 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*.
- [30] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training gans. In *Proceedings of the Advances in Neural Information Processing Systems*.
- [31] Zeyang Sha, Zheng Li, Ning Yu, and Yang Zhang. 2022. DE-FAKE: Detection and Attribution of Fake Images Generated by Text-to-Image Diffusion Models. *arXiv preprint arXiv:2210.06998* (2022).
- [32] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *Proceedings of the International Conference on Learning Representations*.
- [33] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [34] Antonio Torralba and Alexei A Efros. 2011. Unbiased look at dataset bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [35] 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*.
- [36] Xin Wang, Hui Guo, Shu Hu, Ming-Ching Chang, and Siwei Lyu. 2022. Gan-generated faces detection: A survey and new perspectives. *arXiv preprint arXiv:2202.07145* (2022).
- [37] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365* (2015).
- [38] 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*.
- [39] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. 2017. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing* (2017).
- [40] Mingxu Zhang, Hongxia Wang, Peisong He, Asad Malik, and Hanqing Liu. 2022. Improving GAN-generated image detection generalization using unsupervised domain adaptation. In *Proceedings of the IEEE International Conference on Multimedia and Expo*.
- [41] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. 2019. Detecting and simulating artifacts in gan fake images. In *Proceedings of the IEEE international Workshop on Information Forensics and Security*.