



11-6-2023

The Age of Synthetic Realities: Challenges and Opportunities

João Phillipe Cardenuto
State University of Campinas, Brazil

Jing Yang
State University of Campinas, Brazil

Rafael Padilha
State University of Campinas, Brazil

Renjie Wan
Hong Kong Baptist University

Daniel Moreira
Loyola University Chicago, dmoreira1@luc.edu

Follow this and additional works at https://ecommons.luc.edu/cs_facpubs



Part of the [Computer Sciences Commons](#)

Recommended Citation

Cardenuto, João Phillipe; Yang, Jing; Padilha, Rafael; Wan, Renjie; Moreira, Daniel; Li, Haoliang; Wang, Shiqi; Andaló, Fernanda; Marcel, Sébastien; and Rocha, Anderson. The Age of Synthetic Realities: Challenges and Opportunities. *APSIPA Transactions on Signal and Information Processing*, 12, 1: 1–62, 2023. Retrieved from Loyola eCommons, Computer Science: Faculty Publications and Other Works, <http://dx.doi.org/10.1561/116.00000138>

This Article is brought to you for free and open access by the Faculty Publications and Other Works by Department at Loyola eCommons. It has been accepted for inclusion in Computer Science: Faculty Publications and Other Works by an authorized administrator of Loyola eCommons. For more information, please contact ecommons@luc.edu.



This work is licensed under a [Creative Commons Attribution 4.0 International License](#).
© The Authors, 2023.

Authors

João Phillipe Cardenuto, Jing Yang, Rafael Padilha, Renjie Wan, Daniel Moreira, Haoliang Li, Shiqi Wang, Fernanda Andaló, Sébastien Marcel, and Anderson Rocha

Overview Paper

The Age of Synthetic Realities: Challenges and Opportunities

João Phillipe Cardenuto^{1*}, Jing Yang¹, Rafael Padilha¹, Renjie Wan², Daniel Moreira³, Haoliang Li⁴, Shiqi Wang⁵, Fernanda Andaló¹, Sébastien Marcel^{6,7} and Anderson Rocha¹

¹*Artificial Intelligence Lab., [Recod.ai](https://www.recod.ai), Institute of Computing, Universidade Estadual de Campinas, Campinas, SP, Brazil*

²*Department of Computer Science, Hong Kong Baptist University, Hong Kong*

³*Department of Computer Science, Loyola University Chicago, USA*

⁴*Department of Electrical Engineering, City University of Hong Kong, Hong Kong*

⁵*Department of Computer Science, City University of Hong Kong, Hong Kong*

⁶*Idiap Research Institute, Martigny, Switzerland*

⁷*University of Lausanne, Lausanne, Switzerland*

ABSTRACT

Synthetic realities are digital creations or augmentations that are contextually generated through the use of Artificial Intelligence (AI) methods, leveraging extensive amounts of data to construct new narratives or realities, regardless of the intent to deceive. In this paper, we delve into the concept of synthetic realities and their implications for Digital Forensics and society at large within the rapidly advancing field of AI. We highlight the crucial need for the development of forensic techniques capable of identifying harmful synthetic creations and distinguishing them from reality. This is especially important in scenarios involving the creation and dissemination of fake news, disinformation, and misinformation. Our focus extends to various forms of media, such as images, videos, audio, and text, as we examine how synthetic realities are crafted and explore approaches to detecting these malicious creations. Additionally,

*Corresponding author: João Phillipe Cardenuto, phillipe.cardenuto@ic.unicamp.br.

we shed light on the key research challenges that lie ahead in this area. This study is of paramount importance due to the rapid progress of AI generative techniques and their impact on the fundamental principles of Forensic Science.

1 Introduction

In the last decade, there has been a growing expectation that Artificial Intelligence (AI) companies and researchers would dedicate their efforts to integrating humans into the digital realm or the so-called Metaverse. However, while technologies like Augmented Reality (AR) and Virtual Reality (VR) have been topics of discussion for a long time, it is only recently that significant technological advancements have made it possible to materialize such systems.

This version of a *synthetic reality* has been technically discussed at least since the 90s [69], and this was the main vision of where AI and related technologies would take us. However, what came as a less expected outcome is that these very technologies would inundate our physical world with content and creations, profoundly transforming our interactions with the virtual realm and reshaping how we engage with one another.

This more complex notion of *synthetic reality* has been a topic of discussion by the greatest minds of our time [272]. On the one hand, some hold a positive perspective, recognizing the immense advantages it can bring in domains such as automation, healthcare, and innovation. On the other, a group expresses concerns about the potential perils posed by AI, such as the generation of propaganda and untruth. They even advocate for temporary halts in AI experimentation within laboratories [12] to allow time for legal and ethical considerations to align with the pace of progress.

We can adopt a more pragmatic perspective and carefully embrace this emerging paradigm. This entails rapidly adapting ourselves and our societies and understanding and revitalizing our scientific endeavors. Hence, we redefine the term “synthetic realities” herein as any contextual digital creation or augmentation enabled by artificial intelligence methods. These techniques/models draw upon massive amounts of data leading to a new “reality” or narrative regardless of its intention to deceive the individual interacting with it. When the synthetic creation harms individuals, minorities, human rights, or the rule of law, it is paramount to devise forensic techniques to pinpoint such creations and separate what is real from what is synthetic. As an example, consider the creation of a fake news piece. Someone could fabricate a story from scratch using a chatbot, illustrate it with a synthetic image and a video, and then broadcast it to the world as if it were real via social media.

Consequently, Forensic Science has been continually adapting to these evolving circumstances. Rooted in the foundational principle that “every contact leaves a trace”, coined by researcher Edmond Locard [34], Forensic Science asserts that every interaction between individuals, objects, and places leaves behind a trail of evidence. While this concept was initially centered around physical traces like fingerprints, footprints, and blood, it has recently expanded to encompass digital counterparts such as photos, audio, video, and social media posts [190].

The revolving question around digital evidence is: “Are they fake or not?”, as manipulating these multimedia assets can be quickly done with simple and inexpensive tools. Moreover, credible manipulations can be used to fabricate more believable multimedia stories. Research in Multimedia Forensics has yielded important approaches to detecting altered media [67]. More specifically, progress has been made in analyzing digital media (image, video, and audio), identifying manipulations [2, 16, 42, 45, 123], tracing provenance [11, 20, 28, 47, 249], and establishing links with other digital evidence [15, 146, 178].

However, the emerging concept of *synthetic reality* paints an even more unsettling scenario: around 90% of the digital content will be synthetic in the upcoming years, meaning that almost all content will be generated synthetically by definition [57]. The distinction between what is genuine and what is fake takes on a new meaning. This phenomenon becomes evident in various domains, including movies [109], social media [98, 248], marketing [214], and education [169].

Notably, companies are now exploring the adoption of AI-generated models to promote their products [214] or employing AI to simulate eye contact in video conferencing software, enhancing the sense of connection during remote interactions [98]. Schools worldwide are banning chatbots and the use of generative AI on their networks in response to concerns about students submitting unauthentic and potentially plagiarized work [169]. The examples are many when thinking about how AI is shaping our reality.

Therefore, Forensic Science has to adapt yet again to this new reality. To expose synthetic content and tell apart malicious from harmless manipulations or even positive creations, there is one new key element: context. Contextual information can be leveraged to understand the semantics behind media objects. This can empower fact-checking solutions that mitigate the effect of falsified news, misinformation, and false political propaganda. To that end, and following a cognitive science interpretation, we can view digital objects through three perspectives: technological artifacts, sources of information, and platforms to convey ideas.

Traditional forensic techniques thus far have primarily focused on the first perspective. Analysts examined an asset as a digital signal and aimed to detect any possible artifact related to pixel-level or physical-level inconsistencies (e.g., concerning compression, sensor noise, illumination, shadows) to establish

its authenticity. The second perspective pertains to standard fact-checking procedures going beyond multimedia forensics. When treating an object as a source of information, it is essential to identify (or know) the acquisition device and to identify the time and location where it was produced as basic steps towards a fact-checking effort. The third perspective considers a digital object as a platform for conveying ideas. Determining the intentional goal of the asset leads to answering the question of why something happened. The answer to this question and the result of forensic analyses from the other two perspectives can reveal the ultimate goal of the falsified information. For example, it can help identify if there is an ongoing campaign to bias public opinion, influence the mood of a social group, or even incite a group to articulate plans for violent acts.

The latest advancements in AI have compelled forensic techniques to navigate the intricacies between these perspectives. Taking this into consideration, we focus on studying synthetic realities in different forms of media: images, videos, audio, and text. In the remainder of this paper, we discuss how synthetic media is created, considering each of these modalities, and the implications for Digital Forensics when such creations intend to harm third parties in different ways. We explore how to detect such malicious creations and pinpoint key research challenges that lie ahead. This is particularly significant due to the remarkable progress of AI generative techniques in generating realistic content and effectively concealing the typical artifacts left behind during the creation process. Each new generative method aims at creating ever-more-believable realities, thus directly colliding with Locard's principle, the cornerstone of forensics.

2 Synthetic Images

The proliferation of sophisticated synthetic and manipulated media has captured people's attention worldwide. For forensic researchers, this surge in synthetic reality evokes the daunting scenario reminiscent of the early days of Digital Forensics, where image editing was recognized as a powerful tool capable of altering reality [179]. While traditional manual image-editing software like Photoshop and GIMP continue to improve, they are being overshadowed (or enhanced, in the particular case of Photoshop [251]) by the emergence of powerful AI-based generative techniques. Today, it has become remarkably effortless to transform a simple concept or idea into a realistic image, with no requirement for drawing or painting skills to produce stunning, high-quality results.

Amid this rapidly evolving landscape lies generative models. Generative images have outstanding widespread applications in entertainment, as reviving legendary artists [215]; healthcare, as aiding surgeons in developing new

abilities [19]; and accessible tools, as serving people with disabilities [164]. However, many other harmful uses have been reported, such as nonconsensual DeepFake porn [185], misinformation generation [276], and sophisticated types of scams [253]. Scientific integrity researchers are also concerned that such technology would create fraudulent synthetic images in science [78, 198] and the medical area [160, 174], in particular.

Given the alarming potential for misuse of generative models in creating synthetic realities, this section delves into state-of-the-art generative models and the detection of AI-generated images, providing perspectives on the future of synthetic images.

2.1 Image Synthesis

The accelerated research on generative approaches in recent years gave birth to a plethora of AI models and techniques that are developed and open-sourced to the community. They are coupled with easy-to-use environments and applications [54, 172], allowing users to freely explore and share their creations. The increased accessibility further expanded the hype in image synthesis, fostering novel use cases and commercial applications that range from *outpainting* famous art pieces [140] (i.e., synthetically extending the borders of an image) to designing political campaign ads [235]. With the increased interest in the topic from research and industry communities and the rapid development of techniques, one can safely assume that not all synthetic images are born equally. We can categorize existing approaches by how the generation task is conditioned and what family of AI models they rely on.

The generation task defines the goal of the method and, consequently, how it learns to map the expected input to a synthetic output image. Additionally, the input data modality conditions the generation process into expressing particular visual concepts and characteristics desired by the user [294]. The most common tasks fall into *text-to-image* or *image-to-image* generation. Popularized by recent applications such as MidJourney [172] and DreamStudio [54], *text-to-image* generation involves a natural language prompt describing the desired image. This often includes the object or concepts that should be created, the desired artistic style, and the feeling the composition should convey. Whereas, in *image-to-image* generation, guidance is provided through visual cues like images, semantic maps, or body pose keypoints. These cues facilitate the generation process with information that can be challenging to convey through natural language prompts. For instance, they help capture relative element positioning. Examples of conditioning tasks and modalities are shown in Figure 2, including tasks like inpainting (reconstructing missing parts of an image), image translation (converting representations while preserving content), outpainting (generating beyond image boundaries), style transfer (blending




Figure 1: Examples of synthetic realities. (a) Synthetic faces generated by StyleGAN2 [107] (reproduced from [223]); (b) Synthetic art created by MidJourney [172], DALL·E 2 [205], and Stable Diffusion [211] using the prompt “*The Age of Synthetic Realities*”; and (c) Examples of video Deepfakes (reproduced from [139]).

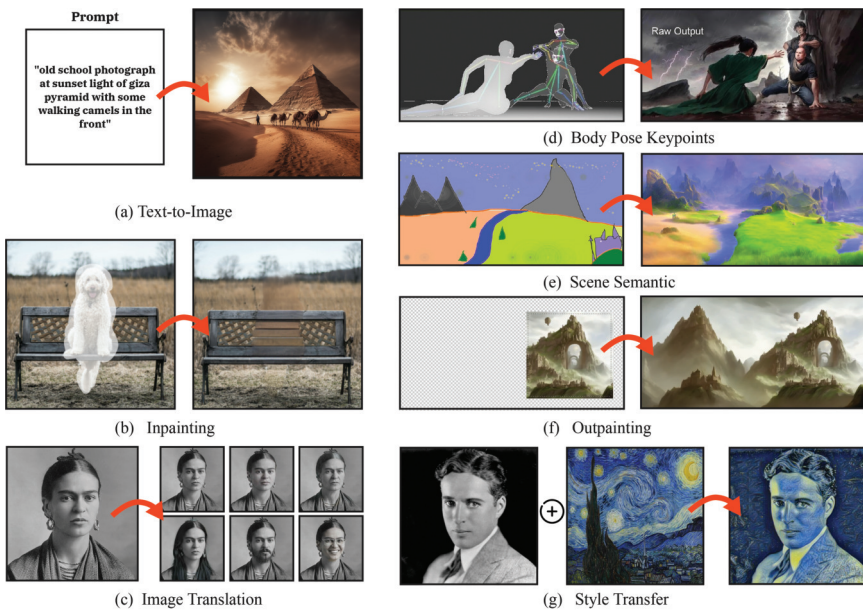


Figure 2: Examples of image synthesis conditioning. The input modality and expected output define the type of generation task performed by the model and help express desired characteristics in the synthesized creation. Examples created with and/or reproduced from [60, 172, 233, 234, 287].

artistic style), scene semantic image synthesis (from a semantic map), and body pose keypoint (content generated constrained to a body pose).

The families of AI models utilized in image generation techniques comprise several types of architectures, with the most common being Generative Ad-

versarial Networks (GANs), Diffusion models, and Variational Autoencoders (VAE).

GANs [75, 192] are built on two components: a generator and a discriminator. The generator learns the underlying distribution of real examples to generate new data, while the discriminator decides whether the input is from the real data space. Through an adversarial training process, the discriminator learns to identify synthetic images, while the generator progressively improves its ability to produce high-quality images that can deceive its counterpart. Numerous GAN variants have been developed in recent years to enhance the performance and stability of image generation [104, 202, 295, 297]. Among them, StyleGAN [105, 106, 108] allowed for intuitive control over the generated image attributes by modulating the convolutional kernels at different levels of the generator, instead of directly controlling the network input. Figure 1(a) shows examples of synthetic faces generated by StyleGAN2 [107]. Its successor, StyleGAN-T [219], builds upon its architecture for text-guided image synthesis. It leverages Contrastive Language-Image Pre-Training [201] (CLIP), a powerful text encoder that aligns textual descriptions with corresponding images. Other approaches [242, 306] follow a similar path, relying on CLIP to integrate natural language understanding into the image synthesis process.

As an alternative to the min-max optimization game of GANs, **Diffusion models** [43, 87] are trained to revert a stochastic diffusion process that progressively adds noise to a target image. To generate new images, the model iteratively denoises the perturbed image at each step, until a high-quality picture is reconstructed. By relying on a deterministic denoising function instead of adversarial learning, their training is more stable and easier to control than GANs. On the other hand, diffusion models rely on multiple network passes to reconstruct samples, constituting a considerably more computationally expensive method than adversarial networks. To improve efficiency, Stable Diffusion [211] operates on compressed latent representations instead of pixel space, mapping the denoising function to smaller manifolds. When considering textual prompts, Imagen [217] leverages text encoders, such as CLIP, to combine them with multiple cascaded diffusion models to generate high-resolution outputs from text. Similarly, DALL-E 2 presented by Ramesh *et al.* [205] trains a diffusion decoder that produces images from CLIP embeddings extracted from textual prompts.

Another prominent family of models is **Variational Autoencoders** (VAE) [116]. Autoencoders follow an encoder-decoder architecture that projects the input data into a low-dimensional latent space and learns to reconstruct the original input from it. VAE, in turn, extends upon autoencoders by adding a probabilistic component to the latent representation. Instead of learning a deterministic encoding for each input, the network learns the parameters of a probabilistic distribution that models the latent space. By doing so, it can sample from the latent space distribution to generate new samples.

Constraining the low-dimensional latent space further, VQ-VAE [186, 208] uses vector quantization to learn discrete latent variables, which improves the interpretability of the learned concepts and allows for easily manipulating them when generating new compositions. Text-to-image Transformer models like DALL-E [206] and CogView [49] employ VQ-VAE modules in their first stages to encode input images into enriched visual token representations. These tokens, selected from a codebook with a limited set of values (e.g., 8192 in CogView [49]), are then fed into an autoregressive Transformer [254]. The Transformer models the joint distribution of textual and visual components, ultimately yielding coherent and semantically meaningful synthesized creations.

Each of these AI model families has its pros and cons. GANs excel at generating sharp and visually compelling images, but they may suffer from mode collapse and training instability issues when used in large and diverse datasets [192]. Diffusion models offer a powerful approach to generating high-quality visual data, but they can be computationally expensive due to their iterative nature. VAEs provide a more straightforward training process with a clear optimization objective, but they may generate less sharp images than GANs and diffusion models. Nonetheless, all of them made significant advancements in the field of image generation, enabling the synthesis of realistic and diverse pictures. As the interest in this area increases, more advances will come, and the realism gap between real and synthetic data will shorten to the point that distinguishing between them will be challenging. This poses numerous problems in assessing the reliability and authenticity of visual content in an increasingly digital world. With this in mind, we discuss existing forensic approaches that may help to identify synthetic creations in the next section.

2.2 Synthetic Images Detection

In contrast to old-fashioned types of image manipulation, modern synthetic images take their realism to higher standards. Figure 3 compares classic Digital Forensics cases with those created by generative models, showcasing the remarkable advancement achieved. The level of refinement and potential harm associated with these synthetic images raise concerns about the capability of Digital Forensics to identify such content. However, we anticipate that **Locard’s exchange principle still holds for synthetic imagery**, with forensic traces taking the form of visual inconsistencies and artifacts left by the generation process. Nevertheless, as these models continue to evolve and new forms of counter-forensics attacks emerge, the validity of this claim may be challenged.

In this section, our analysis focuses on examining possible traces left by GANs and other generative models. We categorize our study based on the types of forensic evidence utilized for detection, namely **Visual Artifacts** and **Noise Fingerprints**.



Figure 3: Classic manual image manipulation versus modern image synthesis. The image on the left (a) represents a well-known case of state-level manual image manipulation, which was misleadingly published by multiple news websites as genuine in 2008. In contrast, we generated the image on the right (b) with MidJourney by using a prompt as simple as “missile test”. Cases like the former one currently represent an even greater challenge to authenticity verification.




Figure 4: Examples of generated image inconsistencies. All images were generated with version 5.1 of the MidJourney model – the latest one released at the writing of this article. In (a), an unnatural synthetic hand with six fingers. In (b), a synthetic wheelchair with inconsistent design; the seat orientation does not match the wheels’ position. In (c), a synthetic billboard with text that makes no sense and presents aberrant letters. In (d), a synthetic paisage with cloudy skies and mountains by a lake; the highlighted cloud is not congruently reflected on the lake’s surface. To generate these images, we used the following prompts: (a) “lady’s hand with a ring on it”, (b) “wheelchair in a hospital”, (c) “outdoor sign with a religious statement on it”, and (d) “realistic photo; mountains with a lake at the bottom”.

2.2.1 Visual Artifacts

Despite the impressive realism of cutting-edge synthetic images, a closer look reveals various aberrational results and visual inconsistencies. Borji [21] has presented several image failures that can occur when generating synthetic content, even with recent generative models such as DALL-E 2 [205], MidJourney [172], and StableDiffusion [233]. These failures may occur in the background, reflections, lighting, shadows, text, body parts, and objects, as depicted by Figure 4. This figure illustrates clues that have been explored by Digital Forensics researchers to detect synthetic content.

Farid, for instance, analyzed the 3D illumination [62] and geometric [63] consistency of structures and objects in a photograph generated by state-of-the-art generative models. His analysis employed classic digital forensics techniques for on-scene illumination and 3D geometric analysis, similar to those used in previous classic forensic works [13, 27, 189]. By doing so, Farid showed that while the local structures in the photo may be globally consistent, they exhibit local inconsistencies that serve as valuable clues for forensic analysts.

However, as generative models continue to evolve, it is expected that the visual inconsistencies and artifacts observed in synthetic images will eventually become rarer or imperceptible, as demonstrated in the case of the AI-synthetic faces examined by Nightingale and Farid [183]. This situation asks for other strategies that rely on other types of forensic clues, such as the noise left by the generation processes.

2.2.2 Noise Fingerprint

As image synthesis is rapidly improving, it is crucial to employ a variety of alternative detectors that explicitly exploit different characteristics of synthetic images. Therefore, as an alternative to visual artifacts, noise-based detectors have been a promising path to expose synthetic content.

In this direction, Marra *et al.* [162] investigated statistics-based techniques to detect potential noise fingerprints left by GANs on their generated content. By utilizing photo response non-uniformity (PRNU) analysis, similar to camera attribution methods, they discovered a correlation between residual noise patterns and specific GAN models. Furthermore, Marra *et al.* demonstrated the feasibility of differentiating between distinct GAN models used for image synthesis through residual noise analysis, enabling GAN model attribution.

In a similar study, Mandelli *et al.* [159] revealed that comparable residual noise patterns could be leveraged to identify GAN-generated scientific images, indicating the potential extension of this approach to other image types beyond natural images. These findings emphasize the applicability of residual noise analysis in detecting and identifying synthetic images, contributing to the field of Digital Forensics. Noise signatures allied with visual clues were explored in tandem by Kong *et al.* [118], showing that combining different evidence might be the way forward in dealing with the challenges of synthetic realities detection.

In a more recent investigation, Corvi *et al.* [40] examined the presence of fingerprints left by state-of-the-art generative models, including GAN-based and Diffusion-based models. Their findings reveal that **no generative model appears to be completely artifact-free** at present. Both GAN-generated and Diffusion-generated images exhibit anomalous periodic patterns in the Fourier spatial domain.

However, as highlighted by Gragnaniello *et al.* [76], such artifacts may be challenging to detect when post-processing operations are applied, such as image resizing and compression. These operations are frequently employed on social media platforms to save storage and speed up sharing, further complicating the identification of these visual irregularities. Besides investigating how post-processing operations impact synthetic image detectors, researchers have also identified the potential for synthetic image detectors to be deceived through counter-forensics attacks.

2.2.3 Counter Forensics

In a recent study that challenges noise-based detection methods, Osakabe *et al.* [188] developed a GAN model that can generate images without “checkerboard artifacts”, a specific type of artifact in the Fourier domain that is common in synthetic images. They achieved this by incorporating a fixed convolutional layer into every upsampling and downsampling layer of the GAN architecture. Remarkably, their model successfully fooled a detector that previously identified fake images with 92% accuracy, reducing the accuracy to a mere 12%.

Similarly, Cozzolino *et al.* [41] demonstrated that synthetic image detectors could be fooled by transferring residual noise fingerprints from real cameras onto GAN-generated images. This process produces a spoofed image that can avoid accurate GAN detectors and camera-model identifiers, causing the image to be misattributed as originating from the transferred camera model.

As counter-forensic attacks indicate, synthetic image detection presents multiple research challenges to ensure media integrity and prevent images from being used by a malicious actor.

2.3 Challenges and Directions

As with classic edited images, creating synthetic realities using cutting-edge generative models has sparked ethical debates and raised concerns about their use. Once again, Digital Forensics plays a crucial role in this debate by investigating ways to detect the traces of artificial intelligence techniques left in these images. While most synthetic images can be identified through a close look into scene inconsistencies and object aberrations, as depicted in Figure 4, the advancement of image synthesis will inevitably render these visual incongruences invisible to the naked eye.

Consequently, researchers also rely on new types of fingerprints inherently left by the generation process, such as specific patterns on the Fourier spatial spectrum and residual noise analysis. Some of these artifacts can be compared to the PRNU noise left by camera sensors. They can aid forensic analysis not only in detecting fake images but also in identifying the specific generative model used to render them (source attribution). However, traces alone

provide a vulnerable target for synthetic image detection, as they can easily be manipulated to deceive accurate fake image detectors. Such attacks may involve common post-processing operations or sophisticated techniques like camera noise transference. Therefore, it is imperative for forensic researchers to develop robust techniques capable of detecting and distinguishing these artifacts, even in the presence of common post-processing operations or more sophisticated attacks.

A more challenging and socially responsible aspect of forensics involves preventing the harmful applications of synthetic images. Given the ease with which such content can be created and shared, it is essential for forensics researchers to design traceable techniques whose synthetic images can be readily distinguished. Traceable evidence would assist analysts in swiftly identifying the source and author of such content, thereby preventing its widespread dissemination. In this vein, researchers have developed deep learning-based watermarking approaches [3, 65, 151] to identify synthesized content. These methods use encoder layers to imbue watermark information in the image pixels without perceptually altering its content. For successfully marking creations, researchers aim to be robust to most online alterations, such as compression, cropping, and intensity changes. Unfortunately, most approaches act from the generator’s perspective, either adapting existing models or adding external components in the generation process to enable watermarking. This might be viable for well-established and commercial applications (e.g., Midjourney [172]) but will be hardly used with open-sourced models that are trained and distributed by the community.

3 Synthetic Videos

The ability to generate realistic and useful videos holds immense value across various application domains such as entertainment, virtual reality, and education [207, 257]. Undoubtedly, video generation techniques have made significant positive contributions in these domains. The advancements have opened up new avenues for creativity, synthetic realities, and immersive experiences. However, it is essential to acknowledge that along with their benefits, these techniques also raise potential security concerns. Synthesize realistic videos can be exploited for malicious purposes, such as financial fraud and the dissemination of fake news. Consequently, ensuring the integrity and authenticity of digital content becomes increasingly critical.

Generally speaking, video synthesis can be divided into video generation and text-to-video synthesis. Previous methods for video generation mainly employ GANs [244, 250, 257] and VAEs [128, 281] to generate videos. But with the advent of diffusion models, recent methods explore them to generate more realistic videos [18, 205, 226, 256, 283]. On the other hand, text-to-video

synthesis incorporates text information to guide the model in generating video content that is responsive to specific demands. Analogously, previous text-to-video synthesis methods have resorted to GANs [136, 191] and diffusion models [6, 59, 84, 86, 111], achieving exceptional generation quality in terms of fidelity, resolution, and temporal consistency.

In this section, we provide a review of the methods for video generation, text-to-video synthesis, synthetic video detection and discuss their challenges. We further outline possible future research directions for synthetic video generation and detection techniques.

3.1 Video Generation

In the pursuit of advancing video synthesis, previous research has extensively explored diverse generative models, including Generative Adversarial Networks (GANs) [218, 257], autoregressive models [232, 281], and implicit neural representations [227, 292]. However, recent attention has been drawn to the exceptional achievements of diffusion models in visual data synthesis. Several notable works propose outstanding video generation methods and investigate their practical applications [18, 92, 256, 273]. For instance, a pioneering work on diffusion video generation [88] primarily focuses on network architecture modifications to extend image synthesis to video. The 3D U-Net is adopted [38] and achieves outstanding generation results in two cases, including unconditional and text-conditional video generation. For longer video generation, they apply an autoregressive approach, where subsequent video segments are conditioned on the preceding ones. Another example of a diffusion video generation work [283] adopts frame-by-frame video generation models. To evaluate different prediction strategies, the authors conduct an ablation study to determine whether predicting the residual of the next frame yields superior results compared to predicting the actual frame. Besides, MM-Diffusion [213] introduces a technique that employs a pair of denoising autoencoders to facilitate the joint generation of audio and video content. VideoFusion [153] presents an innovative approach involving a decomposed diffusion process, which separates the noise associated with each frame into a shared base noise across all frames and a variable residual noise along the temporal axis. TaleCrafter [74] utilizes the Low Rank Adaptation (LoRA) module to devise a personalized image generation system. This system adeptly accommodates various new characters, supporting layout adjustments and localized structural edits.

Furthermore, Hoppe *et al.* [92] introduced the Random Mask Video Diffusion (RaMViD) technique, which can be utilized for both video generation and infilling tasks. The unmasked frames are used to enforce conditions on the diffusion process, while the masked frames undergo diffusion through the forward process. By employing this training strategy, RaMViD demonstrates

outstanding video generation quality. These recent advancements in diffusion-based video generation highlight the potential of this family of models to push the boundaries of video synthesis, addressing the challenges of generating realistic and diverse video content in synthetic realities.

3.2 Text-to-video Synthesis

Text-to-video models are highly data-hungry, which require massive amounts of data to learn caption relatedness, frame photorealism, and temporal dynamics [72]. However, video data resources are comparatively more limited in terms of style, volume, and quality. This scarcity of video data poses significant challenges for training text-to-video generation models. To overcome these challenges, additional controls are often incorporated to enhance the responsiveness of generated videos to user demands [166, 191, 267].

Early text-to-video generation models heavily relied on convolutional GAN models combined with Recurrent Neural Networks (RNNs) to capture temporal dynamics [136, 191]. Despite the introduction of complex architectures and auxiliary losses, GAN-based models exhibit limitations in generating videos beyond simplistic scenes involving digit movements or close-up actions. To that end, recent advancements in the field have aimed to extend text-to-video generation to more diverse domains using large-scale transformers [288] or diffusion models [86]. Furthermore, Follow-Your-Pose [158] designs a novel training scheme that can utilize readily obtained datasets and the pre-trained text-to-image model to generate character videos with controllable poses. Tune-A-Video [273] proposes One-Shot Video Tuning, presenting a tailored spatiotemporal attention mechanism and an efficient one-shot tuning strategy. During inference, it employs DDIM inversion to provide structure guidance for sampling. These approaches provide promising directions for generating more complex and realistic video content by leveraging the expressive power of these advanced network architectures.

However, modeling high-dimensional videos and addressing the scarcity of text-video datasets present considerable challenges in training text-to-video generation models from scratch. To tackle this issue, most approaches adopt a transfer learning paradigm, leveraging pre-trained text-to-image models to acquire knowledge and improve performance. For instance, CogVideo [91] builds upon the pre-trained text-to-image model CogView2 [50], while Imagen Video [86] and Phenaki [255] employ joint image-video training techniques to leverage pre-existing visual representations. In contrast, Make-A-Video [226] focuses on learning motion solely from video data, reducing the reliance on text-video pairs for training.

Another key consideration in video synthesis is the high computational cost associated with generating high-quality videos. To mitigate this issue, latent diffusion has emerged as a popular technique for video generation, as it offers

a computationally efficient alternative [6, 18, 59, 84]. Various powerful but computational-efficient methods, such as MagicVideo [303], which introduces a simple adaptor after the 2D convolution layer, and Latent-Shift [6], which incorporates a parameter-free temporal shift module, have successfully utilized latent diffusion for video synthesis. Additionally, PDVM [291] adopts a novel approach of projecting the 3D video latent space into three 2D image-like latent spaces, further optimizing the computational cost of the video generation process.

Despite the active research in text-to-video generation, existing studies have predominantly overlooked the interplay and intrinsic correlation between spatial and temporal modules. These modules play crucial roles in understanding the complex dynamics of videos and ensuring coherent and realistic video generation.

3.3 Synthetic Video Detection

Again, detecting synthetic videos relies on the fingerprints left by generative models, which have been explored extensively in the context of synthetic image detection. Existing research in synthetic image detection has shown promise by identifying inconsistencies in illumination and geometric structure [62, 63], as well as specific noise patterns in the Fourier domain [40]. However, the extension of image-based detection techniques to videos is still in its early stages.

One straightforward approach for extending image-based detection to videos is through frame-level voting, where each frame is individually analyzed and classified as real or synthetic. However, exploiting temporal information, such as temporal coherence, presents a significant challenge. The temporal domain contains valuable cues that can aid in distinguishing synthetic videos from real ones. For instance, temporal coherence refers to the consistent motion and flow of objects across frames in a real video. Detecting such temporal fingerprints could provide valuable insights into the authenticity of a video. The temporal coherence, which is a significant challenge in the field of video synthesis [268], shall also be vital for synthetic video detection. Consequently, exploiting the temporal inconsistency can be utilized to identify generated videos. Currently, there is no existing method specifically designed for detecting video synthesis. However, there have been relevant work on detecting Deepfake videos using generic architectures. For instance, 3DCNN [296], RNN [5], LSTM [33], and temporal transformer [177] have been widely employed for video-level Deepfake detection. Additionally, some studies focus on detecting deepfake videos by examining specific temporal artifacts like lip movement [82], rPPG artifacts [199], and head pose inconsistency [284]. Thus, how to explore the temporal artifacts (*e.g.*, unnatural motions) in generated videos to detect video synthesis is an intriguing task.

This temporal aspect of video analysis introduces additional complexities compared to image analysis. Generated video content can exhibit various artifacts due to the processing techniques employed, including both handcrafted designs and deep neural networks. These artifacts, including blur, compression artifacts, and noise, can be intentionally or unintentionally injected during the video generation process. Consequently, these artifacts may pose significant obstacles to the detection of synthetic videos, particularly when the detection model is trained solely on high-quality video data.

To address this challenge, detection models need to be highly generalized to handle data from different domains. The models must be capable of recognizing and adapting to various levels of quality, distortion types, and content sources. This requirement calls for the application of domain generalization techniques, which enable the model to generalize well beyond the training data distribution. By training the model on a diverse range of video data, encompassing different quality levels, distortion types, and content sources, the detection system can become more robust and effective in identifying synthetic videos across a wide range of scenarios.

Overall, the extension of image-based detection techniques to videos presents a meaningful yet challenging direction for research. Leveraging temporal fingerprints and addressing the presence of artifacts in generated videos require novel approaches and further exploration. Developing detection models that can effectively analyze and distinguish synthetic videos while being adaptable to various domains will play a crucial role in combating the increasing threat of synthetic videos in today’s digital landscape.

3.4 Challenges and Directions

Recent advancements in diffusion models have revolutionized text-to-video synthesis, achieving remarkable capabilities that surpass previous state-of-the-art approaches and deliver unprecedented generative performance. This breakthrough has significantly enhanced the quality and fidelity of generated videos. However, as we delve deeper into this domain, it becomes evident that there is a need for further research and development.

Despite the demonstrated success in generated image content in the past few years, video generation is still in its infancy. As we have discussed, the challenges of video generation mainly lie in the following three aspects: (1) lacking large-scale, diverse, and in-the-wild video datasets; (2) demanding computational costs; and (3) unstable in synthesizing coherent content from both spatial and temporal perspectives.

In the realm of synthetic video detection, existing methods predominantly focus on identifying face forgery, where the manipulation targets primarily involve facial features and expressions. However, to address the evolving landscape of synthetic videos and their potential threats, it is imperative

to explore detection techniques that encompass a broader range of scenes and contexts. Detecting synthetic videos with diverse scenes, objects, and backgrounds poses an interesting avenue for future advancements in forensic research. By expanding the scope of detection techniques, we can develop robust and comprehensive methods that effectively identify and mitigate the risks associated with the increasing sophistication of synthetic videos.

4 Synthetic Audio

Synthetic realities in audio are an emerging technology transforming how we experience sound. From augmented and virtual reality to interactive audio installations, synthetic realities offer a new dimension to our auditory senses. These immersive audio experiences create a simulated environment that can transport listeners to different worlds, trigger emotions, and enhance storytelling. With the advancements in audio technology and the increasing demand for immersive experiences, synthetic realities are poised to revolutionize the entertainment, gaming, and education industries. While the rise of synthetic audio technology has brought about significant benefits to various fields, it also presents a considerable threat to the integrity of our society. One of the most concerning implications is the potential misuse of this technology by malicious actors, who can exploit it for nefarious purposes such as telecommunication fraud. Cai *et al.* [25] have highlighted the dangers of using generative models to create fake audio that can deceive individuals and organizations, leading to financial losses and reputational damage. The use of synthetic audio in such fraudulent activities underscores the urgent need for developing robust and reliable methods for detecting and mitigating the harms of this technology. In this section, we propose to survey synthetic realities in audio and dive into the possibilities and challenges of this rapidly evolving field.

4.1 Synthetic Audio Generation

Audio synthesis is a vital and rapidly evolving research area with a wide range of applications, including text-to-speech (TTS), speech enhancement, voice conversion, and binaural audio synthesis. In the field of TTS, previous works have extensively utilized deep learning-based architectures such as WaveNet [187] and Clarinet [194], as well as transformer models like FastSpeech [209] and Neural TTS [134], and variational autoencoder (VAE) approaches such as MultiSpeech [79] and Hierarchical VAE [94]. Recently, diffusion models have gained prominence in addressing TTS problems, with notable contributions from WaveGrad [29], DiffWave [120], Gradient Flow [195], and Diffusion TTS [100]. Recently, PromptTTS [81] introduces an innovative approach for speech synthesis. By utilizing a prompt containing both style and content attributes, the system successfully generates speech with remarkable accuracy in

terms of style manipulation and overall speech quality. DiffVoice [147] involves a two-step process. Initially, speech signals are encoded into a latent representation at the phoneme rate, employing a variational autoencoder that benefits from adversarial training. Subsequently, a joint modeling of duration and latent representation is achieved through the implementation of a diffusion model.

Speech enhancement techniques aim to improve speech recognition system performance by mitigating the impacts of ambient noise. The advancement of generative models has led to the development of various approaches for speech enhancement. These include GAN-based methods like MetricGAN [71] and Speech Enhancement GAN [144], as well as diffusion-based models such as Storm [130], Conditional Diffusion [149], and Cold Filter [285]. These models have exhibited general and robust speech enhancement performance.

Voice conversion, another critical task in speech synthesis, aims to transform the voice of one speaker into that of another. Different approaches have been explored for voice conversion, including transformer-based models (Voice-Filter [96]), GAN-based methods (CycleGAN-VC [103], VoiceGAN [93]), and VAE-based techniques (ACVAE [102] and Neural Voice Cloning [35]). These methods facilitate the manipulation of speaker characteristics while maintaining the linguistic content of the speech. In recent study, FreeVC [132] focuses on achieving superior waveform reconstruction quality. This is accomplished through the integration of strategies aimed at extracting pristine content information from data, all without the necessity of relying on textual annotations.

Lastly, binaural audio synthesis [131, 210] focuses on transforming mono audio signals into binaural audio, which enables accurate sound localization and immersive auditory experiences. NFS [129] is trained to predict the temporal delays and amplitude scales associated with diverse early reflections. It achieves more efficient memory utilization and computational expenses. By simulating the perception of sound through two ears, binaural audio synthesis contributes to creating a more realistic and interactive auditory environment.

Overall, the continuous advancements in deep learning, generative models, and various synthesis techniques have significantly expanded the possibilities and applications of audio synthesis, enhancing the quality, naturalness, and versatility of synthesized speech and audio.

4.2 Synthetic Audio Detection

Existing methods for synthetic audio detection can be categorized into two different streams: feature-based and image-based methods [14]. Feature-based approaches describe the audio through signal features such as Mel frequency cepstral coefficient (MFCC) and constant Q cepstral coefficient (CQCC) [245, 282]. These features are then fed into typical classifiers (e.g., support vector machines) [275] and deep neural networks [170]), which are trained to detect synthetic audio. Image-based methods, on the other hand, utilize either

spectrogram images [64, 142] computed from the audio signal and use them as the input for deep neural networks to extract discriminative information for synthetic audio detection. The aforementioned techniques have also been widely applied to deepfake detection related to synthetic audio [31] (more detail can be found in Section 7.3.2).

4.3 Challenges and Directions

Despite the progress made in synthetic audio detection, there are still challenges to overcome. One of the critical challenges is the availability of large datasets of synthetic audio that can be used to train detection models effectively [241]. Generating a large dataset of synthetic audio can be time-consuming and resource-intensive. Additionally, the increasing complexity and sophistication of synthetic audio algorithms may require more advanced detection methods that can keep up with these advancements.

While there are challenges to overcome, such as the availability of large datasets and the need for more advanced detection methods, recent research has shown promising results in developing more effective detection techniques. As synthetic audio technology advances, it is essential to continue developing and improving detection methods to prevent the misuse of deepfake audio and ensure that this technology is used safely and responsibly.

5 Synthetic Text

Large Language Models (LLMs) have revolutionized artificial intelligence, marking a significant milestone in the field. Since the rise of GPT models [22, 203], competitors from other companies like Google and Microsoft have also developed their own LLM, including Gopher [204], GLaM [55], LaMDA [243], Megatron-Turing NLG [228], Chinchilla [90], PaLM [37], BloombergGPT [274]. Some open-source models have also been released, including Jurassic-1 [141], GPT-J [264], GPT-NeoX [17], BLOOM [220], OPT [298], LLaMA [246], and OpenAssistant [121]. Although there are many variants of LLMs models, most of them have the same architecture of the GPT-2 model [203], which is an auto-regressive model (decoder-only transformer) that generates the next word based on previous inputs.

Figure 5 illustrates how these models have reached a quality comparable to humans, even in highly complex tasks such as poem writing, as both poems seem to be written by an excellent poet. The widespread adoption of LLMs has been observed across diverse domains, including Medicine [237], Journalism [161], and Science [247]. These powerful tools possess immense potential to enhance human capabilities in various areas, ranging from code

What kind of thoughts now, do you carry
 In your travels day by day
 Are they bright and lofty visions,
 Or neglected, gone astray?

Matters not how great in fancy,
 Or what deeds of skill you've wrought;
 Man, though high may be his station,
 Is no better than his thoughts.

Catch your thoughts and hold them tightly,
 Let each one an honor be;
 Purge them, scourge them, burnish brightly,
 Then in love set each one free.

(a) Poem "Thoughts" - Human Version.

In the realm of silent contemplation,
 Thoughts weave their intricate creation.
 Whispers of ideas, ethereal and vast,
 A symphony of musings from the past.

They flutter like butterflies in the mind,
 Seeking connections, truths to find.
 Embrace the thoughts that come your way,
 For within their essence, you may find,
 The power to shape your world, your mind.

In the realm of thoughts, oh, so grand,
 Unleash your creativity, take a stand.
 Let them flow, let them soar,
 And in their symphony, discover more.

(b) Poem "Thoughts" - LLM Version.

Figure 5: Human versus Machine poem writing. (a) is a public-domain poem entitled Thoughts from Myra Viola Wilds; (b) is a poem written by chatGPT when asked to write a poem entitled 'Thoughts'.

development (e.g., GitHub Copilot¹) to combating online hate speech and harassment (e.g., Cohere Classify ²).

While they offer numerous benefits, there is also a risk of them being used to produce harmful content, either deliberately by malicious actors or inadvertently due to their inherent flaws. In this section, we will investigate potential flaws and damaging applications of LLMs through a forensic lens. Our exploration will encompass the emergence of threats, machine-generated text detectors, and the underlying research challenges.

5.1 Large Language Models Threats

Text generation, like any form of machine-generated content, possesses inherent scalability, granting it the power to be employed in both beneficial and

¹<https://github.com/features/copilot>

²<https://txt.cohere.com/content-moderation-classify/>

detrimental ways. While AI-generated text may still exhibit semantic flaws or hallucinations [4], it has become increasingly difficult to differentiate between human and machine-generated text. This convergence of quality between human and AI-generated text poses a significant concern, particularly in the hands of malicious actors.

In a comprehensive study about computer-generated text threat modeling, Crothers *et al.* [44] categorize various types of attacks facilitated by large language models (LLMs). They group these attacks into four primary threats: (1) Facilitating Malware and Social Engineering; (2) Spam and Harassment; (3) Online Influence Campaigns; and (4) Exploiting AI authorship. While we will enumerate some of these threats within this section, it is worth noticing that LLMs have opened up a wide range of possibilities for malicious actors, extending beyond the scope of our enumerated list.

5.1.1 Facilitating Malware and Social Engineering

This threat makes use of LLMs for facilitating scalable and customizable scams, making them a significant threat in the realm of malware and social engineering [73]. By leveraging techniques like fine-tuning and prompt engineering, malicious actors can generate tailored scams that are highly convincing and appealing to specific targets or communities. For instance, by incorporating social media data from a target’s profile, such as their interests, lifestyle, and social connections, LLMs can create more sophisticated and personalized scams that manipulate individuals into taking harmful actions or providing sensitive information.

Another notable threat within this category is *Data Poisoning*. It involves the injection of exploitable data into the training process of LLMs or fine-tuning them with malicious intent. Schuster *et al.* [222] demonstrated a possible attack that poisons code-completion models (e.g., GitHub Co-pilot), making them include code vulnerabilities in their output, which attackers can later exploit. Such attacks open an important discussion about training datasets, as they are often gathered from the web without any rigorous curation.

5.1.2 Spam and Harassment

This threat weaponizes LLMs through trolls and hateful communities to propagate toxic content, disseminate misinformation, and target specific communities for harassment. An example of such an attack is the creation of GPT-4chan, as highlighted by Yannic Kilcher in his video “This is the worst AI ever” [114]. Kilcher fine-tuned GPT-J using data collected from the /pol/ channel on 4chan, a controversial online platform forum channel. The resulting model was used to interact with users on the same channel, encapsulating the offen-

sive, nihilistic, and trolling nature that characterizes many /pol/ posts [114]. Kilcher’s experiment raised the alarm to the scientific community on the ease with which LLMs can be misused and the potential consequences of such actions [126]. It emphasized the need for careful consideration and ethical responsibility when deploying and sharing LLMs, as they can be harnessed to amplify harm and propagate hateful ideologies.

5.1.3 Online Influence Campaigns

The utilization of LLMs for spreading fake news and manipulating public opinion has emerged as a significant concern. Political campaigns, in particular, could be a perilous case through the use of LLMs. Bai *et al.* [7] have demonstrated the susceptibility of individuals to persuasion on political matters when exposed to tailored messages generated by LLMs. Malicious actors could use this phenomenon in a devastating scenario to influence elections and other democratic processes.

5.1.4 Exploiting AI authorship

An intriguing threat posed by LLMs is the generation of academic articles. One can recall the case of SCIGen (2005), in which MIT graduate students developed a system that could automatically generate computer science papers, demonstrating the vulnerability of academic conferences to such submissions [236]. There is a growing concern that LLMs could be exploited to generate much more sophisticated fake articles than SCIGen, compromising scientific integrity. Research integrity experts fear that paper mills³ will improve their production in quality and quantity by using LLMs [247].

5.2 Detection Methods

A few detection methods have been proposed for LLMs generated content. One of the pioneering approaches is GROVER, proposed by Zellers *et al.* [293]. GROVER was capable of generating fluent and highly realistic fake articles using LLMs, which motivated the authors to explore detection techniques for such content. Zellers *et al.* found that employing the same model used for generating the text achieved higher detection accuracy compared to using a different one. Their results demonstrated an impressive accuracy rate of 92% in detecting LLM-generated fake articles.

However, over the past few years, models’ ability to generate text has significantly advanced, raising a bigger challenge. One recent approach, DetectGPT,

³Potentially illegal organizations that offer ghostwritten fraudulent or fabricated manuscripts [23].

introduced by Mitchell *et al.* [175], aims to address this challenge by detecting whether a given passage is generated by a specific model. The method is based on the hypothesis that AI-generated text exhibits a more negative log probability curvature compared to human-written text. To validate this hypothesis, Mitchell *et al.* proposed an approximation method for estimating the Hessian trace of the log probability function for both model-generated and human-written text, yielding promising results. However, a limitation of their approach is the requirement of knowing the specific generator model, which may not always be feasible in practice.

In recent research efforts, there has been a specific focus on ChatGPT-generated text due to its global attention. In [176], Mitrović *et al.* focused on detecting short texts such as online reviews. They employed a transformer-based model and applied an explanation method (SHAP [150]) to gain insights into distinguishing between human-written and machine-generated text. They found that detecting machine-generated text becomes more challenging when it is paraphrased from human text, where a human provides the initial text and asks the model to improve it. Additionally, the authors noted that ChatGPT tends to use uncommon words, exhibits politeness and impersonality, and lacks human-like emotional expressions. Another related study [259] explored the AI ability in text paraphrasing using GPT-3 and T5 models. They generated machine-paraphrased text and evaluated human performance in detecting these generated texts. The study showed that humans could not accurately detect GPT-3 paraphrased text, with accuracy only slightly above random (53%).

In response to concerns about AI-generated text, several proprietary tools have emerged to address the detection of AI-authored content. One such tool is GPTZero, which has gained attention in the media as a promising method for identifying AI-generated text [239]. However, we were unable to locate the source code or a scientific article detailing their approach. Similarly, numerous applications have been developed claiming to detect AI-generated text, such as GPTkit,⁴ Illuminary,⁵ OpenAI's AI Text Classifier,⁶ and AICheatCheck.⁷ However, many of these tools lack comprehensive studies on the reliability of their detection methods.

In an effort to facilitate the detection of ChatGPT-generated content, Yu *et al.* [290] released a large dataset specifically designed for the identification of ChatGPT-written abstracts. This dataset includes over 35,000 synthetic abstracts generated by ChatGPT, comprising fully generated texts, polished outputs, and mixtures of human-written and machine-generated abstracts. Additionally, the dataset contains more than 15,000 human-written abstracts

⁴<https://gptkit.ai/>.

⁵<https://illuminarty.ai/en/text/ai-generated-text-detection.html>.

⁶<https://platform.openai.com/ai-text-classifier>.

⁷<https://www.aicheatcheck.com/>.

for comparison. The results of their detection experiments demonstrated the ability to identify content that was entirely generated by ChatGPT. However, the task becomes more challenging when the generated text is mixed with human-written content. This work provides a valuable dataset into the complexities of detecting machine-generated text, particularly in scenarios involving a combination of human and AI-authored content.

A potential solution to address the misuse of LLMs is the use of text watermarks [77, 117]. Grinbaum and Adomaitis [77] argue that machine-generated long texts should include a watermark to indicate their source and ensure transparency. In [117], Kirchenbauer *et al.* propose embedding watermarks by modifying the sampling rules of next-word prediction. They use a hash function and pseudo-random generator to assign random colors (green and red) to words in the vocabulary. During next-word prediction, words from the red list are prohibited from appearing. However, they acknowledge the difficulty of a watermarking low-entropy text, as substituting a prohibited red word in such cases could result in poor quality output with high perplexity. To address this, they suggest a soft rule that encourages substituting red words in a high-entropy text. A third party familiar with the hash function and random generator can easily determine the colors of words by computing them. This detection method does not require knowledge of the specific generation model, making it a cheaper and more straightforward approach.

Although watermarking shows promise as a solution, it is important to consider that it modifies the output text. The method proposed in [117] evaluates quality based on perplexity, but there is a possibility that the meaning and semantics of the output may be altered due to the watermarking process.

5.3 Challenges and Directions

Detecting machine-generated text versus human-written text poses increasing challenges as LLMs continue to improve their ability to mimic human language [8]. Several challenges in this regard are highlighted below:

- **Generalization:** Detection methods often lack generalizability, meaning that a method developed to detect text generated by one specific model may not easily transfer to detecting text generated by another model. However, in real-world scenarios, prior knowledge about the specific model generating the text is typically unavailable.
- **Mixed reality:** Existing detection methods struggle when it comes to identifying machine-generated text that is mixed with human-written text. The combination of both types makes it more difficult to differentiate between them.

- **Adaptability:** LLMs demonstrate high adaptability to given prompts, making it challenging for methods that rely on finding patterns in the generated text. Models can exhibit different personalities⁸ and respond differently based on prompts, such as ChatGPT’s ability to adopt various tones depending on the prompt (e.g., DAN⁹), even faking emotions and swear words¹⁰. This adaptability further complicates detection efforts.

Despite these challenges, the remarkable capabilities of LLMs also present great research opportunities for detecting synthetic text. Some potential directions include:

- **Differentiating machine-generated from human-written text:** Achieving this requires collaborative efforts between humans and machines. Humans can contribute their technical knowledge of how models operate and how humans typically express themselves, facilitating the development of detection methods.
- **Attributing the source model of generated text:** Just as humans exhibit distinct writing styles, different language models may possess unique traits when generating text. By differentiating text generated by different models, researchers can gain insights into the behaviors of each model and identify each model’s fingerprint.
- **Fact-checking machine-generated text:** LLMs often struggle with generating factual content due to limited training data and the prevalence of fictional stories in their training corpus. In addition, after models are trained, the knowledge stored in these models can quickly become outdated. Conducting fact-checking on machine-generated text is crucial to ensure the reliability of AI-generated information.

6 NERFs and Metaverse

Neural Radiance Fields (NeRF) have emerged as an effective method for implicit volumetric scene representation, enabling learning from multiple viewing angles [173]. NeRF has been successfully applied in various domains, including transparent object grasping [99], scene understanding and reasoning [258, 302], and clear representations in challenging scenarios [80, 97, 157]. Recent NeRF variants, such as NeRF-W [163] and Ha-NeRF [32], have demonstrated their ability to reconstruct scenes from input with various perturbations. We can also see these developments as an example of synthetic realities, especially when we consider the possibility of totally synthesizing new “worlds”.

⁸<https://learn.microsoft.com/en-us/azure/cognitive-services/openai/how-to/chatgpt>

⁹<https://www.mlyearning.org/dan-chatgpt-prompt/>

¹⁰<https://metaroids.com/learn/jailbreaking-chatgpt-everything-you-need-to-know/>

6.1 NeRF for Metaverse Applications

In the context of metaverse applications, NeRF has been utilized in virtual concerts [48] and metaverse platforms for architecture and urban planning [10]. However, challenges remain, such as the need for real-time rendering of complex scenes with multiple dynamic objects and efficient methods to handle large and challenging scenes. Further research is necessary to address these challenges and to explore new use cases for NeRF in metaverse development.

One significant challenge to NeRF's applicability in metaverse development is its reliance on pre-computed camera parameters for scene representation. Several methods, such as NeRF [271] and BARF [143], have been proposed to optimize camera parameters and scene representation. However, avoiding interference from undesired scenes during camera parameter optimization remains an unsolved problem.

A potential research area for the future is building an occlusion-free scene reconstruction based on inaccurate or even unknown camera parameters, enabling greater flexibility in the use of NeRF for scene representation, leading to more effective applications in computer vision and graphics [173]. In a metaverse, where scenes are typically composed of multiple dynamic objects, avatars, and user interactions, occlusion-free scene reconstruction based on NeRF would enable a more thorough scene representation, resulting in more immersive and realistic virtual environments.

Moreover, an occlusion-free scene reconstruction based on NeRF that does not rely on accurate camera parameters would enable greater flexibility in metaverse development, allowing designers to create and share virtual spaces more efficiently [10]. Finally, optimizing NeRF-based methods for real-time rendering of complex scenes with multiple dynamic objects would enable seamless user interactions in a metaverse, leading to a more responsive and interactive virtual environment [48].

In summary, an occlusion-free scene reconstruction based on NeRF has the potential to significantly benefit metaverse applications by enabling more thorough scene representation, flexibility in scene creation and sharing, and real-time performance and scalability. Further research in this area could lead to even more effective applications of NeRF in the context of metaverse development.

6.2 Challenges and Limitations

Despite the progress achieved thus far, there are also some challenges and limitations to consider when using NeRF in metaverse development. One significant challenge is the reliance on pre-computed camera parameters for scene representation [308], making it infeasible when the pre-computation is not possible. This can limit the flexibility of scene creation and sharing. Although

some solutions [143, 271] have been proposed to optimize camera parameters along with scene representation, avoiding interference from undesired scenes during camera parameter optimization remains an unsolved problem [307].

Another limitation of NeRF is its computational cost. NeRF-based methods require significant computational resources and can suffer from slow rendering times [180], limiting their real-time performance for complex scenes. This can be a considerable challenge for metaverse applications, where real-time performance is critical for a seamless user experience.

Moreover, NeRF-based methods may not be suitable for all types of scenes. Scenes with complex geometry, occlusions [308], and dynamic objects [197] can pose challenges for NeRF-based methods, leading to incomplete scene representation and rendering. Though several image restoration methods [156, 260–263, 269] have been proposed, they are far from being practical solutions for NeRF and its variants. Therefore, it is essential to carefully evaluate the suitability of NeRF-based methods for a given scene and application.

Besides, there are also some potential adverse impacts regarding its broader societal implications. One potential concern is the potential for NeRF-based metaverse applications to become addictive and negatively impact mental health. The immersive and interactive nature of metaverse environments, combined with the potential for NeRF to create highly realistic and detailed scenes, could create a compelling and addictive experience for users. This could negatively impact mental health [252], including addiction, social isolation, and other adverse effects associated with prolonged use of virtual environments.

Another potential concern is the impact of NeRF-based metaverse applications on social dynamics and inequality [154]. NeRF-based metaverse applications could potentially exacerbate existing social inequalities and create new ones. For example, access to high-quality hardware and internet connectivity could become a barrier to participation in these environments, further marginalizing disadvantaged communities.

Last, using NeRF-based metaverse applications has raised significant concerns from both privacy and forensic perspectives [277]. These applications have the potential to collect and store vast amounts of personal data, which could be exploited for targeted advertising, surveillance, and other forms of data mining, leading to further erosion of individual autonomy and privacy. This could result in new forms of digital inequality and harm. Furthermore, the difficulty of collecting and preserving evidence in the NeRF and Metaverse contexts has been discussed in recent literature [110, 270]. Traditional forensic techniques may not be applicable in these virtual environments, where data are decentralized, and ownership is often unclear, making the determination of the chain of custody for digital assets within the Metaverse a complex and challenging task. Additionally, the potential for manipulating digital evidence within these environments raises concerns about the reliability and authenticity of such evidence, particularly with the use of deepfakes and synthetic media [277].

Therefore, there is a pressing need to develop new forensic techniques and tools to address these challenges and ensure the integrity and reliability of digital evidence in the NeRF and Metaverse contexts.

7 DeepFakes

In the context of synthetic realities, one particular example of utmost attention is deepfakes. Deepfakes are synthetic media that are digitally manipulated to replace one person’s identity or personal traits convincingly with that of another. Therefore, when synthetic media comprises the replacement of someone’s biometric traits, we are referring to a deepfake. It is typically present in images, audio samples, and videos.

7.1 Deepfake Images

7.1.1 Deepfake Image Generation

The issue of falsified image contents has been a long-standing problem in the image forensics area. With the emergence of deep learning, numerous powerful learning-based models are able to generate the so-called deepfake images with a high level of realism. In recent years, various deepfake techniques have been proposed, including image inpainting/removal, image composition, entire image synthesis, image translation, and text-to-image. Image inpainting/removal is used to fill in image regions with convincing content. Meanwhile, image composition, which encompasses object placement, image blending, image harmonization, and shadow generation, involves cutting out the foreground from one image and pasting it onto another image. Entire image synthesis involves the generation of images entirely by generative models such as GAN [75], VAE [115, 116], and diffusion models [229, 231]. Image translation, on the other hand, enables the transfer of an image’s style, such as converting a sketch image to a colored image. With the rapid development of diffusion models, the images generated based on text prompts are becoming increasingly realistic. Despite their remarkable quality, deepfake images can be misused for malicious purposes, leading to various security issues such as fake news and fraud.

7.1.2 Deepfake Image Detection

Deepfake image detection methods can be broadly classified into two categories: image-level and pixel-level detection. While image-level methods aim to identify the authenticity of the entire input image, pixel-level methods localize the manipulated regions. Traditional detection methods for detecting image manipulation in image inpainting/removal and image composition, rely on

capturing artifacts based on prior knowledge, such as lens distortions [167], CFA artifacts [66], noise patterns [155], compression artifacts [61], etc. Learning-based methods have improved the detection performance by capturing noise prints [304], JPEG features [265], High-frequency (HF) artifacts [310], and forgery boundary [53]. Additionally, detecting manipulated images generated through entire image synthesis, image translation, and text-to-image is another challenging problem. Various methods propose to extract visual artifacts [165], color artifacts [168], specific GAN fingerprints [289], and spectral features [58] for generated image detection. Nevertheless, these methods have limitations in generalizing across different GANs. To address this issue, more general methods such as CNN and generalization methods have been proposed [266, 278]. As image manipulation technology continues to advance, deepfake image detection is an essential field of research to prevent the spread of misinformation and protect the integrity of visual media.

7.1.3 Challenges and Future Work

Despite significant progress in deepfake image detection, there are still several challenges that need to be addressed. One of the main challenges is the generalization of deepfake detection models to unseen datasets and scenarios, which is crucial for practical applications. Another challenge is the robustness of these models against anti-forensics techniques such as recapturing and adversarial attacks. Moreover, the industry is now somewhat ahead of academia in terms of deploying deepfake detection technologies in real-world settings (e.g., ChatGPT). This gap can be narrowed by updating and creating more up-to-date deepfake databases, as most existing ones are somewhat outdated in the research community. Additionally, many deepfake detection models are not explainable, making it difficult to understand how they make decisions. Future work should aim to develop explainable models that can provide clear and interpretable justifications for their decisions. Overall, addressing these challenges can lead to more reliable and effective deepfake image detection systems in the future.

7.2 Deepfake Video

7.2.1 Deepfake Video Generation

Recently, deepfake videos typically refer to manipulated face videos. Face information plays a vital role in human communication [70]. However, the spread of deepfake videos on social media platforms can result in significant security concerns due to the potential dissemination of disinformation and misinformation, posing tangible and pressing security concerns. Generally speaking, there are four primary categories of deepfake videos, which are identity swap, face

reenactment, attribute manipulation, and entire synthesis [119]. These videos are generated using powerful generative models such as GAN [75], VAE [115, 116], and diffusion models [229, 231], which are capable of producing highly sophisticated videos. Identity swap replaces the original face regions with target faces, while face reenactment transfers the source facial expression to the target one. Attribute manipulation can alter specific facial features like hair, eyeglasses, nose, etc. With the advent of foundation models, entire synthesized videos can be generated. Powerful deep learning tools have been used to create sophisticated deepfake video datasets like UADFV [284], DF-TIMIT [122], FaceForensics++ [212], DFD [56], DFDC [51], Celeb-DF [139], DF-Forensics-1.0 [101], ForgeryNet [83], FFIW [305], KoDF [127], and FakeAVCeleb [112]. As deepfake techniques continue to evolve, it is crucial to develop effective methods for detecting deepfake videos and preventing their malicious use.

7.2.2 Deepfake Video Detection

To counteract malicious deepfake attacks, many detection methods have been proposed. Traditional methods mainly focus on hand-crafted features, such as lack of eye-blinking [137] and warping artifacts [138]. However, these methods are not accurate enough. Learning-based methods such as convolutional neural networks (CNN) [1, 36, 182, 240], recurrent neural networks (RNN) [216], and vision transformer (ViT) [85], have been proposed to achieve more promising detection performance. Afchar *et al.* [1] designed MesoNet and MesoInception4 to detect Deepfake and Face2Face videos automatically. Besides, some generic networks such as Xception Net [36], Efficient Net [240], and Capsule Net [182] have been demonstrated effective on deepfake detection tasks. Subsequent works have employed RNN [216] and ViT [85] to further improve forgery detection accuracy. Other methods capture spatial artifacts [30, 124, 135, 184, 224, 300], frequency artifacts [118, 171, 200], and biological signals [39, 199] to perform deepfake detection. Follow-up works [26, 133, 145, 152, 225, 301, 309] focus on improving the generalization capability and robustness of the model. Temporal information has also been exploited in many deepfake video detection methods based on typical generic networks such as 3DCNN [296], LSTM [5, 89], RNN [33, 216], and ViT [113]. Combining spatial and temporal information can achieve more reliable detection and improve the model's generalization capability.

7.2.3 Challenges and Future Work

Deepfake video creation and detection have seen significant success in recent years, but many issues remain unresolved. While accurate and secure, deepfake detectors lack interpretability, limiting their applications in practical scenarios.

Localizing forgery regions and forgery frames is also a crucial yet understudied task. Additionally, the two-player nature of face forgery and forgery detection means that attack techniques will continue to become more powerful, thereby calling for more general detection methods. Furthermore, deepfake videos often involve audio manipulation, which is largely overlooked in existing methods. Therefore, more visual-audio joint datasets and multi-modal detectors are expected in future works.

7.3 Deepfake Audio

7.3.1 Deepfake Audio Generation

Deepfake audio refers to manipulated or synthetic audio created using deep learning techniques. The aim of deepfake audio is to impersonate the speaker’s speech characteristics, such as accent, timbre, and intonation, by learning from target voice resources. Traditional methods for audio manipulation involve removing, duplicating, copying within an audio sample, or pasting and inserting fragments into other audios. Deep learning-based speech synthesis makes the generated audio more realistic and difficult to distinguish from real ones. Subsequent models based on likelihood algorithms, such as WaveNet [187] and WaveGlow [196], have been developed to perform audio generation. However, these methods often require conditional information and may fail to generate long signal sequences. Recent waveform generative models, such as GAN [52, 125, 280] and VAE [193], take advantage of various auxiliary losses, thereby achieving superior generation performance. On the other hand, recent diffusion models (*e.g.*, DiffWave [120]), have exhibited remarkable generation performance even in challenging unconditional and class-conditional waveform generation scenarios. In the context of text-to-speech tasks, diffusion models can be classified into: acoustic model (*e.g.*, Diff-TTS [100]), vocoder (*e.g.*, DiffWave [120]), and end-to-end framework (*e.g.*, FastDiff [95]). The promising results of diffusion models indicate their potential to revolutionize the field of audio generation and synthesis.

7.3.2 Deepfake Audio Detection

Automatic Speaker Verification (ASV) systems [181] currently detect manipulated audio through three tasks: logical access (LA), physical access (PA), and speech deepfake (DF) [279]. The LA task focuses on detecting synthetic speech injected into a communication system, while the PA task includes acoustic propagation and real physical factors. The DF task aims to detect deepfake speech circulating on social media platforms. Traditional methods for detecting deepfake audio involve analyzing the spectrogram of the audio and exposing audio inconsistencies, such as abrupt changes in frequency or

amplitude. These methods are based upon the assumption that synthetic audios have unique frequency and amplitude patterns. However, recent deep learning deepfake techniques raise the difficulty in identifying authenticity and call for the design of deep learning countermeasures. Generally speaking, deep learning methods can be categorized into feature-based, image-based, and waveform-based [24, 238]. Feature-based methods utilize critical digital signal features, such as Mel-frequency cepstral coefficients (MFCCs) [148], constant Q cepstral coefficient (CQCC) [282], and energy, to detect deepfake audio. Image-based methods apply the spectrogram image of the signal to conduct inconsistency detection. Waveform-based methods aim to analyze the raw waveform of the audio signal instead. To facilitate the development of deepfake audio detection models, numerous databases, such as M-AILABS Speech [230], GAN based synthesized audio dataset [299], Half-Truth [286], and H-Voice [9] have been created. Overall, deepfake audio detection is a challenging but important task that requires the constant development and refinement of detection methods.

7.3.3 Challenges and Future Work

Despite recent advancements in synthetic audio generation, there are still several challenges that need to be addressed. As generation techniques continue to evolve, deepfake audio will become increasingly difficult to distinguish by both human and AI-based detectors. Even worse, existing detection methods have shown poor robustness to compression, encoding, and noise. Additionally, current deepfake audio detection methods suffer from inefficient training datasets and overfitting issues [279], resulting in limited generalization capability. Moreover, most detection methods extract specific features (such as MFCC, CQCC, and energy) to conduct deepfake detection. However, it is challenging to extract appropriate features for specific detection tasks. How to effectively combine various features for more robust detection opens an important research path forward. Last but not least, it is crucial to conduct further research on the ethical implications of deepfake audio, such as its potential for misuse and its impacts on audio professionals.

8 Conclusion and Final Thoughts

Trust plays a fundamental role in our society. Citizens entrust infrastructures, services (including education and health), media (including social networks nowadays), the judiciary system (including law enforcement), and political decision-making in general. Democracies are endangered when citizens no longer trust the system and their elected representatives.

Unfortunately, trust can be tampered with through influence or disinformation. Although disinformation has probably always existed in human history (e.g., spreading rumors to influence elections), the message's quality and scale were low, restricting its impact. However, in our contemporary digital world, disinformation (also coined fake news) with greater realistic content spreads at an unprecedented scale on the Internet through alternative media without any filtering by the traditional mainstream channels. With the advance of Artificial Intelligence (AI) technologies, all sorts of media (text, images, and audio) can be synthetically generated. More particularly, recent generative AI models trained on very large datasets can produce more plausible and realistic content.

Distinguishing truth from falsity is becoming even more difficult, and the difference between reality and fiction is getting thinner daily. We are now facing the **Era of Synthetic Realities**. Disinformation exploits cognitive biases (e.g., anchoring bias, third-person effect, authority bias, bandwagon effect, to mention a few) [46], which are systematic errors in judgment that humans can make, and because of this, synthetic realities represent a threat to our society.

Synthetic realities are now generated by criminals and hostile agents for various malicious operations, including: political disinformation and state espionage (e.g., fake social network profiles), national security (e.g., facilitating a military coup), financial fraud (e.g., CEO scam impersonation), blackmail (e.g., ransomfake), defamation (e.g., revengeporn), plausible deniability of forensic evidence.

As human beings, because of cognitive biases, citizens will never stop falling for disinformation. A way to fight head on is to create tools to analyze digital content prior authentication by human experts. We anticipate some factors of utmost importance when developing new solutions.

The first one involves exploring the context of a digital asset as much as possible, even with limited training data. The second one involves efforts on robustness and interpretability, as decisions must be intelligible to human beings. The final one consists in being conscious of the incompleteness of individual methods and orchestrating decision-making fusion methods to combine different telltales for final detection.

1. **Limited training data.** Data-driven approaches often rely upon large amounts of training data. This problem may become critical in the rapidly evolving scenario of fake information. New forms of falsification, unknown to the forensic analyst, are proposed daily, preventing the timely collection of all relevant training data. We discuss this issue by posing the problem as an open-set recognition problem [221]; that is, we need to define a suitable model for pristine data and analyze false information by looking for inconsistencies concerning this model. If

possible, researchers also need to consider few-shot learning approaches that only require a tiny amount of labeled data to update to new threats.

2. **Robustness and Interpretability.** Being robust to wide-spectrum and unforeseen conditions is a basic system requirement, but it becomes central in a forensic environment featuring two active players. Robustness to adversarial attacks is paramount nowadays, especially when data-driven methods are applied, in light of the many literature findings that emphasize their vulnerability. Besides using basic solutions, we pose that reliability by relying on interpretable machine learning is necessary. We advocate for strategies that help us understand why a learning-based system behaves a certain way and provides the observed answers. Methods should also include semantics and context to support the entire decision-making process.
3. **Fusion.** Combining different methods toward a unified detection framework is very promising. As discussed in prior art for image forgery detection [68], fusion in different learning stages (early, middle, or even late-stage) plays a fundamental role within dynamic and adversarial setups. We envision learning strategies combining different telltales as a promising way forward.

Therefore, some driving research questions involve challenges in:

- **Detection:** is it possible to detect plausible and realistic digital content (e.g., text generated by large-scale language models, synthetic images, and voices generated by generative models)?
- **Attribution:** is it possible to accomplish source attribution by assigning manipulated digital content to a known type of attack vector?
- **Explainability:** is it possible to automatically uncover cues or inconsistencies in digital content to corroborate falsity, as discussed above?
- **Context and fusion:** how to incorporate context? How to combine different telltales?

Many other challenges will play out in the coming years as synthetic realities become ever more realistic, directly affecting fundamental pillars of our society, such as democratic values, individual freedom, and social tolerance. In this paper, we strived to discuss some of these challenges and what lies ahead, but it was just the tip of the iceberg.

Only an orchestrated effort of government representatives, society at large, and researchers will be able to curb such threats. We believe possible explorations might lie in regulatory acts, education investments, and scientific research for more powerful detection methods.

Acknowledgement

The Swiss, Brazilian and U.S. Governments are authorized to reproduce and distribute reprints for Governmental purposes, notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Swiss and Brazil governments or of DARPA, AFRL or the U.S. Government.

Financial Support

This work was supported by the São Paulo Research Foundation (FAPESP) (J.P.C., grant numbers 2020/02211-2, 2017/12646-3), (J.Y., grant numbers 2019/04053-8, 2022/05002-0), (A.R., grant number 2017/12646-3); Blue Sky Research Fund of HKBU (R.W., grant number BSRF/21-22/16); Guangdong Basic and Applied Basic Research Foundation (R.W., grant number 2022A1515110692); Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) (D.M., grant number FA8750-20-2-1004); and the Research Grant Council (RGC) of Hong Kong through Early Career Scheme (ECS) (H.L., grant number 21200522). Also, part of this work within the Idiap Research Institute received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI), the Swiss State of Valais and the City of Martigny.

Biographies

João Phillipe Cardenuto received a degree in Bachelor in Computer Engineering and Computer Science at the Universidade Estadual de Campinas (UNICAMP) (2019). Currently, he is pursuing a Ph.D. degree at the UNICAMP, Brazil. Due to his research work, he was honored with the Google Latin America Research Award (LARA) 2021. He also had the opportunity to collaborate with Google Research as an intern in 2022. His latest works have focused on image provenance analysis and the development of forensic algorithms aimed at detecting doctored scientific images. His research interests include media forensics, computer vision, machine learning, and scientific integrity.

Jing Yang is a Ph.D. student of the Artificial Intelligence Lab., Recod.ai, and is currently doing a research internship at the Ubiquitous Knowledge Processing (UKP) Lab. She received her Master's degree in Computer Science at Hunan University, China, and her Bachelor's in Information and Computing

Science at Hubei University of Technology, China. Jing’s research interests include natural language understanding, fact-checking, and forensics.

Rafael Padilha is a researcher associated with the Artificial Intelligence Lab. (Recod.ai) at the Institute of Computing, Universidade Estadual de Campinas (UNICAMP), Brazil. He received his Ph.D. in 2022 from the same university, with a joint research internship at the University of Kentucky, USA. He works for Microsoft Research, innovating on agriculture and sustainability under the Research for Industry team. His research interests lie in computer vision, machine learning, and digital forensics.

Renjie Wan received his BEng degree from the University of Electronic Science and Technology of China in 2012 and the Ph.D. degree from Nanyang Technological University, Singapore, in 2019. He is currently an Assistant Professor at Hong Kong Baptist University, Hong Kong. He is the outstanding reviewer of ICCV 2019 and the recipient of the Microsoft CRSF Award, VCIP 2020 Best Paper Award, and the Wallenberg-NTU Presidential Postdoctoral Fellowship.

Daniel Moreira received a Ph.D. degree in computer science from the Universidade Estadual de Campinas (UNICAMP), Brazil, in 2016. After working four years as a systems analyst with the Brazilian Federal Data Processing Service (SERPRO), he joined the University of Notre Dame for six years, first as a post-doctoral fellow and later as an assistant research professor. He is currently an assistant professor in the Department of Computer Science at Loyola University Chicago. He is also a member of the IEEE Information Forensics and Security Technical Committee (IFS-TC), 2021–2023 term, IEEE Signal Processing Society Education Center Editorial Board, 2022–2023 term, and associate editor of IEEE Transactions on Information Forensics and Security (T-IFS) and Elsevier Pattern Recognition journals. His research interests include media forensics, machine learning, computer vision, and biometrics.

Haoliang Li received his Ph.D. degree from Nanyang Technological University (NTU), Singapore in 2018. He is currently an assistant professor in Department of Electrical Engineering, City University of Hong Kong. His research mainly focuses on AI security, multimedia forensics and transfer learning. He received the Wallenberg-NTU presidential postdoc fellowship in 2019, doctoral innovation award in 2019, VCIP best paper award in 2020, Top 50 Chinese Young Scholars in AI+X 2022, and Stanford’s top 2% most highly cited scientists in 2022.

Shiqi Wang received the B.S. degree in computer science from the Harbin Institute of Technology in 2008 and the Ph.D. degree in computer application technology from Peking University in 2014. From 2014 to 2016, he was a Post-Doctoral Fellow with the Department of Electrical and Computer Engineering,

University of Waterloo, Waterloo, ON, Canada. From 2016 to 2017, he was a Research Fellow with the Rapid-Rich Object Search Laboratory, Nanyang Technological University, Singapore. He is currently an Assistant Professor with the Department of Computer Science, City University of Hong Kong. He has authored or coauthored more than 200 refereed journal articles/conference papers. His research interests include video compression, image/video quality assessment, and image/video search and analysis. He received the Best Paper Award from IEEE VCIP 2019, ICME 2019, IEEE Multimedia 2018, and PCM 2017. His coauthored article received the Best Student Paper Award in the IEEE ICIP 2018.

Fernanda Andaló is a researcher associated with the Artificial Intelligence Lab. (Recod.ai) at the Institute of Computing, Universidade Estadual de Campinas (UNICAMP), Brazil. Andaló received a Ph.D. in Computer Science from the same university in 2012, during which she was a research fellow at Brown University. She worked for Samsung as a researcher and was a postdoctoral researcher in collaboration with Motorola, from 2014 to 2018. Currently, she works at The LEGO Group devising machine learning solutions for digital products. She was the 2016–2017 Chair of the IEEE Women in Engineering South Brazil Section, and is an elected member of the IEEE Information Forensics and Security Technical Committee. Her research interests include machine learning and computer vision.

Sébastien Marcel (IEEE Senior member) is a senior researcher at the Idiap Research Institute (Switzerland), he heads the Biometrics Security and Privacy group and conducts research on face recognition, speaker recognition, vein recognition, attack detection (presentation attacks, morphing attacks, deepfakes) and template protection. He is also Professor at the University de Lausanne (UNIL) at the School of Criminal Justice and lecturer at the Ecole Polytechnique Fédérale de Lausanne (EPFL). He received his Ph.D. degree in signal processing from Université de Rennes I in France (2000) at CNET, the research center of France Telecom (now Orange Labs). He is also the Director of the Swiss Center for Biometrics Research and Testing, which conducts certifications of biometric products. He is Associate Editor of IEEE Transactions on Biometrics and Identity Science. He was Associate Editor of IEEE Signal Processing Letters, Associate Editor of IEEE Transactions on Information Forensics and Security, a Guest Editor of the IEEE Transactions on Information Forensics and Security Special Issue on “Biometric Spoofing and Countermeasures”, and Co-editor of the IEEE Signal Processing Magazine Special Issue on “Biometric Security and Privacy”. He is also the lead Editor of the Springer Handbook of Biometrics Anti-Spoofing (Editions 1, 2 and 3).

Anderson Rocha is a full-professor for Artificial Intelligence and Digital Forensics at the Institute of Computing, Universidade Estadual de Campinas

(UNICAMP), Brazil. He is the Director of the Artificial Intelligence Lab., Recod.ai, and was the Director of the Institute of Computing for the 2019–2023 term. He is an elected affiliate member of the Brazilian Academy of Sciences (ABC) and the Brazilian Academy of Forensic Sciences (ABC). He is a two-term elected member of the IEEE Information Forensics and Security Technical Committee (IFS-TC) and its chair for the 2019–2020 term. He is a Microsoft Research and a Google Research Faculty Fellow. In addition, in 2016, he has been awarded the Tan Chin Tuan (TCT) Fellowship, a recognition promoted by the Tan Chin Tuan Foundation in Singapore. Finally, he is ranked Top-2% among the most influential scientists worldwide, according to recent studies from Research.com and Standford/PlosOne.

References

- [1] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen, “MesoNet: a compact facial video forgery detection network,” in *IEEE International Workshop on Information Forensics and Security*, 2018, 1–7.
- [2] S. Agarwal and H. Farid, “Photo forensics from JPEG dimples,” in *IEEE International Workshop on Information Forensics and Security*, 2017, 1–6.
- [3] M. Ahmadi, A. Norouzi, N. Karimi, S. Samavi, and A. Emami, “ReD-Mark: Framework for residual diffusion watermarking based on deep networks,” *Expert Systems with Applications*, 146, 2020, 113157.
- [4] H. Alkaissi and S. I. McFarlane, “Artificial Hallucinations in ChatGPT: Implications in Scientific Writing,” *Cureus*, 15(2), 2023.
- [5] I. Amerini and R. Caldelli, “Exploiting prediction error inconsistencies through LSTM-based classifiers to detect deepfake videos,” in *ACM Workshop on Information Hiding and Multimedia Security*, 2020, 97–102.
- [6] J. An, S. Zhang, H. Yang, S. Gupta, J.-B. Huang, J. Luo, and X. Yin, “Latent-Shift: Latent Diffusion with Temporal Shift for Efficient Text-to-Video Generation,” *arXiv preprint*, arXiv:2304.08477, 2023.
- [7] H. Bai, J. G. Voelkel, Johannes Christopher Eichstaedt, and R. Willer, “Artificial Intelligence Can Persuade Humans on Political Issues,” *OSF Preprints*, 2023.
- [8] C. Bail, L. Pinheiro, and J. Royer, “Difficulty Of Detecting AI Content Poses Legal Challenges,” <https://www.law360.com/articles/1593766/difficulty-of-detecting-ai-content-poses-legal-challenges>, (Accessed on 06/06/2023), 2023.
- [9] D. M. Ballesteros, Y. Rodriguez, and D. Renza, “A dataset of histograms of original and fake voice recordings (H-Voice),” *Data in brief*, 29, 2020, 105331.

- [10] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, "Mip-NeRF 360: Unbounded anti-aliased neural radiance fields," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 5470–9.
- [11] S. Bayram, H. T. Sencar, and N. Memon, "Sensor fingerprint identification through composite fingerprints and group testing," *IEEE Transactions on Information Forensics and Security*, 10(3), 2014, 597–612.
- [12] Y. Bengio *et al.*, "Pause Giant AI Experiments: An Open Letter," <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>, (Accessed on 05/16/2023), 2023.
- [13] M. Bertamini, A. Spooner, and H. Hecht, "Naive optics: Predicting and perceiving reflections in mirrors," *Journal of Experimental Psychology: Human Perception and Performance*, 29(5), 2003, 982–1002.
- [14] K. Bhagtani, A. K. S. Yadav, E. R. Bartusiak, Z. Xiang, R. Shao, S. Baireddy, and E. J. Delp, "An overview of recent work in media forensics: Methods and threats," *arXiv preprint*, arXiv:2204.12067, 2022.
- [15] A. Bharati, D. Moreira, J. Brogan, P. Hale, K. Bowyer, P. Flynn, A. Rocha, and W. Scheirer, "Beyond pixels: Image provenance analysis leveraging metadata," in *IEEE Winter Conference on Applications of Computer Vision*, 2019, 1692–702.
- [16] T. Bianchi and A. Piva, "Image forgery localization via block-grained analysis of JPEG artifacts," *IEEE Transactions on Information Forensics and Security*, 7(3), 2012, 1003–17.
- [17] S. Black, S. Biderman, E. Hallahan, Q. Anthony, L. Gao, L. Golding, H. He, C. Leahy, K. McDonell, J. Phang, *et al.*, "GPT-NeoX-20B: An Open-Source Autoregressive Language Model," in *ACL Workshop on Challenges & Perspectives in Creating Large Language Models*, 2022, 95–136.
- [18] A. Blattmann, R. Rombach, H. Ling, T. Dockhorn, S. W. Kim, S. Fidler, and K. Kreis, "Align your Latents: High-Resolution Video Synthesis with Latent Diffusion Models," *arXiv preprint*, arXiv:2304.08818, 2023.
- [19] A. Bohr and K. Memarzadeh, "The rise of artificial intelligence in healthcare applications," in *Artificial Intelligence in Healthcare*, Elsevier, 2020, 25–60.
- [20] L. Bondi, L. Baroffio, D. Güera, P. Bestagini, E. J. Delp, and S. Tubaro, "First steps toward camera model identification with convolutional neural networks," *IEEE Signal Processing Letters*, 24(3), 2016, 259–63.
- [21] A. Borji, "Qualitative Failures of Image Generation Models and Their Application in Detecting Deepfakes," *arXiv preprint*, arXiv:2304.06470, 2023.

- [22] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, *et al.*, “Language models are few-shot learners,” in *Advances in Neural Information Processing Systems*, Vol. 33, 2020, 1877–901.
- [23] J. A. Byrne and J. Christopher, “Digital magic, or the dark arts of the 21st century—how can journals and peer reviewers detect manuscripts and publications from paper mills?” *FEBS Letters*, 594(4), 2020, 583–9.
- [24] Z. Cai, W. Wang, and M. Li, “Waveform Boundary Detection for Partially Spoofed Audio,” in *International Conference on Acoustics, Speech and Signal Processing*, 2023, 1–5.
- [25] Z. Cai, Z. Xiong, H. Xu, P. Wang, W. Li, and Y. Pan, “Generative adversarial networks: A survey toward private and secure applications,” *ACM Computing Surveys*, 54(6), 2021, 1–38.
- [26] J. Cao, C. Ma, T. Yao, S. Chen, S. Ding, and X. Yang, “End-to-End Reconstruction-Classification Learning for Face Forgery Detection,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 4113–22.
- [27] T. J. de Carvalho, C. Riess, E. Angelopoulou, H. Pedrini, and A. de Rezende Rocha, “Exposing Digital Image Forgeries by Illumination Color Classification,” *IEEE Transactions on Information Forensics and Security*, 8(7), 2013, 1182–94.
- [28] M. Chen, J. Fridrich, M. Goljan, and J. Lukás, “Determining image origin and integrity using sensor noise,” *IEEE Transactions on Information Forensics and Security*, 3(1), 2008, 74–90.
- [29] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, “WaveGrad: Estimating gradients for waveform generation,” *arXiv preprint*, arXiv:2009.00713, 2020.
- [30] S. Chen, T. Yao, Y. Chen, S. Ding, J. Li, and R. Ji, “Local relation learning for face forgery detection,” in *AAAI Conference on Artificial Intelligence*, Vol. 35, No. 2, 2021, 1081–8.
- [31] T. Chen, A. Kumar, P. Nagarsheth, G. Sivaraman, and E. Khoury, “Generalization of Audio Deepfake Detection,” in *The Speaker and Language Recognition Workshop (Odyssey)*, 2020, 132–7.
- [32] X. Chen, Q. Zhang, X. Li, Y. Chen, Y. Feng, X. Wang, and J. Wang, “Hallucinated neural radiance fields in the wild,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 12933–42.
- [33] A. Chintha, B. Thai, S. J. Sohrawardi, K. Bhatt, A. Hickerson, M. Wright, and R. Ptucha, “Recurrent convolutional structures for audio spoof and video deepfake detection,” *IEEE Journal of Selected Topics in Signal Processing*, 14(5), 2020, 1024–37.
- [34] W. J. Chisum and B. Turvey, “Evidence dynamics: Locard’s exchange principle & crime reconstruction,” *Journal of Behavioral Profiling*, 1(1), 2000, 1–15.

- [35] H.-S. Choi, J. Lee, W. Kim, J. Lee, H. Heo, and K. Lee, “Neural analysis and synthesis: Reconstructing speech from self-supervised representations,” in *Advances in Neural Information Processing Systems*, Vol. 34, 2021, 16251–65.
- [36] F. Chollet, “Xception: Deep learning with depthwise separable convolutions,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, 1251–8.
- [37] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, *et al.*, “PaLM: Scaling language modeling with pathways,” *arXiv preprint*, arXiv:2204.02311, 2022.
- [38] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, “3D U-Net: Learning dense volumetric segmentation from sparse annotation,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2016, 424–32.
- [39] U. A. Ciftci, I. Demir, and L. Yin, “FakeCatcher: Detection of synthetic portrait videos using biological signals,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020, 1–1.
- [40] R. Corvi, D. Cozzolino, G. Poggi, K. Nagano, and L. Verdoliva, “Intriguing properties of synthetic images: from generative adversarial networks to diffusion models,” *arXiv preprint*, arXiv:2304.06408, 2023.
- [41] D. Cozzolino, J. Thies, A. Rössler, M. Nießner, and L. Verdoliva, “SpoC: Spoofing Camera Fingerprints,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, 990–1000.
- [42] D. Cozzolino and L. Verdoliva, “Noiseprint: a CNN-based camera model fingerprint,” *IEEE Transactions on Information Forensics and Security*, 15, 2019, 144–59.
- [43] F.-A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah, “Diffusion models in vision: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023, 1–20.
- [44] E. Crothers, N. Japkowicz, and H. L. Viktor, “Machine Generated Text: A Comprehensive Survey of Threat Models and Detection Methods,” *arXiv preprint*, arXiv:2210.07321, 2022.
- [45] D. D’Avino, D. Cozzolino, G. Poggi, and L. Verdoliva, “Autoencoder with recurrent neural networks for video forgery detection,” *IS&T Electronic Imaging (EI)*, 29(7), 2017, 92–6.
- [46] K. Daniel, *Thinking, fast and slow*, Penguin, 2012.
- [47] M. Darvish Morshedi Hosseini and M. Goljan, “Camera identification from HDR images,” in *ACM Workshop on Information Hiding and Multimedia Security*, 2019, 69–76.

- [48] L. De Luigi, D. Bolognini, F. Domeniconi, D. De Gregorio, M. Poggi, and L. Di Stefano, “ScanNeRF: a Scalable Benchmark for Neural Radiance Fields,” in *IEEE Winter Conference on Applications of Computer Vision*, 2023, 816–25.
- [49] M. Ding, Z. Yang, W. Hong, W. Zheng, C. Zhou, D. Yin, J. Lin, X. Zou, Z. Shao, H. Yang, and J. Tang, “CogView: Mastering text-to-image generation via transformers,” in *Advances in Neural Information Processing Systems*, Vol. 34, 2021, 19822–35.
- [50] M. Ding, W. Zheng, W. Hong, and J. Tang, “CogView2: Faster and better text-to-image generation via hierarchical transformers,” *arXiv preprint*, arXiv:2204.14217, 2022.
- [51] B. Dolhansky, R. Howes, B. Pflaum, N. Baram, and C. C. Ferrer, “The deepfake detection challenge (DFDC) preview dataset,” *arXiv preprint*, arXiv:1910.08854, 2019.
- [52] J. Donahue, S. Dieleman, M. Bińkowski, E. Elsen, and K. Simonyan, “End-to-end adversarial text-to-speech,” *arXiv preprint*, arXiv:2006.03575, 2020.
- [53] C. Dong, X. Chen, R. Hu, J. Cao, and X. Li, “MVSS-Net: Multi-view multi-scale supervised networks for image manipulation detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3), 2022, 3539–53.
- [54] “DreamStudio,” <https://dreamstudio.ai/>, (Accessed on 05/22/2023), 2023.
- [55] N. Du, Y. Huang, A. M. Dai, S. Tong, D. Lepikhin, Y. Xu, M. Krikun, Y. Zhou, A. W. Yu, O. Firat, *et al.*, “GLaM: Efficient scaling of language models with mixture-of-experts,” in *International Conference on Machine Learning*, 2022, 5547–69.
- [56] N. Dufour and A. Gully, “Contributing Data to Deepfake Detection Research,” <https://ai.googleblog.com/2019/09/contributing-data-to-deepfake-detection.html>, (Accessed on 06/06/2023), 2019.
- [57] M. den Dunnen, “Synthetic Reality & Deep fakes Impact on Police Work,” <https://enlets.eu/wp-content/uploads/2021/11/Final-Synthetic-Reality-Deep-fakes-Impact-on-Police-Work-04.11.21.pdf>, (Accessed on 05/16/2023), 2021.
- [58] R. Durall, M. Keuper, and J. Keuper, “Watch your up-convolution: CNN based generative deep neural networks are failing to reproduce spectral distributions,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, 7890–9.
- [59] P. Esser, J. Chiu, P. Atighehchian, J. Granskog, and A. Germanidis, “Structure and content-guided video synthesis with diffusion models,” *arXiv preprint*, arXiv:2302.03011, 2023.
- [60] “FaceApp,” <https://www.faceapp.com/>, (Accessed on 05/22/2023), 2023.

- [61] Z. Fan and R. L. De Queiroz, "Identification of bitmap compression history: JPEG detection and quantizer estimation," *IEEE Transactions on Image Processing*, 12(2), 2003, 230–5.
- [62] H. Farid, "Lighting (In)consistency of Paint by Text," *arXiv preprint*, arXiv:2207.13744, 2022.
- [63] H. Farid, "Perspective (In)consistency of Paint by Text," *arXiv preprint*, arXiv:2206.14617, 2022.
- [64] A. Fathan, J. Alam, and W. H. Kang, "Mel-Spectrogram Image-Based End-to-End Audio Deepfake Detection Under Channel-Mismatched Conditions," in *IEEE International Conference on Multimedia and Expo*, 2022, 1–6.
- [65] P. Fernandez, G. Couairon, H. Jégou, M. Douze, and T. Furon, "The Stable Signature: Rooting Watermarks in Latent Diffusion Models," *arXiv preprint*, arXiv:2303.15435, 2023.
- [66] P. Ferrara, T. Bianchi, A. De Rosa, and A. Piva, "Image forgery localization via fine-grained analysis of CFA artifacts," *IEEE Transactions on Information Forensics and Security*, 7(5), 2012, 1566–77.
- [67] A. M. Ferreira, T. Carvalho, F. A. Andalo, and A. Rocha, "Counteracting the contemporaneous proliferation of digital forgeries and fake news," *Anais da Academia Brasileira de Ciências (AABC)*, 91(suppl. 1), 2019, e20180149.
- [68] A. Ferreira, S. C. Felipussi, C. Alfaro, P. Fonseca, J. E. Vargas-Munoz, J. A. Dos Santos, and A. Rocha, "Behavior knowledge space-based fusion for copy-move forgery detection," *IEEE Transactions on Image Processing*, 25(10), 2016, 4729–42.
- [69] J. Franchi, "Digest 95-5," http://phys.bspu.unibel.by/static/met/guides/eric_it/digests/virtual.html, (Accessed on 05/16/2023), 1995.
- [70] C. Frith, "Role of facial expressions in social interactions," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1535), 2009, 3453–8.
- [71] S.-W. Fu, C.-F. Liao, Y. Tsao, and S.-D. Lin, "MetricGAN: Generative adversarial networks based black-box metric scores optimization for speech enhancement," in *International Conference on Machine Learning*, 2019, 2031–41.
- [72] S. Ge, S. Nah, G. Liu, T. Poon, A. Tao, B. Catanzaro, D. Jacobs, J.-B. Huang, M.-Y. Liu, and Y. Balaji, "Preserve Your Own Correlation: A Noise Prior for Video Diffusion Models," *arXiv preprint*, arXiv:2305.10474, 2023.
- [73] A. Giarretta and N. Dragonì, "Community Targeted Phishing," in *International Conference in Software Engineering for Defence Applications*, 2020, 86–93.

- [74] Y. Gong, Y. Pang, X. Cun, M. Xia, H. Chen, L. Wang, Y. Zhang, X. Wang, Y. Shan, and Y. Yang, “TaleCrafter: Interactive Story Visualization with Multiple Characters,” *arXiv preprint arXiv:2305.18247*, 2023.
- [75] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio, “Generative Adversarial Nets,” in *Advances in Neural Information Processing Systems*, Vol. 27, 2014.
- [76] D. Gragnaniello, D. Cozzolino, F. Marra, G. Poggi, and L. Verdoliva, “Are GAN Generated Images Easy to Detect? A Critical Analysis of the State-Of-The-Art,” in *IEEE International Conference on Multimedia and Expo*, 2021, 1–6.
- [77] A. Grinbaum and L. Adomaitis, “The Ethical Need for Watermarks in Machine-Generated Language,” *arXiv preprint*, arXiv:2209.03118, 2022.
- [78] J. Gu, X. Wang, C. Li, J. Zhao, W. Fu, G. Liang, and J. Qiu, “AI-enabled image fraud in scientific publications,” *Patterns*, 3(7), 2022, 100511.
- [79] H. Guo, F. Xie, F. K. Soong, X. Wu, and H. Meng, “A Multi-Stage Multi-Codebook VQ-VAE Approach to High-Performance Neural TTS,” *arXiv preprint*, arXiv:2209.10887, 2022.
- [80] Y.-C. Guo, D. Kang, L. Bao, Y. He, and S.-H. Zhang, “NeRFReN: Neural Radiance Fields With Reflections,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 18388–97.
- [81] Z. Guo, Y. Leng, Y. Wu, S. Zhao, and X. Tan, “PromptTTS: Controllable text-to-speech with text descriptions,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2023, 1–5.
- [82] A. Haliassos, K. Vougioukas, S. Petridis, and M. Pantic, “Lips don’t lie: A generalisable and robust approach to face forgery detection,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 5039–49.
- [83] Y. He, B. Gan, S. Chen, Y. Zhou, G. Yin, L. Song, L. Sheng, J. Shao, and Z. Liu, “ForgeryNet: A versatile benchmark for comprehensive forgery analysis,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 4360–9.
- [84] Y. He, T. Yang, Y. Zhang, Y. Shan, and Q. Chen, “Latent Video Diffusion Models for High-Fidelity Video Generation with Arbitrary Lengths,” *arXiv preprint*, arXiv:2211.13221, 2022.
- [85] Y.-J. Heo, Y.-J. Choi, Y.-W. Lee, and B.-G. Kim, “Deepfake detection scheme based on vision transformer and distillation,” *arXiv preprint*, arXiv:2104.01353, 2021.

- [86] J. Ho, W. Chan, C. Saharia, J. Whang, R. Gao, A. Gritsenko, D. P. Kingma, B. Poole, M. Norouzi, D. J. Fleet, and T. Salimans, “Imagen video: High definition video generation with diffusion models,” *arXiv preprint*, arXiv:2210.02303, 2022.
- [87] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, 33, 2020, 6840–51.
- [88] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet, “Video diffusion models,” *arXiv preprint*, arXiv:2204.03458, 2022.
- [89] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, 9(8), 1997, 1735–80.
- [90] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. d. L. Casas, L. A. Hendricks, J. Welbl, A. Clark, *et al.*, “Training compute-optimal large language models,” *arXiv preprint*, arXiv:2203.15556, 2022.
- [91] W. Hong, M. Ding, W. Zheng, X. Liu, and J. Tang, “CogVideo: Large-scale pretraining for text-to-video generation via transformers,” *arXiv preprint*, arXiv:2205.15868, 2022.
- [92] T. Höpfe, A. Mehrjou, S. Bauer, D. Nielsen, and A. Dittadi, “Diffusion models for video prediction and infilling,” *arXiv preprint*, arXiv:2206.07696, 2022.
- [93] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, “Voice conversion from unaligned corpora using variational autoencoding Wasserstein generative adversarial networks,” *arXiv preprint*, arXiv:1704.00849, 2017.
- [94] W. Hsu, Y. Zhang, R. J. Weiss, H. Zen, Y. Wu, Y. Wang, Y. Cao, Y. Jia, Z. Chen, J. Shen, P. Nguyen, and R. Pang, “Hierarchical generative modeling for controllable speech synthesis,” *arXiv preprint*, arXiv:1810.07217, 2018.
- [95] R. Huang, M. W. Lam, J. Wang, D. Su, D. Yu, Y. Ren, and Z. Zhao, “FastDiff: A fast conditional diffusion model for high-quality speech synthesis,” in *International Joint Conference on Artificial Intelligence*, 2022, 4157–63.
- [96] W.-C. Huang, T. Hayashi, Y.-C. Wu, H. Kameoka, and T. Toda, “Voice transformer network: Sequence-to-sequence voice conversion using transformer with text-to-speech pretraining,” *arXiv preprint*, arXiv:1912.06813, 2019.
- [97] X. Huang, Q. Zhang, Y. Feng, H. Li, X. Wang, and Q. Wang, “HDR-NeRF: High Dynamic Range Neural Radiance Fields,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 18377–87.

- [98] M. Humphries, “Nvidia Uses AI to Make Our Eyes Always Look at the Camera,” <https://uk.pcmag.com/video-conferencing-software/144893/nvidia-uses-ai-to-make-our-eyes-always-look-at-the-camera>, (Accessed on 05/16/2023), 2023.
- [99] J. Ichnowski, Y. Avigal, J. Kerr, and K. Goldberg, “Dex-NeRF: Using a neural radiance field to grasp transparent objects,” in *Conference on Robot Learning (CoRL)*, 2021.
- [100] M. Jeong, H. Kim, S. J. Cheon, B. J. Choi, and N. S. Kim, “Diff-TTS: A denoising diffusion model for text-to-speech,” *arXiv preprint*, arXiv:2104.01409, 2021.
- [101] L. Jiang, R. Li, W. Wu, C. Qian, and C. C. Loy, “DeeperForensics-1.0: A large-scale dataset for real-world face forgery detection,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, 2889–98.
- [102] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, “ACVAE-VC: Non-parallel voice conversion with auxiliary classifier variational autoencoder,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(9), 2019, 1432–43.
- [103] T. Kaneko, H. Kameoka, K. Tanaka, and N. Hojo, “CycleGAN-VC2: Improved cyclegan-based non-parallel voice conversion,” in *International Conference on Acoustics, Speech and Signal Processing*, 2019, 6820–4.
- [104] T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive Growing of GANs for Improved Quality, Stability, and Variation,” in *International Conference on Learning Representations*, 2018.
- [105] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila, “Alias-free generative adversarial networks,” in *Advances in Neural Information Processing Systems*, Vol. 34, 2021, 852–63.
- [106] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, 4401–10.
- [107] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, “Analyzing and Improving the Image Quality of StyleGAN,” in *Proc. CVPR*, 2020.
- [108] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, “Analyzing and improving the image quality of StyleGAN,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, 8110–9.
- [109] J. Kell, “How A.I. is reshaping the way movies are made,” <https://fortune.com/2023/02/14/tech-forward-everyday-ai-hollywood-movies/>, (Accessed on 06/05/2023), 2023.
- [110] N.-A. Le-Khac and K.-K. R. Choo, “Databases in Digital Forensics,” in, *A Practical Hands-on Approach to Database Forensics*, Springer International Publishing, 2022, 1–2.

- [111] L. Khachatryan, A. Movsisyan, V. Tadevosyan, R. Henschel, Z. Wang, S. Navasardyan, and H. Shi, “Text2Video-Zero: Text-to-image diffusion models are zero-shot video generators,” *arXiv preprint*, arXiv:2303.13439, 2023.
- [112] H. Khalid, S. Tariq, M. Kim, and S. S. Woo, “FakeAVCeleb: a novel audio-video multimodal deepfake dataset,” *arXiv preprint*, arXiv:2108.05080, 2021.
- [113] S. A. Khan and H. Dai, “Video transformer for deepfake detection with incremental learning,” in *ACM International Conference on Multimedia*, 2021, 1821–8.
- [114] Y. Kilcher, “GPT-4chan: This is the worst AI ever,” <https://www.youtube.com/watch?v=efPrtcLdcdM>, (Accessed on 06/06/2023), 2022.
- [115] D. P. Kingma and M. Welling, “An introduction to variational autoencoders,” *Foundations and Trends® in Machine Learning*, 12(4), 2019, 307–92.
- [116] D. P. Kingma and M. Welling, “Auto-encoding variational Bayes,” *arXiv preprint*, arXiv:1312.6114, 2013.
- [117] J. Kirchenbauer, J. Geiping, Y. Wen, J. Katz, I. Miers, and T. Goldstein, “A watermark for large language models,” *arXiv preprint*, arXiv:2301.10226, 2023.
- [118] C. Kong, B. Chen, H. Li, S. Wang, A. Rocha, and S. Kwong, “Detect and locate: Exposing face manipulation by semantic-and noise-level telltales,” *IEEE Transactions on Information Forensics and Security*, 17, 2022, 1741–56.
- [119] C. Kong, S. Wang, and H. Li, “Digital and Physical Face Attacks: Reviewing and One Step Further,” *arXiv preprint*, arXiv:2209.14692, 2022.
- [120] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, “DiffWave: A versatile diffusion model for audio synthesis,” *arXiv preprint*, arXiv:2009.09761, 2020.
- [121] A. Köpf, Y. Kilcher, D. von Rütte, S. Anagnostidis, Z.-R. Tam, K. Stevens, A. Barhoum, N. M. Duc, O. Stanley, R. Nagyfi, *et al.*, “OpenAssistant Conversations – Democratizing Large Language Model Alignment,” *arXiv preprint*, arXiv:2304.07327, 2023.
- [122] P. Korshunov and S. Marcel, “Deepfakes: a new threat to face recognition? Assessment and detection,” *arXiv preprint*, arXiv:1812.08685, 2018.
- [123] P. Korus and J. Huang, “Multi-scale analysis strategies in PRNU-based tampering localization,” *IEEE Transactions on Information Forensics and Security*, 12(4), 2016, 809–24.
- [124] A. Kumar, A. Bhavsar, and R. Verma, “Detecting deepfakes with metric learning,” in *International Workshop on Biometrics and Forensics (IWBF)*, 2020, 1–6.

- [125] K. Kumar, R. Kumar, T. De Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, “MelGAN: Generative adversarial networks for conditional waveform synthesis,” in *Advances in Neural Information Processing Systems*, Vol. 32, 2019.
- [126] A. Kurenkov, “Lessons from the GPT-4Chan Controversy,” <https://thegradient.pub/gpt-4chan-lessons/>, (Accessed on 06/06/2023), 2022.
- [127] P. Kwon, J. You, G. Nam, S. Park, and G. Chae, “KoDF: A large-scale korean deepfake detection dataset,” in *IEEE International Conference on Computer Vision*, 2021, 10744–53.
- [128] G. Le Moing, J. Ponce, and C. Schmid, “CCVS: context-aware controllable video synthesis,” in *Advances in Neural Information Processing Systems*, Vol. 34, 2021, 14042–55.
- [129] J. W. Lee and K. Lee, “Neural fourier shift for binaural speech rendering,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2023, 1–5.
- [130] J.-M. Lemercier, J. Richter, S. Welker, and T. Gerkmann, “StoRM: A Diffusion-based Stochastic Regeneration Model for Speech Enhancement and Dereverberation,” *arXiv preprint*, arXiv:2212.11851, 2022.
- [131] Y. Leng, Z. Chen, J. Guo, H. Liu, J. Chen, X. Tan, D. Mandic, L. He, X.-Y. Li, T. Qin, S. Zhao, and T.-Y. Liu, “BinauralGrad: A two-stage conditional diffusion probabilistic model for binaural audio synthesis,” in *Advances in Neural Information Processing Systems*, Vol. 35, 2022, 23689–700.
- [132] J. Li, W. Tu, and L. Xiao, “Freevc: Towards High-Quality Text-Free One-Shot Voice Conversion,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2023, 1–5.
- [133] L. Li, J. Bao, T. Zhang, H. Yang, D. Chen, F. Wen, and B. Guo, “Face X-ray for more general face forgery detection,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, 5001–10.
- [134] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, “Neural speech synthesis with transformer network,” in *AAAI Conference on Artificial Intelligence*, Vol. 33, No. 1, 2019, 6706–13.
- [135] X. Li, Y. Lang, Y. Chen, X. Mao, Y. He, S. Wang, H. Xue, and Q. Lu, “Sharp multiple instance learning for deepfake video detection,” in *ACM International Conference on Multimedia*, 2020, 1864–72.
- [136] Y. Li, M. Min, D. Shen, D. Carlson, and L. Carin, “Video generation from text,” in *AAAI Conference on Artificial Intelligence*, Vol. 32, No. 1, 2018.
- [137] Y. Li, M.-C. Chang, and S. Lyu, “In Ictu Oculi: Exposing AI created fake videos by detecting eye blinking,” in *IEEE International Workshop on Information Forensics and Security*, 2018, 1–7.

- [138] Y. Li and S. Lyu, “Exposing deepfake videos by detecting face warping artifacts,” *arXiv preprint*, arXiv:1811.00656, 2018.
- [139] Y. Li, X. Yang, P. Sun, H. Qi, and S. Lyu, “Celeb-DF: A large-scale challenging dataset for deepfake forensics,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, 3204–13.
- [140] J. Liang, C. Wu, X. Hu, Z. Gan, J. Wang, L. Wang, Z. Liu, Y. Fang, and N. Duan, “NUWA-Infinity: Autoregressive over Autoregressive Generation for Infinite Visual Synthesis,” *Advances in Neural Information Processing Systems*, 35, 2022, 15420–32.
- [141] O. Lieber, O. Sharir, B. Lenz, and Y. Shoham, “Jurassic-1: Technical details and evaluation,” *AI21*, 2021.
- [142] S.-Y. Lim, D.-K. Chae, and S.-C. Lee, “Detecting Deepfake Voice Using Explainable Deep Learning Techniques,” *Applied Sciences*, 12(8), 2022, 3926.
- [143] C.-H. Lin, W.-C. Ma, A. Torralba, and S. Lucey, “BARF: Bundle-adjusting neural radiance fields,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 5721–31.
- [144] J. Lin, S. Niu, Z. Wei, X. Lan, A. J. Wijnngaarden, M. C. Smith, and K.-C. Wang, “Speech enhancement using forked generative adversarial networks with spectral subtraction,” in *Interspeech*, 2019, 3163–7.
- [145] H. Liu, X. Li, W. Zhou, Y. Chen, Y. He, H. Xue, W. Zhang, and N. Yu, “Spatial-phase shallow learning: rethinking face forgery detection in frequency domain,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 772–81.
- [146] Y. Liu, X. Zhu, X. Zhao, and Y. Cao, “Adversarial learning for constrained image splicing detection and localization based on atrous convolution,” *IEEE Transactions on Information Forensics and Security*, 14(10), 2019, 2551–66.
- [147] Z. Liu, Y. Guo, and K. Yu, “DiffVoice: Text-to-Speech with Latent Diffusion,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2023, 1–5.
- [148] B. Logan, “Mel frequency cepstral coefficients for music modeling,” in *International Society for Music Information Retrieval Conference*, 2000.
- [149] Y.-J. Lu, Z.-Q. Wang, S. Watanabe, A. Richard, C. Yu, and Y. Tsao, “Conditional diffusion probabilistic model for speech enhancement,” in *International Conference on Acoustics, Speech and Signal Processing*, 2022, 7402–6.
- [150] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Advances in Neural Information Processing Systems*, Vol. 30, 2017.

- [151] X. Luo, R. Zhan, H. Chang, F. Yang, and P. Milanfar, “Distortion agnostic deep watermarking,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, 13548–57.
- [152] Y. Luo, Y. Zhang, J. Yan, and W. Liu, “Generalizing face forgery detection with high-frequency features,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 16317–26.
- [153] Z. Luo, D. Chen, Y. Zhang, Y. Huang, L. Wang, Y. Shen, D. Zhao, J. Zhou, and T. Tan, “VideoFusion: Decomposed Diffusion Models for High-Quality Video Generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, 10209–18.
- [154] C. Lutz, “Digital inequalities in the age of artificial intelligence and big data,” *Human Behavior and Emerging Technologies*, 1(2), 2019, 141–8.
- [155] S. Lyu, X. Pan, and X. Zhang, “Exposing region splicing forgeries with blind local noise estimation,” *International Journal of Computer Vision*, 110, 2014, 202–21.
- [156] D. Ma, R. Wan, B. Shi, A. C. Kot, and L.-Y. Duan, “Learning to Jointly Generate and Separate Reflections,” in *IEEE International Conference on Computer Vision*, 2019, 2444–52.
- [157] L. Ma, X. Li, J. Liao, Q. Zhang, X. Wang, J. Wang, and P. V. Sander, “Deblur-NeRF: Neural Radiance Fields from Blurry Images,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 12851–60.
- [158] Y. Ma, Y. He, X. Cun, X. Wang, Y. Shan, X. Li, and Q. Chen, “Follow Your Pose: Pose-Guided Text-to-Video Generation using Pose-Free Videos,” *arXiv preprint arXiv:2304.01186*, 2023.
- [159] S. Mandelli, D. Cozzolino, E. D. Cannas, J. P. Cardenuto, D. Moreira, P. Bestagini, W. J. Scheirer, A. Rocha, L. Verdoliva, S. Tubaro, and E. J. Delp, “Forensic Analysis of Synthetically Generated Western Blot Images,” *IEEE Access*, 10, 2022, 59919–32.
- [160] N. Mangaokar, J. Pu, P. Bhattacharya, C. K. Reddy, and B. Viswanath, “Jekyll: Attacking Medical Image Diagnostics using Deep Generative Models,” in *IEEE European Symposium on Security and Privacy (EuroSecP)*, 2020, 139–57.
- [161] F. Manjoo, “ChatGPT Is Already Changing How I Do My Job,” <https://www.nytimes.com/2023/04/21/opinion/chatgpt-journalism.html>, (Accessed on 06/06/2023), 2023.
- [162] F. Marra, D. Gagnaniello, L. Verdoliva, and G. Poggi, “Do GANs Leave Artificial Fingerprints?” In *IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2019, 506–11.

- [163] R. Martin-Brualla, N. Radwan, M. S. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth, “NeRF in the Wild: Neural radiance fields for unconstrained photo collections,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 7210–9.
- [164] C. Martinez, “Artificial Intelligence and Accessibility: Examples of a Technology that Serves People with Disabilities,” <https://www.inclusivecitymaker.com/artificial-intelligence-accessibility-examples-technology-serves-people-disabilities/>, (Accessed on 06/05/2023), 2021.
- [165] F. Matern, C. Riess, and M. Stamminger, “Exploiting visual artifacts to expose deepfakes and face manipulations,” in *IEEE Winter Conference on Applications of Computer Vision Workshops*, 2019, 83–92.
- [166] M. Mathieu, C. Couprie, and Y. LeCun, “Deep multi-scale video prediction beyond mean square error,” *arXiv preprint*, arXiv:1511.05440, 2015.
- [167] O. Mayer and M. C. Stamm, “Accurate and efficient image forgery detection using lateral chromatic aberration,” *IEEE Transactions on Information Forensics and Security*, 13(7), 2018, 1762–77.
- [168] S. McCloskey and M. Albright, “Detecting GAN-generated imagery using saturation cues,” in *IEEE International Conference on Image Processing*, 2019, 4584–8.
- [169] L. Mearian, “Schools look to ban ChatGPT, students use it anyway,” <https://www.computerworld.com/article/3694195/schools-look-to-ban-chatgpt-students-use-it-anyway.html>, (Accessed on 05/16/2023), 2023.
- [170] A. Mehrish, N. Majumder, R. Bharadwaj, R. Mihalcea, and S. Poria, “A review of deep learning techniques for speech processing,” *Information Fusion*, 2023, 101869.
- [171] C. Miao, Z. Tan, Q. Chu, N. Yu, and G. Guo, “Hierarchical Frequency-Assisted Interactive Networks for Face Manipulation Detection,” *IEEE Transactions on Information Forensics and Security*, 17, 2022, 3008–21.
- [172] “Midjourney,” <https://www.midjourney.com>, (Accessed on 05/22/2023), 2023.
- [173] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis,” in *European Conference on Computer Vision*, 2020, 405–21.
- [174] Y. Mirsky, T. Mahler, I. Shelef, and Y. Elovici, “CT-GAN: Malicious Tampering of 3D Medical Imagery Using Deep Learning,” in *USENIX Conference on Security Symposium*, 2019, 461–78.
- [175] E. Mitchell, Y. Lee, A. Khazatsky, C. D. Manning, and C. Finn, “DetectGPT: Zero-shot machine-generated text detection using probability curvature,” *arXiv preprint*, arXiv:2301.11305, 2023.

- [176] S. Mitrović, D. Andreoletti, and O. Ayoub, “ChatGPT or human? detect and explain. explaining decisions of machine learning model for detecting short ChatGPT-generated text,” *arXiv preprint*, arXiv:2301.13852, 2023.
- [177] T. Mittal, U. Bhattacharya, R. Chandra, A. Bera, and D. Manocha, “Emotions don’t lie: An audio-visual deepfake detection method using affective cues,” in *ACM International Conference on Multimedia*, 2020, 2823–32.
- [178] D. Moreira, A. Bharati, J. Brogan, A. Pinto, M. Parowski, K. W. Bowyer, P. J. Flynn, A. Rocha, and W. J. Scheirer, “Image provenance analysis at scale,” *IEEE Transactions on Image Processing*, 27(12), 2018, 6109–23.
- [179] E. Morris, “Photography as a Weapon,” <https://archive.nytimes.com/opinionator.blogs.nytimes.com/2008/08/11/photography-as-a-weapon/>, (Accessed on 06/05/2023), 2008.
- [180] T. Müller, A. Evans, C. Schied, and A. Keller, “Instant neural graphics primitives with a multiresolution hash encoding,” *ACM Transactions on Graphics*, 41(4), 2022, 1–15.
- [181] R. Naika, “An overview of automatic speaker verification system,” in *International Conference on Intelligent Computing and Information and Communication (ICICC)*, 2018, 603–10.
- [182] H. H. Nguyen, J. Yamagishi, and I. Echizen, “Capsule-Forensics: Using capsule networks to detect forged images and videos,” in *International Conference on Acoustics, Speech and Signal Processing*, 2019, 2307–11.
- [183] S. J. Nightingale and H. Farid, “AI-synthesized faces are indistinguishable from real faces and more trustworthy,” *Proceedings of the National Academy of Sciences*, 119(8), 2022.
- [184] Y. Nirkin, L. Wolf, Y. Keller, and T. Hassner, “DeepFake detection based on discrepancies between faces and their context,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10), 2021, 6111–21.
- [185] D. O’Sullivan, “Nonconsensual deepfake porn puts AI in spotlight,” <https://edition.cnn.com/2023/02/16/tech/nonconsensual-deepfake-porn/index.html>, (Accessed on 06/05/2023), 2023.
- [186] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, “Neural discrete representation learning,” in *International Conference on Neural Information Processing Systems*, Vol. 30, 2017, 6309–18.
- [187] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “WaveNet: A generative model for raw audio,” *arXiv preprint*, arXiv:1609.03499, 2016.

- [188] T. Osakabe, M. Tanaka, Y. Kinoshita, and H. Kiya, “CycleGAN without checkerboard artifacts for counter-forensics of fake-image detection,” in *International Workshop on Advanced Imaging Technology (IWAIT)*, 2021, 1176609.
- [189] Y. Ostrovsky, P. Cavanagh, and P. Sinha, “Perceiving Illumination Inconsistencies in Scenes,” *Perception*, 34(11), 2005, 1301–14.
- [190] R. Padilha, A. Theófilo, F. A. Andaló, D. A. Vega-Oliveros, J. P. Cardenuto, G. Bertocco, J. Nascimento, J. Yang, and A. Rocha, “A Inteligência Artificial e os desafios da Ciência Forense Digital no século XXI,” *Estudos Avançados*, 35(101), 2021, 113–38.
- [191] Y. Pan, Z. Qiu, T. Yao, H. Li, and T. Mei, “To create what you tell: Generating videos from captions,” in *ACM International Conference on Multimedia*, 2017, 1789–98.
- [192] Z. Pan, W. Yu, X. Yi, A. Khan, F. Yuan, and Y. Zheng, “Recent progress on generative adversarial networks (GANs): A survey,” *IEEE Access*, 7, 2019, 36322–33.
- [193] K. Peng, W. Ping, Z. Song, and K. Zhao, “Non-autoregressive neural text-to-speech,” in *International Conference on Machine Learning*, 2020, 7586–98.
- [194] W. Ping, K. Peng, and J. Chen, “ClariNet: Parallel wave generation in end-to-end text-to-speech,” *arXiv preprint*, arXiv:1807.07281, 2018.
- [195] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov, “Grad-TTS: A diffusion probabilistic model for text-to-speech,” in *International Conference on Machine Learning*, 2021, 8599–608.
- [196] R. Prenger, R. Valle, and B. Catanzaro, “WaveGlow: A flow-based generative network for speech synthesis,” in *International Conference on Acoustics, Speech and Signal Processing*, 2019, 3617–21.
- [197] A. Pumarola, E. Corona, G. Pons-Moll, and F. Moreno-Noguer, “D-NeRF: Neural radiance fields for dynamic scenes,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 10318–27.
- [198] C. Qi, J. Zhang, and P. Luo, “Emerging Concern of Scientific Fraud: Deep Learning and Image Manipulation,” *bioRxiv*, 2021.
- [199] H. Qi, Q. Guo, F. Juefei-Xu, X. Xie, L. Ma, W. Feng, Y. Liu, and J. Zhao, “DeepRhythm: Exposing deepfakes with attentional visual heartbeat rhythms,” in *ACM International Conference on Multimedia*, 2020, 4318–27.
- [200] Y. Qian, G. Yin, L. Sheng, Z. Chen, and J. Shao, “Thinking in frequency: Face forgery detection by mining frequency-aware clues,” in *European Conference on Computer Vision*, 2020, 86–103.
- [201] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning transferable visual models from natural language supervision,” in *International Conference on Machine Learning*, 2021, 8748–63.

- [202] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” *arXiv preprint*, arXiv:1511.06434, 2015.
- [203] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, “Language models are unsupervised multitask learners,” *OpenAI*, 2019.
- [204] J. W. Rae, S. Borgeaud, T. Cai, K. Millican, J. Hoffmann, H. F. Song, J. Aslanides, S. Henderson, R. Ring, S. Young, E. Rutherford, *et al.*, “Scaling language models: Methods, analysis & insights from training Gopher,” *arXiv preprint*, arXiv:2112.11446, 2021.
- [205] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, “Hierarchical text-conditional image generation with clip latents,” *arXiv preprint*, arXiv:2204.06125, 2022.
- [206] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, “Zero-shot text-to-image generation,” in *International Conference on Machine Learning*, 2021, 8821–31.
- [207] M. Ranzato, A. Szlam, J. Bruna, M. Mathieu, R. Collobert, and S. Chopra, “Video (language) modeling: a baseline for generative models of natural videos,” *arXiv preprint*, arXiv:1412.6604, 2014.
- [208] A. Razavi, A. Van den Oord, and O. Vinyals, “Generating diverse high-fidelity images with VQ-VAE-2,” in *Advances in Neural Information Processing Systems*, Vol. 32, 2019.
- [209] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “FastSpeech: Fast, robust and controllable text to speech,” in *Advances in Neural Information Processing Systems*, Vol. 32, 2019.
- [210] A. Richard, D. Markovic, I. D. Gebru, S. Krenn, G. A. Butler, F. Torre, and Y. Sheikh, “Neural synthesis of binaural speech from mono audio,” in *International Conference on Learning Representations*, 2021.
- [211] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-Resolution Image Synthesis With Latent Diffusion Models,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 10684–95.
- [212] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Nießner, “Faceforensics++: Learning to detect manipulated facial images,” in *IEEE International Conference on Computer Vision*, 2019, 1–11.
- [213] L. Ruan, Y. Ma, H. Yang, H. He, B. Liu, J. Fu, N. J. Yuan, Q. Jin, and B. Guo, “Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, 10219–28.

- [214] S. Ruberg, “Backlash against AI supermodels triggers wider fears in fashion workforce,” <https://www.nbcnews.com/business/business-news/ai-models-levis-controversy-backlash-rcna77280>, (Accessed on 05/16/2023), 2023.
- [215] M. Ruggier, “Whitney Houston died 10 year ago, Vegas hologram show captures legacy,” <https://eu.usatoday.com/story/entertainment/music/2022/02/11/whitney-houston-died-10-year-ago-las-vegas-hologram-show-captures-legacy/6669768001/>, (Accessed on 06/05/2023), 2022.
- [216] E. Sabir, J. Cheng, A. Jaiswal, W. AbdAlmageed, I. Masi, and P. Natarajan, “Recurrent convolutional strategies for face manipulation detection in videos,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, 80–7.
- [217] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. Gontijo Lopes, B. Karagol Ayan, T. Salimans, J. Ho, D. J. Fleet, and M. Norouzi, “Photorealistic text-to-image diffusion models with deep language understanding,” in *Advances in Neural Information Processing Systems*, Vol. 35, 2022, 36479–94.
- [218] M. Saito, E. Matsumoto, and S. Saito, “Temporal generative adversarial nets with singular value clipping,” in *IEEE International Conference on Computer Vision*, 2017, 2830–9.
- [219] A. Sauer, T. Karras, S. Laine, A. Geiger, and T. Aila, “StyleGAN-T: Unlocking the Power of GANs for Fast Large-Scale Text-to-Image Synthesis,” in *International Conference on Machine Learning*, 2023.
- [220] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé, *et al.*, “BLOOM: A 176b-parameter open-access multilingual language model,” *arXiv preprint*, arXiv:2211.05100, 2022.
- [221] W. J. Scheirer, A. de Rezende Rocha, A. Sapkota, and T. E. Boult, “Toward open set recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7), 2012, 1757–72.
- [222] R. Schuster, C. Song, E. Tromer, and V. Shmatikov, “You Autocomplete Me: Poisoning Vulnerabilities in Neural Code Completion,” in *USENIX Security Symposium*, 2021, 1559–75.
- [223] K. Seonghyeon, “Github Repository of Style GAN2 Pytorch,” <https://github.com/rosinality/stylegan2-pytorch>, (Accessed on 05/22/2023), 2021.
- [224] Z. Shang, H. Xie, Z. Zha, L. Yu, Y. Li, and Y. Zhang, “PRRNet: Pixel-Region relation network for face forgery detection,” *Pattern Recognition*, 116, 2021, 107950.
- [225] K. Shiohara and T. Yamasaki, “Detecting Deepfakes with Self-Blended Images,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 18720–9.

- [226] U. Singer, A. Polyak, T. Hayes, X. Yin, J. An, S. Zhang, Q. Hu, H. Yang, O. Ashual, O. Gafni, D. Parikh, S. Gupta, and Y. Taigman, “Make-a-video: Text-to-video generation without text-video data,” in *International Conference on Learning Representations*, 2023.
- [227] I. Skorokhodov, S. Tulyakov, and M. Elhoseiny, “StyleGAN-V: A continuous video generator with the price, image quality and perks of StyleGAN2,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 3626–36.
- [228] S. Smith, M. Patwary, B. Norick, P. LeGresley, S. Rajbhandari, J. Casper, Z. Liu, S. Prabhunoye, G. Zerveas, V. Korthikanti, et al., “Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, A Large-Scale Generative Language Model,” *arXiv preprint*, arXiv:2201.11990, 2022.
- [229] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *International Conference on Machine Learning*, 2015, 2256–65.
- [230] I. Solak, “The M-AILABS Speech Dataset,” <https://www.caito.de/2019/01/03/the-m-ailabs-speech-dataset/>, (Accessed on 06/06/2023), 2019.
- [231] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” *arXiv preprint*, arXiv:2011.13456, 2020.
- [232] N. Srivastava, E. Mansimov, and R. Salakhudinov, “Unsupervised learning of video representations using LSTMs,” in *International Conference on Machine Learning*, 2015, 843–52.
- [233] “Stable Diffusion GitHub Repository,” <https://github.com/CompVis/stable-diffusion>, (Accessed on 05/22/2023), 2023.
- [234] “Stable Diffusion web UI,” <https://github.com/AUTOMATIC1111/stable-diffusion-webui>, (Accessed on 05/22/2023), 2023.
- [235] I. Stanley-Becker and J. Wagner, “RNC counters Biden announcement with dystopian, AI-aided video,” <https://www.washingtonpost.com/politics/2023/04/25/rnc-biden-ad-ai/>, (Accessed on 06/05/2023), 2023.
- [236] J. Stribling, M. Krohn, and D. Aguayo, “SCIgen - An Automatic CS Paper Generator,” <https://pdos.csail.mit.edu/archive/scigen/>, (Accessed on 06/06/2023), 2005.
- [237] N. Subbaraman, “ChatGPT Will See You Now: Doctors Using AI to Answer Patient Questions,” <https://www.wsj.com/articles/dr-chatgpt-physicians-are-sending-patients-advice-using-ai-945cf60b>, (Accessed on 06/06/2023), 2023.
- [238] C. Sun, S. Jia, S. Hou, and S. Lyu, “AI-Synthesized Voice Detection Using Neural Vocoder Artifacts,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2023, 904–12.

- [239] S. Svrluga, “Was that essay written by AI? A student made an app that might tell you,” <https://www.washingtonpost.com/education/2023/01/12/gptzero-chatgpt-detector-ai/>, (Accessed on 06/06/2023), 2023.
- [240] M. Tan and Q. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in *International Conference on Machine Learning*, 2019, 6105–14.
- [241] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, “A survey on neural speech synthesis,” *arXiv preprint*, arXiv:2106.15561, 2021.
- [242] M. Tao, B.-K. Bao, H. Tang, and C. Xu, “GALIP: Generative Adversarial CLIPs for Text-to-Image Synthesis,” *arXiv preprint*, arXiv:2301.12959, 2023.
- [243] R. Thoppilan, D. de Freitas, J. Hall, N. Shazeer, A. Kulshreshtha, H. Cheng, A. Jin, T. Bos, L. Baker, Y. Du, *et al.*, “LaMDA: Language models for dialog applications,” *arXiv preprint*, arXiv:2201.08239, 2022.
- [244] Y. Tian, J. Ren, M. Chai, K. Olszewski, X. Peng, D. N. Metaxas, and S. Tulyakov, “A good image generator is what you need for high-resolution video synthesis,” *arXiv preprint*, arXiv:2104.15069, 2021.
- [245] M. Todisco, H. Delgado, and N. W. Evans, “A New Feature for Automatic Speaker Verification Anti-Spoofing: Constant Q Cepstral Coefficients,” in *The Speaker and Language Recognition Workshop (Odyssey)*, 2016, 283–90.
- [246] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, *et al.*, “LLaMA: Open and efficient foundation language models,” *arXiv preprint*, arXiv:2302.13971, 2023.
- [247] T. H. Tran, “A Doctor Published Several Research Papers with Breakneck Speed. ChatGPT Wrote Them All.,” <https://www.thedailybeast.com/how-this-doctor-wrote-dozens-of-science-papers-with-chatgpt>, (Accessed on 06/06/2023), 2023.
- [248] J. Traylor, “AI-generated synthetic media is starting to permeate the Internet,” <https://www.nbcnews.com/tech/tech-news/ai-generated-synthetic-media-future-content-rcna72958>, (Accessed on 06/05/2023), 2023.
- [249] A. Tuama, F. Comby, and M. Chaumont, “Camera model identification with the use of deep convolutional neural networks,” in *IEEE International Workshop on Information Forensics and Security*, 2016, 1–6.
- [250] S. Tulyakov, M.-Y. Liu, X. Yang, and J. Kautz, “MoCoGAN: Decomposing motion and content for video generation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, 1526–35.

- [251] L. Ulanof, “Photoshop AI Generative Fill is so powerful it might change photo editing forever,” <https://www.techradar.com/features/photoshop-ai-generative-fill-is-so-powerful-it-might-change-photo-editing-forever>, (Accessed on 06/05/2023), 2023.
- [252] S. S. Usmani, M. Sharath, and M. Mehendale, “Future of mental health in the metaverse,” *General Psychiatry*, 35(4), 2022, e100825.
- [253] J. Vainilavičius, “Deepfakes could facilitate real estate fraud, experts warn,” <https://cybernews.com/security/deepfakes-could-facilitate-real-estate-fraud/>, (Accessed on 06/05/2023), 2022.
- [254] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, Vol. 30, 2017.
- [255] R. Villegas, M. Babaeizadeh, P.-J. Kindermans, H. Moraldo, H. Zhang, M. T. Saffar, S. Castro, J. Kunze, and D. Erhan, “Phenaki: Variable length video generation from open domain textual description,” *arXiv preprint*, arXiv:2210.02399, 2022.
- [256] V. Voleti, A. Jolicoeur-Martineau, and C. Pal, “Masked conditional video diffusion for prediction, generation, and interpolation,” *arXiv preprint*, arXiv:2205.09853, 2022.
- [257] C. Vondrick, H. Pirsiavash, and A. Torralba, “Generating videos with scene dynamics,” in *Advances in Neural Information Processing Systems*, Vol. 29, 2016.
- [258] S. Vora, N. Radwan, K. Greff, H. Meyer, K. Genova, M. S. Sajjadi, E. Pot, A. Tagliasacchi, and D. Duckworth, “NeSF: Neural semantic fields for generalizable semantic segmentation of 3D scenes,” *arXiv preprint*, arXiv:2111.13260, 2021.
- [259] J. P. Wahle, T. Ruas, F. Kirstein, and B. Gipp, “How Large Language Models are Transforming Machine-Paraphrased Plagiarism,” *arXiv preprint*, arXiv:2210.03568, 2022.
- [260] R. Wan, B. Shi, H. Li, L.-Y. Duan, and A. C. Kot, “Face image reflection removal,” *International Journal of Computer Vision*, 129, 2021, 385–99.
- [261] R. Wan, B. Shi, H. Li, L.-Y. Duan, and A. C. Kot, “Reflection scene separation from a single image,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, 2398–406.
- [262] R. Wan, B. Shi, H. Li, L.-Y. Duan, A.-H. Tan, and A. C. Kot, “CoRRN: Cooperative reflection removal network,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(12), 2019, 2969–82.
- [263] R. Wan, B. Shi, W. Yang, B. Wen, L.-Y. Duan, and A. C. Kot, “Purifying Low-light Images via Near-Infrared Enlightened Image,” *IEEE Transactions on Multimedia*, 2022, 1–13.
- [264] B. Wang and A. Komatsuzaki, “GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model,” <https://github.com/kingoflolz/mesh-transformer-jax>, (Accessed on 05/16/2023), 2021.

- [265] M. Wang, X. Fu, J. Liu, and Z.-J. Zha, “JPEG Compression-aware Image Forgery Localization,” in *ACM International Conference on Multimedia*, 2022, 5871–9.
- [266] S.-Y. Wang, O. Wang, R. Zhang, A. Owens, and A. A. Efros, “CNN-generated images are surprisingly easy to spot... for now,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, 8695–704.
- [267] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, G. Liu, A. Tao, J. Kautz, and B. Catanzaro, “Video-to-video synthesis,” *arXiv preprint*, arXiv:1808.06601, 2018.
- [268] W. Wang, H. Yang, Z. Tuo, H. He, J. Zhu, J. Fu, and J. Liu, “VideoFactory: Swap Attention in Spatiotemporal Diffusions for Text-to-Video Generation,” *arXiv preprint*, arXiv:2305.10874, 2023.
- [269] Y. Wang, R. Wan, W. Yang, H. Li, L.-P. Chau, and A. Kot, “Low-light image enhancement with normalizing flow,” in *AAAI Conference on Artificial Intelligence*, Vol. 36, No. 3, 2022, 2604–12.
- [270] Y. Wang, Z. Su, N. Zhang, R. Xing, D. Liu, T. H. Luan, and X. Shen, “A survey on metaverse: Fundamentals, security, and privacy,” *IEEE Communications Surveys and Tutorials*, 25(1), 2022, 319–52.
- [271] Z. Wang, S. Wu, W. Xie, M. Chen, and V. A. Prisacariu, “NeRF—: Neural Radiance Fields Without Known Camera Parameters,” *arXiv preprint*, arXiv:2102.07064, 2021.
- [272] K. W. Wong, “Titans of AI Andrew Ng and Yann LeCun oppose call for pause on powerful AI systems,” <https://venturebeat.com/ai/titans-of-ai-industry-andrew-ng-and-yann-lecun-oppose-call-for-pause-on-powerful-ai-systems/>, (Accessed on 05/16/2023), 2023.
- [273] J. Z. Wu, Y. Ge, X. Wang, W. Lei, Y. Gu, W. Hsu, Y. Shan, X. Qie, and M. Z. Shou, “Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation,” *arXiv preprint arXiv:2212.11565*, 2022.
- [274] S. Wu, O. Irsoy, S. Lu, V. Dabrovolski, M. Dredze, S. Gehrmann, P. Kambadur, D. Rosenberg, and G. Mann, “BloombergGPT: A large language model for finance,” *arXiv preprint*, arXiv:2303.17564, 2023.
- [275] Z. Wu, T. Kinnunen, N. Evans, J. Yamagishi, C. Hanilçi, M. Sahidullah, and A. Sizov, “ASVspoof 2015: the first automatic speaker verification spoofing and countermeasures challenge,” in *Interspeech*, 2015.
- [276] C. Xiang, “People Are Creating Records of Fake Historical Events Using AI,” <https://www.vice.com/en/article/k7zqdw/people-are-creating-records-of-fake-historical-events-using-ai>, (Accessed on 06/05/2023), 2023.
- [277] Y. Xing, W. He, J. Z. Zhang, and G. Cao, “AI privacy opinions between US and chinese people,” *Journal of Computer Information Systems*, 63(3), 2023, 492–506.

- [278] X. Xuan, B. Peng, W. Wang, and J. Dong, “On the generalization of GAN image forensics,” in *Chinese Conference on Biometric Recognition (CCBR)*, 2019, 134–41.
- [279] J. Yamagishi, X. Wang, M. Todisco, M. Sahidullah, J. Patino, A. Nautsch, X. Liu, K. A. Lee, T. Kinnunen, N. Evans, *et al.*, “ASVspoof 2021: accelerating progress in spoofed and deepfake speech detection,” in *Automatic Speaker Verification and Spoofing Countermeasures Challenge*, 2021.
- [280] R. Yamamoto, E. Song, and J.-M. Kim, “Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram,” in *International Conference on Acoustics, Speech and Signal Processing*, 2020, 6199–203.
- [281] W. Yan, Y. Zhang, P. Abbeel, and A. Srinivas, “VideoGPT: Video generation using VQ-VAE and transformers,” *arXiv preprint*, arXiv:2104.10157, 2021.
- [282] J. Yang, R. K. Das, and H. Li, “Extended constant-Q cepstral coefficients for detection of spoofing attacks,” in *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2018, 1024–9.
- [283] R. Yang, P. Srivastava, and S. Mandt, “Diffusion probabilistic modeling for video generation,” *arXiv preprint*, arXiv:2203.09481, 2022.
- [284] X. Yang, Y. Li, and S. Lyu, “Exposing deep fakes using inconsistent head poses,” in *International Conference on Acoustics, Speech and Signal Processing*, 2019, 8261–5.
- [285] H. Yen, F. G. Germain, G. Wichern, and J. Le Roux, “Cold Diffusion for Speech Enhancement,” in *International Conference on Acoustics, Speech and Signal Processing*, 2023, 1–5.
- [286] J. Yi, Y. Bai, J. Tao, Z. Tian, C. Wang, T. Wang, and R. Fu, “Half-truth: A partially fake audio detection dataset,” *arXiv preprint*, arXiv:2104.03617, 2021.
- [287] D. Yingying, T. Fan, D. Weiming, M. Chongyang, P. Xingjia, W. Lei, and X. Changsheng, “StyTr2: Image Style Transfer with Transformers,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [288] L. Yu, Y. Cheng, K. Sohn, J. Lezama, H. Zhang, H. Chang, A. G. Hauptmann, M.-H. Yang, Y. Hao, I. Essa, and L. Jiang, “MAGVIT: Masked generative video transformer,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2023, 10459–69.
- [289] N. Yu, L. S. Davis, and M. Fritz, “Attributing fake images to GANs: Learning and analyzing GAN fingerprints,” in *IEEE International Conference on Computer Vision*, 2019, 7556–66.

- [290] P. Yu, J. Chen, X. Feng, and Z. Xia, “CHEAT: A Large-scale Dataset for Detecting ChatGPT-writtEn AbsTracts,” *arXiv preprint*, arXiv:2304.12008, 2023.
- [291] S. Yu, K. Sohn, S. Kim, and J. Shin, “Video probabilistic diffusion models in projected latent space,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2023, 18456–66.
- [292] S. Yu, J. Tack, S. Mo, H. Kim, J. Kim, J.-W. Ha, and J. Shin, “Generating videos with dynamics-aware implicit generative adversarial networks,” *arXiv preprint*, arXiv:2202.10571, 2022.
- [293] R. Zellers, A. Holtzman, H. Rashkin, Y. Bisk, A. Farhadi, F. Roesner, and Y. Choi, “Defending against neural fake news,” in *Advances in Neural Information Processing Systems*, Vol. 32, 2019.
- [294] F. Zhan, Y. Yu, R. Wu, J. Zhang, S. Lu, L. Liu, A. Kortylewski, C. Theobalt, and E. Xing, “Multimodal image synthesis and editing: A survey,” *arXiv preprint*, arXiv:2112.13592, 2021.
- [295] B. Zhang, S. Gu, B. Zhang, J. Bao, D. Chen, F. Wen, Y. Wang, and B. Guo, “StyleSwin: Transformer-based GAN for High-resolution Image Generation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 11304–14.
- [296] D. Zhang, C. Li, F. Lin, D. Zeng, and S. Ge, “Detecting Deepfake Videos with Temporal Dropout 3DCNN,” in *International Joint Conference on Artificial Intelligence*, 2021, 1288–94.
- [297] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. Metaxas, “StackGAN: Text to Photo-realistic Image Synthesis with Stacked Generative Adversarial Networks,” in *IEEE International Conference on Computer Vision*, 2017, 5907–15.
- [298] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, *et al.*, “OPT: Open pre-trained transformer language models,” *arXiv preprint*, arXiv:2205.01068, 2022.
- [299] Z. Zhang, X. Yi, and X. Zhao, “GAN based synthesized audio dataset,” <https://ieee-dataport.org/documents/gan-based-synthesized-audio-dataset>, (Accessed on 06/06/2023), 2019.
- [300] H. Zhao, W. Zhou, D. Chen, T. Wei, W. Zhang, and N. Yu, “Multi-attentional deepfake detection,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 2185–94.
- [301] T. Zhao, X. Xu, M. Xu, H. Ding, Y. Xiong, and W. Xia, “Learning self-consistency for deepfake detection,” in *IEEE International Conference on Computer Vision*, 2021, 15023–33.
- [302] S. Zhi, T. Laidlow, S. Leutenegger, and A. J. Davison, “In-place scene labelling and understanding with implicit scene representation,” in *IEEE International Conference on Computer Vision*, 2021, 15818–27.

- [303] D. Zhou, W. Wang, H. Yan, W. Lv, Y. Zhu, and J. Feng, “MagicVideo: Efficient video generation with latent diffusion models,” *arXiv preprint*, arXiv:2211.11018, 2022.
- [304] P. Zhou, X. Han, V. I. Morariu, and L. S. Davis, “Learning rich features for image manipulation detection,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, 1053–61.
- [305] T. Zhou, W. Wang, Z. Liang, and J. Shen, “Face forensics in the wild,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 5778–88.
- [306] Y. Zhou, R. Zhang, C. Chen, C. Li, C. Tensmeyer, T. Yu, J. Gu, J. Xu, and T. Sun, “Towards language-free training for text-to-image generation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, 17907–17.
- [307] C. Zhu, R. Wan, and B. Shi, “Neural Transmitted Radiance Fields,” in *Advances in Neural Information Processing Systems*, Vol. 35, 2022, 38994–9006.
- [308] C. Zhu, R. Wan, Y. Tang, and B. Shi, “Occlusion-Free Scene Recovery via Neural Radiance Fields,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2023, 20722–31.
- [309] X. Zhu, H. Wang, H. Fei, Z. Lei, and S. Z. Li, “Face forgery detection by 3D decomposition,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, 2929–39.
- [310] L. Zhuo, S. Tan, B. Li, and J. Huang, “Self-adversarial training incorporating forgery attention for image forgery localization,” *IEEE Transactions on Information Forensics and Security*, 17, 2022, 819–34.