Robust Deepfake On Unrestricted Media: Generation And Detection

Trung-Nghia Le⋆, Huy H Nguyen, Junichi Yamagishi, and Isao Echizen

National Institute of Informatics, Tokyo, Japan
The Graduate University for Advanced Studies (SOKENDAI), Kanagawa, Japan
University of Tokyo, Tokyo, Japan
`{ltnghia,nhhuy,jyamagis,iechizen}@nii.ac.jp`

Abstract. Recent advances in deep learning have led to substantial improvements in deepfake generation, resulting in fake media with a more realistic appearance. Although deepfake media have potential application in a wide range of areas and are drawing much attention from both the academic and industrial communities, it also leads to serious social and criminal concerns. This chapter explores the evolution of and challenges in deepfake generation and detection. It also discusses possible ways to improve the robustness of deepfake detection for a wide variety of media (*e.g.*, in-the-wild images and videos). Finally, it suggests a focus for future fake media research.

Keywords: deepfake generation, deepfake detection, robustness, in-the-wild

1 Introduction

Deep learning has been successfully used to power various applications such as big data analysis, natural language processing, signal processing, computer vision, human-computer interaction, medical imaging, and media forensics. Recent advances in deep learning have led to substantial improvements in deepfake generation ("deepfake" for short) (*e.g.*, deep learning-based face forgery, AI-based face forgery) that can change the target person's identity [1–4]. In addition, emerging technologies such as autoencoders (AEs) and generative adversarial networks (GANs) enable transferring one person's face to another person while retaining the target person's facial expression and head pose [5–8].

The realistic appearances of images and videos synthesized using deepfake techniques have recently drawn much attention in the computer vision and computer graphics fields. Moreover, they have potential applications in a wide range of areas, such as education [9], pattern and design creation, film and art creation (*i.e.*, digital avatars [10, 11], and beauty filters [12–14]).

However, the ease of creating falsified AI-synthesized images and videos has led to serious concerns about individual harassment, criminal deception, fake news, hoaxes, and financial fraud. Public figures such as celebrities and politicians are easy targets

---

⋆ Corresponding author. This article will appear as one chapter of the book Fake Media Generation and Detection, edited by Mahdi Khosravy, Isao Echizen, Noboru Babaguchi.
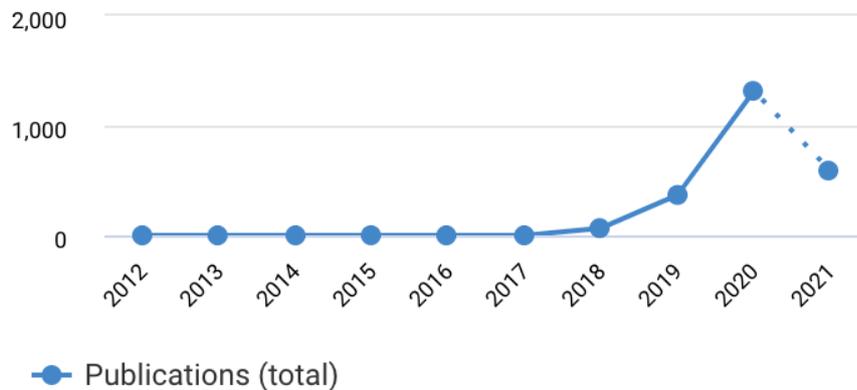
Fig. 1: Number of deepfake papers published 2012-2021.

of deepfake attacks. Their faces can be swapped onto the faces of porn stars, for example [15]. Fake speeches by world leaders can be crafted for falsification purposes, threatening world security. Deepfakes can also be a threat to ordinary people, such as through money scamming [16]. DeepNude software [17] is a particularly disturbing threat as it can transform anyone into a non-consensual porn actor. Fake news can cause political and/or religious tensions between countries, fool the public and thereby affect the results of elections, and create chaos in financial markets. These forms of falsification are a massive threat to privacy and identity and can affect many aspects of human lives.

To address this threat, it is essential to develop countermeasures against face forgeries in digital media. The research community has thus made a great effort to accelerate the development of means for detecting facial manipulation in images and videos. The number of media forensics workshops and conferences has been rapidly increasing. As shown in Figure 1, the number of deepfake detection papers[1] has increased substantially in the last few years. In addition, competitions in both academia (NIST's Open Media Forensics Challenge[2] and NTU's DeeperForensics Challenge[3]) and industry (Facebook's Deepfake Detection Challenge[4]) are held regularly to address the threat of face-swapping.

This chapter provides an in-depth review of deepfake generation and detection methods from the viewpoint of computer vision. It also highlights deepfake datasets and critical benchmarks for different tasks related to face manipulation. In addition, it insightfully discusses potential ways of improving the robustness of deepfake detection methods.

---

[1] https://app.dimensions.ai

[2] https://mfc.nist.gov/

[3] https://competitions.codalab.org/competitions/25228

[4] https://www.kaggle.com/c/deepfake-detection-challenge

Table 1: Evolution of deepfake generation methods 2016–2021.

2016	2017	2018	2019	2020	2021
Face2Face [5]	Deepfakes [2]	FaceSwapGAN [3]	ZAO [14]	Reface [13]	Zhu et al. [18]
	FaceApp [12]	ProGAN [19]	StyleGAN [20]	DeepFaceLab [21]	Le et al. [22]
	Korshunova et al. [23]	StarGAN [24]	NeuralTextures [6]	StyleGANv2 [25]	
		ReenactGAN [26]	GANimation [27]	StarGANv2 [28]	
		RsGAN [29]	FSGAN [7]	InterFaceGAN [30]	
		X2Face [31]	Zhang et al. [32]	StyleALAE [33]	
		FSNet [34]	Egor et al. [11]	FaceShifter [4]	
		Zhixin et al. [8]		ICface [35]	
		Kim et al. [36]		FaR-GAN [37]	

The remainder of this chapter is organized as follows. Section 2 introduces deepfake generation methods. Next, Section 3 presents and discusses deepfake detection methods. Finally, Section 4 summarizes what is covered in this chapter.

2   Deepfake Generation

2.1   Deepfake Evolution

Table 1 illustrates the evolution of deepfakes from 2016 to 2021. Traditional manipulation methods based on hand-crafted features were developed to add (splice), remove (inpaint), and replicate (copy/move) objects within images. Recent advancements in deep learning support various automated manipulation approaches that achieve realistic appearances, namely deepfake generation. The first deepfake generation method developed was based on a simple AE [2]. Subsequently, GANs [3] started to be used to create realistic fake media. GAN-based methods are more generalized than AE-based ones in synthesizing realistic manipulated faces because they work without being explicitly trained on subject-specific images. Now, AE- and GAN-based methods are the two most widely used deepfake generation methods.

Deepfake generators are usually categorized by their application (*e.g.*, identity swapping, expression reenactment, face synthesis, and attribute manipulation). Identity swapping replaces the target identity with the source identity; expression reenactment manipulates facial expressions; face synthesis creates fake facial images; attribute manipulation creates attribute fabrication. In addition, here we also differentiate methods by the evolution of the techniques they use. Furthermore, we analyze their pros and cons. We argue that modern deep learning techniques can be used in a variety of areas.

Paired-training techniques [2], first proposed in 2017, have the highest ratio of usage in deepfake generators. They focus on specific individuals and require a substantial amount of data on both the target and source individual. Moreover, the use of paired training for deepfake generation is a time-consuming process. The evolution of deep learning techniques has recently enabled the development of few-/one-/zero-shot learning approaches that overcome these disadvantages.

2.2   Identity Swapping

In identity swapping, or face replacement, a person's face in a source image is replaced with another person's face. This manipulation approach uses a deep generator (*i.e.*, AE or GAN) to give a victim's face someone else's features while preserving the original facial expression.

The first deep learning method for identity swapping, Deepfakes [2], uses AEs and was developed in 2017. It has been used, for example, to replace the faces of actors in pornographic videos with those of celebrities. Two AEs (*i.e.*, encoder-decoder pairs) are used in which the encoder is used to extract the latent features of a face and the decoder is used to reconstruct the face. After each encoder-decoder pair is trained on a person's video, the decoders are swapped to regenerate the target face with the features of the source face.

Korshunova *et al.* [23] used a fully convolutional network to transfer the facial appearance of the source person to that of another person in another image. Utilizing the style transfer technique, they combined style loss, content loss, light loss, and total variation regularization to produce realistic images. Open-source projects (*e.g.*, Deep-FaceLab [21] and FaceSwap-GAN [3]) with tutorials available on GitHub have opened the door for anyone to create fake images and videos. However, these approaches require a large amount of data on both the target and source individual for paired training, resulting in a time-consuming process.

The time-consuming limitation has been partially overcome by recent developments that have made the generation of deepfakes more efficient. Zakharov *et al.* [11] proposed using GAN-based few-/one-shot learning to generate a realistic talking-head video from an image, but this requires lengthy meta-learning on a large-scale video dataset. Li *et al.* [4] trained a two-stage framework in a zero-shot learning manner (*i.e.*, self-supervised learning) for high fidelity and occluded face-swapping. Zhu *et al.* [18] extended the latent space to maintain more facial details and then used StyleGAN2 [25] to generate high-resolution swapped facial images. Natsume *et al.* [34] combined separately encoded face and facial landmarks to generate a fake identity. These methods can be used to swap any two faces without retraining. Commercial deepfake applications launched since 2017 (*e.g.*, FaceApp [12], Reface [13], ZAO [14]) enable users without technical knowledge to easily swap faces in images and videos. However, they cannot handle faces with pose variations or that are partly occluded.

2.3   Expression Reenactment

Expression reenactment, also known as face reenactment or emotion synchronization, manipulates facial expressions by transferring facial expressions, gestures, and head movements of the source person to the target person while keeping the identity of the target person. This kind of manipulation is usually aimed at altering the person's facial expressions and synchronizing lip movements to create fabricated content.

Facial expression reenactment techniques are mainly based on 3D face reconstruction and a GAN architecture. The 3D facial modeling-based approach can accurately capture the geometry and movement of the head, resulting in a photorealistic reenacted face. Thies *et al.* [5, 6] introduced the use of 3D facial modeling combined with image

rendering for real-time transfer of the facial expressions of a person talking in front of a commodity webcam into a target face video.

Although GAN-based methods can generate photorealistic images, they require a large amount of training data in order to achieve photorealistic reenactment for unknown identities. Wu *et al*. [26] proposed encoding the source face into a boundary latent space and using a target-specific decoder to transfer the latent features onto the target face. Shu *et al*. [8] combined the disentangling of shapes in an unsupervised manner with AE models to morph expressions. Kim *et al*. [36] proposed using space-time encoding for temporally coherent synthesis in combination with a conditional GAN (cGAN) to synthesize video portraits of a target individual from a still image. This enables full control over the target by transferring a rigid head pose, facial expression, and eye motion with a high level of photorealism. Pumarola *et al*. [27] trained emotion action units based on a dual-cGAN in a fully unsupervised manner for use in generating facial animation from a single image. In addition to transferring expressions, the head pose can be controlled by using a recurrent neural network to enhance naturalness [7] by using different modalities [31] and by using human interpretable attributes and actions [35].

Few-/one-shot facial expression reenactment methods have recently been introduced to overcome the disadvantage of training on large-scale data by using a few- or even a single-source image. Ha *et al*. [38] proposed using a few-shot face reenactment framework that includes image attention analysis, target feature alignment, and landmark transformation to prevent quality degradation in unseen identity mismatch situations. These methods do not need additional fine-tuning phases for identity adaptation, yielding deepfake images that can be usefully deployed in the wild. SPADE blocks [39], used in one-shot face reenactment, can effectively generate a new source's expressed appearance by learning the latent representation of the target's facial landmarks [32, 37]. The problem has also been extended to multimedia-based reenactment. Fried *et al*. [40] developed a text-based editing approach to generating talking-head videos—the dialogue is modified to match the corresponding head movement of the speaker.

2.4 Face Synthesis

Face synthesis has been applied in many fields, such as video games and 3D modeling. Almost all face synthesis methods utilize powerful GAN models to generate entirely new facial images. GAN models have been developed that have progressively improved resolution, image quality, and realism. StyleGAN [20] improved the image resolution (1024×1024) of ProGAN [19] (128×128) while StyleGAN2 [25] further improved the image quality of StyleGAN by removing unwanted artifacts.

GAN-based face synthesis methods have been used in various applications, such as for facial attribute translation [20, 24, 25, 28], identity-attribute combination [41], identified characteristics removal [42], and interactive semantic manipulation [43, 44]. In addition, some GAN-based methods can be used in virtual try-on applications that enable consumers to try on various cosmetics in a virtual environment [45, 46]. These methods can be extended to the entire body. For example, DeepNude uses the Pix2PixHD GAN model [47] to inpaint clothing regions to generate fake nude images.

2.5   Attribute Manipulation

Facial attribute manipulation, also known as face editing or face retouching, modifies facial attributes such as hair color, hairstyle, skin color, gender, age, smile, eyeglasses, and makeup. This type of manipulation can be considered partially conditional face synthesis. Thus, GAN techniques for face synthesis are usually used for facial attribute manipulation. Choi *et al.* [24, 28] proposed using a unified model architecture that enables simultaneous training of multiple datasets with different domains to transfer various facial attributes and expressions at the same time. Disentangled facial features can be interpreted in different latent spaces, resulting in more precise control of attribute manipulation in face editing [20, 25, 30, 33]. These GAN-based methods are agnostic and trained on a huge number of images in order to work without being explicitly trained on subject-specific images. However, they tend to generate faces with new identities if the target face is not in the training distribution. In addition, they cannot handle faces with pose variations or occlusions.

2.6   Hybrid Applications

We argue that current deep learning methods can combine different deepfakes to create hybrid applications. Nirkin *et al.* [7] introduced a GAN model for both face-swapping and reenactment in real time that follows the reenact and blend strategy. Natsume *et al.* [29] encoded latent vectors for face and hair as two separate variational AEs (VAEs) and then conditionally swapped or edited the identity of the target. Le *et al.* [22] proposed a framework for generating fake identities by combining facial attribute editing and face-swapping. These methods are subject agnostic and can be applied to any pair of faces without retraining. Commercial FaceApp mobile applications [12] enable users to swap faces as well as change the age, gender, smile, and hairstyle of the faces.

2.7   Limitations

The biggest limitation of deepfake methods is the **lack of generalizability**, which can lead to such problems as low resolution, low quality, strange artifacts, facial pose variations, and occlusions (*c.f.* Fig. 2).

Images generated using conventional deepfake methods usually have **low resolution and low quality**. While some recently developed methods have improved resolution and image quality, they require a large amount of training data (*i.e.*, more than 1.5 million images), leading to time-consumed training.

Strange artifacts in generated images are mainly caused by ineffective blending and facial pose variations. Although deepfake methods for generating fake faces have steadily improved, the mismatch between the synthesized regions and original regions mostly appears due to the limitations of the blending algorithms. Hence, analysis of undesired colors and patterns in blending boundaries around the face can lead to strong deepfake detection [48], even for unknown deepfake generators.

Facial pose variations, which generally affect in-the-wild images, usually cannot be handled with non-paired-training methods. Existing deepfake methods are generally trained on frontal view facial images; they thus fail to preserve the face when there are

Fig. 2: Problems with current deepfake generation methods. From left to right: low resolution, low quality, strange artifacts due to wearable items, and facial pose variations.

large pose variations. Although the improved blending with recently developed methods [4, 22] helps to reduce artifacts, it does not completely solve the problem.

Occlusions caused by hair, glasses, or other objects cannot be handled by conventional deepfake generation methods. This leads to visual artifacts in the synthesized fake faces. A recently developed heuristic error acknowledging refinement network [4] helps to remove facial inconsistencies.

3 Deepfake Detection

Although recently developed deepfake generation methods generate realistic images that can spoof machines and fool people, visual anomalies in the images, such as low quality, low resolution, and strange artifacts, are still present. This means that differences between the genuine and manipulated areas can be parameterized and used for deepfake detection.

Conventional deepfake detection methods can only classify a given cropped face as real or fake[5]. Recently developed methods can not only recognize the authenticity of faces but also localize manipulated areas at the bounding-box and pixel levels, namely detection and segmentation, respectively (*c.f.* Fig. 3).

This section describes the two basic approaches to deepfake detection, the conventional approach and the end-to-end approach (*c.f.* Fig. 4 and Table 2). It also analyzes the robustness of deepfake detection methods and suggests future research directions.

---

[5] In this chapter, we use "deepfake detection" following existing work, but the actual task is "recognition/classification.

a) Classification   b) Detection   b) Segmentation
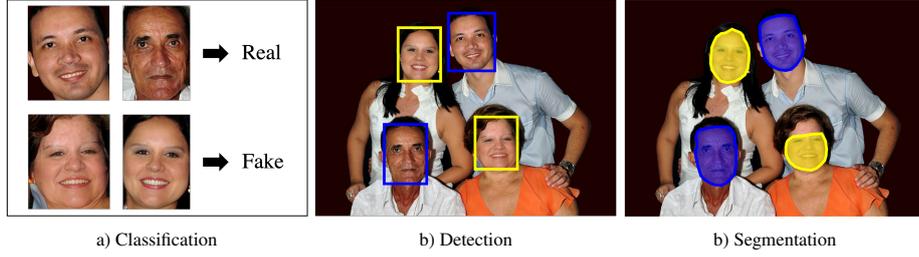
Fig. 3: Deepfake tasks (from left to right): conventional classification, end-to-end detection, and segmentation.

a) Conventional approach

b) End-to-end approach
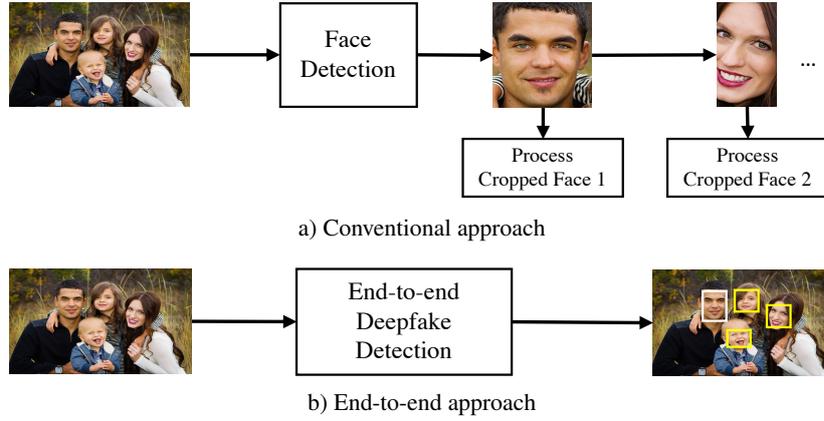
Fig. 4: Two basic approaches to deepfake detection: methods based on conventional approach use sequential processing while those based on end-to-end approach use parallel processing.

Table 2: Categorization of deepfake detection methods. ★, ♣, and ♠ indicate classification, detection, and segmentation tasks, respectively. Methods that can process multiple faces are shown in blue.

Conventional Approach				End-to-end Approach
Data-Driven	**Visual Artifacts**	**Frequency Domain**	**Attention & Segmentation**	
Zhou *et al.* [49] ★	Li *et al.* [50] ★	Li *et al.* [51] ★	Nguyen *et al.* [52] ★, ♠	Zhou *et al.* [53] ★, ♣
Afchar *et al.* [54] ★	Matern *et al.* [55] ★	Liu *et al.* [56] ★	Dang *et al.* [57] ★	Le *et al.* [22] ★, ♣, ♠
Nguyen *et al.* [58] ★	Yuezun *et al.* [59] ★		Li *et al.* [48] ★, ♠	
Rossler *et al.* [60] ★	Yang *et al.* [61] ★		Bojia *et al.* [62] ★	
Wang *et al.* [63] ★	Haliassos *et al.* [64] ★		Wang *et al.* [65] ★	
	Zhu *et al.* [66] ★			

3.1   Conventional Deepfake Detection

Face forgery classification, *i.e.*, conventional deepfake detection, is aimed at classifying facial images as real or fake. Face forgery classification methods require the input of

identified face regions (*c.f.* Fig. 4a). They do not have a face localization ability, so their performance greatly depends on the accuracy of the independent face detection method used.

Conventional deepfake classifiers are data-driven; their deep networks are directly trained on real and fake images and videos without relying on specific artifact analysis, such as a two-stream neural network [49], MesoNet [54], CapsuleNet [58], Xception-Net [60], and FakeSpotter [63].

Some methods exploit inconsistencies created by visual artifacts in deepfake images and videos; they analyze biological clues such as eye blinking [50], head pose [61], skin texture [67], and iris and teeth color [55]. Other methods use artifacts in affine face warping [59] or in mouth movements [64] to distinguish real and fake faces. Another method [66] decomposes the facial image to extract facial details (*i.e.*, direct light and identity texture) and combines them with the original face through a two-stream network to find critical forgery clues.

The frequency domain has been used in a few methods, which improves their transferability. Li *et al*. [51] developed an adaptive frequency feature generation module to mine frequency clues in combination with metric learning for improved separability in the embedding space. Liu *et al*. [56] combined the spatial image and phase spectrum to capture up-sampling artifacts and thereby improve the transferability of deepfake detection.

Most current deepfake classifiers are based on attention or segmentation masking. Wang *et al*. [65] presented an attention-based data augmentation method for guiding the detector to refine and enlarge its attention mask by occluding sensitive facial regions. Zi *et al*. [62] fed both the facial image and attention mask (*i.e.*, blurred facial landmark regions) into an attention-based face fusion network for improved classification. Dang *et al*. [57] used an attention mechanism to highlight informative regions and thereby further improve feature maps used for the classification task. Li *et al*. [48] focused on blending regions to enable forged faces to be recognized without relying on knowledge of specific face manipulation techniques. Nguyen *et al*. [52] designed a multi-task learning network for simultaneously locating manipulated regions in fake facial images.

Most existing methods aim to identify the authenticity of faces without categorizing deepfake generation types. The main reasons are the lack of data and ambiguity between deepfake categories. Recently introduced methods, which are trained on datasets [60] with multiple deepfake types, generally have the ability to recognize deepfake generation types [51,56–58,60,63,64,66]. However, these methods lack generalization due to overfitting to few generation methods in the training process and thus cannot be used in practical contexts.

3.2   End-to-end Deepfake Detection

Detecting forged faces among multiple faces in multi-person scenes, in which only a small subset of them have been manipulated, is still challenging for traditional deepfake detection methods. Several researchers have recently begun to target multi-person in-the-wild images (*c.f.* Fig. 4b). Zhou *et al*. [53] trained an attention framework to detect face forgeries in multi-person scenes. Le *et al*. [22] used transfer learning to build

end-to-end models for multiple face forgery detection and segmentation for in-the-wild images.

3.3 Robust Deepfake Detection

Vulnerability to Adversarial Attacks Adversarial attacks can modify images and video by adding specific noises in such a way that machine learning models misclassify them. CNNs are vulnerable to such attacks, and their accuracy may drop to near zero [68]. Indeed, the addition of insignificant noise can cause remarkable changes in the prediction and even completely confuse deep learning classifiers. Hussain *et al*. [69] adversarially modified fake videos and presented pipelines in both white-box and black-box attack scenarios that can fool deepfake classifiers into classifying fake videos as real. Huang *et al*. [70] showed the existence of both individual and universal adversarial perturbations that can cause well-performing deepfake classifiers to misbehave. Li *et al*. [71] demonstrated that fake facial images generated using adversarial points on a face manifold can defeat two strong forensic classifiers. Even methods that won the Deepfake Detection Challenge (DFDC) [72] were easily bypassed in a practical attack scenario using transferable and accessible adversarial attacks [73].

Object detectors are more robust than classifiers due to their complex network architecture, but there have been few reports on attacking such detectors. Treu *et al*. [74] developed a robust adversarial example method that degrades human detectors by targeting clothing regions.

Defense against adversarial attacks is a critical step toward developing robust solutions for biometrics verification (*i.e.*, deepfake detection). Many solutions have been proposed for increasing the robustness of CNNs against adversarial attacks. Most of them are independent solutions and thus do not require retraining on the victim images. Naseer *et al*. [75] introduced a self-supervised countermeasure against unseen adversarial attacks that can be deployed as a plug-and-play solution to protect different vision systems, such as those for classification, segmentation, and detection. Salman *et al*. [76] presented a defense method called "denoised smoothing" that converts existing pretrained classifiers into provably robust ones without any retraining or fine-tuning. This method can be easily integrated into public vision API providers (*i.e.*, Azure, Google, AWS, and Clarifai APIs). To defend black-box face biometrics classifiers against adversarial attacks, Theagarajan *et al*. [77] used an ensemble of iterative adversarial image purifiers for which the performance is continuously validated in a loop using Bayesian uncertainties. Akhil *et al*. [78] presented a parameter-free defense layer that is plugged into a CNN to prevent gradient- and optimization-based adversarial attacks in black-box and gray-box settings. In addition, AE- and GAN-based denoising methods [79] have been used to reconstruct the original images.

Nonrobustness on Unseen Scenarios Most conventional deepfake detection methods greatly depend on the training scenario. They are generally trained on a few of the more widely used deepfake generation methods, which are applied to a limited number of images and videos. This training of deepfake detection models on similar data can lead to overfitting poor performance in the real world. Furthermore, conventional

deepfake detection methods are not robust against media processing, such as compression, noise addition, color variation, and light variation. Indeed, although strong deep learning models have attained very high accuracy [48, 58] in laboratory environments with a simple background and a single clear front-view face [61, 80], they generally do not work well in diverse contexts that do not share a close distribution with the training dataset. Le *et al.* [22] showed that even end-to-end deepfake detectors are easily degraded on unseen image scenarios, *i.e.*, ones distant from the training distribution.

To deal with unseen forgery types, Haliassos *et al.* [64] targeted inconsistencies in mouth movements by leveraging rich representations learned via lipreading. Li *et al.* [48] focused on blending regions to recognize forged faces without relying on knowledge of specific face manipulation techniques. Sun *et al.* [81] developed a temporal method for geometrically modeling discriminative features to improve robustness for highly compressed or noise corrupted videos.

Data augmentation has also been shown to be effective for training robust deep networks. Chen *et al.* [82] augmented the process of training general object detectors. They explored model-dependent data augmentation by dynamically selecting more substantial and relevant adversarial images sourced from classification and localization branches, which can be generalized to different object detectors. In addition, improvements have been made to a couple of existing datasets by using simple perturbations, which have increased their size, such as the DFDC [72] and DeeperForensics [83] datasets. Le *et al.* [22] imitated real-world data distributions through diverse perturbations, such as color and edge manipulation, block-wise distortion, image corruption, and weather effects. Hence, to deal with unseen scenarios, recently developed augmentation libraries can be used to better train robust deepfake detection methods. In addition, if the new contextual images are still distant from the training distribution, a budget-wise solution is to use CycleGAN [84] to simulate unseen real-world contexts for training robust deep learning models.

3.4 Deepfake Datasets and Benchmarks

This section introduces deepfake datasets and analyzes methods that attained a top ranking in recent deepfake challenges.

Deepfake Datasets Table 3 shows the basic information for the main deepfake datasets. They can be categorized into two groups: those that support only image-wise or video-wise deepfake classification and those that support face-wise classification and localization tasks.

The DF-TIMIT dataset [80] has 640 fake videos crafted from the Vid-TIMIT dataset [88] using Faceswap-GAN [3]. The UADFV dataset [61] consists of 98 videos, half of which are fake, created using FaceApp [12]. The FaceForensics++ dataset [60] contains 1000 pristine videos from YouTube and 4000 synthetic videos manipulated using deepfake methods [1, 2, 5, 6]. The Google DFD dataset [86] includes 3068 fake videos. The Facebook DFDC dataset [72] contains 128K original and manipulated videos created using various deepfake and augmentation methods [7, 11, 20, 21, 89]. The Celeb-DF dataset [87] comprises 590 YouTube celebrity videos and 5,639 fake videos. The

Fig. 5: Images synthesized using robust deepfake detection methods that were trained using augmentation libraries. From top to bottom, original images, images generated using AugLy [85]. and images generated using OpenForensics [22].

Table 3: Basic information about deepfake datasets. "Cls.," "Det.," and "Seg." stand for classification, detection, and segmentation, respectively. Pristine scenarios are originally collected images/videos used to generate fake data. Unique fake scenarios are fake images/videos ignoring perturbations.

Dataset	Year	Type	Task	Fake Identity	Multi-Face	Face Occlusion	No. of Pristine Scenario	No. of Unique Fake Scenario	Data Augmentation
DF-TIMIT [80]	2018	Video	Cls.	Other videos	✗	✗	320	320	✗
UADFV [61]	2019	Video	Cls.	Other videos	✗	✗	49	49	✗
FaceForensics++ [60]	2019	Video	Cls.	Other videos	✗	✗	1,000	4,000	✗
Google DFD [86]	2019	Video	Cls.	Other videos	✗	✗	363	3,068	✗
Facebook DFDC [72]	2020	Video	Cls.	Other videos	✗	✗	48,190	104,500	✓
Celeb-DF [87]	2020	Video	Cls.	Other videos	✗	✗	590	5,639	✗
DeeperForensics [83]	2020	Video	Cls.	Hired actors	✗	✗	1,000	1,000	✓
WildDeepfake [62]	2020	Image	Cls.	N/A	✗	✗	0	707	✗
FFIW [53]	2021	Video	Det.	Other videos	✓	✗	12,000	10,000	✗
OpenForensics [22]	2021	Image	Det. / Seg.	GAN	✓	✓	45,473	70,325	✓

DeeperForensics dataset [83] consists of 10,000 videos manipulated using a deepfake VAE and augmentations on 1000 original videos in the FaceForensics++ dataset. The WildDeepfake dataset [62] contains face sequences extracted from 707 deepfake videos collected from the Internet.

The FFIW dataset [53] includes 10K videos, in which each video frame has three faces on average. Faces were randomly forged using three deepfake methods [1, 7, 21]. This dataset contains face-wise annotations for both deepfake classification and detection. The OpenForensics dataset [22] has 115K images with more than 334K faces. Different from the other datasets, the forged faces in the OpenForensics dataset were

Table 4: Top-ranked teams in DFDC Challenge 2020 [72].

Ranking	Team	Overall Log Loss
1	Selim Seferbekov	0.4279
2	WM	0.4284
3	NTechLab	0.4345
4	Eighteen Years Old	0.4347
5	The Medics	0.4371

infinitely synthesized using GAN models [30, 33]. The OpenForensics dataset also contains rich annotations that support multiple tasks.

DFDC Challenge 2020 The DFDC Challenge 2020 [72][6] was launched by Facebook in partnership with industrial and academic leaders to support the development of innovative technologies for detecting deepfakes and manipulated media. Participants had to identify whether a video contained a manipulated face. The DFDC Challenge was performed on 100,000 videos, making it the largest deepfake database to date. The challenge had 2,265 submissions on the leaderboard[7].

Table 4 lists the top-ranked teams. The top-ranked team[8] used a multi-task cascaded convolutional neural network (MTCNN) [90] for face detection and then the EfficientNet-B7 for deepfake classification. A noisy student model pre-trained on the ImageNet dataset was fine-tuned using heavy augmentations. The prediction results for video frames were combined through a heuristic fusion module.

The second-ranked team[9] used an ensemble of the Xception classifier and two weakly supervised models for data augmentation to compute frame-by-frame results and then take the average of all frames.

The third-ranked team[10] used an ensemble of EfficientNet-B7 models. One model ran on frame sequences using 3D convolution layers, and the others ran frame-by-frame, differing in the size of the face crops and augmentations. A noisy student model pre-trained on the ImageNet dataset was fine-tuned using mixup augmentation [91].

DeeperForensics Challenge 2020 The DeeperForensics Challenge 2020 [92][11] was hosted by Nanyang Technological University in conjunction with ECCV 2020. The challenge was aimed at promoting face forgery detection methods through evaluation of their performances on 60,000 videos generated by face-swapping frameworks. Of the

---

[6] https://ai.facebook.com/blog/deepfake-detection-challenge-results-an-open-initiative-to-advance-ai/

[7] https://www.kaggle.com/c/deepfake-detection-challenge

[8] https://github.com/selimsef/dfdc_deepfake_challenge

[9] https://github.com/cuihaoleo/kaggle-dfdc

[10] https://github.com/NTech-Lab/deepfake-detection-challenge

[11] https://competitions.codalab.org/competitions/25228

Table 5: Top-ranked teams in DeeperForensics Challenge 2020 [92].

Ranking	Team	Binary Cross-Entropy Loss
1	Forensics	0.2674
2	RealFace	0.3699
3	VISG	0.4060
4	jiashangplus	0.4064
5	Miao	0.4132

115 teams registered for the competition, 25 teams made valid submissions and were on the leaderboard[12].

Table 5 lists the top-ranked teams. The top-ranked team[13] used an MTCNN [90] for face detection and an ensemble of the EfficientNet-B0, EfficientNet-B1, and EfficientNet-B2 models for face-wise classification. These deep-learning models were fine-tuned using many data augmentations.

The second-ranked team used RetinaFace [93] to detect faces and then averaged the predicted results by using a video-based model and an image-based model. The image-based model took the median of the frame-wise prediction probabilities while the video-based model used an attention mechanism to fuse the temporal information between frames. Noisy student EfficientNet-B5 models pre-trained on the ImageNet dataset were fine-tuned using various data augmentations.

The third-ranked team combined all facial images detected using an MTCNN [90] into a sequence and then predicted the results using a 3D CNN. Different networks (*i.e.*, I3D [94], 3D ResNet [95], and R(2+1)D [96]) trained on diverse data augmentations were used.

OpenForensics Benchmark Different from existing deepfake datasets, which contain only a face in an image, the OpenForensics dataset [22] is aimed at identifying forged faces among many real faces in an image. This aim is achieved by providing face-wise ground truths for all classification, detection, and segmentation tasks. In addition, the dataset provides two kinds of images, ones created using seen deepfake generation methods and ones created using unseen methods.

Table 6 shows the benchmark evaluation results for multi-face forgery detection and segmentation tasks on the OpenForensics dataset [22]. For standard evaluation, Blend-Mask [103], a modern single-stage method with an attention mechanism, achieved the best performance on both the detection and segmentation tasks. The other modern single-stage methods (*i.e.*, CenterMask [102], PolarMask [104], and CondInst [106]) also had high performance.

As also shown in Table 6, YOLACT++ [101] and BlendMask [103] were the most robust against unseen images. There was a substantial drop in performance for all methods, especially CenterMask [102], for which the accuracy dropped to nearly zero. Even

---

[12] https://competitions.codalab.org/competitions/25228#results
[13] https://github.com/beibuwandeluori/DeeperForensicsChallengeSolution

Table 6: OpenForensics benchmark for multi-face forgery detection and segmentation [22]. Seen/Unseen stand for seen/unseen manipulated image generation methods.

Method	Year	Detection		Segmentation	
		Seen Images	Unseen Images	Seen Images	Unseen Images
Mask R-CNN [97]	2017	79.2	42.1	83.6	43.7
MS R-CNN [98]	2019	79.0	42.2	85.1	43.3
RetinaMask [99]	2019	80.0	48.5	82.8	48.0
YOLACT [100]	2019	68.1	49.4	72.5	51.8
YOLACT++ [101]	2020	72.9	53.7	77.3	54.7
CenterMask [102]	2020	85.5	0.03	87.2	0.02
BlendMask [103]	2020	87.0	53.9	89.2	54.0
PolarMask [104]	2020	85.0	51.7	85.0	52.7
MEInst [105]	2020	82.8	46.1	82.2	46.0
CondInst [106]	2020	84.0	52.7	87.7	54.1

the leading detection methods remain limited and cannot yet effectively address unseen images, *i.e.*, those beyond the training set's distribution.

To complete the evaluation of all deepfake detection tasks on the OpenForensics dataset, we conducted two additional experiments on recognizing the authenticity of cropped faces on both the attack and defense sides, which were not described in the original report [22].

In the attack experiment, we evaluated the performance of two widely used and high-performance deepfake classifiers (*i.e.*, XceptionNet [60] and EfficientNet-B4 [72]) against three commonly used gradient-based adversarial attacks [107] (the Fast Gradient Sign Method (FGSM) [108], the Basic Iterative Method (BIM) [109], and the Projected Gradient Descent (PGD) method [110]). FGSM is a simple yet effective method for generating adversarial images. BIM is an improved version of FGSM—FGSM is repeated multiple times with small changes in the parameters to make the attack more effective. This improved method is called "iterative-FGSM." PGD [110] is also an iterative version of FGSM—random starts are used to generate adversarial disturbances that are less perceptible.

We trained XceptionNet [60] and EfficientNet-B4 [72] on the training set and evaluated them on the test-development set of the OpenForensics dataset using their public training parameters. Table 7 shows the results for the attack methods. The performance of XceptionNet dropped more than 20% when it was attacked by FGSM and by more about 90% when it was attacked by BIM and PGD. Although EfficientNet achieved about the same accuracy as XceptionNet on original images, it was not as robust as XceptionNet when it was attacked by any of the adversarial methods: its accuracy was greatly degraded by 80% to more than 90%. Hence, deepfake classifiers are quite vulnerable to adversarial attacks.

Figure 6 shows example images produced by the compared adversarial attack methods. FGSM was the worst in terms of both performance and realism. Both BIM and PGD had similar performance that was much better than that of FGSM.

Fig. 6: Example images produced by compared adversarial attack methods. From left to right: original images followed by images produced by FGSM [108], BIM [109], and PGD [110]. Best viewed in color with zoom.

Table 7: Performance of deepfake classifiers against adversarial attacks evaluated on OpenForensics dataset [22].

Adversarial Example	Victim Model	
Attack Method	XceptionNet [60]	EfficientNet-B4 [72]
None	99.22	99.58
FGSM [108]	76.98	19.76
BIM [109]	10.92	8.39
PSD [110]	8.19	12.96

In the defense experiment, we evaluated the performance of defense solutions against seen and unseen scenarios. Particularly, we evaluated deepfake classifiers on the test-challenge set of the OpenForensics dataset, which was created by unknown generators. We trained XceptionNet [60], EfficientNet-B4 [72], and CapsuleNet-R50 [58] on standard images in the training set using their public training parameters.

As shown in Table 8, we evaluated them in two situations. First, under the assumption that the generators are seen, we re-used these methods to augment standard images

Table 8: Performance of defense solutions for deepfake classifiers against seen/unseen scenarios in OpenForensics dataset [22].

Victim Model	Defense Solution		
	None	For Seen Generators	For Unseen Generators
XceptionNet [60]	75.60	98.68	78.93
EfficientNet-B4 [72]	83.89	93.67	85.99
CapsuleNet-R50 [58]	73.07	95.99	76.30

Fig. 7: Predicted evolution of deepfake generation. From left to right, deepfake techniques will enable manipulation of human face, human body, and any visual content in the future. Note that DeepNude [17] result (middle column) has been censored.

in the training set for use in fine-tuning the deepfake classifiers. Using seen generators greatly improved the performance of the classifiers, to nearly 100%. Second, for unseen generators, we reconstructed their effects through unsupervised learning. In particular, we trained CycleGAN [84] on both standard images in the training set and unseen images in the test-challenge set, using public parameters. We then deployed CycleGAN on the training set to synthesize pseudo data for fine-tuning the deepfake classifiers to make them more robust. This increased the accuracy of all the classifiers by about 3%. This demonstrates that data augmentation can be used to train deepfake classifiers to make them more robust in both seen and unseen situations.

3.5 Future Directions

Deepfake generation methods are currently focused on the human face, so current deepfake detection methods are based on the assumption that the target objects to be manipulated are always faces. However, we cannot know which objects will be manipulated by attackers in the future because deepfake techniques can be applied to different visual objects. Indeed, DeepNude [17] slightly modifies the Pix2PixHD GAN model [47] to inpaint clothing areas with human skin, thus enabling anyone to be transformed non-consensually into a porn star. Thus, we believe that **deepfake techniques will be able to manipulate any visual content in the future**. We argue that this future manipulation will differ from conventional hand-crafted-based manipulation (*i.e.*, adding, removing, and cloning objects). Figure 7 shows examples of different kinds of deepfakes. Hence, future generations of deepfake detection methods should have the **ability to identify** not only human faces but also **various manipulated visual contents**.

Existing deepfake detectors [22, 48, 52] **lack convincibility** because they show only possible forged regions in the images or video frames without any evidence, such as
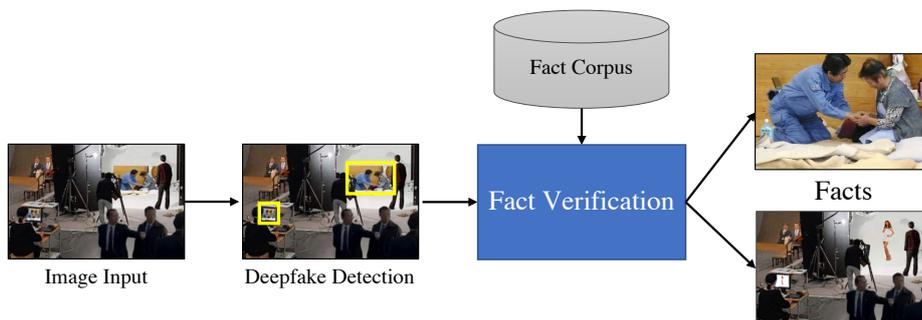
Fig. 8: Suggested workflow for fact verification.

showing the original images before manipulation, namely fact verification. The problem is more challenging when the manipulation is adversarial and done using image filters in the digital world, so promoting efforts to detect deepfakes is very tricky for both people and machines.

Verifying facts can help reduce false-positive results from deepfake detectors (authentic images or videos are classified as forged) and make them more reliable for real-world usage [111] (e.g., journalism [112]). Thus, image-based fact verification (*c.f.* Fig. 8) should be the focus in the next generation of deepfake detection methods, aiming towards trusted AI. However, to the best of our knowledge, current methods are unable to show the original data if they are manipulated.

Fact verification is a new and extremely challenging target for today's technologies. We argue that image-based face verification can be simulated via visual similarity search, *i.e.*, copy detection [113] and near-duplicate detection [114]). In particular, assuming that a huge fact database can be crawled from the Internet, we can use a similarity search to find original-like images/videos in a fact database if they are manipulated. This research field has recently drawn attention from the research community through newly developed datasets [115] and competitions [113]. However, it still is a new problem without a clear state-of-the-art solution It is already particularly challenging to search through the huge volume of content currently being generated on social media; the volume of images and videos will likely reach into the millions or even billions.

4   Summary

This chapter overviewed deepfake generation and detection methods from the viewpoint of technical evolution in computer vision. In particular, it described deepfake generation methods in different categories and analyzed their limitations. In addition, it clarified the different tasks of deepfake detection, from conventional classification to modern end-to-end detection and segmentation. It further discussed the limitations of deepfake detection methods and suggested solutions for improving the robustness of deepfake detection. Finally, it suggested a future direction for deepfake detection.

Acknowledgments

References

1. "3d face swapping," https://github.com/MarekKowalski/FaceSwap, 2016, [Online; accessed 18-Feb-2021].
2. "Deepfakes software for all," https://github.com/deepfakes/faceswap, 2017, [Online; accessed 18-Feb-2021].
3. "A denoising autoencoder, adversarial losses and attention mechanisms for face swapping," https://github.com/shaoanlu/faceswap-GAN, 2018, [Online; accessed 18-Feb-2021].
4. L. Li, J. Bao, H. Yang, D. Chen, and F. Wen, "Faceshifter: Towards high fidelity and occlusion aware face swapping," in *Conference on Computer Vision and Pattern Recognition*, 2019.
5. J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner, "Face2face: Real-time face capture and reenactment of rgb videos," in *Computer Vision and Pattern Recognition*, 2016.
6. J. Thies, M. Zollhöfer, and M. Nießner, "Deferred neural rendering: Image synthesis using neural textures," *ACM Transactions on Graphics*, vol. 38, no. 4, 2019.
7. Y. Nirkin, Y. Keller, and T. Hassner, "FSGAN: Subject agnostic face swapping and reenactment," in *International Conference on Computer Vision*, 2019, pp. 7184–7193.
8. Z. Shu, M. Sahasrabudhe, R. A. Guler, D. Samaras, N. Paragios, and I. Kokkinos, "Deforming autoencoders: Unsupervised disentangling of shape and appearance," in *European Conference on Computer Vision*, September 2018.
9. M. O. Learning, "Tackling the misinformation epidemic with "in event of moon disaster"," https://news.mit.edu/2020/mit-tackles-misinformation-in-event-of-moon-disaster-0720, 2020, [Online; accessed 18-Feb-2021].
10. H. Kim, P. Garrido, A. Tewari, W. Xu, J. Thies, M. Niessner, P. Pérez, C. Richardt, M. Zollhöfer, and C. Theobalt, "Deep video portraits," *ACM Transactions on Graphics*, vol. 37, no. 4, 2018.
11. E. Zakharov, A. Shysheya, E. Burkov, and V. Lempitsky, "Few-shot adversarial learning of realistic neural talking head models," in *International Conference on Computer Vision*, 2019.
12. "Faceapp - face editor, makeover and beauty app," https://www.faceapp.com/, 2017, [Online; accessed 18-Feb-2021].
13. "Reface," https://hey.reface.ai/, 2020, [Online; accessed 18-Feb-2021].
14. "Zao app deepfake," https://zaodownload.com/, 2019, [Online; accessed 20-July-2021].
15. J. Brandon, "Terrifying high-tech porn: Creepy 'deepfake' videos are on the rise," https://www.foxnews.com/tech/terrifying-high-tech-porn-creepy-deepfake-videos-are-on-the-rise, 2018, [Online; accessed 18-Feb-2021].
16. J. Damiani, "A voice deepfake was used to scam a ceo out of \$243,000," https://www.forbes.com/sites/jessedamiani/2019/09/03/a-voice-deepfake-was-used-to-scam-a-ceo-out-of-243000/, 2019, [Online; accessed 18-Feb-2021].
17. S. Samuel, "A guy made a deepfake app to turn photos of women into nudes. it didn't go well." https://www.vox.com/2019/6/27/18761639/ai-deepfake-deepnude-app-nude-women-porn, 2019, [Online; accessed 18-Feb-2021].

18. Y. Zhu, Q. Li, J. Wang, C.-Z. Xu, and Z. Sun, "One shot face swapping on megapixels," in *Conference on Computer Vision and Pattern Recognition*, June 2021, pp. 4834–4844.
19. 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.
20. T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4401–4410.
21. I. Perov, D. Gao, N. Chervoniy, K. Liu, S. Marangonda, C. Umé, M. Dpfks, C. S. Facenheim, L. RP, J. Jiang, S. Zhang, P. Wu, B. Zhou, and W. Zhang, "Deepfacelab: A simple, flexible and extensible face swapping framework," *Arxiv Pre-print Arxiv:2005.05535*, 2020.
22. T.-N. Le, H. H. Nguyen, J. Yamagishi, and I. Echizen, "Openforensics: Large-scale challenging dataset for multi-face forgery detection and segmentation in-the-wild," in *International Conference on Computer Vision*, 2021.
23. I. Korshunova, W. Shi, J. Dambre, and L. Theis, "Fast face-swap using convolutional neural networks," in *International Conference on Computer Vision*, 2017, pp. 3677–3685.
24. Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "Stargan: Unified generative adversarial networks for multi-domain image-to-image translation," in *Conference on Computer Vision and Pattern Recognition*, 2018.
25. T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of StyleGAN," in *Conference on Computer Vision and Pattern Recognition*, 2020.
26. W. Wu, Y. Zhang, C. Li, C. Qian, and C. C. Loy, "Reenactgan: Learning to reenact faces via boundary transfer," in *European Conference on Computer Vision*, 2018.
27. A. Pumarola, A. Agudo, A. Martinez, A. Sanfeliu, and F. Moreno-Noguer, "Ganimation: One-shot anatomically consistent facial animation," *International Journal of Computer Vision*, 2019.
28. Y. Choi, Y. Uh, J. Yoo, and J.-W. Ha, "Stargan v2: Diverse image synthesis for multiple domains," in *Conference on Computer Vision and Pattern Recognition*, 2020.
29. R. Natsume, T. Yatagawa, and S. Morishima, "Rsgan: face swapping and editing using face and hair representation in latent spaces," in *SIGGRAPH*, 2018.
30. Y. Shen, J. Gu, X. Tang, and B. Zhou, "Interpreting the latent space of gans for semantic face editing," in *Conference on Computer Vision and Pattern Recognition*, 2020.
31. O. Wiles, A. Koepke, and A. Zisserman, "X2face: A network for controlling face generation by using images, audio, and pose codes," in *European Conference on Computer Vision*, 2018.
32. Y. Zhang, S. Zhang, Y. He, C. Li, C. C. Loy, and Z. Liu, "One-shot face reenactment," in *British Machine Vision Conference*, 2019.
33. S. Pidhorskyi, D. A. Adjeroh, and G. Doretto, "Adversarial latent autoencoders," in *Conference on Computer Vision and Pattern Recognition*, 2020.
34. R. Natsume, T. Yatagawa, and S. Morishima, "Fsnet: An identity-aware generative model for image-based face swapping," in *Asian Conference on Computer Vision*, 2018.
35. S. Tripathy, J. Kannala, and E. Rahtu, "Icface: Interpretable and controllable face reenactment using gans," in *Winter Conference on Applications of Computer Vision*, 2020.
36. H. Kim, P. Garrido, A. Tewari, W. Xu, J. Thies, M. Niessner, P. Pérez, C. Richardt, M. Zollhöfer, and C. Theobalt, "Deep video portraits," *ACM Trans. Graph.*, vol. 37, no. 4, 2018.
37. H. Hao, S. Baireddy, A. R. Reibman, and E. J. Delp, "Far-gan for one-shot face reenactment," in *Conference on Computer Vision and Pattern Recognition Workshops*, 2020.
38. S. Ha, M. Kersner, B. Kim, S. Seo, and D. Kim, "Marionette: Few-shot face reenactment preserving identity of unseen targets," in *AAAI Conference on Artificial Intelligence*, 2020.

39. T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *Conference on Computer Vision and Pattern Recognition*, 2019.

40. O. Fried, A. Tewari, M. Zollhöfer, A. Finkelstein, E. Shechtman, D. B. Goldman, K. Genova, Z. Jin, C. Theobalt, and M. Agrawala, "Text-based editing of talking-head video," *ACM Transactions on Graphics*, vol. 38, no. 4, 2019.

41. J. Bao, D. Chen, F. Wen, H. Li, and G. Hua, "Towards open-set identity preserving face synthesis," in *Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6713–6722.

42. M. Maximov, I. Elezi, and L. Leal-Taixe, "Ciagan: Conditional identity anonymization generative adversarial networks," in *Conference on Computer Vision and Pattern Recognition*, June 2020.

43. C.-H. Lee, Z. Liu, L. Wu, and P. Luo, "Maskgan: Towards diverse and interactive facial image manipulation," in *Conference on Computer Vision and Pattern Recognition*, 2020.

44. P. Zhu, R. Abdal, Y. Qin, and P. Wonka, "Sean: Image synthesis with semantic region-adaptive normalization," in *Conference on Computer Vision and Pattern Recognition*, June 2020.

45. W. Jiang, S. Liu, C. Gao, J. Cao, R. He, J. Feng, and S. Yan, "Psgan: Pose and expression robust spatial-aware gan for customizable makeup transfer," in *Conference on Computer Vision and Pattern Recognition*, June 2020.

46. T. Nguyen, A. Tran, and M. Hoai, "Lipstick ain't enough: Beyond color matching for in-the-wild makeup transfer," in *Conference on Computer Vision and Pattern Recognition*, 2021.

47. T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, "High-resolution image synthesis and semantic manipulation with conditional gans," in *Conference on Computer Vision and Pattern Recognition*, 2018.

48. 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 *Conference on Computer Vision and Pattern Recognition*, June 2020.

49. P. Zhou, X. Han, V. I. Morariu, and L. S. Davis, "Two-stream neural networks for tampered face detection," in *Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 1831–1839.

50. Y. Li, M. Chang, and S. Lyu, "In ictu oculi: Exposing ai created fake videos by detecting eye blinking," in *International Workshop on Information Forensics and Security*, 2018, pp. 1–7.

51. J. Li, H. Xie, J. Li, Z. Wang, and Y. Zhang, "Frequency-aware discriminative feature learning supervised by single-center loss for face forgery detection," in *Conference on Computer Vision and Pattern Recognition*, June 2021, pp. 6458–6467.

52. H. H. Nguyen, F. Fang, J. Yamagishi, and I. Echizen, "Multi-task learning for detecting and segmenting manipulated facial images and videos," in *International Conference on Biometrics: Theory, Applications and Systems*, 2019.

53. T. Zhou, W. Wang, Z. Liang, and J. Shen, "Face forensics in the wild," in *Conference on Computer Vision and Pattern Recognition*, June 2021, pp. 5778–5788.

54. D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen, "Mesonet: a compact facial video forgery detection network," in *International Workshop on Information Forensics and Security*, 2018, pp. 1–7.

55. F. Matern, C. Riess, and M. Stamminger, "Exploiting visual artifacts to expose deepfakes and face manipulations," in *Winter Applications of Computer Vision Workshops*, 2019, pp. 83–92.

56. 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 *Conference on Computer Vision and Pattern Recognition*, 2021, pp. 772–781.

57. H. Dang, F. Liu, J. Stehouwer, X. Liu, and A. K. Jain, "On the detection of digital face manipulation," in *Conference on Computer Vision and Pattern recognition*, 2020, pp. 5781–5790.

58. 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.

59. Y. Li and S. Lyu, "Exposing deepfake videos by detecting face warping artifacts," in *Conference on Computer Vision and Pattern Recognition Workshops*, 2019.

60. A. Rössler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Niessner, "Faceforensics++: Learning to detect manipulated facial images," in *International Conference on Computer Vision*, Oct 2019, pp. 1–11.

61. X. Yang, Y. Li, and S. Lyu, "Exposing deep fakes using inconsistent head poses," in *International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 8261–8265.

62. B. Zi, M. Chang, J. Chen, X. Ma, and Y.-G. Jiang, "Wilddeepfake: A challenging real-world dataset for deepfake detection," in *International Conference on Multimedia*, 2020, p. 2382–2390.

63. R. Wang, F. Juefei-Xu, L. Ma, X. Xie, Y. Huang, J. Wang, and Y. Liu, "Fakespotter: A simple yet robust baseline for spotting ai-synthesized fake faces," in *International Joint Conference on Artificial Intelligence*, 2020, pp. 3444–3451.

64. A. Haliassos, K. Vougioukas, S. Petridis, and M. Pantic, "Lips don't lie: A generalisable and robust approach to face forgery detection," in *Conference on Computer Vision and Pattern Recognition*, 2021, pp. 5039–5049.

65. C. Wang and W. Deng, "Representative forgery mining for fake face detection," in *Conference on Computer Vision and Pattern Recognition*, June 2021, pp. 14 923–14 932.

66. X. Zhu, H. Wang, H. Fei, Z. Lei, and S. Z. Li, "Face forgery detection by 3d decomposition," in *Conference on Computer Vision and Pattern Recognition*, June 2021, pp. 2929–2939.

67. Z. Liu, X. Qi, and P. H. Torr, "Global texture enhancement for fake face detection in the wild," in *Conference on Computer Vision and Pattern Recognition*, June 2020.

68. N. Carlini and H. Farid, "Evading deepfake-image detectors with white-and black-box attacks," in *Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 658–659.

69. S. Hussain, P. Neekhara, M. Jere, F. Koushanfar, and J. McAuley, "Adversarial deepfakes: Evaluating vulnerability of deepfake detectors to adversarial examples," in *Winter Conference on Applications of Computer Vision*, 2021, pp. 3348–3357.

70. R. Huang, F. Fang, H. H. Nguyen, J. Yamagishi, and I. Echizen, "Security of facial forensics models against adversarial attacks," in *International Conference on Image Processing*, 2020, pp. 2236–2240.

71. D. Li, W. Wang, H. Fan, and J. Dong, "Exploring adversarial fake images on face manifold," in *Conference on Computer Vision and Pattern Recognition*, 2021, pp. 5789–5798.

72. B. Dolhansky, J. Bitton, B. Pflaum, J. Lu, R. Howes, M. Wang, and C. C. Ferrer, "The deepfake detection challenge dataset," *arXiv preprint arXiv:2006.07397*, 2020.

73. P. Neekhara, B. Dolhansky, J. Bitton, and C. C. Ferrer, "Adversarial threats to deepfake detection: A practical perspective," in *Conference on Computer Vision and Pattern Recognition Workshops*, 2021, pp. 923–932.

74. M. Treu, T.-N. Le, H. H. Nguyen, J. Yamagishi, and I. Echizen, "Fashion-guided adversarial attack on person segmentation," in *Conference on Computer Vision and Pattern Recognition Workshops*, 2021.

75. M. Naseer, S. Khan, M. Hayat, F. S. Khan, and F. Porikli, "A self-supervised approach for adversarial robustness," in *Conference on Computer Vision and Pattern Recognition*, 2020, pp. 262–271.

76. H. Salman, M. Sun, G. Yang, A. Kapoor, and J. Z. Kolter, "Denoised smoothing: A provable defense for pretrained classifiers," in *Conference on Neural Information Processing Systems*, 2020.

77. R. Theagarajan and B. Bhanu, "Defending black box facial recognition classifiers against adversarial attacks," in *Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 3537–3547.

78. A. Goel, A. Agarwal, M. Vatsa, R. Singh, and N. K. Ratha, "Dndnet: Reconfiguring cnn for adversarial robustness," in *Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 103–110.

79. N. Moran, D. Schmidt, Y. Zhong, and P. Coady, "Noisier2noise: Learning to denoise from unpaired noisy data," in *Conference on Computer Vision and Pattern Recognition*, June 2020.

80. P. Korshunov and S. Marcel, "Deepfakes: a new threat to face recognition? assessment and detection," *arXiv preprint arXiv:1812.08685*, 2018.

81. Z. Sun, Y. Han, Z. Hua, N. Ruan, and W. Jia, "Improving the efficiency and robustness of deepfakes detection through precise geometric features," in *Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3609–3618.

82. X. Chen, C. Xie, M. Tan, L. Zhang, C.-J. Hsieh, and B. Gong, "Robust and accurate object detection via adversarial learning," in *Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16 622–16 631.

83. 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 *Conference on Computer Vision and Pattern Recognition*, June 2020.

84. J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *International Conference on Computer Vision*, 2017.

85. J. Bitton and Z. Papakipos, "Augly: A data augmentations library for audio, image, text, and video." https://github.com/facebookresearch/AugLy, 2021.

86. N. Dufour, A. Gully, P. Karlsson, A. V. Vorbyov, T. Leung, J. Childs, and C. Bregler, "Deepfakes detection dataset by google & jigsaw," 2019.

87. Y. Li, X. Yang, P. Sun, H. Qi, and S. Lyu, "Celeb-df: A large-scale challenging dataset for deepfake forensics," in *Conference on Computer Vision and Pattern Recognition*, June 2020.

88. C. Sanderson and B. C. Lovell, "Multi-region probabilistic histograms for robust and scalable identity inference," in *International Conference on Advances in Biometrics*, 2009, p. 199–208.

89. D. Huang and F. De la Torre, "Facial action transfer with personalized bilinear regression," in *European Conference on Computer Vision*, 2012.

90. K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.

91. H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *International Conference on Learning Representations*, 2018.

92. L. Jiang, Z. Guo, W. Wu, Z. Liu, Z. Liu, C. C. Loy, S. Yang, Y. Xiong, W. Xia, B. Chen, P. Zhuang, S. Li, S. Chen, T. Yao, S. Ding, J. Li, F. Huang, L. Cao, R. Ji, C. Lu, and G. Tan, "DeeperForensics Challenge 2020 on real-world face forgery detection: Methods and results," *arXiv preprint arXiv:2102.09471*, 2021.

93. J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, "Retinaface: Single-shot multi-level face localisation in the wild," in *Conference on Computer Vision and Pattern Recognition*, June 2020.

94. J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6299–6308.

95. K. Hara, H. Kataoka, and Y. Satoh, "Learning spatio-temporal features with 3d residual networks for action recognition," in *International Conference on Computer Vision Workshops*, 2017, pp. 3154–3160.

96. D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, "A closer look at spatiotemporal convolutions for action recognition," in *Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6450–6459.

97. K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *International Conference on Computer Vision*, 2017, pp. 2980–2988.

98. Z. Huang, L. Huang, Y. Gong, C. Huang, and X. Wang, "Mask Scoring R-CNN," in *Conference on Computer Vision and Pattern Recognition*, 2019.

99. C.-Y. Fu, M. Shvets, and A. C. Berg, "RetinaMask: Learning to predict masks improves state-of-the-art single-shot detection for free," in *arXiv preprint arXiv:1901.03353*, 2019.

100. D. Bolya, C. Zhou, F. Xiao, and Y. J. Lee, "Yolact: Real-time instance segmentation," in *International Conference on Computer Vision*, 2019.

101. D. Bolya, C. Zhou, F. Xiao, and Y. J. Lee, "Yolact++: Better real-time instance segmentation," *Transactions on Pattern Analysis and Machine Intelligence*, 2020.

102. Y. Lee and J. Park, "Centermask: Real-time anchor-free instance segmentation," in *Conference on Computer Vision and Pattern Recognition*, 2020.

103. H. Chen, K. Sun, Z. Tian, C. Shen, Y. Huang, and Y. Yan, "Blendmask: Top-down meets bottom-up for instance segmentation," in *Conference on Computer Vision and Pattern Recognition*, 2020.

104. E. Xie, P. Sun, X. Song, W. Wang, X. Liu, D. Liang, C. Shen, and P. Luo, "Polarmask: Single shot instance segmentation with polar representation," in *Conference on Computer Vision and Pattern Recognition*, 2020.

105. R. Zhang, Z. Tian, C. Shen, M. You, and Y. Yan, "Mask encoding for single shot instance segmentation," in *Conference on Computer Vision and Pattern Recognition*, 2020.

106. Z. Tian, C. Shen, and H. Chen, "Conditional convolutions for instance segmentation," in *European Conference on Computer Vision*, 2020.

107. J. Rauber, W. Brendel, and M. Bethge, "Foolbox: A python toolbox to benchmark the robustness of machine learning models," in *Reliable Machine Learning in the Wild Workshop, 34th International Conference on Machine Learning*, 2017. [Online]. Available: http://arxiv.org/abs/1707.04131

108. I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *International Conference on Learning Representations*, 2015.

109. A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," in *International Conference on Learning Representations Workshop*, 2016.

110. A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *International Conference on Learning Representations*, 2018.

111. C. Kruengkrai, J. Yamagishi, and X. Wang, "A multi-level attention model for evidence-based fact checking," in *ACL*, 2021.

112. N. Vo and K. Lee, "Where are the facts? searching for fact-checked information to alleviate the spread of fake news," *arXiv preprint arXiv:2010.03159*, 2020.

113. M. Douze, G. Tolias, E. Pizzi, Z. Papakipos, L. Chanussot, F. Radenovic, T. Jenicek, M. Maximov, L. Leal-Taixé, I. Elezi *et al.*, "The 2021 image similarity dataset and challenge," *arXiv preprint arXiv:2106.09672*, 2021.

114. Y. Ke, R. Sukthankar, L. Huston, Y. Ke, and R. Sukthankar, "Efficient near-duplicate detection and sub-image retrieval," in *ACM multimedia*, vol. 4, no. 1, 2004, p. 5.

115. S. Heller, L. Rossetto, and H. Schuldt, "The ps-battles dataset-an image collection for image manipulation detection," *arXiv preprint arXiv:1804.04866*, 2018.

Appendix

This appendix lists links to the deepfake datasets and widely used methods described in this chapter.

Name	Link
OpenForensics [22]	https://sites.google.com/view/ltnghia/research/openforensics
FaceForensics++ [60]	https://github.com/ondyari/FaceForensics
FF++ Leaderboard	http://kaldir.vc.in.tum.de/faceforensics_benchmark
DFDC [72]	https://ai.facebook.com/datasets/dfdc
DeeperForensics-1.0 [83]	https://github.com/EndlessSora/DeeperForensics-1.0
FFIW [53]	https://github.com/tfzhou/FFIW
XceptionNet [60]	https://github.com/ondyari/FaceForensics
EfficientNet-B4 [72]	https://github.com/lukemelas/EfficientNet-PyTorch
CapsuleNet-R50 [58]	https://github.com/nii-yamagishilab/Capsule-Forensics-v2
CycleGAN [84]	https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix
FoolBox [107]	https://github.com/bethgelab/foolbox

This chapter aims to introduce general knowledge about deepfake for beginners and/or students. A deeper understanding of the contents can be obtained by a review of the following materials for beginners. They should be helpful in obtaining basic knowledge about computer vision and biometrics, which is necessary for understanding deepfake.

Name	Information
CVPR 2018 Tutorial on GANs	https://sites.google.com/view/cvpr2018tutorialongans
DCGAN Tutorial - Pytorch	https://pytorch.org/tutorials/beginner/dcgan_faces_tutorial.html
DCGAN Tutorial - Tensorflow	https://www.tensorflow.org/hub/tutorials/tf_hub_generative_image_module
DeepFaceLab Tutorial	https://medium.com/geekculture/creating-deepfake-miracles-with-deepfacelab-tutorial-saehd-model-aa2aa12c08f3
CVPR 2020 Workshop on Media Forensics	https://sites.google.com/view/wmediaforensics2020
Handbook of Biometrics	Jain, Anil K., Flynn, Patrick, Ross, Arun A., "Handbook of Biometrics", Springer, 2008.