

AI-Generated Faces in the Real World: A Large-Scale Case Study of Twitter Profile Images

Jonas Ricker
Ruhr University Bochum
Bochum, Germany

Dennis Assenmacher
GESIS - Leibniz Institute for the
Social Sciences
Cologne, Germany

Thorsten Holz*
CISPA Helmholtz Center for
Information Security
Saarbrücken, Germany

Asja Fischer*
Ruhr University Bochum
Bochum, Germany

Erwin Quiring*
Ruhr University Bochum
Bochum, Germany
ICSI
Berkeley, USA

ABSTRACT

Recent advances in the field of generative artificial intelligence (AI) have blurred the lines between authentic and machine-generated content, making it almost impossible for humans to distinguish between such media. One notable consequence is the use of AI-generated images for fake profiles on social media. While several types of disinformation campaigns and similar incidents have been reported in the past, a systematic analysis has been lacking.

In this work, we conduct the first large-scale investigation of the prevalence of AI-generated profile pictures on Twitter. We tackle the challenges of a real-world measurement study by carefully integrating various data sources and designing a multi-stage detection pipeline. Our analysis of nearly 15 million Twitter profile pictures shows that 0.052 % were artificially generated, confirming their notable presence on the platform. We comprehensively examine the characteristics of these accounts and their tweet content, and uncover patterns of coordinated inauthentic behavior. The results also reveal several motives, including spamming and political amplification campaigns. Our research reaffirms the need for effective detection and mitigation strategies to cope with the potential negative effects of generative AI in the future.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Information systems** → **Social networks**; • **Security and privacy** → **Human and societal aspects of security and privacy**.

KEYWORDS

AI-Generated Content, Fake Image Detection, Social Networks

ACM Reference Format:

Jonas Ricker, Dennis Assenmacher, Thorsten Holz, Asja Fischer, and Erwin Quiring. 2024. AI-Generated Faces in the Real World: A Large-Scale Case

*Equal supervision.



This work is licensed under a Creative Commons Attribution International 4.0 License.

RAID 2024, September 30–October 02, 2024, Padua, Italy

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0959-3/24/09

<https://doi.org/10.1145/3678890.3678922>

Study of Twitter Profile Images. In *The 27th International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2024)*, September 30–October 02, 2024, Padua, Italy. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3678890.3678922>

1 INTRODUCTION

The emergence of generative artificial intelligence (AI) has revolutionized content creation, enabling us to produce highly authentic and diverse outputs, such as images, videos, texts, and music that bear a striking resemblance to human-created media. These AI-driven systems have become ubiquitous in various areas of our society and provide reliable support in numerous applications. Among other use cases, they streamline the writing of emails and texts or enhance programming with advanced code completion tools. However, alongside its impressive benefits, generative AI also has the potential for significant detrimental effects. A pressing problem is the ability to generate compellingly realistic but false content, which can be used as a way to spread misinformation, manipulate people, and influence public opinion.

In a significant action in late 2019, Facebook dismantled an extensive network of over 900 accounts, pages, and groups that had collectively spent more than 9 million USD on advertisements promoting Donald Trump, potentially impacting the 2020 US presidential election [74]. A notable feature of this network was the use of AI-generated profile images, possibly taken from the website *thispersondoesnotexist.com* (TPDNE) which became operational in February 2019. Using NVIDIA’s StyleGAN [55], TPDNE generates a new facial image every time the page is refreshed, making it easily accessible to everyone. Since this incident, the use of AI-synthesized faces in disinformation campaigns has been on the rise, likely because such images reduce the risk of detection through reverse image searches [38]. Investigations have revealed that many of these deceptive clusters operate with state interests in mind, seeking to bolster specific narratives [73, 76, 98, 101] or interfere in the domestic policies of foreign states [42, 43, 75]. Additionally, there are efforts to influence public opinion [99, 100] or establish connections with unsuspecting social media users [37, 106]. The FBI and Europol have expressed concerns that the trend of using AI-generated content in cybercrime and foreign influence operations is expected to grow steadily [25, 29]. Given these examples, it

is essential to understand the detection possibility, prevalence, and usage of AI-generated images in the wild instead of a lab setting.

In this work, we tackle this challenge by concentrating on the phenomenon of AI-generated images in social media. At the time of writing, it is becoming increasingly difficult for humans to differentiate these machine-generated media from authentic photographs, as evidenced by recent studies [34, 51, 59, 71, 72, 95, 103]. Although the detection of generated images has been explored extensively in lab settings, there is a surprising lack of comprehensive research addressing their identification and widespread use on social media platforms in real-world contexts. In this paper, we provide the first systematic and large-scale study of AI-generated profile images on Twitter¹. Our research is founded on three main pillars.

First, we develop a fast and effective detection pipeline tailored to the identification of AI-generated images in real-world scenarios. This task presents unique challenges, including the lack of a definitive ground truth and the diversity of possible image manipulations. To solve these problems, we carefully design a detection pipeline step by step. We consider different dataset types, apply a pre-filter to discard images with too small or no faces, and adapt a state-of-the-art classification model specifically targeting synthetic profile images on Twitter. As mentioned above, observations suggest that the majority of AI-generated profile images originate from TPDNE, which is why we tailor our detection pipeline to this kind of fake faces. Finally, we integrate various tools that help with the manual labeling that is required to estimate error rates on unlabeled in-the-wild data. We study each component of our system in controlled setups and show that the pipeline is capable of accurately recognizing AI-generated images.

Second, we analyze a large collection of 14 989 385 Twitter profile pictures to determine how prevalent AI-generated profile pictures are on the platform. We identify 7 723 accounts that use such images, which corresponds to a prevalence rate of 0.052 %. This result indicates a notable presence of generated profile images on Twitter. We also assess the accuracy and reliability of our findings by estimating error rates. We estimate the false negative rate (FNR)—the fraction of mislabeled fake images—of our approach to lie between 2.75 % and 3.03 %, and the false discovery rate (FDR)—the fraction of real images among all images classified as fake—to be 1.4 %. The results suggest a low error rate of our method.

Third, we contextualize the use of AI-generated profile pictures on Twitter by examining the corresponding accounts and their tweets. Our results show clear differences between the two types of accounts: accounts with fake images tend to have lower social engagement as well as fewer followers and followed accounts. Despite the generally lower activity, some accounts with fake images are very active, suggesting possible involvement in spam campaigns. In addition, fake accounts are often newer and are suspended more frequently by Twitter, indicating inauthentic behavior. A significant portion of accounts was created in bulk shortly before our data collection, which is a common pattern for accounts created for message amplification, disinformation campaigns, or similar disruptive activity. This impression is confirmed by our textual analysis of the accounts’ tweets. We identify large clusters spamming very similar

contents, frequently referring to giveaways, cryptocurrencies, and pornography. Notably, we also observe accounts that engage in contentious or political topics, such as the war in Ukraine, debates on COVID-19 and vaccinations, and election-related discourse.

Contributions. We make the following key contributions:

- (1) *Detection Pipeline.* We propose a multi-step pipeline for detecting AI-generated profile images on social media. We evaluate each stage in a controlled setup and demonstrate the pipeline’s suitability for real-world settings.
- (2) *Prevalence Study on Twitter.* We apply our pipeline to a total of 14 989 385 authentic profile images to systematically study the prevalence of AI-generated faces on Twitter. We identify 7 723 accounts with generated profile images, corresponding to a prevalence rate of 0.052 %.
- (3) *Account and Tweet Analysis.* We analyze the user metrics and tweets of accounts using AI-generated profile images to learn more about their intended purpose. We identify prevalent topics and find a significant number of accounts to apparently participate in coordinated inauthentic behavior.

We release our code and data at github.com/jonasricker/twitter-ai-faces.

2 BACKGROUND

We start by providing a short primer on the creation and detection of AI-generated images.

AI-Generated Content (AIGC). AIGC, sometimes also referred to as “deepfakes”, is content that appears authentic to humans but is synthesized or altered using a deep neural network. It is most prominently associated with manipulated videos in which the face of a person is replaced with a different one [69], but also encompasses other types of media including images, audio, and text. While AIGC offers great creative potential, it is also used for malicious purposes, including defamatory images and videos [14], voice cloning [20, 35], fake customer reviews [115], and machine-generated posts on social media [27, 39].

Image Synthesis. Learning a probability distribution from samples in order to generate novel samples is a longstanding challenge, especially in the high-dimensional image domain. Besides variational autoencoders (VAEs) [58] and autoregressive models [104, 105], generative adversarial networks (GANs) [40] have proven to be effective in synthesizing high-quality images [12, 52–56, 91, 122]. The StyleGAN family [54–56] received special attention due to their ability to generate faces that are practically indistinguishable from real ones [72]. Recently, it has been shown that diffusion models (DMs) [21, 49, 97] are able to match and even surpass the visual quality of GAN-generated images.

Generated Image Detection. There is a continuing arms race for effective detection techniques and newer generations of image synthesis algorithms. Broadly speaking, generated image detection techniques can be divided into two categories: methods that rely on handcrafted features and learning-based methods. Methods from the first category either exploit visual defects (e.g., facial inconsistencies [65], impossible reflections [50], irregular pupil shapes [44])

¹Since we collected our main dataset before Twitter’s rebranding to X, we use the platform’s former name throughout this work.

or “invisible” characteristics such as frequency artifacts [10, 11, 23, 33, 92, 121], pixel statistics [66, 70], or model-specific properties [64, 87, 118]. Learning-based methods, on the other hand, use neural networks to learn a suitable feature representation to distinguish fake from real images [9, 15, 16, 41, 51, 63, 79, 108].

3 METHODOLOGY

A large-scale study on generated images in the wild comes with multiple challenges. First, we do not know the ground truth. As a result, it is difficult to estimate the amount of overlooked generated images (*false negatives*) and to be sure that an image detected as generated is actually generated (*precision*). Finally, studying millions of images comes with a computational overhead so that the detection method has to be efficient, too. We discuss further challenges and limitations of our study in Section 8. To deal with all these challenges, we carefully design a multi-step detection pipeline. The following is a step-by-step description of this pipeline. Note that while the presented approach is applied to Twitter, our method can be adapted to any other social network. We provide implementation details in Appendix A.

3.1 Data Collection

We describe the four types of datasets that we use for studying generated images, with Twitter being our use case. Table 1 summarizes our notation.

In-The-Wild Dataset \mathcal{D}_W^{\times} . To estimate the prevalence of generated images on a social network, it is important to obtain a mostly unconditional sample. In the case of Twitter, this can be achieved by using the API endpoint that provides real-time access to a random 1 % subset of all publicly posted tweets. We download each author’s profile image together with their profile metadata (cf. Appendix A.1 for an overview). Note that this approach only enables us to obtain profile images from users who write posts during the data collection period. Additionally, we omit users who have not set a profile image, that is, who are using Twitter’s default profile image. From March 7 to March 15 2023, we collected 14 989 385 profile images.

Labeled Datasets $\mathcal{D}_R/\mathcal{D}_F$ and Variations. We continue with labeled datasets of fake and real images which can be used to train a detector. As discussed in Section 1, existing observations suggest that the vast majority of generated profile images on Twitter are taken from TPDNE, which generates images with StyleGAN2 [56] trained on the FFHQ [55] dataset². We therefore decide to focus on this specific kind of fake faces and use 10 000 images from TPDNE as our fake-labeled dataset (denoted by \mathcal{D}_F) and correspondingly 10 000 images from FFHQ as our real-labeled dataset (denoted by \mathcal{D}_R). We discuss this limitation of focusing on TPDNE in Section 8. As prior work shows that processing operations like resizing and compression can affect the detection [61, 80], we consider two dataset variations:

- \mathcal{D}_R^{\times} and \mathcal{D}_F^{\times} . To obtain profile images with the social network’s processing steps, we adapt the approach from Boato et al. [7]. We upload both \mathcal{D}_R and \mathcal{D}_F to Twitter, set each

Table 1: Dataset notation. The symbol \times indicates that images were processed by Twitter.

SYMBOL	DESCRIPTION
\mathcal{D}_W^{\times}	Unlabeled dataset of Twitter profile images.
\mathcal{D}_F	Labeled dataset of fake images.
\mathcal{D}_F^{\times}	Labeled dataset of fake images uploaded as profile image and downloaded afterward.
$\mathcal{D}_F^{\times'}$	Version of \mathcal{D}_F^{\times} where images are zoomed into during upload.
\mathcal{D}_R	Labeled dataset of real images.
\mathcal{D}_R^{\times}	Labeled dataset of real images uploaded as profile image and downloaded afterward.
$\mathcal{D}_R^{\times'}$	Version of \mathcal{D}_R^{\times} where images are zoomed into during upload.
\mathcal{D}_P^{\times}	Proxy-labeled dataset of supposedly real Twitter profile images.
\mathcal{D}_D^{\times}	Labeled dataset of documented fake Twitter profile images.

image as profile image, and then download all images again. We denote these processed images by \mathcal{D}_R^{\times} and \mathcal{D}_F^{\times} .

- $\mathcal{D}_R^{\times'}$ and $\mathcal{D}_F^{\times'}$. We additionally simulate a user who zooms into the profile image during the upload, as it is common for social media platforms. We denote these images by $\mathcal{D}_R^{\times'}$ and $\mathcal{D}_F^{\times'}$, respectively.

We confirm in Section 7.2 that considering the preprocessing indeed improves the detection performance under realistic conditions.

Proxy-Labeled Real Dataset \mathcal{D}_P^{\times} . Social media platforms often have very popular users with a lot of followers. These popular users are rather unlikely to use deceptive fake images. Hence, we can build a proxy-labeled dataset with presumably *real* images. In particular, we select 10 000 profile images from the accounts in \mathcal{D}_W^{\times} with the highest numbers of followers that also pass our pre-filter (which is presented in the next section). We denote the so-created proxy-labeled dataset of real profile images by \mathcal{D}_P^{\times} .

Documented Fakes Dataset \mathcal{D}_D^{\times} . Finally, there are documented cases of generated profile images that were discovered manually. For example, blog posts regularly report such images when analyzing inauthentic Twitter accounts [78]. These cases can be used to build a labeled dataset of fake images in the wild, which we denote by \mathcal{D}_D^{\times} . Such a dataset is not free of bias, but provides a good means to finally check the performance of our classifier on an independent source. For our study, we use a dataset of 1 420 generated Twitter profile images that were manually collected between November 2022 and January 2024 [113].

3.2 Detection

Equipped with these different datasets, we can proceed with the detection of generated profile images. Here, we propose a two-stage procedure to improve the accuracy and the efficiency.

Pre-Filter ϕ . We start with a pre-filter ϕ to discard irrelevant samples. In our case, we can discard images without any face or where the face is too small. We use the efficient BlazeFace [4] face detector to detect faces and locate facial landmarks. An image passes ϕ if at least one face is detected and the Euclidean distance between the normalized coordinates of both eyes is greater or equal to 0.1.

²When published in 2019, TPDNE used the original StyleGAN [55], but switched to StyleGAN2 [56] shortly after its release.

The pre-filter serves two purposes: First, the overall computational complexity decreases by reducing the number of analyzed candidates in the subsequent, more demanding detection stage. Second, the detection stage is trained on facial images, so that other types of profile images, such as logos or monochrome images, could be wrongly classified as fake. Filtering irrelevant images can therefore decrease the false positive rate (FPR).

Classifier C. To automatically label a profile image as real or fake, we use a state-of-the-art CNN detector based on ResNet-50 [47]. Previous work [16, 17, 41, 61, 108] has demonstrated that this model is able to effectively distinguish real from generated images and that it provides good generalization capabilities. We initially attempted to use pre-trained fake image detectors, however, we found that the heavy pre-processing performed by Twitter makes it necessary to train our own detector (cf. Section 7.1). In particular, we train on the combination of \mathcal{D}_R^{\times} and \mathcal{D}_P^{\times} for real images, and \mathcal{D}_F^{\times} for fake images. The resulting final classifier is denoted by $C_{R^{\times}, P^{\times} | F^{\times}}$. Note that we experiment with using other dataset variations to train a classifier in our ablation study in Section 7.2. Yet, using processed real, fake, and proxy-labeled real images provides the highest performance for processed and zoomed inputs.

3.3 Assistance for Manual Labeling

To estimate error rates of our detection scheme on unlabeled in-the-wild data, it is necessary to manually label these images as real or fake. As generated images have reached a level of quality which makes them almost indistinguishable from real images [34, 51, 59, 71, 72, 95, 103], we use two tools to facilitate this process.

Alignment. Faces generated by StyleGAN2 [56] are characterized by being almost perfectly aligned with respect to their facial landmarks, caused by the alignment of the training dataset FFHQ. By superimposing multiple images, this characteristic has been leveraged to visually identify clusters of fake accounts in social networks [37, 42, 43, 73, 75, 99, 101]. We automate this manual process by extracting facial landmarks with BlazeFace [4] and computing the deviation from a reference. For each landmark L_1, \dots, L_{12} (x - and y -coordinates of eyes, ears, mouth, and nose), we compute its mean μ_i and standard deviation σ_i over a reference dataset. In our study, we use the training subset of \mathcal{D}_F^{\times} as reference. We define an image x as being aligned, if the condition

$$|L_i(x) - \mu_i| < k\sigma_i \quad \forall i \in \{1, \dots, 12\} \quad (1)$$

holds, where $L_i(x)$ are the landmarks extracted from the image x and $k \in \mathbb{Z}$ controls the maximum deviation from the reference. We set $k = 7$. During our evaluation in Section 4, we find that this is the lowest value at which all generated images in the validation set of \mathcal{D}_F are aligned. While a close alignment hints towards a face generated by StyleGAN2, it is ineffective if the image has been cropped or geometrically transformed.

Inversion. Additionally, we leverage GAN inversion [111] as an assistance tool. For a given input image, this method finds the latent code which reconstructs the original input when passed through the generator. We use the implementation by Karras et al. [56] to invert images using StyleGAN2. Previous work has shown that generated images can be reconstructed more successfully than real

Table 2: Evaluation of our pre-filter ϕ . We separately analyze its two conditions, which are the presence of a face and its sufficient size.

Dataset	With Face		Without Face	
	Face Detected	Size Check	Face Detected	Size Check
\mathcal{D}_F^{\times}	100 %	100 %	—	—
\mathcal{D}_W^{\times}	92.47 %	58.16 %	42.53 %	30.27 %

images [1, 56, 81] (we provide a visual example in Appendix B). Note that inversion also relies on facial alignment. If an adversary uses a cropped version of a fake face, the inversion result will be distorted. We therefore only use inversion as labeling assistance if the image is aligned.

4 EVALUATION

In this section, we proceed with an evaluation of our proposed methodology in a controlled setting with labeled data. This allows us to verify the components of our detection pipeline before studying generated faces in the wild on Twitter in Section 5 and analyzing the corresponding profiles and tweets in Section 6.

Dataset Splits. We randomly split \mathcal{D}_R , \mathcal{D}_F , and \mathcal{D}_P^{\times} into 8 500 train, 500 validation, and 1 000 test images, respectively. \mathcal{D}_R^{\times} and \mathcal{D}_F^{\times} are split in the same manner. As we use \mathcal{D}_R^{\times} and \mathcal{D}_F^{\times} only for evaluation, they only contain the corresponding 1 000 test images, respectively.

Pre-Filter. We start with the pre-filter ϕ that should discard irrelevant images, but keep potentially generated images. An image passes ϕ if a face (a) is detected and (b) has a sufficient size (see Section 3.2). In the following, we apply ϕ to the test set from \mathcal{D}_F^{\times} and to 1 000 randomly sampled images from \mathcal{D}_W^{\times} . For the latter subset, we manually label each image whether it (partly) contains a human face. Our experiment here has three goals: we want to verify that all generated images from \mathcal{D}_F^{\times} pass ϕ , confirm that the face detector works reliably on the in-the-wild images from \mathcal{D}_W^{\times} , and finally get an estimate of the number of kept in-the-wild images passed to the next stage.

Table 2 shows the results for our evaluation of ϕ . All generated images from \mathcal{D}_F^{\times} pass ϕ , fulfilling our first goal. Among the sampled in-the-wild images from \mathcal{D}_W^{\times} with a face, the face detector correctly identifies 92.47 %. We manually look through the undetected faces. In most cases, the face is either very small, obstructed (e.g., by masks or smartphones), or partly outside the frame. The face in these images is not prevalent, so that we consider it acceptable to skip them. For in-the-wild images without a face, the face detector mistakenly locates a face in 42.53 % of the cases. We manually inspect the mislabeled images. The vast majority contains faces, but they are drawn, digitally created, or belong to animals or statues. Only very few detections are obviously “wrong”, such as images with Twitter’s former default profile image. Since these images are just passed to the next stage, having some false positives is not critical. Based on this analysis, we can conclude that the

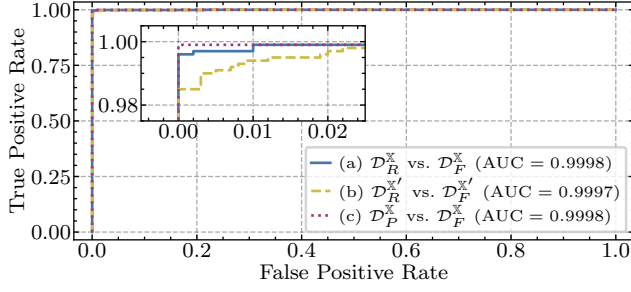


Figure 1: Evaluation of our classifier $C_{R^x, P^x/F^x}$. We show the ROC curve under different conditions.

face detector reliably works, fulfilling our second goal. Finally, we measure how many images from the subset of \mathcal{D}_W^x additionally pass the size check and therefore ϕ . In only 58.16 % of the face images and 30.27 % of the non-face images, the face is considered large enough, considerably reducing the number of images passed to the next stage. Overall, we conclude that our pre-filter allows us to skip irrelevant images efficiently, without mistakenly discarding generated faces.

Classification. Next, we verify that our classifier $C_{R^x, P^x/F^x}$ is capable of spotting generated images in realistic settings. We evaluate the performance of our classifier under three conditions: (a) processed images (\mathcal{D}_R^x vs. \mathcal{D}_F^x), (b) zoomed images ($\mathcal{D}_R^{x'}$ vs. $\mathcal{D}_F^{x'}$), and (c) proxy-labeled real and fake images (\mathcal{D}_P^x vs. \mathcal{D}_F^x). We use the test set from each dataset.

Figure 1 shows the respective ROC curves. Our classifier has an almost perfect detection rate with an AUC value close to 1.0. Note that the setup on zoomed data is slightly more challenging, because there are no examples of zoomed images in the detector’s training data. Still, the error rate remains very small. Due to the strong class imbalance on real Twitter data, a small error rate is required to avoid an excessive amount of false positives.

Assistance Tools. We finally verify our methods that allow us to better label images for the error-rate estimation later.

Alignment. Using the test set from \mathcal{D}_F^x , we confirm that with $k = 7$, all fake images are correctly labeled as being aligned. From the 1 000 randomly sampled images from \mathcal{D}_W^x , only 35 are aligned.

Inversion. We first verify that generated images can be inverted more accurately than real images. We invert 500 images from \mathcal{D}_R^x and \mathcal{D}_F^x , respectively, and compute the LPIPS [120] distance between original and reconstructed images. This distance metric measures the perceptual similarity between two images and has been previously used to estimate the reconstruction quality [56]. The histograms in Figure 2 show that the reconstructions from \mathcal{D}_F^x are perceptually more similar to the originals compared to the reconstructions from \mathcal{D}_R^x . A classification based on the LPIPS scores results in an AUC of 0.97.

As second experiment, we check that inversion is helpful for manual labeling. We divide the 1 000 images (500 real, 500 fake) into 900 training and 100 test images. For each image we construct a side-by-side view with the original, its reconstruction obtained




Figure 2: Evaluation of GAN inversion. The lower the LPIPS distance between the original image and its reconstruction, the more similar they are.

by inversion, and the distance measured in LPIPS and MSE (cf. Appendix B). Using the training set, one annotator practices the manual classification. We then evaluate the performance based on the held-out test images. 99 out of 100 images are correctly assigned, demonstrating a feasible manual inspection.

We emphasize that images from \mathcal{D}_R^x are very similar to images from \mathcal{D}_F^x . In contrast, most in-the-wild profile images are visually different, leading to even worse reconstructions (despite being aligned) and thus to comparatively high LPIPS values. Hence, we expect the actual manual labeling process to be easier than in the controlled setting.

Summary. Our evaluation indicates a valid detection pipeline. The pre-filter allows skipping irrelevant images while the classifier allows detecting generated images. The assistance tools can help with the manual labeling process.

5 DETECTING GENERATED IMAGES IN THE WILD

Equipped with a valid detection pipeline, we can now explore the prevalence of generated images on Twitter. To this end, we first need to calibrate the detection to the real-world setup before we can present the final results.

Manual Labeling. To begin with, we have to label a subset of \mathcal{D}_W^x . First, this allows us to get a detection threshold, so that we are able to classify individual images as real or generated. Note that in the controlled setup before, we evaluated the overall performance of $C_{R^x, P^x/F^x}$ with the AUC metric that takes into account all possible thresholds and thus does not require picking a specific value. Second, a separate labeled set is necessary to estimate error rates.

Unfortunately, manual labeling of all samples in \mathcal{D}_W^x is unfeasible due to the sheer volume of samples within the dataset. Thus, we resort to a random subset of 1 498 938 images, corresponding to 10 % of all samples. We then sort these images based on their score (from $C_{R^x, P^x/F^x}$) from low (real) to high (fake) and select the top 1 000 images that pass ϕ for manual labeling. We acknowledge that choosing the subset based on the classifier that we are trying to evaluate introduces an unwanted bias: there could be fake images with very low scores that are overlooked. However, we argue that this approach strikes a balance between practicability and a sound




Figure 3: Precision-recall curve of $C_{R^x, P^x/F^x}$ on the validation set of the manually labeled images. The circle marks the selected threshold, which maximizes the F1-score.

estimation. Selecting the subset by pure chance would require an enormous amount of manual labeling to gather a sufficient number of fake images. Moreover, the scores of our subset range from 1.0 to 0.33. From the test set of \mathcal{D}_F^x , only 3 out of 1000 images get a score below 0.33. We therefore assume that only a very small number of false negatives is potentially overlooked. Our evaluation on the independent dataset \mathcal{D}_D^x (at the end of this section) also confirms that our approach provides a reliable estimate of the error rates.

We carefully inspect each image and, if it is aligned, its reconstruction from GAN inversion. We label an image as real if the framing and pose do not match with that of \mathcal{D}_F , if it contains a complex and meaningful background, or if the reconstruction deviates significantly from the original. In contrast, images are labeled as fake if they contain diffuse backgrounds, asymmetries (eyes, ear-rings), unnatural clothing, color artifacts, and/or an almost perfect reconstruction. By doing so, we obtain 185 images labeled as “Real”, 725 images labeled as “Fake”, and 90 images labeled as “Unsure”. Most images with label “Unsure” resemble images from TPDNE, but do not contain clear artifacts or were strongly edited. We also assigned this label if we suspect that an image was generated using a different kind of generative model. We randomly split the 910 images labeled as “Real” or “Fake” into a validation set (for calibrating the threshold) and a test set (for estimating the error rates) of equal size, maintaining the label ratio in both splits.

Choosing a Threshold. Due to the high imbalance between real and generated images in \mathcal{D}_W^x , choosing an appropriate threshold is not trivial. A too high threshold leads to many overlooked fake images (low recall), while a too low threshold leads to many real images classified as fake (low precision). Figure 3 depicts the precision-recall curve based on the predictions of $C_{R^x, P^x/F^x}$ on the validation set of our manually labeled images. As recall and precision are equally relevant in our setting, we follow the common practice to select the threshold based on the F1-score. The best F1-score (0.9832) is achieved using a threshold of 0.9899361. Such a high threshold might appear counterintuitive. Figure 4 shows that most fake images are confidently classified as fake—with scores very close or equal to 1. The scores of real images, however, have greater variation. Thus, choosing a relatively high threshold gives the best performance. Note that the scores of real images in this subset are not representative for all real images, since we purposely selected images with high scores.




Figure 4: Score distribution of manually labeled images.




Figure 5: Examples of fake images falsely classified as real, together with their classification score.

Estimating Error Rates. Equipped with our selected threshold, we can now estimate the error rates of our detector. We consider two metrics to be most relevant in our setting: the false negative rate (FNR), which is the fraction of mislabeled fake images among all fake images, and the false discovery rate (FDR), which is the fraction of mislabeled real images among *all images classified as fake*. Note that the FDR is not to be confused with the FPR, which is the fraction of mislabeled real images among *all real images*.

Based on the test set of our manually labeled subset, the FNR is 3.03% and the FDR is 1.4%. To understand the errors, we take a closer look on the misclassified images. Figure 5 shows the false negatives within the test set together with their scores. Although the majority actually gets a high score and is only classified as real due to the high threshold, three images have a considerably lower score. We cannot identify a pattern which causes their misclassification. Neither do we observe any characteristics that would explain the real images classified as fake (false discoveries). As these profile images are real users, we cannot provide visual examples here.

For a second estimate, we leverage the independent dataset \mathcal{D}_D^x of fake profile images that were spotted by users on the web before. We obtain a low FNR of 2.75%, that is, 39 out of 1420 fake profile images are incorrectly labeled as real. This result is similar to our previously estimated FNR of 3.03%. Note that all images in \mathcal{D}_D^x pass our prefilter ϕ . False negatives therefore only depend on the classifier’s score.

Overall, we can confirm the performance of our detector on two different test sets. While the error rates are not zero, they are small enough to draw conclusions in our analysis in the next section.

Prevalence of Fake Profile Images on Twitter. We are now ready for the final step. We apply our detection scheme on the entire in-the-wild dataset \mathcal{D}_W^X . The pre-filter ϕ discards 8 535 322 images, reducing the number of images by 56.94%. Next, using our detector, we classify 7 723 profile images as fake. This is 0.052% of the full dataset. In the next section, we analyze the profiles behind these images and their tweets in more detail.

6 ANALYSIS

Our goal in this section is to understand the context where the generated profile images are used. To this end, we first perform an analysis of the *accounts* behind these images (Section 6.1). Then, we thoroughly analyze the *content of the tweets* that were sent from these accounts (Section 6.2). For simplicity, we refer to accounts using generated profile images as “fake-image accounts” as opposed to “real-image accounts” in the following.

6.1 User Metrics

We begin by analyzing the difference between fake-image and real-image accounts regarding social connections, account activity, as well as account creation and status.

Social Connections. On Twitter, social interactions are primarily measured in the number of followers an account has and the number of other accounts it follows. Figures 6a and 6b visualize the distribution of these metrics for real- and fake-image accounts at the time of data collection. We find that fake-image accounts have fewer followers (mean: 393.35, median: 60) compared to real-image accounts (mean: 5 086.38, median: 165) in our dataset. This is in accordance with previous analyses of fake social media profiles [5, 19, 45, 96]. 1 997 (25.86%) of all fake-image accounts have 9 or fewer followers and 1 063 (13.76%) have exactly zero followers. We notice that 1 996 fake-image accounts (25.84%) have exactly 106 followers. Our content analysis in Section 6.2 reveals that these accounts belong to a large cluster of fake accounts involved in coordinated inauthentic behavior.

We find that fake-image accounts also follow fewer other accounts (mean: 283.18, median: 21) compared to real-image accounts (mean: 759.83, median: 262). Interestingly, 2 175 fake-image accounts (28.16%) follow exactly two other accounts. In contrast to the number of followers, a relatively small number of fake-image accounts (163, 2.11%) follows exactly zero other accounts.

Activity. Figure 6c shows that fake-image accounts do participate in Twitter based on the number of tweets. Yet, they are overall less active than real-image accounts. On average, fake-image accounts posted 3 158.9 (median: 112) tweets, as opposed to 17 096.39 (median: 3450) tweets from real-image accounts. 1 948 (25.22%) of all fake-image accounts have 10 or fewer tweets. In addition, Figure 6d shows the average number of tweets per day, calculated by dividing the total number of tweets by the number of days the account exists. Based on the median, fake-image accounts are still less active than real-image accounts (0.95 vs. 3.68 tweets per day). However, a large fraction of fake-image accounts posts exceptionally many tweets per day, causing a higher mean (12.89 vs. 11.47 tweets per day). In particular, there are 503 fake-image accounts (6.51%) that submitted more than 50 tweets per day.

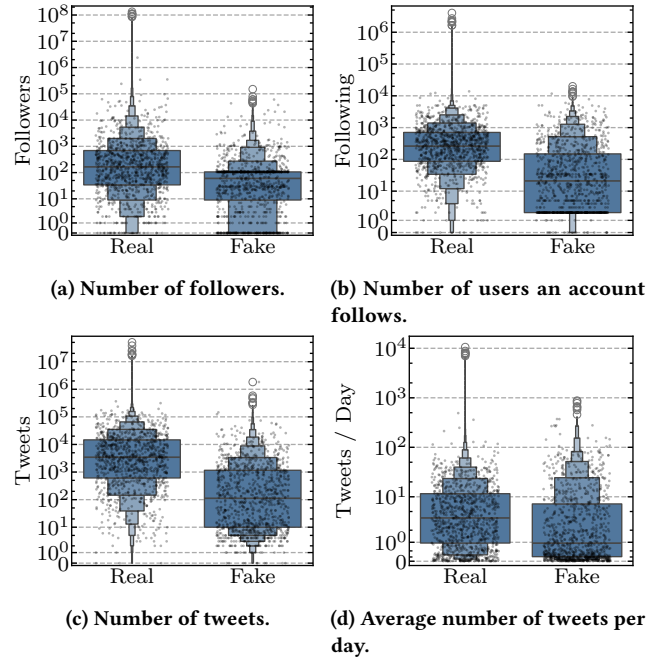


Figure 6: Distributions of user metrics from real- and fake-image accounts. The points depict 1 000 randomly selected samples from each class, respectively.

Account Creation and Status. Figure 7a compares the times of account creation. Fake-image accounts are considerably “younger”, with more than half of them (52.38%) being created in 2023 (note that our data collection happened in March 2023). In contrast, only 6.22% of real-image accounts have been created in this period.

In addition to the creation date, we also examine the account status after a certain period of time. We checked the status of all 7 723 fake-image accounts nine months after data collection by querying the respective profile page. As a reference, we did the same for an equal number of randomly sampled real-image accounts. Accounts can be either active, deactivated (by the user), or suspended (by Twitter). Figure 8 illustrates that more than half of the fake-image accounts (52.07%) have been suspended. In contrast, only 5.01% of real-image accounts in the reference set have been suspended. The high number of suspended fake-image accounts suggests that they were violating Twitter’s rules.

In Figures 7b and 7c, we analyze the account creation of fake-image accounts given their status. We observe various suspended accounts that were created in bulk just shortly before our data collection, especially in the middle of February. Note that we do not know when these accounts were suspended, so that we cannot determine the effective lifetime of these accounts.

Takeaways. Our analysis shows that real and fake-image accounts notably differ. Fake-image accounts have fewer social interactions, both regarding the number of followers and the number of accounts they follow. While these metrics are distributed evenly for real-image accounts, we observe patterns with fake-image accounts.

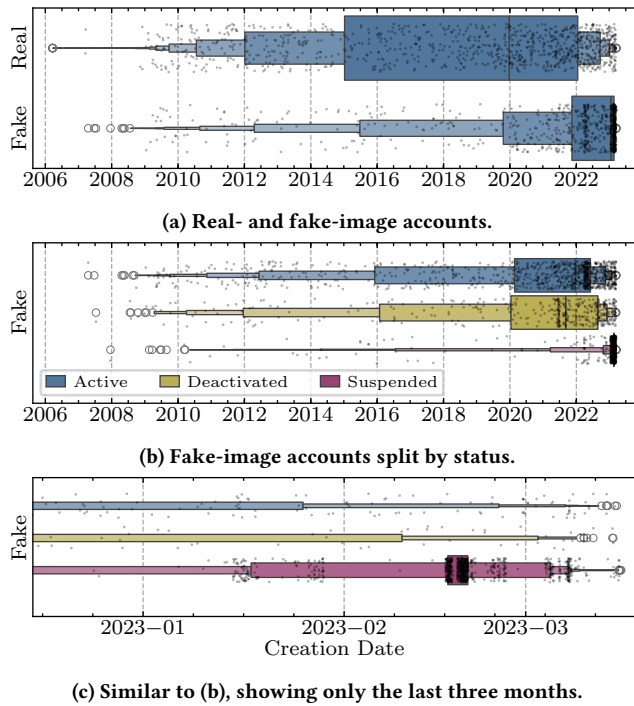


Figure 7: Distributions of account creation times from real- and fake-image accounts. In (b) and (c) we differentiate fake-image accounts by their status nine months after data collection. The points depict up to 1000 randomly selected samples for each label and status, respectively.




Figure 8: Account status nine months after data collection.

There are large groups with identical values, indicating an orchestrated network of inauthentic users. Moreover, fake-image accounts are not passive, they considerably participate in Twitter based on the number of tweets. Although they are in general less active than real-image accounts, there are several fake-image accounts that post very frequently, hinting towards spamming attacks. Finally, fake-image accounts have a more limited lifetime. They are usually created more recently than real-image accounts, and they are also disproportionately often suspended by Twitter. This suggests inauthentic behavior. Moreover, a substantial number was created in bulk just before our data collection period started. This *bulk creation* (or batch creation) is a common pattern for inauthentic behavior, used, for example, to amplify messages or to participate in spamming or trolling activities [31, 45].




Figure 9: Number of tweets (top) and corresponding accounts (bottom) for the 20 most represented languages. We differentiate accounts (and their tweets) by their status nine months after data collection, combining deactivated and suspended accounts as inactive accounts. Note that a single account can be associated with multiple languages.

6.2 Content Analysis

To evaluate the purpose of the identified fake-image accounts, we proceed to analyze their tweets (original as well as retweets) posted in 2023. We utilize data collected in the context of a large-scale Twitter stream archiving effort [26] based on Twitter’s 1% sampled stream (the same we used to create \mathcal{D}_W^X). This allows us to access information about the activity of the profiles before and after the profile collection week (until Twitter restricted access to its API in June 2023). In total, we have access to 111 165 tweets from the 7 723 fake-image accounts in our collection.

We begin our analysis with the *language and account status*. The upper half of Figure 9 shows a breakdown of the number of tweets per language. Using the accounts’ status nine months after our data collection (cf. Section 6.1), we can also calculate the fraction of tweets stemming from accounts that have been deactivated or suspended. For the remaining part of this section we refer to this group as inactive accounts. Overall, 49.6% of all tweets come from inactive accounts. Interestingly, Turkish and Arabic stand out as languages with significantly higher rates of inactive accounts (87.95% and 95.59%, respectively) than other languages. The number of unique accounts that created the tweets in each language are reported in the lower half of Figure 9. It shows that Turkish tweets, for instance, stem from a relatively small number of users.

We proceed with a *textual analysis*. To identify structural patterns, we employ state-of-the-art sentence embeddings [86] to group the tweet texts into semantically related clusters. We utilize the cosine similarity between the sentence embeddings to determine cluster belonging. A new observation (tweet) is assigned to an existing cluster if a certain similarity threshold (in our case 0.6) is reached. Otherwise, a new cluster will be generated. Furthermore, we limit our analysis on clusters that exhibit a minimum cluster size of 50 (i.e., at least 50 tweets should be in one cluster). This approach allows us to identify dominant trends. Note that it does not provide

a distribution of topics, because not every tweet is assigned to a cluster. For the purpose of visualization, we use UMAP [67] as a dimensionality reduction technique to generate a two-dimensional representation of the clustering outcome (cf. Figure 10). For each cluster, we calculate the class-based term frequency–inverse document frequency (TF-IDF) terms to determine representative class tokens. In a subsequent step, we conduct a manual qualitative review of all clusters to identify and describe common themes, which are detailed in the following paragraphs. For translating Turkish and Arabic tweets we use machine translation and double check unclear phrases with native speakers. We describe the general cluster contents and provide representative examples for important topics. We also analyze the metadata of accounts within cluster and report unusual characteristics.

English (inactive accounts). The clustering for English content posted by users that are not on the platform anymore reveals a notable pattern: we observe a single, extremely large cluster that encompasses 49.67 % of all tweets from this group of accounts. Despite the variability of the actual content, these tweets all share a common structure. Each tweet begins by mentioning a specific Twitter user, followed by a short sequence of English terms. Interestingly, these sequences do not form logical sentences, so they are neither semantically nor syntactically correct. Each of these sentences is then followed by a specific Chinese hashtag that can be translated to: “This is really useful”. Unfortunately, we can only speculate about their purpose. Our hypothesis is that the embedded hyperlinks within the tweets may have directed users to malicious external websites. As the links are no longer functional, we cannot verify this hypothesis.

The accounts’ metadata corroborate the assumption that the 1579 accounts within this cluster were part of an organized network. All but three were created between February 16 and February 20, which is consistent with our observations in Figure 7c. Up to 754 accounts were created on a single day. We also find that this cluster contributes to the large number of accounts with identical social connections (cf. Figure 6a and Figure 6b). 94.93 % have exactly 106 followers and 95.31 % follow exactly two other accounts. The usernames (Twitter handles) appear to be constructed from a list of German-sounding first names and last names (or initials), and optionally one or multiple digits (e.g., @GuntherForstner86). 67.44 % of all accounts share the same Arabic display name that can be translated to “Noon Namshi Sivvi discount code is strong and effective” (Noon and Namshi are e-commerce platforms operating in the Arabic region). These accounts also have their location set to “KSA” (Kingdom of Saudi Arabia). Moreover, the accounts contain nonsense descriptions like “Personal west service street laugh small.”. We hypothesize that these were automatically generated or translated. Again, we can only speculate about the reason, especially about the mixed use of English, German, Arabic, and Chinese language. We also notice that 99.18 % of all accounts in this cluster use profile images that are duplicates within our dataset of fake-image accounts. Appendix D elaborates our method for identifying duplicate images.

The remaining clusters mainly focus on giveaways, often related to cryptocurrencies, with tweets like

*\$50 (2 winners x \$25) 24 hours - like, follow)
i will #giveaway 100 usdt worth of \$loop as we celebrate our 10k milestone*

or the promotion of illegal content such as links to broadcasting streams of soccer matches, e.g.,

live stream arouca vs benfica live [link]

Another trend is the distribution of links to websites and Telegram groups containing explicit content, e.g.,

follow for more [link]

English (active accounts). The clustering of tweets from users who were still active after nine months reveals similarities and differences. Figure 10 depicts a visual representation of the top clusters with their representative text tokens. A significant portion of all clusters is again related to various forms of cryptocurrency, stocks, and giveaways, e.g.,

*drop your #tezos #nft if you need it sold!
15000\$ in \$eth — 5 lucky winners!*

Additionally, we find a significant share of adult content/porn related clusters, actively advertising explicit content, also through dedicated patterns like

beautiful/charming/etc. [profile of porn actress] [link]

Compared to inactive users, we observe that active accounts also engage in discussions on contentious or political issues. These include, for example, the war in Ukraine, election-related discourse, and debates on COVID and vaccinations:

*welcome to nazi ukraine #russia
desantis racks up wins while trump, potential 2024
opponents take swipes at florida governor
albos crocodile tears: watch this video, that the mainstream media refuses to show.
someone needs to find an antidote for the vaxxx*

Turkish (inactive accounts). For the Turkish accounts, we restrict our analysis to content posted by inactive users, since this is the majority of the dataset. Our findings indicate that almost all of this content is related to pornography or escort services. The primary distinction among the clusters are the cities mentioned within the posts. Most tweets also contain links to other websites, which are no longer functional. Upon examining the metadata of all 932 accounts, we again identify the systematic pattern for usernames that we already observed in the large cluster of English tweets. However, first and last names appear to be of Turkish descent. Moreover, almost all accounts have their location set to a real Turkish city. 46.78 % of all accounts again use duplicate profile images and 85.52 % were created within one month. These findings again indicate that at least some systematic approach (automatic or semi-automatic) is used to generate the accounts.

Arabic (inactive accounts). For Arabic tweets, we again only consider accounts that have been deactivated or suspended. All clustered tweets appear to be related to literature, with individual clusters being characterized by mentions of certain authors, countries, or topics—all related to the Arabic region. These tweets make up 72.34 % of all Arabic tweets from inactive accounts. Surprisingly, the tweets share a common structure with those from the large




Figure 10: UMAP representation of English tweets posted by active users. Distinct (groups of) clusters are annotated by their most representative tokens. Different clusters are separated by color.

cluster of English tweets: they contain the specific Chinese hashtag, an external link, and an incoherent sentence. Our metadata analysis suggests that the 1 806 accounts indeed belong to the same cluster, despite the different language. Almost all accounts were created between February 16 and February 20, with 892 being created on a single day. We observe the same anomalies regarding the (German) usernames, locations, descriptions, and social connections. Given the book-related content and the frequently occurring username that promotes a discount code, we hypothesize that the external links might have referred to the respective shopping platforms.

Takeaways. Our content analysis reveals that English, Turkish, and Arabic are the dominant languages used by the fake-image accounts in our collection. We identify large networks of fake-image accounts that were probably automatically created and that participated in large-scale spamming attacks. We observe recurring patterns as part of the automation. Accounts are created in bulk. Tweets, usernames, locations, descriptions, and social connections follow a systematic pattern. Multiple accounts within a network share the same profile image. Furthermore, our analysis shows that frequently occurring topics are cryptocurrencies, giveaways, and content related to pornography and escort services. Fake-image accounts also participate in controversial political discussions. These findings align with prior analyses of inauthentic content on Twitter [18, 77, 83, 85].

6.3 Sample Study on Active Accounts

Finally, we analyze the current behavior of fake-image accounts that are still active at the time of writing (February–March 2024). This gives insights about the use-case of rather long-term fake-accounts. As we cannot use data from Twitter’s API any longer, we




Figure 11: Topic distribution of 1 000 manually inspected active accounts.

randomly select 1 000 active fake-image accounts and visit their Twitter profile manually. Two annotators independently check the most recent tweets and assign a topic to each profile (Cohen’s kappa: 0.84). Accounts where both annotators disagree are revisited. We choose topics from five categories, so that we can get a broad understanding of the prevalent application scenarios.

Figure 11 depicts the distribution of topics. The majority of fake-image accounts participates in the political discourse (36.1%) or shares finance-related content (33.1%), mostly related to cryptocurrencies. 5.1% of the profiles revolve around other websites or products (“Business”), while 4.3% share explicit content or promote escort services (“Sex”). The remaining accounts (21.4%) cover diverse topics or have an empty timeline. Taken together, we observe similar topics as before in our cluster analysis.

6.4 Summary

Our systematic analysis revealed 7 723 Twitter accounts that use AI-generated profile images. By analyzing both their user metrics and the content of their tweets, we identify particular patterns. Some of

Table 3: Evaluation of existing pre-trained detectors. We report the AUCs under different conditions.

Condition	C_{Wang}	C_{Grag}	C_{Ojha}	$C_{R^{\mathbb{X}}, P^{\mathbb{X}}/F^{\mathbb{X}}}$
(a) $\mathcal{D}_R^{\mathbb{X}}$ vs. $\mathcal{D}_F^{\mathbb{X}}$	0.7279	0.9249	0.6405	0.9998
(b) $\mathcal{D}_R^{\mathbb{X}'}$ vs. $\mathcal{D}_F^{\mathbb{X}'}$	0.7243	0.9600	0.6338	0.9997
(c) $\mathcal{D}_P^{\mathbb{X}}$ vs. $\mathcal{D}_F^{\mathbb{X}}$	0.8713	0.9015	0.6922	0.9998
(d) \mathcal{D}_R vs. \mathcal{D}_F	0.9466	1.0000	0.8296	—

these patterns, like the high number of suspended accounts, striking similarities within the accounts’ properties, or the multitude of similar tweets posted by different users, represent strong evidence that a subset of these accounts is part of organized, inauthentic networks. While many accounts amplify content related to cryptocurrencies or pornography, we also observe accounts that express controversial political opinions.

7 ABLATION STUDY

Before finishing our study, we shortly confirm the design choices of the classification methodology proposed in Section 3. In particular, we justify the need to train our own classifier (Section 7.1) and study the impact of training data (Section 7.2).

7.1 Evaluation of Pre-Trained Detectors

Detecting GAN-generated images is a well-researched problem and several pre-trained detectors have been proposed (cf. Section 2). However, we observe that the performance of these detectors suffers from Twitter’s image processing, making it necessary to directly train a classifier on processed profile images.

Setup. We test three existing pre-trained classifiers: C_{Wang} [108] (which is the basis of our classifier), C_{Grag} [41], and C_{Ojha} [79]. Appendix C provides more details on these three classifiers. We evaluate four conditions (see Table 3). The conditions (a)-(c) correspond to those in Figure 1 and all use images processed by Twitter. We additionally test the pre-trained detectors on unprocessed images in condition (d).

Results. Table 3 shows the AUCs of the three classifiers compared to $C_{R^{\mathbb{X}}, P^{\mathbb{X}}/F^{\mathbb{X}}}$. Our trained detector significantly outperforms the pre-trained classifiers under the Twitter conditions (a)-(c). The fact that the latter perform better under the clean condition (d) demonstrates the strong effect of Twitter’s processing. It is therefore not possible to use a pre-trained detector for our study of in-the-wild profile images.

7.2 Effect of Training Data

Our datasets described in Section 3.1 allow for different combinations of training data. In the following, we justify the choice of training our detector on real images, proxy-labeled real images, and fake images.

Setup. We consider three classifier variants and analyze their performance under the three conditions from Figure 1, respectively.

Table 4: Evaluation of three detector variants trained on different datasets. We report the AUCs under different conditions.

Condition	$C_{R/F}$	$C_{R^{\mathbb{X}}/F^{\mathbb{X}}}$	$C_{R^{\mathbb{X}}, P^{\mathbb{X}}/F^{\mathbb{X}}}$
(a) $\mathcal{D}_R^{\mathbb{X}}$ vs. $\mathcal{D}_F^{\mathbb{X}}$	0.9971	0.9998	0.9998
(b) $\mathcal{D}_R^{\mathbb{X}'}$ vs. $\mathcal{D}_F^{\mathbb{X}'}$	0.9953	0.9995	0.9997
(c) $\mathcal{D}_P^{\mathbb{X}}$ vs. $\mathcal{D}_F^{\mathbb{X}}$	0.9983	0.9994	0.9998

The classifier $C_{R/F}$ is trained on \mathcal{D}_R and \mathcal{D}_F and represents the most straightforward option. The images are not processed by Twitter, but we resize them to 400×400 pixels to match the resolution of actual profile images. The second classifier, $C_{R^{\mathbb{X}}/F^{\mathbb{X}}}$, is trained on the same images but *with* Twitter’s processing. Finally, $C_{R^{\mathbb{X}}, P^{\mathbb{X}}/F^{\mathbb{X}}}$ is additionally trained on $\mathcal{D}_P^{\mathbb{X}}$ as real images.

Results. Table 4 shows the AUCs of the three detector variants under the different conditions. Our finally chosen classifier, $C_{R^{\mathbb{X}}, P^{\mathbb{X}}/F^{\mathbb{X}}}$, has the highest performance in all conditions. The classifier $C_{R/F}$ trained on unprocessed images performs worse than both variants trained on processed images, confirming our findings from Section 7.1. Note that, while an AUC > 0.99 is still very high, the small difference can cause a significant increase in false positives, given the size of $\mathcal{D}_W^{\mathbb{X}}$. Overall, our results provide two insights. First, the classifier C should be trained on images that are processed similarly to the target images. Second, including proxy-labeled real images from the target distribution ($\mathcal{D}_P^{\mathbb{X}}$) improves the detection performance. A closer look shows that this causes a better separation of the classifier scores, shifting scores towards either end of the output range. This motivates our choice of $C_{R^{\mathbb{X}}, P^{\mathbb{X}}/F^{\mathbb{X}}}$ for our study.

8 DISCUSSION AND LIMITATIONS

Our work systematically examines the prevalence of generated images on Twitter. Despite great effort, our study has limitations that we discuss in the following.

Sampling Bias. The restricted 1% access to Twitter as well as the limited and randomly chosen collection period can introduce a sampling bias [2] to our study. Especially the presence of several large clusters with seemingly orchestrated accounts in our collected dataset has a significant effect on our analysis. These clusters and their concrete topics are expected to change over time. Nevertheless, the characteristics, such as the bulk creation of accounts, should generally apply. The same holds for high-level tendencies, such as political amplification or spamming. These are also in line with prior observations on Twitter misuse [18, 85]. Finally, we note a possible bias due to the restructuring of Twitter/ \mathbb{X} after the takeover by Elon Musk. It is possible that with the rise of hate speech and bots [48], the prevalence of generated profile images has also increased. Unfortunately, the current API limits impede a replication of our analysis.

Selection Bias in Analysis. A full analysis of all tweets is beyond the scope of our work. Thus, our cluster analysis is not exhaustive and only focuses on the prevalent trends. Still, this allows us to

identify the primary contexts in which generated images are used on Twitter, so that we can draw general conclusions on topics.

Focus on Images Generated from TPDNE. We focus on facial images from TPDNE that are generated by StyleGAN2 [56]. Although we are unable to provide statements regarding the prevalence of other types of generated images on Twitter, we expect to cover the most prevalent type. TPDNE has made it considerably easier to access generated images compared to other generative models. Several reports confirm that GAN-generated faces are in fact used by fake social media accounts [42, 43, 73–76, 98, 99, 101, 110]. Moreover, most alternative models need to be deployed locally. This requires technical knowledge and possibly specialized hardware. Although text-to-image models like Stable Diffusion or Midjourney can be accessed through a browser, generating images at scale may require significant time and additional costs. Achieving good images can require multiple attempts and services like Midjourney require payment. Finally, we note that detecting *all* kinds of generated images, especially in a real-world setting where images are heavily processed, is still an open challenge [15, 41]. Therefore, we focus on one setting where we aim at developing a highly reliable detector.

Likelihood of Overlooked Fake Profiles. As discussed in Section 5, classifying in-the-wild data always requires trading off the number of overlooked fakes against the number of falsely detected real images. While we make our best efforts to evaluate the performance of our detection pipeline under realistic conditions, we cannot exclude that the actual FNR is higher than our estimate. Fake profile images with an unusual processing could potentially bypass our detector. Furthermore, our tool-assisted manual labeling process is not guaranteed to be error-free. However, as the FNR on the independent dataset \mathcal{D}_D^{\times} closely matches our estimate, the likelihood of overlooked generated images should be low.

Applicability for Fake Account Detection. While our proposed detection pipeline can reliably identify accounts using generated profile images, this is naturally not sufficient for mitigating fake social media accounts in general. Not only does a generated profile image not necessarily imply that an account is fake, inauthentic accounts can simply use other image sources. State-of-the-art fake account detection methods usually rely only on account metadata and social interactions [114]. However, we believe that profile images can be an additional factor to decide whether an account is authentic or not. While stock photos or images taken from the Internet can be easily identified using reverse image search, generated images have the advantage (for the adversary) of being unique. Our work can contribute to closing this loophole.

9 RELATED WORK

Studying generated faces on social media touches different research areas. In the following, we examine related methods and concepts.

Detecting Generated Images on Social Media. Despite the plethora of proposed fake image detection methods (cf. Section 2), there exists only little work on the detection in real-world settings. Boato et al. [7] create a synthetic dataset of processed images by

sharing real and generated images on different social media platforms. They find that a classifier trained on “original” images is not able to effectively detect shared images, unless it is fine-tuned. This confirms our results in Section 7. In a related work [62], the same methodology is applied to deepfake videos, yielding similar findings. Sabel and Stiff [90] present an approach to detect generated text and profile images on Twitter. They collect tweets related to controversial topics (e.g., COVID-19) and separately classify the tweet’s text and the corresponding profile image. Their method can detect generated media but is highly sensitive to selected thresholds. High precision thresholds cause a significant decrease in true positives, resulting in many overlooked generated images.

Closest to our work is the concurrent study by Yang et al. [113]. While they share the goal of estimating the prevalence of generated profile images on Twitter, their dataset, methodology, and analysis is more limited. The number of analyzed in-the-wild profile images is considerably smaller (254 275 vs. our 14 989 385 samples). Moreover, our detection approach consists of multiple stages to obtain a reliable and robust detection. In contrast, the detection approach by Yang et al. [113] is solely based on StyleGAN2’s facial alignment, which is not sufficient to effectively identify generated faces. First, some real faces inevitably have the same facial alignment. This explains that their evaluation shows a high FDR of 85.86%. Thus, each detected image needs to be manually checked, which does not scale with a larger dataset. Second, relying exclusively on facial alignment is also vulnerable to simple geometric transformations like scaling or cropping. After applying the approach on \mathcal{D}_W^{\times} , we find that 440 generated images in \mathcal{D}_W^{\times} cannot be detected due to being misaligned. We provide a detailed description of their method and the results of our comparison in Appendix E. Finally, in contrast to Yang et al. [113], we contextualize the use of generated profile images by examining the corresponding accounts and tweets.

Yang et al. [113] estimate a lower bound of 0.021% to 0.044% active Twitter accounts that use generated profile images. Our estimated rate with 0.052% is slightly higher, which we attribute to our higher detection performance and the fact that we discard accounts with Twitter’s default profile image.

Human Perception of Generated Social Media Profiles. Since it is unlikely that artificially generated social media profiles can be prevented completely, studying their effect on humans and our society is crucial. Mink et al. [68] conduct a user study to measure users’ trust towards such profiles in a social engineering context. They find that users are likely to accept a connection request from a LinkedIn profile using generated faces or texts. Even participants that were explicitly informed about the presence of fake accounts had an acceptance rate of 43%. A similar work [88] in which participants were asked to label profiles as real or fake in a Twitter-like environment, shows that human performance is almost equivalent to random guessing (48.9%). These findings emphasize the need for reliable detection methods of generated contents in social networks.

Social Media Studies. Complimentary to our work, a large body of interdisciplinary research has focused on the misuse of social media [18, 32, 116]. For example, the 2016 US elections were marked by accusations of opinion manipulation through automated accounts

on social media, particularly on Twitter, so that researchers investigated these inauthentic and coordinated campaigns [3, 6]. In recent years, research has increasingly focused on the harms caused to online communities and the potential to manipulate public sentiment. Studies have extensively explored the roles of disinformation spread, online conspiracy proliferation, and political interference [60, 94, 107]. Another research direction is the identification of inauthentic behavior in context of financial campaigns [19, 102]. Recently, these research efforts are facing new challenges given the increasing use of AI-generated content by social bots [32, 112].

10 CONCLUSION

Generative AI provides unprecedented capabilities to create deceptively realistic content, be it images, videos, text, or music. Despite the considerable applications for the good, these methods also raise significant concerns about their harmful effects. On social media, generated images can be misused to create seemingly real accounts that spread, for example, political misinformation or spam. While the detection of generated content has been explored extensively in controlled laboratory settings, there has been limited systematic research on the prevalence on social media. In this paper, we provide the first systematic large-scale study of generated profile images on Twitter. To build a reliable detection method, we carefully build a pipeline step by step where we consider different dataset types, pre-filtering, classification, and labeling-assistance methods.

In our dataset of 14 989 385 profile images from Twitter, we classify 7 723 profile images as generated. This is 0.052 % of the dataset, showing that generated profile images are notably present on Twitter. Our analysis of the corresponding accounts and their tweets leads to various insights. Fake-image accounts and real-image accounts differ regarding social connections, account activity, account creation time, and account status. For example, many fake-image accounts are created in batches and have identical metadata, indicating that they are part of an organized network. The tweet analysis shows that frequently occurring topics are cryptocurrencies, giveaways, content related to pornography and escort services, as well as controversial political discussions.

In summary, our work introduces a detection method for studying generated content on social media. Our analysis underlines that generated images are used as profile images for a wide range of applications. Addressing this threat will require several steps. First, platforms can implement detection algorithms to flag generated content, as Meta has announced lately [13]. Second, watermarking methods (e.g., [30]) that integrate a detectable watermarking directly into the generation process can facilitate the detection. Finally, raising more awareness about the existence and impact of generated content will be necessary.

ETHICS STATEMENT AND DATA AVAILABILITY

Working with real-world data from social media carries ethical and privacy-related risks. We take different measures to reduce these risks. In our study, statistics of real accounts are reported in aggregated form. We show personal information, such as profile images and tweet texts, only for accounts using generated images. However, we acknowledge that we cannot completely avoid the risk of falsely labeling a real image as generated.

To foster the development and evaluation of real-world generated image detectors, we plan to share our labeled image datasets. Moreover, to comply with Twitter’s/X’s terms of service (ToS), we will release the IDs of users and tweets from our in-the-wild dataset. Due to the recent changes to Twitter’s API, we are aware that accessing the full dataset based on the IDs is challenging. We therefore invite researchers to contact us for discussing further uses of the dataset and potential collaborations.

ACKNOWLEDGMENTS

We thank our reviewers for their valuable comments and suggestions. This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA - 390781972. Moreover, this work was supported by the Leibniz Association Competition (P101/2020) as well as by the IFI program of the German Academic Exchange Service (DAAD) funded by the Federal Ministry of Education and Research (BMBF).

REFERENCES

- [1] Michael Albright and Scott McCloskey. 2019. Source Generator Attribution via Inversion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- [2] Daniel Arp, Erwin Quiring, Feargus Pendlebury, Alexander Warnecke, Fabio Pierazzi, Christian Wressnegger, Lorenzo Cavallaro, and Konrad Rieck. 2022. Dos and Don’ts of Machine Learning in Computer Security. In *USENIX Security Symposium*.
- [3] Adam Badawy, Emilio Ferrara, and Kristina Lerman. 2018. Analyzing the Digital Traces of Political Manipulation: The 2016 Russian Interference Twitter Campaign. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*.
- [4] Valentin Bazarevsky, Yuri Kartynnik, Andrey Vakunov, Karthik Raveendran, and Matthias Grundmann. 2019. BlazeFace: Sub-millisecond Neural Face Detection on Mobile GPUs. *arXiv preprint* (2019).
- [5] Fabricio Benevenuto, Gabriel Magno, Tiago Rodrigues, and Virgilio Almeida. 2010. Detecting Spammers on Twitter. In *Collaboration, Electronic messaging, Anti-Abuse and Spam Conference (CEAS)*.
- [6] Alessandro Bessi and Emilio Ferrara. 2016. Social Bots Distort the 2016 US Presidential Election Online Discussion. *First Monday* (2016).
- [7] Giulia Boato, Cecilia Pasquini, Antonio L. Stefani, Sebastiano Verde, and Daniele Miorandi. 2022. TrueFace: A Dataset for the Detection of Synthetic Face Images from Social Networks. In *IEEE International Joint Conference on Biometrics (IJCB)*.
- [8] Johannes Buchner. 2024. ImageHash: An image hashing library written in Python. <https://github.com/JohannesBuchner/imagehash>.
- [9] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. 2020. What Makes Fake Images Detectable? Understanding Properties That Generalize. In *European Conference on Computer Vision (ECCV)*.
- [10] Keshigeyan Chandrasegaran, Ngoc-Trung Tran, and Ngai-Man Cheung. 2021. A Closer Look at Fourier Spectrum Discrepancies for CNN-generated Images Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [11] Yuanqi Chen, Ge Li, Cece Jin, Shan Liu, and Thomas Li. 2021. SSD-GAN: Measuring the Realness in the Spatial and Spectral Domains. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [12] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [13] Nick Clegg. 2024. Labeling AI-Generated Images on Facebook, Instagram and Threads. <https://about.fb.com/news/2024/02/labeling-ai-generated-images-on-facebook-instagram-and-threads/>.
- [14] Samantha Cole. 2017. AI-assisted Fake Porn Is Here and We’re All Fucked. <https://www.vice.com/en/article/gydydm/gal-gadot-fake-ai-porn>.
- [15] Riccardo Corvi, Davide Cozzolino, Giada Zingarini, Giovanni Poggi, Koki Nagano, and Luisa Verdoliva. 2023. On the Detection of Synthetic Images Generated by Diffusion Models. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [16] Davide Cozzolino, Diego Garganiello, Giovanni Poggi, and Luisa Verdoliva. 2021. Towards Universal GAN Image Detection. In *International Conference on*

- Visual Communications and Image Processing (VCIP).*
- [17] Davide Cozzolino, Justus Thies, Andreas Rössler, Matthias Niessner, and Luisa Verdoliva. 2021. SpoC: Spoofing Camera Fingerprints. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
 - [18] Stefano Cresci. 2020. A Decade of Social Bot Detection. *Commun. ACM* (2020).
 - [19] Stefano Cresci, Fabrizio Lillo, Daniele Regoli, Serena Tardelli, and Maurizio Tesconi. 2019. Cashtag Piggybacking: Uncovering Spam and Bot Activity in Stock Microblogs on Twitter. *ACM Trans. Web* (2019).
 - [20] Jesse Damiani. 2019. A Voice Deepfake Was Used to Scam a CEO out of \$243,000. <https://www.forbes.com/sites/jessedamiani/2019/09/03/a-voice-deepfake-was-used-to-scam-a-ceo-out-of-243000/>.
 - [21] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
 - [22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2020. An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*.
 - [23] Ricard Durall, Margret Keuper, and Janis Keuper. 2020. Watch Your Up-Convolution: CNN Based Generative Deep Neural Networks Are Failing to Reproduce Spectral Distributions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [24] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. 1996. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *International Conference on Knowledge Discovery and Data Mining (KDD)*.
 - [25] Europol Innovation Lab. 2024. Facing reality? Law enforcement and the challenge of deepfakes. <https://www.europol.europa.eu/publications-events/publications/facing-reality-law-enforcement-and-challenge-of-deepfakes>.
 - [26] Pavlos Fafalios, Vasileios Iosifidis, Eirini Ntoutsi, and Stefan Dietze. 2018. Tweet-sKB: A Public and Large-Scale RDF Corpus of Annotated Tweets. In *European Semantic Web Conference (ESWC)*.
 - [27] Tiziano Fagni, Fabrizio Falchi, Margherita Gambini, Antonio Martella, and Maurizio Tesconi. 2021. TweepFake: About Detecting Deepfake Tweets. *PLoS ONE* (2021).
 - [28] Hany Farid. 2021. An Overview of Perceptual Hashing. *Journal of Online Trust and Safety* (2021).
 - [29] Federal Bureau of Investigation (FBI), Cyber Division. 2021. Malicious Actors Almost Certainly Will Leverage Synthetic Content for Cyber and Foreign Influence Operations. <https://www.ic3.gov/Media/News/2021/210310-2.pdf>.
 - [30] Pierre Fernandez, Guillaume Couairon, Hervé Jégou, Matthijs Douze, and Teddy Furon. 2023. The Stable Signature: Rooting Watermarks in Latent Diffusion Models. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.
 - [31] Emilio Ferrara. 2022. Twitter Spam and False Accounts Prevalence, Detection, and Characterization: A Survey. *First Monday* (2022).
 - [32] Emilio Ferrara. 2023. Social Bot Detection in the Age of ChatGPT: Challenges and Opportunities. *First Monday* (2023).
 - [33] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. 2020. Leveraging Frequency Analysis for Deep Fake Image Recognition. In *International Conference on Machine Learning (ICML)*.
 - [34] Joel Frank, Franziska Herbert, Jonas Ricker, Lea Schönherr, Thorsten Eisenhofer, Asja Fischer, Markus Dürmuth, and Thorsten Holz. 2024. A Representative Study on Human Detection of Artificially Generated Media Across Countries. In *IEEE Symposium on Security and Privacy (S&P)*.
 - [35] Yang Gao, Rita Singh, and Bhiksha Raj. 2018. Voice Impersonation Using Generative Adversarial Networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
 - [36] Kiran Garimella and Dean Eckles. 2020. Images and Misinformation in Political Groups: Evidence from WhatsApp in India. *Harvard Kennedy School (HKS) Misinformation Review* (2020).
 - [37] Josh A. Goldstein and Renée DiResta. 2022. Research Note: This Salesperson Does Not Exist: How Tactics from Political Influence Operations on Social Media Are Deployed for Commercial Lead Generation. *Harvard Kennedy School (HKS) Misinformation Review* (2022).
 - [38] Josh A. Goldstein and Shelby Grossman. 2021. How Disinformation Evolved in 2020. <https://www.brookings.edu/articles/how-disinformation-evolved-in-2020/>.
 - [39] Josh A. Goldstein, Girish Sastry, Micah Musser, Renee DiResta, Matthew Gentzel, and Katerina Sedova. 2023. Generative Language Models and Automated Influence Operations: Emerging Threats and Potential Mitigations. *arXiv preprint* (2023).
 - [40] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems (NeurIPS)*.
 - [41] D. Gagnaniello, D. Cozzolino, F. Marra, G. Poggi, and L. Verdoliva. 2021. Are GAN Generated Images Easy to Detect? A Critical Analysis of the State-of-the-Art. In *IEEE International Conference on Multimedia and Expo (ICME)*.
 - [42] Graphika Team. 2020. Step into My Parler. <https://graphika.com/reports/step-into-my-parler>.
 - [43] Graphika Team. 2021. Fake Cluster Boosts Huawei. <https://graphika.com/reports/fake-cluster-boosts-huawei>.
 - [44] Hui Guo, Shu Hu, Xin Wang, Ming-Ching Chang, and Siwei Lyu. 2022. Eyes Tell All: Irregular Pupil Shapes Reveal GAN-generated Faces. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
 - [45] Supraja Gurajala, Joshua S. White, Brian Hudson, and Jeanna N. Matthews. 2015. Fake Twitter Accounts: Profile Characteristics Obtained Using an Activity-Based Pattern Detection Approach. In *International Conference on Social Media & Society*.
 - [46] Harmon, Joshua Roesslein, and other contributors. 2024. Tweepy. <https://github.com/tweepy/tweepy>.
 - [47] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [48] Daniel Hickey, Matheus Schmitz, Daniel Fessler, Paul E. Smaldino, Goran Muric, and Keith Burghardt. 2023. Auditing Elon Musk’s Impact on Hate Speech and Bots. In *International AAAI Conference on Web and Social Media (ICWSM)*.
 - [49] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
 - [50] Shu Hu, Yuezun Li, and Siwei Lyu. 2021. Exposing GAN-Generated Faces Using Inconsistent Corneal Specular Highlights. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
 - [51] Nils Hulzebosch, Sarah Ibrahim, and Marcel Worring. 2020. Detecting CNN-generated Facial Images in Real-World Scenarios. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
 - [52] Mingku Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. 2023. Scaling up GANs for Text-to-Image Synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [53] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *International Conference on Learning Representations (ICLR)*.
 - [54] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2021. Alias-Free Generative Adversarial Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*.
 - [55] Tero Karras, Samuli Laine, and Timo Aila. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [56] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [57] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*.
 - [58] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*.
 - [59] Federica Lago, Cecilia Pasquini, Rainer Böhme, Hélène Dumont, Valérie Goffaux, and Giulia Boato. 2022. More Real than Real: A Study on Human Visual Perception of Synthetic Faces [Applications Corner]. *IEEE Signal Processing Magazine* (2022).
 - [60] Luca Luceri, Ashok Deb, Silvia Giordano, and Emilio Ferrara. 2019. Evolution of Bot and Human Behavior During Elections. *First Monday* (2019).
 - [61] Sara Mandelli, Nicolò Bonettini, Paolo Bestagini, and Stefano Tubaro. 2020. Training CNNs in Presence of JPEG Compression: Multimedia Forensics vs Computer Vision. In *IEEE International Workshop on Information Forensics and Security (WIFS)*.
 - [62] Federico Marcon, Cecilia Pasquini, and Giulia Boato. 2021. Detection of Manipulated Face Videos over Social Networks: A Large-Scale Study. *Journal of Imaging* (2021).
 - [63] Francesco Marra, Diego Gagnaniello, Davide Cozzolino, and Luisa Verdoliva. 2018. Detection of GAN-generated Fake Images over Social Networks. In *IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*.
 - [64] Francesco Marra, Diego Gagnaniello, Luisa Verdoliva, and Giovanni Poggi. 2019. Do GANs Leave Artificial Fingerprints?. In *IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*.
 - [65] Falko Matern, Christian Riess, and Marc Stamminger. 2019. Exploiting Visual Artifacts to Expose Deepfakes and Face Manipulations. In *IEEE Winter Applications of Computer Vision (WACV) Workshops*.
 - [66] Scott McCloskey and Michael Albright. 2019. Detecting GAN-generated Imagery Using Saturation Cues. In *IEEE International Conference on Image Processing (ICIP)*.
 - [67] Leland McInnes, John Healy, and James Melville. 2020. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv preprint* (2020).
 - [68] Jaron Mink, Licheng Luo, Natã M. Barbosa, Olivia Figueira, Yang Wang, and Gang Wang. 2022. DeepPhish: Understanding User Trust towards Artificially Generated Profiles in Online Social Networks. In *USENIX Security Symposium*.
 - [69] Yisroel Mirsky and Wenke Lee. 2021. The Creation and Detection of Deepfakes: A Survey. *ACM Comput. Surv.* (2021).

- [70] Lakshmanan Nataraj, Tajuddin Manhar Mohammed, B. S. Manjunath, Shivkumar Chandrasekaran, Arjuna Flenner, Jawadul H. Bappy, and Amit K. Roy-Chowdhury. 2019. Detecting GAN Generated Fake Images Using Co-Occurrence Matrices. *IS&T International Symposium on Electronic Imaging: Media Watermarking, Security, and Forensics*.
- [71] Sophie Nightingale, Shruti Agarwal, Erik Härkönen, Jaakko Lehtinen, and Hany Farid. 2021. Synthetic Faces: How Perceptually Convincing Are They? *Journal of Vision* (2021).
- [72] Sophie J. Nightingale and Hany Farid. 2022. AI-synthesized Faces Are Indistinguishable from Real Faces and More Trustworthy. *Proceedings of the National Academy of Sciences* (2022).
- [73] Ben Nimmo, C. Shawn Eib, and Léa Ronzaud. 2020. Operation Naval Gazing. <https://graphika.com/reports/operation-naval-gazing>.
- [74] Ben Nimmo, C. Shawn Eib, L. Tamora, Kate Johnson, Ian Smith, Eto Buziashvili, Alyssa Kann, Kanishk Karan, Esteban Ponce de León Rosas, and Max Rizzuto. 2019. #OperationFFS: Fake Face Swarm. <https://graphika.com/reports/operationffs-fake-face-swarm>.
- [75] Ben Nimmo, Camille François, C. Shawn Eib, and Léa Ronzaud. 2020. IRA Again: Unlucky Thirteen. <https://graphika.com/reports/ira-again-unlucky-thirteen>.
- [76] Ben Nimmo, Camille François, C. Shawn Eib, and Léa Ronzaud. 2020. Spamouflage Goes to America. <https://www.graphika.com/reports/spamouflage-dragon-goes-to-america>.
- [77] Leonardo Nizzoli, Serena Tardelli, Marco Avvenuti, Stefano Cresci, Maurizio Tesconi, and Emilio Ferrara. 2020. Charting the Landscape of Online Cryptocurrency Manipulation. *IEEE Access* (2020).
- [78] Conspirador Norteño. 2024. Conspirador Norteño | Substack. <https://conspirador0.substack.com/>.
- [79] Utkarsh Ojha, Yuheng Li, and Yong Jae Lee. 2023. Towards Universal Fake Image Detectors That Generalize across Generative Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [80] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. 2022. On Aliased Resizing and Surprising Subtleties in GAN Evaluation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [81] Cecilia Pasquini, Francesco Laiti, Davide Lobba, Giovanni Ambrosi, Giulia Boato, and Francesco De Natale. 2023. Identifying Synthetic Faces through GAN Inversion and Biometric Traits Analysis. *Applied Sciences* (2023).
- [82] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research (JMLR)* (2011).
- [83] Jürgen Pfeffer, Daniel Matter, Kokil Jaidka, Onur Varol, Afra Mashhadi, Jana Lasser, Dennis Assenmacher, Siqi Wu, Diyi Yang, Cornelia Brantner, Daniel M. Romero, Jahna Otterbacher, Carsten Schwemmer, Kenneth Joseph, David Garcia, and Fred Morstatter. 2023. Just Another Day on Twitter: A Complete 24 Hours of Twitter Data. In *International AAAI Conference on Web and Social Media (ICWSM)*.
- [84] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models from Natural Language Supervision. In *International Conference on Machine Learning (ICML)*.
- [85] Jacob Ratkiewicz, Michael Conover, Mark Meiss, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. 2011. Detecting and Tracking Political Abuse in Social Media. In *International AAAI Conference on Web and Social Media (ICWSM)*.
- [86] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [87] Jonas Ricker, Denis Lukovnikov, and Asja Fischer. 2024. AEROBLADE: Training-Free Detection of Latent Diffusion Images Using Autoencoder Reconstruction Error. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [88] Sippo Rossi, Youngjin Kwon, Odd Harald Auglend, Raghava Rao Mukkamala, Matti Rossi, and Jason Thatcher. 2023. Are Deep Learning-Generated Social Media Profiles Indistinguishable from Real Profiles?. In *Hawaii International Conference on System Sciences (HICSS)*.
- [89] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* (2015).
- [90] Johan Sabel and Harald Stiff. 2021. Detecting Generated Media: A Case Study on Twitter Data. *NATO STO Meeting Proceedings* (2021).
- [91] Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. 2021. Projected GANs Converge Faster. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [92] Katja Schwarz, Yiyi Liao, and Andreas Geiger. 2021. On the Frequency Bias of Generative Models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [93] Selenium Developers. 2024. Selenium. <https://github.com/SeleniumHQ/selenium>.
- [94] Chengcheng Shao, Giovanni Luca Ciampaglia, Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer. 2018. The Spread of Low-Credibility Content by Social Bots. *Nature Communications* (2018).
- [95] Bingyu Shen, Brandon RichardWebster, Alice O'Toole, Kevin Bowyer, and Walter J. Scheirer. 2021. A Study of the Human Perception of Synthetic Faces. In *IEEE International Conference on Automatic Face and Gesture Recognition (FG)*.
- [96] Naman Singh, Tushar Sharma, Abha Thakral, and Tanupriya Choudhury. 2018. Detection of Fake Profile in Online Social Networks Using Machine Learning. In *International Conference on Advances in Computing and Communication Engineering (ICACCE)*.
- [97] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics. In *International Conference on Machine Learning (ICML)*.
- [98] Stanford Internet Observatory. 2020. Analysis of Twitter Takedowns Linked to Cuba, the Internet Research Agency, Saudi Arabia, and Thailand. <https://cyber.fsi.stanford.edu/io/news/twitter-takedown-october-2020>.
- [99] Stanford Internet Observatory. 2020. Reply-Guys Go Hunting: An Investigation into a U.S. Astroturfing Operation on Facebook, Twitter, and Instagram. <https://purl.stanford.edu/vh222ch4142>.
- [100] Benjamin Strick. 2020. West Papua: New Online Influence Operation Attempts to Sway Independence Debate. <https://www.bellingcat.com/news/2020/11/11/west-papua-new-online-influence-operation-attempts-to-sway-independence-debate/>.
- [101] Benjamin Strick. 2021. Analysis of the Pro-China Propaganda Network Targeting International Narratives. <https://www.info-res.org/post/revealed-coordinated-attempt-to-push-pro-china-anti-western-narratives-on-social-media>.
- [102] Serena Tardelli, Marco Avvenuti, Maurizio Tesconi, and Stefano Cresci. 2022. Detecting Inorganic Financial Campaigns on Twitter. *Information Systems* (2022).
- [103] Raffaele Tucciarelli, Neza Vehar, Shamil Chandaria, and Manos Tsakiris. 2022. On the Realness of People Who Do Not Exist: The Social Processing of Artificial Faces. *iScience* (2022).
- [104] Aaron van den Oord, Nal Kalchbrenner, Lasse Espeholt, Koray Kavukcuoglu, Oriol Vinyals, and Alex Graves. 2016. Conditional Image Generation with Pixel-CNN Decoders. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [105] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. 2016. Pixel Recurrent Neural Networks. In *International Conference on Machine Learning (ICML)*.
- [106] James Vincent. 2019. A Spy Reportedly Used an AI-generated Profile Picture to Connect with Sources on LinkedIn. <https://www.theverge.com/2019/6/13/18677341/ai-generated-fake-faces-spy-linked-in-contacts-associated-press>.
- [107] Emily L. Wang, Luca Luceri, Francesco Pierri, and Emilio Ferrara. 2023. Identifying and Characterizing Behavioral Classes of Radicalization within the QAnon Conspiracy on Twitter. In *International AAAI Conference on Web and Social Media (ICWSM)*.
- [108] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [109] Yuping Wang, Chen Ling, and Gianluca Stringhini. 2023. Understanding the Use of Images to Spread COVID-19 Misinformation on Twitter. *Proc. ACM Hum.-Comput. Interact.* (2023).
- [110] Tyler Williams and Razmik Kozakjian. 2022. Portrait Mode: GAN Collages and Fake Personas. <https://graphika.com/posts/portrait-mode-gan-collages-and-fake-personas>.
- [111] Weihao Xia, Yulun Zhang, Yujia Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. 2023. GAN Inversion: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023).
- [112] Kai-Cheng Yang and Filippo Menczer. 2024. Anatomy of an AI-powered malicious social botnet. *Journal of Quantitative Description: Digital Media* (2024).
- [113] Kai-Cheng Yang, Danishjeet Singh, and Filippo Menczer. 2024. Characteristics and Prevalence of Fake Social Media Profiles with AI-generated Faces. *Journal of Online Trust and Safety* (2024).
- [114] Kai-Cheng Yang, Onur Varol, Pik-Mai Hui, and Filippo Menczer. 2020. Scalable and Generalizable Social Bot Detection through Data Selection. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [115] Yuanshun Yao, Bimal Viswanath, Jenna Cryan, Haitao Zheng, and Ben Y. Zhao. 2017. Automated Crowdturfing Attacks and Defenses in Online Review Systems. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*.
- [116] Sarita Yardi, Daniel Romero, Grant Schoenebeck, and Danah Boyd. 2010. Detecting Spam in a Twitter Network. *First Monday* (2010).
- [117] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. 2016. LSUN: Construction of a Large-Scale Image Dataset Using Deep Learning with Humans in the Loop. *arXiv preprint* (2016).
- [118] Ning Yu, Larry S. Davis, and Mario Fritz. 2019. Attributing Fake Images to GANs: Learning and Analyzing GAN Fingerprints. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.

- [119] Savvas Zannettou, Tristan Caulfield, Jeremy Blackburn, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Guillermo Suarez-Tangil. 2018. On the Origins of Memes by Means of Fringe Web Communities. In *Internet Measurement Conference (IMC)*.
- [120] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [121] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. 2019. Detecting and Simulating Artifacts in GAN Fake Images. In *IEEE International Workshop on Information Forensics and Security (WIFS)*.
- [122] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.

A METHODOLOGY DETAILS

Here, we provide implementation details of our methodology.

A.1 Data Collection

In-the-Wild Dataset \mathcal{D}_W^X . We access the Twitter API using the tweepy [46] Python package and download the profile image of each tweet’s author from the respective `profile_image_url`. Table 5 lists all metadata fields we obtain from the API. The second column denotes how many accounts in \mathcal{D}_W^X have a value in the respective field.

Labeled Datasets $\mathcal{D}_R/\mathcal{D}_F$ and Variations. We collect 10 000 images from TPDNE by repeatedly querying the website, mimicking a user creating a fake profile. We analogously take the first 10 000 real images from the FFHQ dataset. To avoid an unwanted bias based on image processing, we convert the PNG files from FFHQ to JPEG using the same parameters as TPDNE. Then, to obtain processed images as they would appear on Twitter (\mathcal{D}_R^X and \mathcal{D}_F^X), we upload each image as a profile image and download it. We observed a difference in the image processing between API-based and browser-based uploads. Images uploaded with the API kept their resolution, while images uploaded in the browser were resized to 400×400 pixels. As the majority of in-the-wild images has the resized resolution, we select the browser-based approach and automate the upload using the web automation framework Selenium [93]. To obtain the zoomed-in versions ($\mathcal{D}_R^{X'}$ and $\mathcal{D}_F^{X'}$), the automated upload procedure is extended by first zooming into each image by a random amount and then moving the image by a random x- and y-offset. We ensure that the image still looks like a plausible profile image at the maximum zoom rate.

A.2 Pre-Filter

BlazeFace [4] predicts a bounding box as well as the x- and y-positions of six facial landmarks (eyes, ears, mouth, and nose) in normalized coordinates between 0 and 1. If an image contains multiple faces, we select the one with the largest bounding box.

A.3 Classifier

Our architecture and training procedure is adapted from Wang et al. [108]. We follow the common practice of initializing a ResNet-50 [47] with weights from an image classifier trained on ImageNet [89] and replace the final layer to reflect the binary classification setting. During training, we use a batch size of 32 and optimize the model using Adam [57] and binary cross-entropy loss.



(a) Real image from \mathcal{D}_R .



(b) Generated image from \mathcal{D}_F .

Figure 12: Examples of original images (left) and their reconstructions (right).

In the case of $C_{R^X, P^X/F^X}$ we ensure balanced sampling of real/proxy-labeled real and fake samples. The learning rate is reduced by a factor of 10 if the validation loss does not decrease by 0.001 during 5 epochs. We perform early stopping once the learning rate becomes smaller than 10^{-6} . For training $C_{R/F}$, the images in \mathcal{D}_R and \mathcal{D}_F are resized to 400×400 using bilinear interpolation to match the profile image dimensions of Twitter. The training data is augmented using three kinds of perturbations, each applied with probability $p = 0.1$: Gaussian blurring with a kernel size of 9 and σ uniformly sampled from $[0.5, 5.0]$, JPEG compression with quality uniformly sampled from $[30, 100]$, and resizing, with scale and aspect ratio uniformly sampled from $[0.25, 0.75]$ and $[0.8, 1.25]$, respectively. During training, we randomly extract crops of size 224×224 , while we take the center crop of the same size during validation and testing.

B INVERSION EXAMPLES

Figure 12 depicts example images to demonstrate the assisted manual labeling. The left image is the original while the right image is its reconstruction obtained by GAN inversion. For the real image from \mathcal{D}_R , we observe that the background is inaccurate and the face is slightly blurred. In contrast, the generated image from \mathcal{D}_F can be reconstructed very accurately, including the background.

C PRE-TRAINED DETECTORS

Here we provide details on the three existing pre-trained classifiers we evaluate in Section 7.1. C_{Wang} [108] is the model on that our detector is based on. However, it is trained on a diverse set of images generated by ProGAN [53] and corresponding real images from LSUN [117]. We select the version *Blur+JPEG (0.1)* since the authors report a good performance on images generated by StyleGAN2 [56]. C_{Grag} [41] is an improved version of C_{Wang} that avoids

Table 5: Overview of metadata for accounts in our dataset \mathcal{D}_W^X .


Field	Count	Description
id	14 989 385	Unique user identifier.
username	14 989 385	Username (handle).
name	14 989 385	Name shown in profile (display name).
created_at	14 989 385	Account creation time.
location	7 940 863	User-specified location.
description	14 989 385	Profile bio.
url	3 654 749	User-specified URL.
profile_image_url	14 989 385	URL to user’s profile image.
public_metrics.followers_count	14 989 385	Number of followers.
public_metrics.following_count	14 989 385	Number of accounts user is following.
public_metrics.tweet_count	14 989 385	Number of tweets.
public_metrics.listed_count	14 989 385	Number of lists containing user.
protected	14 989 385	Whether account is private.
verified	14 989 385	Whether account is verified.
withheld.country_codes	4 192	Countries where user is not available.
pinned_tweet_id	6 866 224	Identifier of user’s pinned tweet.
entities.url.urls	3 654 749	Details about profile website.
entities.description.mentions	1 611 345	Details about user mentions in description.
entities.description.urls	780 440	Details about URLs in description.
entities.description.hashtags	1 340 520	Details about hashtags in description.
entities.description.cashtags	35 746	Details about cashtags in description.

downsampling in the first layer of the ResNet-50 [47] backbone to preserve high-frequency artifacts (at the cost of a larger model). Besides training on ProGAN [53] images, the authors provide a detector trained on StyleGAN2 [56] images, which we select since it should yield the best results on our dataset. Finally, C_{Ojha} follows a different approach and leverages the feature space of a pre-trained vision transformer (CLIP-ViT [22, 84]). It uses a single linear layer on top (trained on ProGAN [53] images) to predict whether an image is real or fake.

D DUPLICATE IMAGE DETECTION

Despite the trivial access to generated faces using TPDNE, creators of fake account clusters might use the same face for multiple accounts. To identify such duplicates, we need an approach that is robust to subtle differences caused by varying image processing. We adapt the technique used by previous works [36, 109, 119] and cluster images based on their perceptual hashes (pHashes). Perceptual image hashing [28] aims to extract a meaningful representation of an image that does not depend on individual pixel values, but on the perceived content. The algorithm we use [8] achieves this by deriving 64 bits from the DCT coefficients belonging to the lower frequencies of an image. To obtain groups of duplicate images, we apply the DBSCAN [24] clustering algorithm to our calculated pHashes. We use the implementation from scikit-learn [82] and set the minimum number of elements to 2. We empirically find that we obtain meaningful clusters by setting the maximum allowed Hamming distance between two pHashes to 3.

In total, we identify 540 groups of duplicated images with an average size of 4.88 images. The distribution of the sizes is given in Figure 13. About half of all groups consist of only two or three

**Figure 13: Size distribution of duplicate image clusters.**

duplicated images, while the most frequently used faces appeared in 18 profiles.

E EVALUATION OF ALIGNMENT-BASED DETECTION

In the concurrent work by Yang et al. [113], the authors identify GAN-generated faces on Twitter using a method that is related to our concept of alignment (cf. Section 3.3). They define the $GANEyeDistance$ \mathcal{G} as the normalized Euclidean distance between the actual and expected location of each eye. They propose to consider an image to be potentially GAN-generated if $\mathcal{G} < 0.02$. To reach a final decision, they propose to manually classify images based on visual artifacts. While this approach is easy to implement and computationally efficient, we find that is suboptimal regarding (a) the number of false positives (causing a large manual workload)




Figure 14: Examples of fake images that evade alignment-based detection. Below each image we provide its *GANEyeDistance* \mathcal{G} . The reference eye position is highlighted. Note that we only display images which we confidently consider to be fake to avoid disclosing real profile images.

and (b) the number of false negatives (overlooking generated faces that are not aligned).

We test \mathcal{G} with the suggested threshold on 150 000 randomly chosen images from $\mathcal{D}_W^{\mathbb{X}}$ (about 1%, which yields 730 candidate profiles. For all samples in $\mathcal{D}_W^{\mathbb{X}}$, the estimated number of candidates therefore is 73 000. Classifying these images would require an excessive amount of manual effort.

On the other hand, we find 440 images in $\mathcal{D}_W^{\mathbb{X}}$ that are detected as fake by $C_{R^{\mathbb{X}}, P^{\mathbb{X}}/F^{\mathbb{X}}}$ but are overlooked when classifying based on \mathcal{G} . Naturally, it can be assumed that in this subset our classifier has a higher number of false positives, since most generated images are in fact aligned. Still, after manual inspection, we rate 303 of these images to be definitely or very likely generated. Note that manual labeling is more challenging on these images since we cannot resort to GAN inversion. Figure 14 depicts some examples together with their value of \mathcal{G} . One can see that zooming in by a small amount is sufficient to cause a misalignment. We consider it probable that malicious accounts do this on purpose to appear more credible and avoid detection based on facial landmarks.