
Characteristics and Prevalence of Fake Social Media Profiles with AI-generated Faces

Kai-Cheng Yang, Danishjeet Singh, and Filippo Menczer

Abstract. Recent advancements in generative artificial intelligence (AI) have raised concerns about its potential to create convincing fake social media accounts, but empirical evidence is lacking. In this paper, we present a systematic analysis of Twitter (X) accounts using human faces generated by Generative Adversarial Networks (GANs) for their profile pictures. We present a dataset of 1,420 such accounts and show that they are used to spread scams, disseminate spam, and amplify coordinated messages, among other inauthentic activities. Leveraging a feature of GAN-generated faces—consistent eye placement—and supplementing it with human annotation, we devise an effective method for identifying GAN-generated profiles in the wild. Applying this method to a random sample of active Twitter users, we estimate a lower bound for the prevalence of profiles using GAN-generated faces between 0.021% and 0.044%—around 10,000 daily active accounts. These findings underscore the emerging threats posed by multimodal generative AI. We release the source code of our detection method and the data we collect to facilitate further investigation. Additionally, we provide practical heuristics to assist social media users in recognizing such accounts.

1 Introduction

Recent advances in generative artificial intelligence (AI) have resulted in remarkable competencies across various domains (Cao et al. 2023). Large language models, such as ChatGPT,¹ can understand human language and generate realistic text in different contexts (Jakesch, Hancock, and Naaman 2023; Ye et al. 2023). Generative vision models, such as Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) and diffusion models (Song and Ermon 2019; Ho, Jain, and Abbeel 2020), produce photorealistic images that people can no longer distinguish from human-generated images (Lago et

1. <https://chat.openai.com>

al. 2021; Nightingale and Farid 2022; Bray, Johnson, and Kleinberg 2023). State-of-the-art multimodal models, such as Google’s Gemini models (Pichai and Hassabis 2023), can even understand visual and audio inputs and respond with text or images. This evolving landscape of generative AI models underscores a significant leap in AI technology, opening new avenues for research and applications.

These advancements also present opportunities for misuse. For instance, adversarial actors can exploit AI models to create and disseminate disinformation or conduct influence campaigns (Yamin et al. 2021; Guembe et al. 2022; Goldstein et al. 2023; Yang and Menczer 2024), threatening the integrity of online information exchange. In this study, we focus on the fake social media accounts that leverage generative AI for deceptive purposes (Bommasani et al. 2021). In documented cases, individuals have used AI-generated human images to fabricate fake personas online, leading to deception and targeted attacks (O’Sullivan 2020b, 2020a). Major social media platforms, such as Facebook (Martineau 2019) and LinkedIn (Bond 2022), have already encountered groups of fake accounts that use AI-generated faces.

Despite these cases, a systematic analysis of such accounts is still lacking. What are the tactics and goals of these fake accounts? What is their prevalence on social media platforms? These key questions remain unanswered in part because social media platforms are reluctant to share data and detailed analysis results. To address this knowledge gap, here we present an analysis of inauthentic accounts on Twitter² using GAN-generated pictures as their profiles to mimic real human users.

Although there are other generative vision models, we focus on GANs for several reasons. First, the introduction of GAN-based models, particularly the StyleGAN series (Karras, Laine, and Aila 2019; Karras et al. 2020), marks a major milestone in the field of generative AI models (Cao et al. 2023) because they are capable of producing photorealistic images almost indistinguishable from real ones (Nightingale and Farid 2022). Second, GANs are easily accessible. With their source code publicly available online,³ adversarial actors can easily produce fake human faces as needed. Additionally, for those without technical expertise, websites like <https://thispersondoesnotexist.com> make it easy to download GAN-generated images in large quantities. Last but not least, GAN-generated faces share some distinctive features, allowing us to identify them with high confidence.

In this study, we create a novel dataset of 1,420 inauthentic Twitter accounts, each using a GAN-generated face as its profile picture. Upon close examination, these accounts typically feature human-like names, suggesting a deliberate effort to mimic real users. These accounts engage in inauthentic activities, such as disseminating spam, executing scams, and amplifying messages in a coordinated fashion, often through automation. We further estimate the prevalence of such accounts on the platform. Leveraging the

2. Twitter was rebranded as “X” in July 2023, but we refer to the platform as “Twitter,” as most data analyzed in this paper predates the rebranding.

3. See <https://github.com/NVlabs/stylegan> and <https://github.com/NVlabs/stylegan2>.

feature that GANs consistently place eyes at the same locations, together with human annotation, we develop a simple yet effective method to identify GAN-generated profiles in the wild. We apply this method to a random sample of 254,275 active Twitter users, and our analysis indicates that the prevalence of accounts using GAN-generated faces is between 0.021% and 0.044%. Since our method emphasizes precision and might miss accounts leveraging more advanced AI models to generate their profiles, these estimates should be regarded only as lower bounds.

Our findings reveal that generative AI tools, particularly GANs, have been deployed at scale to create fake personas on Twitter. These accounts are involved in various inauthentic activities, posing a challenge to the integrity of online interactions. To support further research and expand the scope of investigation to other cases and platforms, we share our dataset and code. We also discuss practical heuristics that can empower social media users to identify and defend themselves against the manipulative tactics of these AI-powered accounts.

2 Related work

2.1 GAN-generated face detection

Due to the popularity of the StyleGAN models (Karras, Laine, and Aila 2019; Karras et al. 2020) and their broad impact, many studies have aimed to detect GAN-generated pictures in the past few years. In addition to deep learning models that can capture distinctive features of GAN-generated pictures automatically (Nataraj et al. 2019; Dang et al. 2018), many methods rely on physical or physiological features of the generated faces for detection. For example, Hu, Li, and Lyu (2021) leverage corneal specular highlights; Guo et al. (2022) focus on pupil shapes; X. Yang et al. (2019) and Mundra et al. (2023) exploit the locations of face landmarks. For details on these methods, readers can refer to the review by Wang et al. (2023).

Although we also need to detect GAN-generated faces in this study, our case differs from previous studies significantly. The aforementioned methodology papers typically aim to distinguish generated and real human faces from datasets such as FFHQ (Karras, Laine, and Aila 2019) and CelebA (Liu et al. 2015). Our objective, on the other hand, is to identify generated faces among random Twitter profiles, which presents unique challenges. First, only a subset of Twitter profiles feature human faces, and their styles may differ greatly from GAN-generated ones (see Figure 1(c) for examples). Second, the instances of GAN-generated faces may be relatively scarce, necessitating a detection method with high accuracy to address the class imbalance issue and provide a reliable estimate of the prevalence of fake accounts. Given these differences, existing methods are not suitable for our purpose.

2.2 GAN-generated faces on social media

While there is a substantial body of literature on detecting GAN-generated images, only a few studies have specifically targeted social media contexts. Marra et al. (2018) simulate the case of detecting generated pictures compressed by social media platforms by preprocessing the pictures in their datasets, although their focus is not on human faces. Boato et al. (2022) upload GAN-generated human pictures to social media platforms, including Facebook, Twitter, and Telegram, and then download them to examine the effectiveness of various detection methods on compressed images.

Unfortunately, these studies cannot help address our research questions about the fake social media accounts that use GAN-generated faces. Our knowledge of these accounts is largely derived from sporadic journalistic reporting, brief communications from social media platforms, and case studies provided by hobbyists (Norteño 2023). We hope the present study can further our understanding of these accounts and their negative impact.

2.3 Social bots and their detection

This study also encompasses the topic of social bots, which, while related, are distinct from our primary focus. Social bots are accounts partially controlled by algorithms to emulate humans and act automatically following different instructions (Ferrara et al. 2016). To appear authentic, bots sometimes employ realistic images as profile pictures, including those stolen from other accounts (Confessore et al. 2018) or generated using AI models (as we show below). However, accounts using GAN-generated faces are not necessarily bots; sometimes, they are operated by real humans.

The detection of social bots often depends on the content produced by the accounts and their behavioral patterns, such as posting frequency and types of interactions with other accounts (K.-C. Yang et al. 2019; Sayyadiharikandeh et al. 2020). Some recent studies have explored using GANs for bot detection in adversarial settings (Cresci et al. 2021; Najari, Salehi, and Farahbakhsh 2022). However, their focus is mainly on the behavioral patterns and content, not on the profile pictures. Account profile information, such as the user name, description, and number of friends and followers, typically contributes to bot detection (K.-C. Yang et al. 2020), but profile pictures are usually excluded, likely due to image-processing challenges.

Since we aim to identify accounts that use GAN-generated faces in this study, we will mainly rely on profile pictures for detection. We will then analyze the profile information and activities of the identified accounts.

Table 1: Detailed information about different groups of accounts in TwitterGAN.

Group name	Size	Date of Collection	Note
blue_check	11	2022-11-10	Subscribed to the Twitter’s blue-checkmark paid verification program
bart	5	2023-02-02	Used to propagate spammy hashtags, YouTube videos, and Tumblr links
elon_watch	246	2023-02-03	Repeatedly replied to other accounts with scams involving Elon Musk
sunday	18	2023-03-05	Used to amplify the tweets of Twitter user @ilkersenock
astroturf	1,066	2023-04-08	Used to amplify content from a common set of accounts in a coordinated fashion
chatgpt	67	2024-01-31	Used ChatGPT to generate human-like content
ind_acc	7	varies	Individual accounts that do not belong to any coordinated groups
Total	1,420		

3 Data collection

We compile three datasets for our analyses: TwitterGAN, consisting of Twitter accounts collected from the wild that use GAN-generated faces as profile pictures; AcademicGAN, a set of synthetic faces generated by GANs for research purposes; and RandomTwitter, a collection of random Twitter profiles. The key dataset is shared with the public for further investigation and can be accessed at a public data repository.⁴

3.1 TwitterGAN dataset

One of our main contributions is the novel TwitterGAN dataset consisting of a total of 1,420 Twitter accounts that use GAN-generated human faces as their profiles. We collected their profile information and recent tweets for analysis. When Twitter’s V2 API was still available, we used it to collect the data for most of the accounts. After Twitter shut down its free API, we manually collected the data through Twitter’s web interface. We also downloaded the accounts’ profile pictures. Twitter allowed public access to these pictures in various resolutions. We primarily aimed for the 400 × 400 pixel versions, although some variations in resolution and aspect ratios exist.

The dataset contains multiple account groups collected between November 2022 and January 2024. The timeline of data collection and notes can be found in Table 1. The initial identification of some of these accounts was done by the Twitter user @conspirator0, while others were spotted by the authors. Following this, we conducted a thorough inspection to verify the nature of these accounts.

As mentioned earlier, the GAN-generated faces are very realistic, and layman participants

4. <https://zenodo.org/doi/10.5281/zenodo.10436888>




Figure 1: Sample profile pictures from (a) AcademicGAN, (b) TwitterGAN, and (c) RandomTwitter.

in controlled experiments can no longer distinguish them from real human faces (Nightingale and Farid 2022). However, after reading extensive literature and reviewing a large number of GAN-generated human faces, we have identified several distinctive characteristics of these faces that bolster our confidence in their identification. First, GANs consistently show face landmarks, especially the eyes, at the same location in the generated pictures (X. Yang et al. 2019). This feature becomes evident when comparing multiple GAN-generated faces side by side, as illustrated in Figure 1. Second, many GAN-generated pictures have clear defects, which will be discussed in Section 4.3. Third, some profile pictures contain watermarks explicitly revealing their machine-generated origin.

We also examine the profile information and activities of the accounts to triangulate their nature. For instance, there are occasional mismatches between account names and the perceived gender and race of the faces. We further rely on signals of coordinated inauthentic activities, which proved effective in similar contexts (Pacheco et al. 2021; Yang and Menczer 2024). The TwitterGAN accounts in some groups are interconnected through mutual following, similar profile information, and identical activity patterns. (These connections also allow us to identify more members of their groups.) Our conclusions are drawn from a holistic consideration of all these factors.

3.2 AcademicGAN dataset

Since capturing a large number of accounts using GAN-generated profiles in the wild is not trivial, we supplement the study with the AcademicGAN dataset. It contains 10,000 GAN-generated human faces obtained from the <https://generated.photos/datasets> website upon request. The website trains its own GAN models on proprietary human photos to produce more synthetic ones for sale. All pictures in AcademicGAN have the same resolution of 256×256 pixels and do not have the common defects observed in TwitterGAN. Nevertheless, their face landmarks are still placed at the same locations, as demonstrated in Figure 1. This dataset is used as the ground truth in our analysis to compare with TwitterGAN and develop our detection method.

3.3 RandomTwitter dataset

To estimate the prevalence of accounts using GAN-generated profiles on Twitter, we assemble the RandomTwitter dataset, a random sample of diverse active Twitter users. We leverage the streaming endpoint of Twitter’s V2 API, which delivers a 1% sample of all live tweets in real time. The data collection took place on June 1, 2023, before Twitter shut down our API access. This yielded 285,747 tweets from 254,275 unique users. Similar to the TwitterGAN dataset, we also downloaded the 400×400 pixel versions of their profile images whenever possible. Figure 1 shows a sample of 16 of these images.

4 Characterization of TwitterGAN

In this section, we characterize the profiles and activities of accounts in TwitterGAN.

4.1 Profile information




Figure 2: Profile characteristics of accounts in TwitterGAN. We show the distributions of (a) follower count, (b) following (friend) count, (c) tweet count, and (e) year of creation, for astroturf and other accounts.

Let us first focus on the profile information of the accounts in TwitterGAN. The distributions of their follower/following count, tweet count, and creation year are shown in Figure 2. Since the astroturf group has more accounts than the other groups combined and shows distinctive features, we separate it from the rest of the accounts.

On average, accounts in astroturf have 9.95 followers, 55.68 friends, and 14.91 tweets (standard deviation = 24.03, 51.75, and 15.62, respectively), whereas the other accounts

have 67.75 followers, 162.71 friends, and 509.23 tweets (standard deviation = 163.64, 177.17, and 504.02, respectively).

All the differences between astroturf and the other accounts are statistically significant according to Kolmogorov–Smirnov tests ($p < 0.05$). These results suggest that accounts in TwitterGAN are engaging in various activities on Twitter; however, astroturf accounts are less active than others.

Inspecting the creation years of the accounts suggests that most of them joined Twitter in the past two years, and astroturf accounts are more recent than others. This is not surprising since Twitter, both under the old management and while under Elon Musk’s ownership, routinely removes inauthentic accounts identified by their algorithm or reported to them. But this process is imperfect; some accounts in TwitterGAN are still alive on Twitter despite violating the policies. Specifically, all accounts in the astroturf group appeared to be suspended when we checked their status in December 2023, but only 14.9% of the other accounts in TwitterGAN had been suspended.

4.2 Tactics and inauthentic activities

To uncover the tactics and inauthentic activities of accounts in TwitterGAN, we performed human annotation and bot detection. The annotation covered all accounts from blue_check, bart, elon_watch, sunday, chatgpt, and ind_acc. Since the accounts in astroturf share similar patterns, we randomly selected 250 for annotation to reduce the workload. In total, we annotated 604 accounts, and each account was annotated by two authors independently. Our analyses reveal several patterns, and we illustrate the five most noteworthy ones in Figure 3. Some patterns are common across multiple groups, while others are exclusive to specific ones.

Impersonation. The primary goal of using GAN-generated images in the profiles is to impersonate real humans. As illustrated in Figure 3(a), these fake accounts often feature realistic photos, plausible names, and detailed descriptions, creating the illusion that these accounts are operated by real persons. Our analysis suggests that almost all TwitterGAN accounts employ believable names, though some are uncommon. None of these accounts disclose their use of GAN-generated profile images.

Scamming. Our analysis suggests that the accounts in the elon_watch group are used to promote scams. These accounts typically reply to others with screenshots depicting Elon Musk’s Twitter feed, containing deceptive messages like “Just google [keywords] and thank me later.” An illustration can be found in Figure 3(b). The screenshots are fabricated; Elon Musk never posted these messages. Searching the keywords on Google leads to scam websites.

Spamming. Accounts within TwitterGAN frequently engage in spamming. We show two such examples from the bart group in Figure 3(c), characterized by repeated

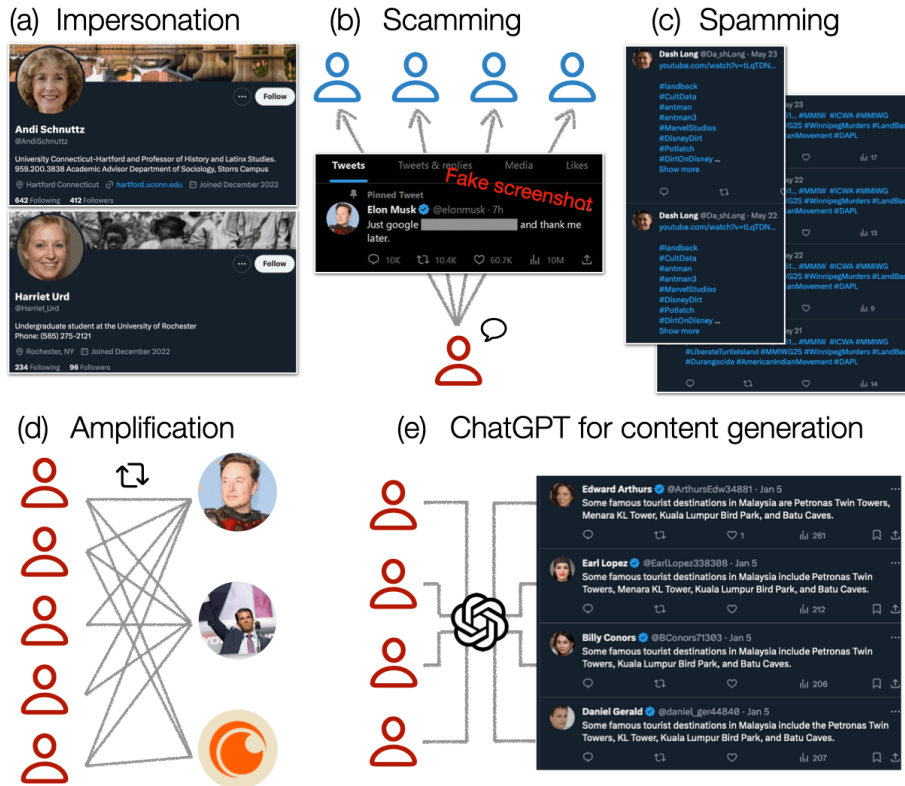


Figure 3: Illustrations of the tactics and inauthentic activities of accounts in TwitterGAN. (a) Fake accounts impersonating people that do not exist through GAN-generated profiles, human-like names, and fake descriptions. (b) Accounts using fake screenshots with scam messages to reply to other accounts (the keywords are redacted to avoid further spreading the scams). (c) Accounts posting spam messages repeatedly. (d) Accounts amplifying the messages from certain accounts by coordinately reposting their tweets. (e) Accounts leveraging ChatGPT to produce human-like replies. The account illustrations colored in red and blue represent inauthentic and authentic ones, respectively.

messages with excessive hashtags and external links to platforms like YouTube and Tumblr.

Coordinated amplification. Some of the TwitterGAN accounts exhibit coordinated inauthentic activity (see Figure 3(d)), a tactic commonly adopted by adversarial actors during influence campaigns (Pacheco et al. 2021). A clear example is the astroturf group, where the members retweet content from a common set of accounts. The five most retweeted accounts are @elonmusk (Elon Musk’s account), @DonaldJTrumpJr (Donald Trump Jr.’s account), @FoxNews (the official account of Fox News), @POTUS (the official account of the US President, currently owned by Joe Biden), and @Crunchyroll (the official account of <https://crunchyroll.com>, an anime distribution company). Their motives remain unclear, given the varied nature of the amplified content.

ChatGPT for content generation. Accounts in the chatgpt group were found to leverage ChatGPT to generate human-like content (see Figure 3(e)). They were identified by

the self-revealing tweets posted by these accounts, which contained phrases such as “as an AI language model.” This method was used to uncover a larger botnet that used the same tactic earlier (Yang and Menczer 2024).

Automation. Many accounts in TwitterGAN exhibit characteristics of bots. For instance, repeatedly sending scams, posting spam messages, and amplifying posts from selected accounts in a coordinated fashion are typical signs of automation. We further apply the BotometerLite model (K.-C. Yang et al. 2020), a widely adopted Twitter bot detection tool, to all accounts in TwitterGAN and find that 75.2% of them are labeled likely bots (using 0.5 as threshold).

“Verification.” Under Musk’s management, the acquisition of a blue checkmark no longer requires going through a verification process. Instead, any account willing to pay a premium fee can obtain the mark to increase its perceived legitimacy. This strategy is not very common among TwitterGAN accounts, possibly due to the associated cost. Only the 11 accounts in blue_check and 44 accounts in chatgpt adopted this strategy.

4.3 Heuristics for identification




Figure 4: Common distinctive defects of GAN-generated faces that can help identify them. (a) Unreal hat. (b) Glass frame blended into the face; surreal background. (c) Unreal collar. (d) Earrings blended into the ears; irregular ear.

During the annotation, we identified some heuristic indicators that can help spot GAN-generated profiles. The most distinctive feature lies in the flaws. Although GANs can generate realistic human faces, other parts of the images are often featured with defects. For instance, GANs sometimes have difficulty producing realistic hats (Figure 4(a)), glasses (Figure 4(b)), and ears (Figure 4(d)). Similarly, users can pay attention to accessories, such as earrings (Figure 4(d)) and necklaces, that GANs often struggle with. Clothing details, especially collars, are also problematic for GANs (Figure 4(c)). Last but

not least, the background sometimes can hint at the artificial nature of a picture (Figure 4(b)). While annotating TwitterGAN accounts, we specifically examined the defects in their profile pictures. Two annotators agreed that 60.1% of them have clear defects (Cohen's Kappa = 0.90).

Profile images lacking obvious defects can still be identified through inconsistencies between images and profile information. For instance, there may be a disconnect between the perceived gender or ethnicity in the image and the associated names. Such mismatches often result from creators randomly pairing images and names without validation. However, relying solely on this criterion can be inaccurate and potentially biased, so it should be used cautiously and complemented with additional verification.

Although it is challenging to determine the nature of suspicious accounts individually, patterns emerge when comparing multiple accounts. Users can start with one suspicious account and observe its connections and interactions with others to search for signs of coordinated inauthentic behavior. These include accounts sharing compatible profile pictures and/or exhibiting similar activities, such as posting similar content and retweeting similar posts. Such analysis can contextualize individual accounts within larger networks, aiding in identifying their true nature.

5 Prevalence on Twitter

In this section, we aim to estimate the prevalence of Twitter accounts using GAN-generated profiles. We first propose a simple yet effective detection method, then apply it to RandomTwitter.

5.1 Eye locations

Our detection method leverages the observation that GANs place the eyes in the same locations in all generated pictures, as demonstrated in Figure 1(a,b). This feature has been reported and exploited to distinguish true and fake human faces (X. Yang et al. 2019). Note that our task differs from previous methodology papers since our goal is to identify GAN-generated profiles from random Twitter profiles. As illustrated in Figure 1(c), many profiles do not feature human faces. And for those containing authentic human faces, we hypothesize that very few of them happen to have their eyes placed in the same locations as the GAN-generated ones.

To test our hypothesis, we need to obtain the eye coordinates of input images. We test different face recognition software and find that the Python package `face_recognition`⁵ provides an effective and efficient solution. The package returns an empty list if no faces are detected in the input images. For images containing faces, the package returns the coordinates of the landmarks, such as the eyes and nose, of each face (Kazemi and

5. https://github.com/ageitgey/face_recognition

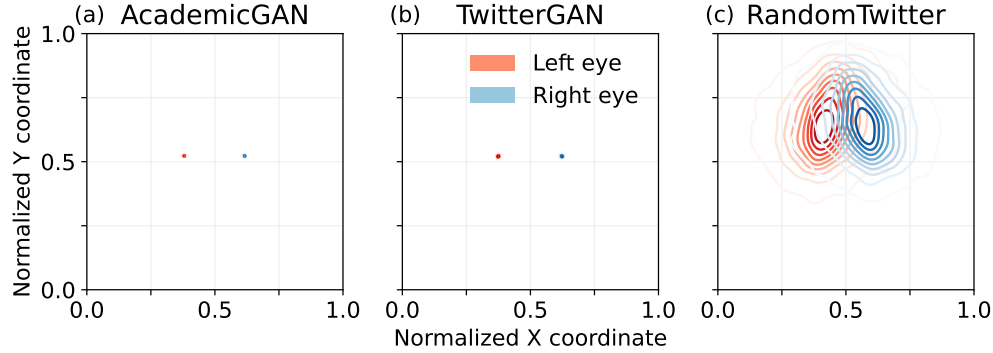


Figure 5: Kernel density estimation distribution of the detected left and right eyes in (a) AcademicGAN, (b) TwitterGAN, and (c) a random sample of 25,000 accounts from RandomTwitter.

Sullivan 2014). We focus on the eyes and obtain the center coordinate of each eye by averaging its landmark points. Since the profile images have different resolutions, we normalize the coordinates using the image height and width so that all the re-scaled coordinates are in the unit interval.

We apply the package to the profiles in TwitterGAN, AcademicGAN, and RandomTwitter. Figure 5 shows the kernel density estimation distributions of the detected eye positions. For the GAN-generated profiles in both TwitterGAN and AcademicGAN, the eye locations are highly concentrated and identical. On the other hand, the eye locations of random profiles in RandomTwitter are scattered across the space. These findings confirm our hypothesis and suggest that TwitterGAN images were likely generated by the same algorithm as those in AcademicGAN. In turn, this indicates that eye locations can be leveraged to distinguish GAN-generated and random Twitter profiles.

Inspecting the results of `face_recognition` yields a few more critical observations. The package can identify exactly one face from each AcademicGAN image and each TwitterGAN image, except for one with clear defects. On the other hand, the package can only identify faces from around 35.7% of the profiles in RandomTwitter. The remaining profiles either do not contain faces or are misclassified by the package. Among the profiles with detected faces, about 92.3% have one face detected, while the others have multiple faces detected. Although `face_recognition` is not perfect in face detection, it works very well on the GAN-generated faces, ensuring a high recall for our objective.

5.2 GANEyeDistance metric

Based on the promising results demonstrated in Figure 5, we propose a metric called GANEyeDistance to quantify the distances between the detected eyes and the expected coordinates of the eyes in GAN-generated pictures.

Let us use $\vec{L} = \{x_l, y_l\}$ and $\vec{R} = \{x_r, y_r\}$ to represent the coordinates of the left and right




Figure 6: (a) Distributions of the GANEyeDistance values in AcademicGAN, TwitterGAN, and RandomTwitter. The results for AcademicGAN and TwitterGAN overlap with each other. (b) A zoom-in of (a) with a focus on the range of [0, 0.1]. The dashed line indicates the threshold we choose for identifying likely GAN-generated profiles in RandomTwitter.

eyes detected by `face_recognition`, where x and y represent the normalized x - and y -coordinates. We further use \vec{L}_{GAN} and \vec{R}_{GAN} to denote the expected locations of left and right eyes in GAN-generated images. The GANEyeDistance for a given profile is defined as:

$$\mathcal{G}(\vec{L}, \vec{R}) = \begin{cases} \frac{\|\vec{L} - \vec{L}_{GAN}\| + \|\vec{R} - \vec{R}_{GAN}\|}{2\sqrt{2}} & \text{if one detected face} \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

where $\|\vec{X} - \vec{Y}\|$ represents the Euclidean distance between \vec{X} and \vec{Y} . Since the longest distance of two points in a unit square is $\sqrt{2}$ and we combine the distance of both eyes, the denominator in Eq. (1) ensures that \mathcal{G} is a number between 0 and 1. A profile with a \mathcal{G} value close to 0 means the eye locations are close to those of GAN-generated faces.

Note that \mathcal{G} is only calculated when one face is detected in the given profile. If `face_recognition` cannot process the input image or cannot detect any faces, we set \mathcal{G} to 1. We also set \mathcal{G} to 1 when multiple faces are detected, since GAN-generated profiles typically have one face in them.

To apply Eq. (1), we also need the ground truth values of \vec{L}_{GAN} and \vec{R}_{GAN} . Inspecting the profiles in TwitterGAN suggests that the profile pictures are subject to modifications such as cropping and stretching, introducing noise. On the other hand, pictures in AcademicGAN have better consistency and quality, so we use the average values of their eye coordinates as our ground truth. Our implementation of GANEyeDistance is publicly available.⁶

To demonstrate the effectiveness of GANEyeDistance, we calculate this value for every image in all three datasets and show their distributions in Figure 6(a). We can see that the profiles in AcademicGAN and TwitterGAN have GANEyeDistance values very close to 0, as expected. The results of RandomTwitter show a bimodal distribution:

6. https://github.com/osome-iu/fake_gan_accounts

the GANEyeDistance values of the profiles with one detected face are distributed across a wide range, and the remaining images have a GANEyeDistance value of 1. The difference in the distributions suggests that GANEyeDistance can effectively distinguish GAN-generated profiles from random Twitter profiles. In Figure 6(b), we zoom in to better inspect the difference among the three datasets. The GANEyeDistance values for AcademicGAN are closest to 0. The results for TwitterGAN deviate from 0 slightly due to noise. The AcademicGAN distribution has a very small overlap with those of the GAN-generated profiles.

Despite the promising results shown in Figure 6, we emphasize again that GANEyeDistance only indicates the distance between the locations of detected eyes and the ground truth. It cannot truly distinguish GAN-generated faces from authentic human faces. This means that although GAN-generated faces have GANEyeDistance close to zero, images with very small GANEyeDistance are not necessarily GAN-generated. In other words, a relatively loose threshold for GANEyeDistance can yield a very high recall for GAN-generated profiles but also introduce false positives.

For instance, if we use 0.02 as the threshold and consider all images with $\mathcal{G} < 0.02$ as positive cases, then the recall for AcademicGAN and TwitterGAN is 100% and 99.4%, respectively. At the same time, 0.46% (1,181) of the random profiles are labeled as positive, so further inspection is needed to determine their nature.

5.3 Prevalence estimation

In this study, our goal is not to build a perfect algorithm to detect GAN-generated profiles. Instead, we aim to estimate the prevalence of accounts using GAN-generated profiles on Twitter, which can be approximated by the prevalence in RandomTwitter. Therefore, the high recall nature of GANEyeDistance makes it a suitable metric to help us narrow down the candidate profiles for further inspection.

Here, we settle on 0.02 as the threshold because it yields a very high recall for TwitterGAN. Increasing the threshold does not boost the recall significantly but introduces many more false positives. As mentioned above, this threshold leads to 1,181 candidate profiles from RandomTwitter.

To reduce the false positives, where the algorithm flags authentic human photos as GAN-generated, two authors manually classify these profiles into three categories: (1) highly likely GAN-generated, (2) likely GAN-generated, and (3) not GAN-generated. Profiles in the first category demonstrate clear and unique features of GAN-generated human faces, such as the defects illustrated in Figure 4. Profiles in the second category are very similar to GAN-generated faces; however, no definitive evidence can be identified.

The Cohen's Kappa for the ratings is 0.85, indicating high agreement between the two annotators. Fifty-four profiles are coded as highly likely generated by GANs by both annotators, and 113 are believed to be likely or highly likely generated by GANs by

both annotators. Profiles are considered not GAN-generated if at least one annotator believes so. Considering that these profiles are selected from the 254,275 accounts in RandomTwitter and our method might miss some positive cases, we estimate a lower bound on the prevalence of active Twitter users using GAN-generated profiles to be 0.021%–0.044%. These accounts were responsible for 0.022–0.053% of the tweets in RandomTwitter.

6 Discussion

In summary, we present a systematic analysis of fake accounts using GAN-generated faces on Twitter. We show that they impersonate humans with realistic profiles and are involved in various inauthentic activities, ranging from spamming and scamming to coordinated amplification of messages. We also provide some heuristics that social media users can use to identify the accounts they encounter.

Using a random sample, we further estimate the prevalence of such accounts on Twitter. We first build a heuristic method exploiting the common feature of GAN-generated faces to narrow down the potential GAN-generated profiles, then perform manual annotation to identify those likely to be synthetic. Our results suggest that the percentage of active Twitter accounts using GAN-generated profiles is between 0.021% and 0.044%. To put the percentages into perspective, let us consider the total number of daily active Twitter users. Pfeffer et al. (2023) collected all 375 million tweets posted on September 21, 2022, and found 40,199,195 active users. Using this number, we estimate that at least 8,537–17,864 daily active accounts on Twitter use GAN-generated profiles.

Since the recall of our machine learning model is not 100%, and GAN-generated images that are heavily stretched or tilted can evade our method, we might miss some positive cases in our analysis. Additionally, the labeling criteria used in our annotation process are very conservative—we only consider a profile positive when two annotators are both highly certain about the label. Therefore, these estimates should be considered as a lower bound. In fact, a recent report by Ricker et al. (2024) produced a similar but slightly higher estimate of 0.052% using different detection methods and excluding accounts with default profile images in the denominator.

Our analysis and estimation have some limitations. We only include a few groups of accounts in TwitterGAN, so their behavior patterns might not generalize to other cases. Although we have proposed an effective method to capture more accounts using GAN-generated profiles in the wild, further analysis is not feasible since Twitter ended free access to its API. Also, since we only focus on one platform, the findings might not be generalized to other platforms, such as Facebook and LinkedIn, that have been reported to suffer from the same issue. Unfortunately, limited access to their data makes it impossible for researchers to study inauthentic accounts on those platforms. The Digital

Services Act by the European Union mandates that large social media platforms, including Twitter and Meta, provide researchers with data access, bringing hope.⁷ Nonetheless, it remains uncertain whether researchers can obtain the necessary data to expand on this research.

The heuristics and detection methods introduced in our study do not necessarily work on other generative AI models. For example, diffusion models (Song and Ermon 2019; Ho, Jain, and Abbeel 2020) yield better image-generation outcomes than GANs (Dhariwal and Nichol 2021; L. Yang et al. 2023). This architecture has been adopted by many AI companies, including OpenAI (Ramesh et al. 2022), Midjourney,⁸ and Stability AI,⁹ to create advanced multimodal models. The images generated by these models are astonishing and free from the issues of GAN-generated ones. Preliminary studies on detecting diffusion model-generated images suggest that the techniques that work on GAN-generated images are no longer effective (Ricker et al. 2022), requiring new detection paradigms.

Despite these limitations, our findings provide empirical evidence of the threats posed by generative AI. We show that GAN-generated images have already been deployed at scale to create large amounts of convincing fake accounts. Many more social media accounts likely already use more advanced AI models to generate more realistic fake profiles. We even uncovered a group of accounts that also leveraged ChatGPT to generate human-like text and conduct realistic interactions (Yang and Menczer 2024). Combining these AI models can lead to more convincing and realistic fake accounts on social media. In fact, recent research suggests that advanced AI models enable the creation of autonomous agents that can ingest multimodal inputs, process the data, reason, and then leverage various tools to achieve their objectives (Xi et al. 2023). Despite the usefulness of these intelligent agents, the risk of adversarial actors abusing them is high.

The looming threats call for new intervention strategies. First, we need effective models that can detect AI-generated images. A recent study finds that assistance from such models can improve people's performance in identifying GAN-generated images (Boyd et al. 2023). However, the research community might struggle to catch up, as these cutting-edge generative models evolve quickly. Another approach requires regulations on the use of generative AI models to create inauthentic accounts and content. For instance, social media platforms can require an account to provide evidence of their authenticity before exposing them and their content to a broader population (Menczer et al. 2023). Finally, we believe that teaching social media users AI literacy is crucial. Understanding the capability and weaknesses of state-of-the-art generative AI models can help social media users better defend themselves from manipulation by AI-powered fake accounts. However, such initiatives should be carried out carefully to avoid unintended consequences (Yan et al. 2023).

7. https://ec.europa.eu/commission/presscorner/detail/en/IP_23_2413

8. <https://docs.midjourney.com>

9. <https://stability.ai/stable-image>

References

- Boato, Giulia, 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 *2022 IEEE International Joint Conference on Biometrics*, 1–7. IEEE. <https://doi.org/10.1109/IJCB54206.2022.10007988>.
- Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, et al. 2021. *On the Opportunities and Risks of Foundation Models*, August 16, 2021. <https://doi.org/10.48550/arXiv.2108.07258>. arXiv: 2108.07258 [cs.LG].
- Bond, Shannon. 2022. "That Smiling LinkedIn Profile Face Might Be a Computer-generated Fake." *NPR*, March 27, 2022. <https://www.npr.org/2022/03/27/1088140809/fake-linkedin-profiles>.
- Boyd, Aidan, Patrick Tinsley, Kevin Bowyer, and Adam Czajka. 2023. "The Value of AI Guidance in Human Examination of Synthetically-Generated Faces." *Proceedings of the AAAI Conference on Artificial Intelligence* 37, no. 5 (June): 5930–38. <https://doi.org/10.1609/aaai.v37i5.25734>.
- Bray, Sergi D., Shane D. Johnson, and Bennett Kleinberg. 2023. "Testing Human Ability to Detect 'Deepfake' Images of Human Faces." *Journal of Cybersecurity* 9, no. 1 (June): tyad011. <https://doi.org/10.1093/cybsec/tyad011>.
- Cao, Yihan, Siyu Li, Yixin Liu, Zhiling Yan, Yutong Dai, Philip S. Yu, and Lichao Sun. 2023. *A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT*, March 7, 2023. <https://doi.org/10.48550/arXiv.2303.04226>. arXiv: 2303.04226 [cs.AI].
- Confessore, Nicholas, Gabriel J.X. Dance, Richard Harris, and Mark Hansen. 2018. "The Follower Factory." *The New York Times*, January 27, 2018. <https://www.nytimes.com/interactive/2018/01/27/technology/social-media-bots.html>.
- Cresci, Stefano, Marinella Petrocchi, Angelo Spognardi, and Stefano Tognazzi. 2021. "The Coming Age of Adversarial Social Bot Detection." *First Monday* 26, no. 7 (May). <https://doi.org/10.5210/fm.v26i7.11474>.
- Dang, L. Minh, Syed Ibrahim Hassan, Suhyeon Im, Jaecheol Lee, Sujin Lee, and Hyeonjoon Moon. 2018. "Deep Learning Based Computer Generated Face Identification Using Convolutional Neural Network." *Applied Sciences* 8 (12): 2610. <https://doi.org/10.3390/app8122610>.
- Dhariwal, Prafulla, and Alexander Nichol. 2021. "Diffusion Models Beat GANs on Image Synthesis." *Advances in Neural Information Processing Systems* 34:8780–94. https://proceedings.neurips.cc/paper_files/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf.

- Ferrara, Emilio, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. "The Rise of Social Bots." *Communications of the ACM* 59 (7): 96–104. <https://doi.org/10.1145/2818717>.
- Goldstein, Josh A., Girish Sastry, Micah Musser, Renée DiResta, Matthew Gentzel, and Katerina Sedova. 2023. *Generative Language Models and Automated Influence Operations: Emerging Threats and Potential Mitigations*, January 10, 2023. <https://doi.org/10.48550/arXiv.2301.04246>. arXiv: 2301.04246 [cs.CY].
- Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. "Generative Adversarial Nets." *Advances in Neural Information Processing Systems* 27. https://proceedings.neurips.cc/paper_files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf.
- Guembe, Blessing, Ambrose Azeta, Sanjay Misra, Victor Chukwudi Osamor, Luis Fernandez-Sanz, and Vera Pospelova. 2022. "The Emerging Threat of AI-driven Cyber Attacks: A Review." *Applied Artificial Intelligence* 36 (1): 2037254. <https://doi.org/10.1080/08839514.2022.2037254>.
- Guo, Hui, Shu Hu, Xin Wang, Ming-Ching Chang, and Siwei Lyu. 2022. "Eyes Tell All: Irregular Pupil Shapes Reveal GAN-generated Faces." In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2904–8. IEEE. <https://doi.org/10.1109/ICASSP43922.2022.9746597>.
- Ho, Jonathan, Ajay Jain, and Pieter Abbeel. 2020. "Denoising Diffusion Probabilistic Models." *Advances in Neural Information Processing Systems* 33:6840–51. https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.
- Hu, Shu, Yuezun Li, and Siwei Lyu. 2021. "Exposing GAN-generated Faces Using Inconsistent Corneal Specular Highlights." In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2500–2504. IEEE. <https://doi.org/10.1109/ICASSP39728.2021.9414582>.
- Jakesch, Maurice, Jeffrey T. Hancock, and Mor Naaman. 2023. "Human Heuristics for AI-generated Language Are Flawed." *Proceedings of the National Academy of Sciences* 120 (11): e2208839120. <https://doi.org/10.1073/pnas.2208839120>.
- Karras, Tero, Samuli Laine, and Timo Aila. 2019. "A Style-based Generator Architecture for Generative Adversarial Networks." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4401–10. <https://doi.org/10.1109/CVPR.2019.00453>.
- Karras, Tero, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. "Analyzing and Improving the Image Quality of StyleGAN." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8110–19. <https://doi.org/10.1109/CVPR42600.2020.00813>.

- Kazemi, Vahid, and Josephine Sullivan. 2014. "One Millisecond Face Alignment with an Ensemble of Regression Trees." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1867–74. <https://doi.org/10.1109/CVPR.2014.241>.
- Lago, Federica, Cecilia Pasquini, Rainer Böhme, Hélène Dumont, Valérie Goffaux, and Giulia Boato. 2021. "More Real Than Real: A Study on Human Visual Perception of Synthetic Faces [Applications Corner]." *IEEE Signal Processing Magazine* 39 (1): 109–16. <https://doi.org/10.1109/MSP.2021.3120982>.
- Liu, Ziwei, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. "Deep Learning Face Attributes in the Wild." In *Proceedings of the IEEE International Conference on Computer Vision*, 3730–38. <https://doi.org/10.1109/ICCV.2015.425>.
- Marra, Francesco, Diego Gagnaniello, Davide Cozzolino, and Luisa Verdoliva. 2018. "Detection of GAN-generated Fake Images over Social Networks." In *2018 IEEE Conference on Multimedia Information Processing and Retrieval*, 384–89. IEEE. <https://doi.org/10.1109/MIPR.2018.00084>.
- Martineau, Paris. 2019. "Facebook Removes Accounts with AI-Generated Profile Photos." *Wired*, December 28, 2019. <https://www.wired.com/story/facebook-removes-accounts-ai-generated-photos>.
- Menczer, Filippo, David Crandall, Yong-Yeol Ahn, and Apu Kapadia. 2023. "Addressing the Harms of AI-generated Inauthentic Content." *Nature Machine Intelligence* 5:679–80. <https://doi.org/10.1038/s42256-023-00690-w>.
- Mundra, Shivansh, Gonzalo J. Aniano Porcile, Smit Marvaniya, James R. Verbus, and Hany Farid. 2023. "Exposing GAN-Generated Profile Photos from Compact Embeddings." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 884–92. <https://doi.org/10.1109/CVPRW59228.2023.00095>.
- Najari, Shaghayegh, Mostafa Salehi, and Reza Farahbakhsh. 2022. "GANBOT: A GAN-based Framework for Social Bot Detection." *Social Network Analysis and Mining* 12:1–11. <https://doi.org/10.1007/s13278-021-00800-9>.
- Nataraj, Lakshmanan, Tajuddin Manhar Mohammed, Shivkumar Chandrasekaran, Arjuna Flenner, Jawadul H. Bappy, Amit K. Roy-Chowdhury, and B. S. Manjunath. 2019. *Detecting GAN Generated Fake Images Using Co-occurrence Matrices*, March 15, 2019. <https://doi.org/10.48550/arXiv.1903.06836>. arXiv: 1903.06836 [cs.CV].
- Nightingale, Sophie J., and Hany Farid. 2022. "AI-synthesized Faces are Indistinguishable from Real Faces and More Trustworthy." *Proceedings of the National Academy of Sciences* 119 (8): e2120481119. <https://doi.org/10.1073/pnas.2120481119>.
- Norteño, Conspirador. 2023. "A Brief and Partial History of Generative AI on Social Media." *Conspirador Norteño* (blog), November 19, 2023. <https://conspirador0.substack.com/p/a-brief-and-partial-history-of-generative>.

- O'Sullivan, Donie. 2020a. "A High School Student Created a Fake 2020 Candidate. Twitter Verified It." *CNN*, February 28, 2020. <https://www.cnn.com/2020/02/28/tech/fake-twitter-candidate-2020>.
- . 2020b. "How Fake Faces Are Being Weaponized Online." *CNN*, February 20, 2020. <https://www.cnn.com/2020/02/20/tech/fake-faces-deepfake>.
- Pacheco, Diogo, Pik-Mai Hui, Christopher Torres-Lugo, Bao Tran Truong, Alessandro Flammini, and Filippo Menczer. 2021. "Uncovering Coordinated Networks on Social Media: Methods and Case Studies." In *Proceedings of the International AAAI Conference on Web and Social Media*, 15:455–66. <https://doi.org/10.1609/icwsm.v15i1.18075>.
- Pfeffer, Jürgen, Daniel Matter, Kokil Jaidka, Onur Varol, Afra Mashhadi, Jana Lasser, Dennis Assenmacher, Siqi Wu, Diyi Yang, Cornelia Brantner, et al. 2023. "Just Another Day on Twitter: A Complete 24 Hours of Twitter Data." In *Proceedings of the International AAAI Conference on Web and Social Media*, 17:1073–81. <https://doi.org/10.1609/icwsm.v17i1.22215>.
- Pichai, Sundar, and Demis Hassabis. 2023. "Introducing Gemini: Our Largest and Most Capable AI Model." *Google Blog* (blog), November 6, 2023. <https://blog.google/technology/ai/google-gemini-ai>.
- Ramesh, Aditya, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. *Hierarchical Text-Conditional Image Generation with CLIP Latents*, April 13, 2022. <https://doi.org/10.48550/arXiv.2204.06125>. arXiv: 2204.06125 [cs.CV].
- Ricker, Jonas, Dennis Assenmacher, Thorsten Holz, Asja Fischer, and Erwin Quiring. 2024. *AI-Generated Faces in the Real World: A Large-Scale Case Study of Twitter Profile Images*, April 22, 2024. <https://doi.org/10.48550/arXiv.2404.14244>. arXiv: 2404.14244 [cs.CR].
- Ricker, Jonas, Simon Damm, Thorsten Holz, and Asja Fischer. 2022. *Towards the Detection of Diffusion Model Deepfakes*, October 26, 2022. <https://doi.org/10.48550/arXiv.2210.14571>. arXiv: 2210.14571 [cs.CV].
- Sayyadiharikandeh, Mohsen, Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer. 2020. "Detection of Novel Social Bots by Ensembles of Specialized Classifiers." In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2725–32. <https://doi.org/10.1145/3340531.3412698>.
- Song, Yang, and Stefano Ermon. 2019. "Generative Modeling by Estimating Gradients of the Data Distribution." *Advances in Neural Information Processing Systems* 32. https://proceedings.neurips.cc/paper_files/paper/2019/file/3001ef257407d5a371a96dcd947c7d93-Paper.pdf.
- Wang, Xin, Hui Guo, Shu Hu, Ming-Ching Chang, and Siwei Lyu. 2023. *GAN-generated Faces Detection: A Survey and New Perspectives*, February 15, 2023. <https://doi.org/10.48550/arXiv.2202.07145>. arXiv: 2202.07145 [cs.CV].

- Xi, Zhiheng, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, et al. 2023. *The Rise and Potential of Large Language Model Based Agents: A Survey*, September 14, 2023. <https://doi.org/10.48550/arXiv.2309.07864>. arXiv: 2309.07864 [cs.AI].
- Yamin, Muhammad Mudassar, Mohib Ullah, Habib Ullah, and Basel Katt. 2021. "Weaponized AI for Cyber Attacks." *Journal of Information Security and Applications* 57:102722. <https://doi.org/10.1016/j.jisa.2020.102722>.
- Yan, Harry Yaojun, Kai-Cheng Yang, James Shanahan, and Filippo Menczer. 2023. "Exposure to Social Bots Amplifies Perceptual Biases and Regulation Propensity." *Scientific Reports* 13 (1): 20707. <https://doi.org/10.1038/s41598-023-46630-x>.
- Yang, Kai-Cheng, and Filippo Menczer. 2024. "Anatomy of an AI-powered Malicious Social Botnet." *Journal of Quantitative Description: Digital Media* 4 (May). <https://doi.org/10.51685/jqd.2024.icwsm.7>.
- Yang, Kai-Cheng, Onur Varol, Clayton A. Davis, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2019. "Arming the Public with Artificial Intelligence to Counter Social Bots." *Human Behavior and Emerging Technologies* 1 (1): 48–61. <https://doi.org/10.1002/hbe2.115>.
- Yang, Kai-Cheng, Onur Varol, Pik-Mai Hui, and Filippo Menczer. 2020. "Scalable and Generalizable Social Bot Detection through Data Selection." *Proceedings of the AAAI Conference on Artificial Intelligence* 34 (01): 1096–103. ISSN: 2374-3468. <https://doi.org/10.1609/aaai.v34i01.5460>.
- Yang, Ling, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. 2023. "Diffusion Models: A Comprehensive Survey of Methods and Applications." *ACM Computing Surveys* 56 (4): 1–39. <https://doi.org/10.1145/3626235>.
- Yang, Xin, Yuezun Li, Honggang Qi, and Siwei Lyu. 2019. "Exposing GAN-synthesized Faces Using Landmark Locations." In *Proceedings of the ACM Workshop on Information Hiding and Multimedia Security*, 113–18. <https://doi.org/10.1145/3335203.3335724>.
- Ye, Junjie, Xuanting Chen, Nuo Xu, Can Zu, Zekai Shao, Shichun Liu, Yuhan Cui, et al. 2023. *A Comprehensive Capability Analysis of GPT-3 and GPT-3.5 Series Models*, March 18, 2023. <https://doi.org/10.48550/arXiv.2303.10420>. arXiv: 2303.10420 [cs.CL].

Authors

Kai-Cheng Yang (yang3kc@gmail.com) is a postdoctoral researcher at Northeastern University. He obtained his PhD in Informatics from Indiana University Bloomington.

Danishjeet Singh is an undergraduate student majoring in Computer Science and a research assistant at the Observatory on Social Media at Indiana University Bloomington.

Filippo Menczer is the Luddy Distinguished Professor of Informatics and Computer Science at the Luddy School of Informatics, Computing, and Engineering and the director of the Observatory on Social Media at Indiana University Bloomington.

Acknowledgements

We thank <https://generated.photos> for sharing their datasets for our research and Twitter user @conspirator0 for flagging candidate accounts using GAN-generated profile images.

Data availability statement

GAN-generated human faces from AcademicGAN can be obtained from <https://generated.photos/datasets> upon request. Other datasets used in this study are publicly available at <https://zenodo.org/doi/10.5281/zenodo.10436888>.

Funding statement

This work was supported in part by the Volkswagen Foundation as part of the “Bots Building Bridges” project, the Knight Foundation, DARPA (grant HR001121C0169), and the Luddy School of Informatics, Computing, and Engineering at Indiana University.

Ethical standards

This is not a human subject study, so it does not require IRB approval.

Keywords

Fake accounts; deepfake; GAN, social media; social bots; coordinated inauthentic activities.