

Jing Han
ZD.B Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Augsburg, GERMANY

Zixing Zhang
Group on Language, Audio & Music, Imperial College London, London, UK

Nicholas Cummins
ZD.B Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Augsburg, GERMANY

Björn Schuller
Group on Language, Audio & Music, Imperial College London, London, UK; ZD.B Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Augsburg, GERMANY

Adversarial Training in Affective Computing and Sentiment Analysis: Recent Advances and Perspectives

Abstract

Over the past few years, adversarial training has become an extremely active research topic and has been successfully applied to various Artificial Intelligence (AI) domains. As a potentially crucial technique for the development of the next generation of emotional AI systems, we herein provide a comprehensive overview of the application of adversarial training to affective computing and sentiment analysis. Various representative adversarial training algorithms are explained and discussed accordingly, aimed at tackling diverse challenges associated with emotional AI systems. Further, we highlight a range of potential future research directions. We expect that this overview will help facilitate the development of adversarial training for affective computing and sentiment analysis in both the academic and industrial communities.

1. Introduction

Affective computing and sentiment analysis currently play a vital role in transforming current Artificial Intelligent (AI) systems into the next genera-

tion of emotional AI devices [1], [2]. It is a highly interdisciplinary research field spanning psychology, cognitive, and computer science. Its motivations include endowing machines with the ability to detect and understand the emotional states of humans and, in turn, respond accordingly [1]. Both the terms affective computing and sentiment analysis relate to the computational interpretation and generation of human emotion or affect. Whereas the former mainly relates to instantaneous emotional expressions and is more commonly associated with speech or image/video processing, the later mainly relates to longer-term opinions or attitudes and is more commonly associated with natural language processing.

A plethora of applications can benefit from the development of affective computing and sentiment analysis [3]–[8]; examples include natural and friendly human–machine interaction systems, intelligent business and customer service systems, and remote healthcare systems. Thus, affective computing and sentiment analysis attract considerable research attention in both the academic and industrial communities.

From a technical point of view, affective computing and sentiment analysis are associated with a wide range of advancements in machine learning, especially in relation to deep learning technologies. For example, deep Convolutional Neural Networks (CNNs) have been reported to considerably outperform conventional models and non-deep neural networks on two benchmark databases for sentiment analysis [9]. Further, an end-to-end deep learning framework which automatically learns high-level representations from raw audio and video signals has been shown to be effective for emotion recognition [10].

However, when deployed in real-life applications, affective computing and sentiment analysis systems face many challenges. These include the sparsity and unbalance problems of the training data [11], the instability of the emotion recognition models [12], [13], and the poor quality of the generated emotional samples [14], [15]. Despite promising research efforts and advances in leveraging techniques, such as semi-supervised learning and transfer learning [11], finding robust solutions to these challenges is an open and ongoing research challenge.

In 2014, a novel learning algorithm called adversarial training (or adversarial

learning) was proposed by Goodfellow et al. [16], and has attracted widespread research interests across a range of machine learning domains [17], [18], including affective computing and sentiment analysis [19]–[21]. The initial adversarial training framework, Generative Adversarial Networks (GANs), consists of two neural networks—a generator and a discriminator, which contest with each other in a two-player zero-sum game. The generator aims to capture the potential distribution of real samples and generates new samples to ‘cheat’ the discriminator as far as possible, whereas the discriminator, often a binary classifier, distinguishes the sources (i.e., real samples or generated samples) of the inputs as accurately as possible. Since its inception, adversarial training has been frequently demonstrated to be effective in improving the robustness of recognition models and the quality of the simulated samples [16]–[18].

Thus, adversarial training is emerging as an efficient tool to help overcome the aforementioned challenges when building affective computing and sentiment analysis systems. More specifically, on the one hand, GANs have the potential to produce an unlimited amount of realistic emotional samples; on the other hand, various GAN variants have been proposed to learn robust high-level representations. Both of the aspects can improve the performance of emotion recognition systems. Accordingly, over the past three years, the number of related papers has grown exponentially. Motivated by the pronounced improvement achieved by these works and by the belief that adversarial training can further advance more works in the community, we thus feel that, there is a necessity to summarize recent studies, and draw attention to the emerging research trends and directions of adversarial training in affective computing and sentiment analysis.

A plethora of surveys can be found in the relevant literature either focusing on conventional approaches or (non-adversarial) deep-learning approaches for both affective computing [11], [22]–[25] and sentiment analysis [6], [26]–[32], or offering more generic overviews of generative adversarial networks [17],

Adversarial training is emerging as an efficient tool to help overcome the challenges when building affective computing and sentiment analysis systems.

[18]. Differing from these surveys, the present article:

- provides, for the first time, a comprehensive overview of the adversarial training techniques developed for affective computing and sentiment analysis applications;
- summarizes the adversarial training technologies suitable, not only for the emotion recognition and understanding tasks, but more importantly, for the emotion synthesis and conversion tasks, which are arguably far from being regarded as mature;
- reviews a wide array of adversarial training technologies covering the text, speech, image and video modalities;
- highlights an abundance of future research directions for the application of adversarial training in affective computing and sentiment analysis.

The remainder of this article is organized as follows. In Section II, we first introduce the background of this overview, which is then followed by a short description of adversarial training in Section III. We then comprehensively summarize the representative adversarial training approaches for emotion synthesis in Section IV, the approaches for emotion conversion in Section V, and the approaches for emotion perception and understanding in Section VI, respectively. We further highlight some promising research trends in Section VII, before drawing the conclusions in Section VIII.

II. Background

In this section, we first briefly describe three of the main challenges associated with affective computing and sentiment analysis, i.e., the naturalness of generated emotions, the sparsity of collected data, and the robustness of trained models. Concurrently, we analyze the drawbacks and limitations of conventional deep learning approaches, and introduce opportunities for the application of adversarial training. Then, we give a

short discussion about the challenge of performance evaluation when generating or converting emotional data.

A typical emotional AI framework consists of two core components: an emotion perception and understanding unit, and an emotion synthesis and conversion unit (cf. Figure 1). The first component (aka a recognition model) interprets human emotions; whereas the second component (aka a generation model) can generate emotionally nuanced linguistics cues, speech, facial expressions, and even gestures. For the remainder of this article, the term emotion synthesis refers to the artificial generation of an emotional entity from scratch, whereas the term emotion conversion refers to the transformation of an entity from one emotional depiction to another. To build a robust and stable emotional AI system, several challenges have to be overcome as discussed in the following sub-sections.

A. Naturalness of Generated Emotions

Emotion synthesis and conversion go beyond the conventional constructs of Natural Language Generation (NLG), Text-To-Speech (TTS), and image/video transformation techniques. This is due in part to the instinct complexity and uncertainty of the emotions, thereby generating an emotional entity remains an ongoing challenge.

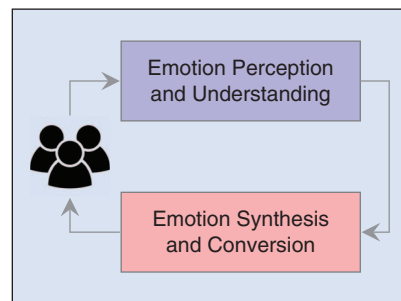


FIGURE 1 The broad framework of a typical emotional artificial intelligence system.

A “classical” GAN consists of a generator and a discriminator as two players in a zero-sum game.

Recently, research has shown the potential of deep-learning based generative models for addressing this challenge. For example, the WaveNet network developed by Oord et al. [33] efficiently synthesizes speech signals, and Pixel Recurrent Neural Networks (Pixel-RNN) and Variational AutoEncoder (VAE), proposed by Oord et al. [34] and Kingma et al. [35] respectively, have been shown to be effective for generating images.

To date, the majority of these studies have not considered emotional information. A small handful of works have been undertaken in this direction [36], [37], however, the generated emotions are far from being considered natural. This is due in part to the highly non-linear nature of emotional expression changes and the variance of individuals [38], [39]. Generative modeling with adversarial training, on the other hand, has frequently been shown to be powerful in regard to generating samples, which are more understandable to humans than the examples simulated by other approaches [17], [18] (see Section IV and Section V for more details).

B. Sparsity of Collected Data

Despite having the possibility to collect massive amounts of unlabeled data through pervasive smart devices and social media, reliably annotated data resources required for emotion analysis are still comparatively scarce [11]. For example, most of the databases currently available for speech emotion recognition contain, at most 10 hours of labeled data [25], [40], which is insufficient for building highly robust models. This issue has become even more pressing in the era of deep learning. The data-sparsity problem mainly lies in the annotation process which is prohibitively expensive and time-consuming [11]. This is especially true in relation to the subjective nature of emotions which dictates the need for several annotators to label the same sam-

ples in order to diminish the effect of personal biases [41].

In tackling this challenge, Kim et al. [42] proposed an unsupervised learning approach to learn the representations across audiovisual modalities for emotion recognition without any labeled data. Similarly, Cummins et al. [43] utilized CNNs pre-trained on large amounts of image data to extract robust feature representations for speech-based emotion recognition. More recently, neural-network-based semi-supervised learning has been introduced to leverage large-scale unlabeled data [13], [44].

Despite the effectiveness of such approaches that distil shared high-level representations between labeled and unlabeled samples, the limited number of labeled data samples means that there is a lack of sufficient resources to extract meaningful and salient representations specific to emotions. In contrast, a generative model with adversarial training has the potential to synthesize an infinite amount of labeled samples to overcome the shortage of conventional deep learning approaches (see Section VI for more details).

C. Robustness of Trained Models

In many scenarios, samples from a target domain are not sufficient or reliable enough to train a robust emotion recognition model. This challenge has motivated researchers to explore transfer learning solutions which leverage related domain (source) samples to aid the target emotion recognition task. This is a highly non-trivial task, the source and target domains are often highly mismatched with respect to the domains in which the data are collected [45], such as different recording environments or websites. For example, in sentiment analysis, the word ‘long’ for evaluating battery life has a positive connotation, whereas when assessing pain it tends to be negative. Moreover, for speech emotion recognition, the source and target samples might have been recorded in distinctive acous-

tic environments and by different speakers [11]. These mismatches have been shown to lead to a performance degradation of models analyzed in real-life settings [27], [45], [46].

In addressing this challenge, Glorot et al. [46] presented a deep neural network based approach to learn the robust representations across different domains for sentiment analysis. Similar approaches have also been proposed by Deng et al. [47] for emotion recognition from speech. Moreover, You et al. [48] successfully transferred the sentiment knowledge from text to predict the sentiment of images. However, it is still unclear if their learned representations are truly domain-invariant or not.

On the other hand, the discriminator of an adversarial training framework has the potential to distinguish from which domain the so-called ‘shared’ representations come from. By doing so, it can help alleviate the robustness problem of an emotion recognition model (see Section VI for more details).

D. Performance Evaluation

Evaluating the performance of the generated or converted emotional samples is essential but challenging in aspects of affective computing and sentiment analysis. Currently, many of the related works directly demonstrate a few appealing samples and evaluate the performance by human judgment [38], [49], [50]. Additionally, a range of metric-based approaches has been proposed to quantitatively evaluate the adversarial training frameworks. For example, the authors in [51] compared the intra-set and inter-set average Euclidean distances between different sets of the generated faces.

Similarly, to quantitatively evaluate models for emotion conversion, other evaluation measurements raised in the literature include BiLingual Evaluation Understudy (BLEU) [52] and Recall-Oriented Understudy for Gisting Evaluation (ROUGE) [53] for text, and a signal-to-noise ratio test for speech [15]. However, the quantitative performance evaluation for emotion perception and understanding is more straightforward. In general, the improvement by implementing adversarial training can be

reported using evaluation metrics such as unweighted accuracy, unweighted average recall, and concordance correlation coefficient [12], [54], [55].

III. Principle of Adversarial Training

In this section, we introduce the basic concepts of adversarial training, so that the interested reader can better understand the design and selection of adversarial networks for a specific task in affective computing and sentiment analysis.

A. Terminology and Notation

With the aim of generating realistic ‘fake’ samples from a complex and high-dimensional true data distribution, the ‘classical’ GAN, consists of two deep neural nets (as two players in a game): a generator (denoted as G) and a discriminator (denoted as D) (cf. Figure 2). During this two-player game, the generator tries to turn input noises from a simple distribution into realistic samples to fool the discriminator, while the discriminator tries to distinguish between true (or ‘real’) and generated (or ‘fake’) data.

Normally, G and D are trained jointly in a minimax fashion. Mathematically, the minimax objective function can be formulated as:

$$\min_{\theta_g} \max_{\theta_d} V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D_{\theta_d}(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D_{\theta_d}(G_{\theta_g}(z)))] \quad (1)$$

where θ_g and θ_d denote the parameters of G and D , respectively; x is a real data instance following the true data distribution $p_{\text{data}}(x)$; whilst z is a vector randomly sampled following a simple distribution (e.g., Gaussian); $G_{\theta_g}(z)$ denotes a generated data given z as the input; and $D_{\theta_d}(\cdot)$ outputs the likelihood of real data given either x or $G_{\theta_g}(z)$ as the input. Note that, the likelihood is in the range of $(0,1)$, indicating to what extent the input is probably a real data instance. Consequently, during training, θ_g is updated to minimize the objective function such that $D_{\theta_d}(G_{\theta_g}(z))$ is close to 1; conversely, θ_d is optimized to maximize the objective such that $D_{\theta_d}(x)$ is close to 1 and $D_{\theta_d}(G_{\theta_g}(z))$ is close to 0. In other words, G and D are trying to optimize a different and oppos-

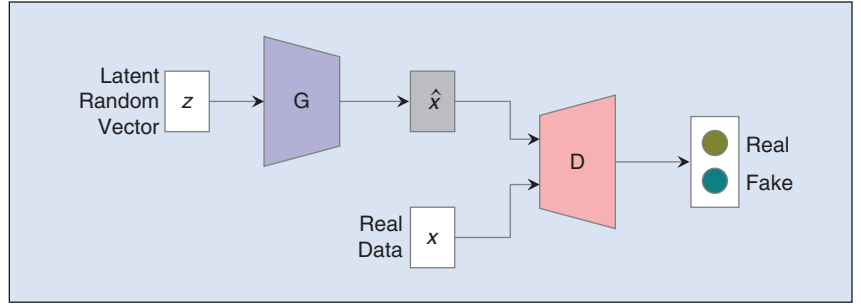


FIGURE 2 Framework of Generative Adversarial Network (GAN).

ing objective function, thus pushing against each other in a zero-sum game. Hence, the strategy is named as adversarial training.

Generally, the training of G and D is done in an iterative manner, i.e., the corresponding neural weights θ_d, θ_g are updated in turns. Once training is completed, the generator G is able to generate more realistic samples, while the discriminator D can distinguish authentic data from fake data. More details of the basic GAN training process can be found in [16].

B. Category of Adversarial Networks

Since the first GAN paradigm was introduced in 2014, numerous variants of the original GAN have been proposed and successfully exploited in many real-life applications. It is roughly estimated that to date, more than 350 variants of GANs have been presented in the literature over the last four years¹, infiltrating into various domains including image, music, speech, and text. For a comprehensive list and other resources of all currently named GANs, interested readers are referred to [56], [57]. Herein, we group these variants into four main categories: optimization-based, structure-based, network-type-based, and task-oriented.

Optimization-based: GANs in this category aim to optimize the minimax objective function to improve the stability and the speed of the adversarial training process. For instance, in the original GAN, the Jensen-Shannon (JS) divergence of the objective function can be a constant, particularly at the start of

the training procedure where there is no overlap between the sampled real data and the generated data. To smooth the training of GANs, the Wasserstein GAN (WGAN) has been proposed by replacing the JS divergence with the earth-mover distance to evaluate the distribution distance between the real and generated data [58].

Other GAN variants in this direction include the Energy-Based GAN (EBGAN) [59], the Least Squares GAN (LSGAN) [60], the Loss-Sensitive GAN (LS-GAN) [61], the Correlational GAN (CorrGAN) [62], and the Mode Regularized GAN (MDGAN) [63], to name but a few.

Structure-based: these GAN variants have been proposed and developed to improve the structure of conventional GAN. For example, the conditional GAN (cGAN) adds auxiliary information to both the generator and discriminator to control the modes of the data being generated [64], while the semi-supervised cGAN (sc-GAN) exploits the labels of real data to guide the learning procedure [65]. Other GAN variants in this category include the BiGAN [66], the CycleGAN [67], the DiscoGAN [68], the InfoGAN [69], and the Triple-GAN [70].

Network-type-based: in addition, several GAN variants have been named after the network topology used in the GAN configuration, such as the DCGAN based on deep convolutional neural networks [19], the AEGAN based on autoencoders [71], the C-RNNGAN based on continuous recurrent neural networks [72], the AttnGAN based on attention mechanisms [73], and the CapsuleGAN based on capsule networks [74].

¹<https://github.com/hindupuravinash/the-gan-zoo/blob/master/gans.tsv>

In contrast to other generative models, GANs directly learn to generate new samples through a two-player game without estimating any explicit density function.

Task-oriented: lastly, there are also a large number of GAN variants that have been designed for a given task, thus serve their own specific interests. Examples, to name just a few, include the Sketch-GAN proposed for sketch retrieval [75], the ArtGAN for artwork synthesis [76], the SEGAN for speech enhancement [77], the WaveGAN for raw audio synthesis [78], and the VoiceGAN for voice impersonation [15].

IV. Emotion Synthesis

As discussed in Section II-A, the most promising generative models, for synthesis, currently include PixelRNN/CNN [34], [79], VAE [35], and GANs [16]. Works undertaken with these models highlight their potential for creating realistic emotional samples. The PixelRNN/CNN approach, for example, can explicitly estimate the likelihood of real data with a tractable density function in order to generate realistic samples. However, the generating procedure is quite slow, as it must be processed sequentially. On the other hand, VAE defines an intractable density function and optimizes a lower bound of the likelihood instead, resulting in a faster-generating speed compared with PixelRNN/CNN. However, it suffers from the generation of low-quality samples.

In contrast to other generative models, GANs directly learn to generate new samples through a two-player game without estimating any explicit density function, and have been shown to obtain state-of-the-art performance for a range of tasks notably in image generation [16], [17], [80]. In particular, GAN-based frameworks can help generate, in theory, an infinite amount of realistic emotional data, including samples with subtle changes which depict more nuanced emotional states.

A. Conditional-GAN-Based Approaches in Image/Video

To synthesize emotions, the most frequently used GAN relates to the conditional GAN (cGAN). In the original cGAN framework, both the generator and discriminator are conditioned on certain extra information c . This extra information can be any kind of auxiliary information, such as the labels or data from other modalities [64]. More specifically, the latent input noise z is concatenated with the condition c as a joint hidden representation for the generator G , in the meanwhile c is combined with either the generated sample \hat{x} or real data x to be fed into the discriminator D , as demonstrated in Figure 3. In this circumstance, the minimax

objective function given in Equation (1) is reformulated:

$$\begin{aligned} \min_{\theta_g} \max_{\theta_d} V(D, G) &= \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D_{\theta_d}(x | c)] \\ &+ \mathbb{E}_{z \sim p_z(z)} [\log(1 - D_{\theta_d}(G_{\theta_g}(z | c)))] \end{aligned} \tag{2}$$

Recently, a collection of works have begun to explore which facial expressions and representations can best be produced via cGAN. The frameworks proposed within these works are either conditioned on attribute vectors including emotion states to generate an image for a given identity [49], or conditioned on various emotions represented by values of features such as facial action unit coefficients to produce dynamic video from a static image [81], or conditioned on arbitrary speech clips to create talking faces synchronized with the given audio sequence [82]. While these approaches can produce faces with convincing realism, they do not fully consider the interpersonal behaviors that are common in social interactions such as mimicry.

In tackling this problem, one novel application was proposed in [51], in which the authors presented a cGAN-based framework to generate valid facial expressions for a virtual agent. The proposed framework consists of two stages: firstly, a person’s facial expressions (in eight emotion classes) are applied as conditions to generate expressive face sketches, then, the generated sketches are leveraged as conditions to synthesize complete face images of a virtual dyad partner. However, this framework does not consider the temporal dependency on faces across various frames, can yield non-smooth facial expressions over time. In light of this, researchers in [50] proposed Conditional Long Short-Term Memory networks (C-LSTMs) to synthesize contextually smooth sequences of video frames in dyadic interactions. Experimental results in [50] demonstrate that the facial expressions in the generated virtual faces reflect appropriate emotional reactions to a person’s behaviors.

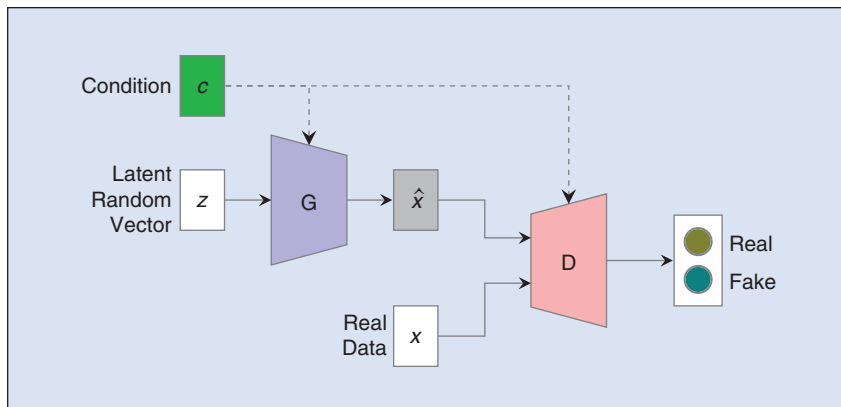


FIGURE 3 Framework of conditional Generative Adversarial Network (cGAN).

B. Other GAN-Based Approaches in Image/Video

In addition to the cGAN-based framework, other GAN variants such as DCGAN [19] and InfoGAN [69] have been investigated for emotional face synthesis. In [19], it is shown that, vector arithmetic operations in the input latent space can yield semantic changes to the image generations. For example, performing vector arithmetic on mean vectors “smiling woman” – “neutral woman” + “neutral man” can create a new image with the visual concept of “smiling man”. The InfoGAN framework, on the other hand, aims to maximize the mutual information between a small subset of the latent variables and the observation, to learn interpretable latent representations which reflect the structured semantic data distribution [69]. For instance, it has been demonstrated that by varying one latent code, the emotions of the generated faces can change from stern to happy [69].

C. Approaches in Other Modalities

As well as the generation of expressive human faces, adversarial training has also been exploited to generate emotional samples in a range of other modalities. For example, in [83] modern artwork images have been automatically generated from an emotion-conditioned GAN. Interestingly, it has been observed that various features, such as colors and shapes, within the artworks are commonly correlated with the emotions which they are conditioned on. Similarly, in [84] plausible motion sequences conditioned by a variety of contextual information (e.g., activity, emotion), have been synthesized by a so-called sequential adversarial autoencoder. More recently, poems conditioned by various sentiment labels (estimated from images) have been created via a multi-adversarial training approach [85].

Correspondingly, adversarial training has also been investigated for both text generation [86], [87] and speech synthesis [78]. In particular, sentence generation conditioned on sentiment (either positive or negative) has been conducted in [14] and [21], but both only on fixed-

Leveraging adversarial training makes the emotion conversion procedure more robust and accurate through the associated game-theoretic approach.

length sequences (11 words in [14] and 40 words in [21]). One example can be the three generated samples with a fixed length (40 words) found in [21] (also shown in Figure 4). Despite the promising nature of these initial works, the performance of such networks is far off when comparing with the quality and naturalness of image generation.

To the best of our knowledge, emotion-integrated synthesis frameworks based on adversarial training has yet to be implemented in speech. Compared with image generation, one main issue we confront in both speech and text is the varied length to generate, which, however, could also be a learned feature in the future.

V. Emotion Conversion

Emotion conversion is a specific style transformation task. In computer vision and speech processing domains, it targets at transforming a source emotion into a target emotion without affecting the identity properties of the subject. Whereas for NLP, sentiment transformation aims to alter the sentiment expressed in the original text while preserving its content. In conventional approaches, paired data are normally required to learn a pairwise transformation function. In this case, the data need to be perfect time aligned to learn an effective model, which is generally achieved by time-warping.

Adversarial training, on the other hand, does away with the need to prepare the paired data as a precondition, as the emotion transformation function can be estimated in an indirect manner. In light of this, adversarial training reshapes conventional emotion conversion procedures and makes the conversion systems simpler to be implemented and used, as time-alignment is not needed. Moreover, leveraging adversarial training makes the emotion conversion procedure more robust and accurate through the associated game-theoretic approach.

A. Paired-Data-Based Approaches

Several adversarial training approaches based on paired training data have been investigated for emotion conversion. For example, in [88], the authors proposed a conditional difference adversarial autoencoder, to learn the difference between the source and target facial expressions of one same person. In this approach, a source face goes through an encoder to generate a latent vector representation, which is then concatenated with the target label to generate the target face through a decoder. Concurrently, two discriminators (trained simultaneously) are used to regularize the latent vector distribution and to help improve the quality of generated faces through an adversarial process.

Moreover, approaches based on facial geometry information have been proposed to guide facial expression conversion

Positive: Follow the Good Earth movie linked Vacation is a comedy that credited against the modern day era yarns which has helpful something to the modern day s best It is an interesting drama based on a story of the famed

Negative: I really can t understand what this movie falls like I was seeing it I m sorry to say that the only reason I watched it was because of the casting of the Emperor I was not expecting anything as

Negative: That s about so much time in time a film that persevered to become cast in a very good way I didn t realize that the book was made during the 70s The story was Manhattan the Allies were to

FIGURE 4 Generated samples on IMDB [21].

By combining seq2seq model with cycleGAN, a chatbot was developed whose response can be transformed from negative to positive.

[38], [89]. In [89], a geometry guided GAN for facial expression transformation was proposed, which is conditioned on facial geometry rather than expression labels. In this way, the facial geometry is directly manipulated, and thus the network ensures a fine-grain control in face editing, which, in general, is not so straightforward in other approaches. In [38], the researchers further disentangled the face encoding and facial geometry (in landmarks) encoding process, which allows the model to perform the facial expression transformations appropriately even for unseen facial expression characteristics.

Another related work is [90], in which the authors focused on voice conversion in natural speech and proposed a variational autoencoding WGAN. Note that, data utilized in [90] are not frame aligned, but still are in pairs. Emotion conversion has not been considered in this work, however, this model could be applied to emotion conversion.

B. Non-Paired-Data-Based Approaches

The methods discussed in the previous section all require pair-wise data of the same subjects in different facial expressions during training. In contrast, Invertible conditional GAN (IcGAN), which consists of a cGAN and two encoders, does not have this constraint [91]. In the

IcGAN framework, the encoders compress a real face image into a latent representation z and a conditional representation c independently. Then, c can be explicitly manipulated to modify the original face with deterministic complex modifications.

Additionally, the ExprGAN framework is a more recent advancement for expression transformation [92], in which the expression intensity can be controlled in a continuous manner from weak to strong. Furthermore, the identity and expression representation learning are disentangled and there is no rigid requirement of paired samples for training [92]. Finally, the authors develop a three-stage incremental learning algorithm to train the model on small datasets [92].

Recently, inspired by the success of the DiscoGAN for style transformation in images, Gao et al. [15] proposed a speech-based style-transfer adversarial training framework, namely VoiceGAN (cf. Figure 5). The VoiceGAN framework consists of two generators/transformers (G_{AB} and G_{BA}) and three discriminators (D_A , D_B , and D_{style}). G_{AB} attempts to transform instances from style A to style B , while G_{BA} does the reverse. D_A and D_B are used to distinguish between the generated data and real data for style A and B , respectively,

and D_{style} is applied to distinguish if all data match the target styles. Importantly, the linguistic information in the speech signals is retained by considering the reconstruction losses of the generated data, and parallel data are not required. To contend with the varied lengths of speech signals, the authors applied a channel-wise pooling to convert variable-sized feature map into a vector of fixed size [15]. Experimental results demonstrate that VoiceGAN is able to transfer the gender of a speaker's voice, and this technique could be easily extended to other stylistic features such as different emotions [15].

More recently, a cycleGAN-based model was proposed to learn sentiment transformation from non-parallel text, with an ultimate goal to automatically adjust the sentiment of a chatbot response [93]. By combining seq2seq model with cycleGAN, the authors developed a chatbot whose response can be transformed from negative to positive.

Compared with the works of adversarial-training-based emotion conversion in the image, it is noticeable that to date, there are only a few related works in video and speech, and only one in text. We believe that the difficulty of applying adversarial training in these domains is threefold: 1) the variable length of corresponding sequential data; 2) the linguistic and language content needed to be maintained during the conversion; and 3) the lack of reliable measurement metrics to rapidly evaluate the performance of such a transformation.

VI. Emotion Perception and Understanding

This section summarizes works which tackle the data-sparsity challenge (see Section VI-A) and the robustness-of-the-emotion-recognizer challenge (see Sections VI-B and VI-C).

A. Data Augmentation

As already discussed, the lack of large amounts of reliable training data is a major issue in the fields of affective computing and sentiment analysis. In this regard, it has been shown that emotion

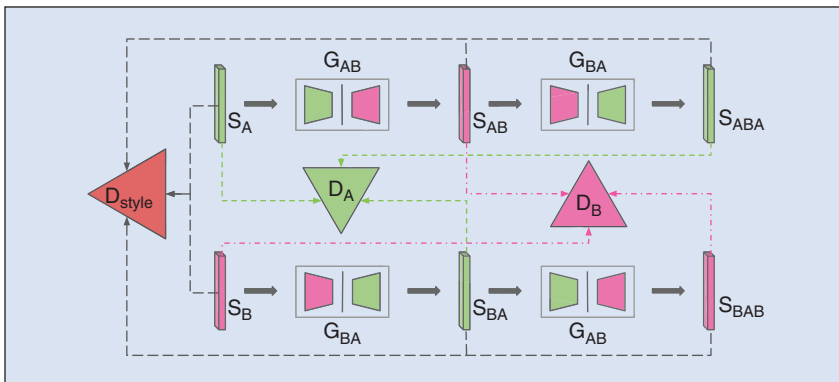


FIGURE 5 Framework of VoiceGAN [15].

recognition performance can be improved with various data augmentation paradigms [94], [95]. Data augmentation is a family of techniques which artificially generate more data to train a more efficient (deep) learning model for a given task.

Conventional data augmentation methods focus on generating data through a series of transformations, such as scaling and rotating an image, or adding noise to speech [95]. However, such perturbations directly on original data are still, to some extent, not efficient to improve overall data distribution estimation. In contrast, as GANs generate realistic data which estimate the distribution of the real data, it is instinctive to apply them to expand the training data required for emotion recognition models. In this regard, some adversarial training based data augmentation frameworks have been proposed in the literature, which aim to supplement the data manifold to approximate the true distribution [96].

For speech emotion recognition, researchers in [97] implemented an adversarial autoencoder model. In this work, high-dimensional feature vectors of real data are encoded into 2-D dimensional representations, and a discriminator is learned to distinguish real 2-D vectors from generated 2-D vectors. The experiments indicate that the 2-D representations of real data can yield suitable margins between different emotion categories. Additionally, when adding the generated data to the original data for training, performance can be marginally increased [97].

Similarly, a cycleGAN has been utilized for face-based emotion recognition [96]. To tackle the data inadequacy and unbalance problems, faces in different emotions have been generated from non-emotion ones, particularly for emotions like disgust and sad, which seemingly have less available samples. Experimental results have demonstrated that, by generating auxiliary data of minority classes for training, not only did the recognition performance of the rare class improve, the average performance over all classes also increased [96].

Ganin et al. first introduced domain adversarial training to tackle the domain shift problem by introducing a gradient reversal layer.

One ongoing research issue relating to GANs is how best to label the generated data. In [97], they adopted a Gaussian mixture model which is built on the original data, whereas the authors in [96] took a set of class-specific GANs to generate images, respectively, which requires no additional annotation process.

In addition to these two approaches, cGAN in a semi-supervised manner (scGAN) can be an interesting alternative worthy of future investigations. The scGAN is an extension of cGAN by forcing the discriminator D to output class labels as well as distinguishing real data from fake data. In this scenario, D acts as both a discriminator and a classifier. More specifically, D classifies the real samples into the first n classes and the generated samples into the $n + 1$ -th class (fake), while G tries to generate the conditioned samples and ‘cheat’ the discriminator to be correctly classified into the first n classes, as illustrated in Figure 6. By taking the class distribution into the objective function, an overall improvement in the quality of the generated samples was observed [65]. Hence, scGAN can be easily adapted for data augmentation in emotion perception and understanding tasks, which to date has yet to be reported in the literature.

Finally, the quality of the generated data is largely overlooked in the works

discussed in this section. It is possible that the generated data might be unreliable, and thus become a form of noise in the training data. In this regard, data filtering approaches should be considered.

B. Domain Adversarial Training

For emotion perception and understanding, numerous domain adaptation approaches have been proposed to date (cf. Section II-C). These approaches seek to extract the most representative features from the mismatched data between the training phase and the test phase, in order to improve the robustness of recognition models (cf. Section II-C). However, it is unclear if the learned representations are truly domain-generative or still domain-specific.

In [98], Ganin et al. first introduced domain adversarial training to tackle this problem. Typically, a feature extractor $G_f(\cdot, \theta_f)$ projects data from two separate domains into high-level representations, which are discriminative for a label predictor $G_y(\cdot, \theta_y)$ and indistinguishable for a domain classifier $G_d(\cdot, \theta_d)$. A typical Domain-Adversarial Neural Network (DANN) is illustrated in Figure 7. Particular to the DANN architecture, a gradient reversal layer is introduced between the domain classifier and the feature extractor, which inverts the sign of the gradient $\partial L_d / \partial \theta_d$ during backward propagation. Moreover, a hyper-parameter

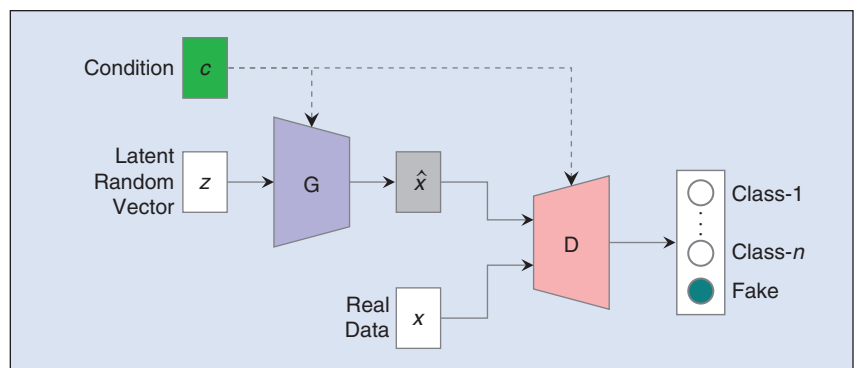


FIGURE 6 Framework of semi-supervised conditional Generative Adversarial Network (scGAN).

Two major open research challenges relating to adversarial training are training instability and mode collapse.

λ is utilized to tune the trade-off between the two branches during the learning process. In this manner, the network attempts to learn domain-invariant feature representations. By this training strategy, the representations learned from different domains cannot be easily distinguished, as demonstrated in Figure 8. Further details on how to train DANN are given in [98].

Using a DANN, a common representation between data from the source and target domains can potentially be learned. This is of relevance in the data mismatch scenario as knowledge learned from the source domain can be applied directly to the target domain [55]. Accordingly, the original DANN paradigm has been adapted to learn domain-

invariant representations for sentiment classification. For example, in [95], [99], attention mechanisms were introduced to give more attention to relevant text when extracting features. In [12], [100], the Wasserstein distance was estimated to guide the optimization of the domain classifier. Moreover, instead of learning common representations between two domains, other research has broadened this concept to tackle the data mismatch issue among multiple probability distributions. In this regard, DANN variants have been proposed for multiple source domain adaptation [101], multi-task learning [102], and multinomial adversarial nets [103].

Finally, DANN has recently been utilized in speech emotion recognition

[55]. These experiments demonstrate that, by aligning the data distributions of the source and target domains (illustrated in Figure 8), an adversarial training approach can yield a large performance gain in the target domain [55].

C. (Virtual) Adversarial Training

Beside factors relating to the quality of the training data, the performance of an emotional AI system is also heavily dependent on its robustness to unseen data. A trivial disturbance on the sample (adversarial examples) might result in an opposite prediction [104], which naturally has to be avoided for a robust recognition model.

Generally speaking, adversarial examples are the examples that are created by making small, but intentionally, perturbations to the input to incur large and significant perturbations in outputs (e.g., incorrect predictions with high confidence) [104]. Adversarial training, however, addresses this vulnerability in recognition models by introducing mechanisms to correctly handle the adversarial examples. In this way, it improves not only robustness to adversarial examples, but also overall generalization for the original examples [104].

Mathematically, adversarial training adds the following term as regularization loss to the original loss function:

$$-\log p(y|x + r_{adv}; \theta), \quad (3)$$

in which x denotes the input, θ denotes the parameters of a classifier, and r_{adv} denotes a worst-case perturbation against the current model $p(y|x; \theta)$, which can be calculated with

$$r_{adv} = \arg \max_{r, \|r\| \leq \epsilon} \log p(y|x + r; \theta). \quad (4)$$

In the context of affective computing and sentiment analysis, the authors in [105] utilized DCGAN and multi-task learning strategies to leverage a large number of unlabeled samples, where the unlabeled samples are considered as adversarial examples. More specifically, the model explores unlabeled data by feeding it through a vanilla

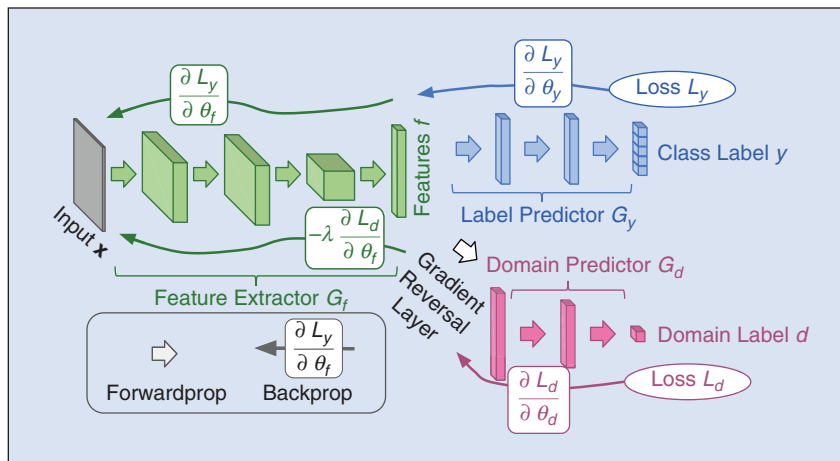


FIGURE 7 Framework of Domain-Adversarial Neural Network (DANN) [98].

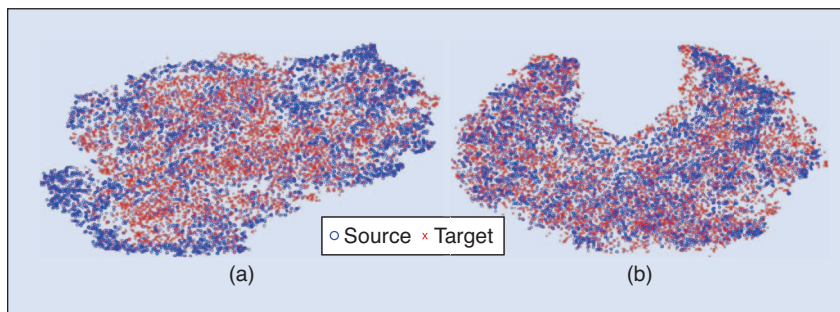


FIGURE 8 Illustration of data distributions from source and target before (a) and after (b) the DANN training [55].

DCGAN, in which a discriminator only learns to classify the input as either real or fake. Hence, no label information is demanded. Note that, the discriminator shares layers with another two classifiers to predict valence and activation simultaneously. This method has been shown to improve generalizability across corpora [105]. A similar approach was conducted in [54] to learn robust representations from emotional speech data for autism detection.

More recently, a cGAN-based framework was proposed for continuous speech emotion recognition in [20], where a predictor and a discriminator are conditioned by acoustic features. In particular, the discriminator is employed to distinguish the joint probability distributions for acoustic features and their corresponding predictions or real annotations. In this way, the predictor is guided to modify the original predictions to achieve a better performance level.

Rather than the above mentioned adversarial training schemes that explicitly rely on the presence of a discriminator network, adversarial training can also be executed in an implicit manner, namely, virtual adversarial training. Virtual adversarial training is conducted by straightforwardly adding an additional regularization term, which is sensitive to the adversarial examples, as a penalty in a loss function.

In virtual adversarial training, first proposed in [106], the regularization loss term of Equation (3) is reformulated without the label y as follows:

$$\mathbf{KL}[p(\cdot|x;\theta) \| p(\cdot|x+r_{\text{adv}};\theta)], \quad (5)$$

where $\mathbf{KL}[\cdot \| \cdot]$ denotes the \mathbf{KL} divergence between two distributions and a worst-case perturbation r_{adv} can be computed by

$$r_{\text{adv}} = \arg \max_{r, \|r\| \leq \epsilon} \mathbf{KL}[p(\cdot|x;\theta) \| p(\cdot|x+r;\theta)], \quad (6)$$

Inspired by these works, authors in [107] reported that, state-of-the-art sentiment classification results can be achieved when adopting (virtual) adversarial training approaches in the text

Generally speaking, not only have adversarial training techniques made great contributions to the development of corresponding generative models, but they are also helpful and instructive for related discriminative models.

domain. In particular, the authors applied perturbations to word embeddings in a recurrent neural network structure, rather than to the original input itself [107]. Following this success, (virtual) adversarial training has also been applied to speech emotion recognition in [108]. Results in [108] demonstrate that, the classification accuracy, as well as the system's overall generalization capability, can be improved.

VII. The Road Ahead

Considerable progress has been made in alleviating some of the challenges related to affective computing and sentiment analysis through the use of adversarial training, for example, synthesizing and transforming image-based emotions through cycleGAN, augmenting data by artificially generating samples, extracting robust representations via domain adversarial training. A detailed summary of these works can be found in Table 1. However, large-scale breakthroughs are still required in both the theorem of adversarial training and its applications to fully realize the potential of this paradigm in affective computing and sentiment analysis.

A. Limitations of Adversarial Training

Arguably, the two major open research challenges relating to adversarial training are training instability and mode collapse. Solving these fundamental concerns will help facilitate its application to affective computing and sentiment analysis.

1) Training Instability

In the adversarial training process, ensuring that there are balance and synchronization between the two adversarial networks plays an important role in obtaining reliable results [16]. That is, the goal optimization of adversarial training

lies in finding a saddle point of, rather than a local minimum between, the two adversarial components. The inherent difficulty in controlling the synchronization of the two adversarial networks increases the risk of instability in the training process.

To date, researchers have made several attempts to address this problem. For example, the implementation of Wasserstein distance rather than the conventional JS divergence partially solves the vanishing gradient problem associated with improvements in the ability of the discriminator to separate the real and generated samples [58]. Furthermore, the convergence of the model and the existence of the equilibrium point have yet to be theoretically proven [109]. Therefore, it remains an open research direction to further optimize the training process.

2) Mode Collapse

Mode collapse occurs when the generator exhibits very limited diversity among generated samples, thus reducing the usefulness of the learned GANs. This effect can be observed as the generated samples can be integrated into a small subset of similar samples (partial collapse), or even a single sample (complete collapse).

Novel approaches dedicated to solving the mode collapse problem are continually emerging. For example, the loss function of the generator can be modified to factor in the diversity of generated samples in batches [110]. Alternatively, the unroll-GAN allows the generator to 'unroll' the updates of the discriminator in a manner which is fully differentiable [111], and the AdaGAN combines an ensemble of GANs in a boosting framework to ensure diversity [112].

TABLE 1 A summary of adversarial training studies in affective computing and sentiment analysis. These studies are listed by their applied tasks (SYN: synthesis, CVS: conversion, DA: data augmentation, DAT: domain adversarial training, AT: adversarial training, VAT: virtual adversarial training), modalities, and published years. GATH: generative adversarial talking head, ASPD: adversarial shared-private model.

PAPER	YEAR	TASK	MODALITY	MODEL	NOTE
RADFOD ET AL. [19]	2016	SYN	IMAGE	DCGAN	VECTOR ARITHMETIC CAN BE DONE IN LATENT VECTOR SPACE, E.G., SMILING WOMAN - NEUTRAL WOMAN + NEUTRAL MAN = SMILING MAN.
CHEN ET AL. [69]	2016	SYN	IMAGE	INFOGAN	LATENT CODE CAN BE INTERPRETED; SUPPORT GRADUAL TRANSFORMATION.
HUANG & KHAN [51]	2017	SYN	IMAGE/VIDEO	DYADGAN	INTERACTION SCENARIO; IDENTITY + ATTRIBUTE FROM THE INTERVIEWEE.
MELIS & AMORES [83]	2017	SYN	IMAGE (ART)	CGAN	GENERATE EMOTIONAL ARTWORK.
BAO ET AL. [49]	2018	SYN	IMAGE	IDENTITY PRESERVING GAN	THE IDENTITY AND ATTRIBUTES OF FACES ARE SEPARATED.
PHAM ET AL. [81]	2018	SYN	IMAGE/VIDEO	GATH	CONDITIONED BY AU; FROM STATIC IMAGE TO DYNAMIC VIDEO.
SONG ET AL. [82]	2018	SYN	IMAGE/VIDEO	CGAN	STATIC IMAGE TO DYNAMIC VIDEO, CONDITIONED BY AUDIO SEQUENCES.
NOJAVANSGHARI ET AL. [50]	2018	SYN	IMAGE	DYADGAN (WITH RNN)	INTERACTION SCENARIO; SMOOTH THE VIDEO SYNTHESIS WITH CONTEXT INFORMATION.
WANG & ARITÈRES [84]	2018	SYN	MOTION	CGAN	TO SIMULATE A LATENT VECTOR WITH SEQ2SEQ AE; CONTROLLED BY EMOTION.
RAJESWAR ET AL. [14]	2017	SYN	TEXT	(W)GAN-GP	GRADIENT PENALTY; GENERATE SEQUENCES WITH FIXED LENGTH.
LIU ET AL. [85]	2018	SYN	TEXT (POETRY)	I2P-GAN	MULTI-ADVERSARIAL TRAINING; GENERATE POETRY FROM IMAGES.
FEDUS ET AL. [21]	2018	SYN	TEXT	MASKGAN	BASED ON ACTOR-CRITIC CGAN; GENERATE SEQUENCES WITH FIXED LENGTH.
PERARNAU ET AL. [91]	2016	CVS	IMAGE	ICGAN	INTERPRETABLE LATENT CODE; SUPPORT GRADUAL TRANSFORMATION.
ZHOU & SHI [88]	2017	CVS	IMAGE	CGAN	LEARN THE DIFFERENCE BETWEEN THE SOURCE AND TARGET EMOTIONS BY ADVERSARIAL AUTOENCODER.
SONG ET AL. [89]	2017	CVS	IMAGE	G2-GAN	GEOMETRY-GUIDED, SIMILAR TO CYCLEGAN.
QIAO ET AL. [38]	2018	CVS	IMAGE	GC-GAN	GEOMETRY-CONTRASTIVE LEARNING; THE ATTRIBUTE AND IDENTITY FEATURES ARE SEPARATED IN THE LEARNING PROCESS.
DING ET AL. [92]	2018	CVS	IMAGE	EXPRGAN	INTENSITY CAN BE CONTROLLED.
LEE ET AL. [93]	2018	CVS	TEXT	CYCLEGAN	NO NEED OF PAIRED DATA; EMOTION SCALABLE.
GAO ET AL. [15]	2018	CVS	SPEECH	VOICEGAN	NO NEED OF PAIRED DATA.
ZHU ET AL. [96]	2018	DA	IMAGE	CYCLEGAN	TRANSFER DATA FROM A TO B; REQUIRE NO FURTHER LABELLING PROCESS ON THE TRANSFERRED DATA.
SAHU ET AL. [97]	2017	DA	SPEECH	ADVERSARIAL AE	USE GMM BUILT ON THE ORIGINAL DATA TO LABEL GENERATED DATA ->NOISY DATA; SENSITIVE TO MODE COLLAPSE.
GANIN ET AL. [98]	2016	DAT	TEXT	DANN	FIRST WORK IN DOMAIN ADVERSARIAL TRAINING.
CHEN ET AL. [12]	2017	DAT	TEXT	DANN	SEMI-SUPERVISED SUPPORTED; WASSERSTEIN DISTANCE USED FOR SMOOTHING TRAINING PROCESS.
ZHANG ET AL. [95]	2017	DAT	TEXT	DANN	ATTENTION SCORING NETWORK IS ADDED FOR DOCUMENT EMBEDDING.
LI ET AL. [99]	2017	DAT	TEXT	DANN	WITH ATTENTION MECHANISMS.
ZHAO ET AL. [101]	2018	DAT	TEXT	MULTISOURCE DANN	EXTENDED FOR MULTIPLE SOURCES.

(Continued)

TABLE 1 A summary of adversarial training studies in affective computing and sentiment analysis. These studies are listed by their applied tasks (SYN: synthesis, CVS: conversion, DA: data augmentation, DAT: domain adversarial training, AT: adversarial training, VAT: virtual adversarial training), modalities, and published years. GATH: generative adversarial talking head, ASPD: adversarial shared-private model. (Continued)

PAPER	YEAR	TASK	MODALITY	MODEL	NOTE
SHEN ET AL. [100]	2018	DAT	TEXT	DANN	WASSERSTEIN DISTANCE GUIDED TO OPTIMIZE DOMAIN DISCRIMINATOR.
LIU ET AL. [102]	2018	DAT	TEXT	ASPD	FOR MULTI-TASK; SEMI-SUPERVISED FRIENDLY.
CHEN & CARDIE [103]	2018	DAT	TEXT	MULTINOMIAL ASPD	MULTINOMIAL DISCRIMINATOR FOR MULTI-DOMAIN.
MOHAMMED & BUSSO [55]	2018	DAT	SPEECH	DANN	ADAPTED FOR SPEECH EMOTION RECOGNITION.
CHANG & SCHERER [105]	2017	AT	SPEECH	DCGAN	SPECTROGRAMS WITH FIXED WIDTH ARE RANDOMLY SELECTED AND CHOPPED FROM A VARIED LENGTH OF AUDIO FILES.
DENG ET AL. [54]	2017	AT	SPEECH	GAN	USE HIDDEN-LAYER REPRESENTATIONS FROM DISCRIMINATORS
HAN ET AL. [20]	2018	AT	SPEECH	CGAN	REGULARIZATION: JOINT DISTRIBUTION.
MIYATO ET AL. [107]	2017	VAT/AT	TEXT	/	FIRST WORK ON VIRTUAL ADVERSARIAL TRAINING.
SAHU ET AL. [108]	2018	VAT/AT	SPEECH	DNN	FIRST WORK FOR SPEECH EMOTION RECOGNITION.

B. Other Ongoing Breakthroughs

In most conditional GAN frameworks, the emotional entity is generated by utilizing discrete emotional categories as the condition controller. However, emotions are more than these basic categories (e.g., Ekman's Six Basic Emotions), and to date, more subtle emotional expressions have been largely overlooked in the literature. While some studies have started addressing this issue in image processing studies (cf. Section IV), it is still one of the major research white spots in speech and text processing. Therefore, using a softer condition to replace the controller remains an open research direction.

To the best of our knowledge, GAN-based emotional speech synthesis has yet to be addressed in the relevant literature. This could be due in part to Speech Emotion Recognition (SER) being a less mature field of research, which leads to a limited capability to distinguish the emotions using speech and thus provides deductible contributions to optimize the generator. However, with the ongoing development of SER and the already discussed success of GANs in conventional speech synthesis [113] and image/video and text generation (cf. Section IV), we strongly believe that major breakthroughs will be made in this area sometime in the near future.

Similarly, current state-of-the-art emotion conversion systems are based on the transformation of static images. However, transforming emotions in the dynamic sequential signals, such as speech, video, and text, remains challenging. This most likely relates to the difficulties associated with sequence-based discriminator and sequence generation. However, the state-of-the-art performance achieved with generative models, such as WaveGAN, indicate that adversarial training can play a key role in helping to break through these barriers.

Additionally, when comparing the performance of different GAN-based models, a fair comparison is vital but not straightforward. In [85], the authors demonstrated that their proposed I2P-GAN outperforms SeqGAN when generating poetry from given images, reporting higher scores on evaluation metrics including BLEU, novelty, and relevance. Also, it has been claimed that InfoGAN converges faster than a conventional GAN framework [69]. However, it should be noted that, a fair experimental comparison of various generative adversarial training models associated with affective computing and sentiment analysis has yet to be reported, to answer questions such as which model is faster, more accurate, or

easier to implement. This absence is mainly due to the lack of benchmark datasets and thoughtfully designed metrics for each specific application (i.e., emotion generation, emotion conversion, and emotion perception and understanding).

Finally, we envisage that, GAN-based end-to-end emotional dialogue systems can succeed the speech-to-text (i.e., ASR) and the text-to-speech (i.e., TTS) processes currently used in conventional dialogue systems. This is motivated by the construct that humans generally do not consciously convert speech into text during conversations [114]. The advantage of this approach is that it avoids the risk of the possible information loss during this internal process. Such an end-to-end emotional framework would further facilitate the next generation of more human-like dialogue systems.

VIII. Conclusion

Motivated by the ongoing success and achievements associated with adversarial training in artificial intelligence, this article summarized the most recent advances of adversarial training in affective computing and sentiment analysis. Covering the audio, image/video, and text modalities, this overview included technologies and paradigms relating to both

emotion synthesis and conversion as well as emotion perception and understanding. Generally speaking, not only have adversarial training techniques made great contributions to the development of corresponding generative models, but they are also helpful and instructive for related discriminative models. We have also drawn attention to further research efforts aimed at leveraging the highlighted advantages of adversarial training. If successfully implemented, such techniques will inspire and foster the new generation of robust affective computing and sentiment analysis technologies that are capable of widespread in-the-wild deployment.

Acknowledgment

This work has been supported by the EU's Horizon 2020 Programme through the Innovation Action No. 645094 (SEWA), the EU's Horizon 2020/EFPIA Innovative Medicines Initiative through GA No. 115902 (RADAR-CNS), and the UK's Economic & Social Research Council through the research Grant No. HJ-253479 (ACLEW).

References

[1] R. Picard, *Affective Computing*. Cambridge, MA: MIT Press, 1997.

[2] M. Minsky, *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*. New York, NY: Simon and Schuster, 2007.

[3] M. Pantic, N. Sebe, J. F. Cohn, and T. Huang, "Affective multimodal human-computer interaction," in *Proc. 13th ACM Int. Conf. Multimedia*, Singapore, 2005, pp. 669–676.

[4] S. Poria, E. Cambria, A. Gelbukh, F. Bisio, and A. Hussain, "Sentiment data flow analysis by means of dynamic linguistic patterns," *IEEE Comput. Intell. Mag.*, vol. 10, no. 4, pp. 26–36, Nov. 2015.

[5] S. Poria, E. Cambria, A. Hussain, and G.-B. Huang, "Towards an intelligent framework for multimodal affective data analysis," *Neural Netw.*, vol. 63, pp. 104–116, Mar. 2015.

[6] E. Cambria, "Affective computing and sentiment analysis," *IEEE Intell. Syst.*, vol. 31, no. 2, pp. 102–107, Mar. 2016.

[7] T. Chen, R. Xu, Y. He, Y. Xia, and X. Wang, "Learning user and product distributed representations using a sequence model for sentiment analysis," *IEEE Comput. Intell. Mag.*, vol. 11, no. 3, pp. 34–44, Aug. 2016.

[8] J. Han, Z. Zhang, N. Cummins, F. Ringeval, and B. Schuller, "Strength modelling for real-world automatic continuous affect recognition from audiovisual signals," *Image Vis. Comput.*, vol. 65, pp. 76–86, Sept. 2017.

[9] C. N. dos Santos and M. Gatti, "Deep convolutional neural networks for sentiment analysis of short texts," in *Proc. 25th Int. Conf. Computational Linguistics*, Dublin, Ireland, 2014, pp. 69–78.

[10] P. Tzirakis, G. Trigeorgis, M. A. Nicolaou, B. Schuller, and S. Zafeiriou, "End-to-end multimodal emotion recognition using deep neural networks," *IEEE J. Sel. Topics Signal Process.*, vol. 11, no. 8, pp. 1301–1309, Dec. 2017.

[11] Z. Zhang, N. Cummins, and B. Schuller, "Advanced data exploitation for speech analysis: An overview," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 107–129, July 2017.

[12] X. Chen, Y. Sun, B. Athiwaratkun, C. Cardie, and K. Weinberger, "Adversarial deep averaging networks for cross-lingual sentiment classification," arXiv Preprint, arXiv:1606.01614, Apr. 2017.

[13] J. Deng, X. Xu, Z. Zhang, S. Frühholz, and B. Schuller, "Semisupervised autoencoders for speech emotion recognition," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 26, no. 1, pp. 31–43, Jan. 2018.

[14] S. Subramanian, S. Rajeswar, F. Dutil, C. Pal, and A. C. Courville, "Adversarial generation of natural language," in *Proc. 2nd Workshop Representation Learning NLP (Rep4NLP@ACL)*, Vancouver, Canada, 2017, pp. 241–251.

[15] Y. Gao, R. Singh, and B. Raj, "Voice impersonation using generative adversarial networks," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Calgary, Canada, 2018, pp. 2506–2510.

[16] I. Goodfellow et al., "Generative adversarial nets," in *Proc. Advances in Neural Information Processing Systems*, Montreal, Canada, 2014, pp. 2672–2680.

[17] K. Wang, C. Gou, Y. Duan, Y. Lin, X. Zheng, and F.-Y. Wang, "Generative adversarial networks: Introduction and outlook," *IEEE/CAA J. Autom. Sinica*, vol. 4, no. 4, pp. 588–598, Sept. 2017.

[18] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, "Generative adversarial networks: An overview," *IEEE Signal Process. Mag.*, vol. 35, no. 1, pp. 53–65, Jan. 2018.

[19] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," San Juan, PR, 2016.

[20] J. Han, Z. Zhang, Z. Ren, F. Ringeval, and B. Schuller, "Towards conditional adversarial training for predicting emotions from speech," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Calgary, Canada, 2018, pp. 6822–6826.

[21] W. Fedus, I. Goodfellow, and A. M. Dai, "MaskGAN: Better text generation via filling in the _," in *Proc. 6th Int. Conf. Learning Representations*, Vancouver, Canada, 2018.

[22] Z. Zeng, M. Pantic, G. I. Roisman, and T. S. Huang, "A survey of affect recognition methods: Audio, visual, and spontaneous expressions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 1, pp. 39–58, Jan. 2009.

[23] R. A. Calvo and S. D'Mello, "Affect detection: An interdisciplinary review of models, methods, and their applications," *IEEE Trans. Affect. Comput.*, vol. 1, no. 1, pp. 18–37, Jan. 2010.

[24] H. Gunes, B. Schuller, M. Pantic, and R. Cowie, "Emotion representation, analysis and synthesis in continuous space: A survey," in *Proc. IEEE Int. Conf. Automatic Face & Gesture Recognition and Workshops*, Santa Barbara, CA, 2011, pp. 827–834.

[25] B. Schuller, "Speech emotion recognition: Two decades in a nutshell, benchmarks, and ongoing trends," *Commun. ACM*, vol. 61, no. 5, pp. 90–99, Apr. 2018.

[26] B. Liu, *Sentiment Analysis and Opinion Mining*. San Rafael, CA: Morgan & Claypool, 2012.

[27] W. Medhat, A. Hassan, and H. Korashy, "Sentiment analysis algorithms and applications: A survey," *Ain Shams Eng. J.*, vol. 5, no. 4, pp. 1093–1113, Dec. 2014.

[28] B. Liu, *Sentiment Analysis: Mining Opinions, Sentiments, and Emotions*. Cambridge, U.K.: Cambridge Univ. Press, June 2015.

[29] E. Cambria, S. Poria, A. Gelbukh, and M. Thelwall, "Sentiment analysis is a big suitcase," *IEEE Intell. Syst.*, vol. 32, no. 6, pp. 74–80, Dec. 2017.

[30] S. Poria, E. Cambria, R. Bajpai, and A. Hussain, "A review of affective computing: From unimodal analysis to multimodal fusion," *Inf. Fusion*, vol. 37, pp. 98–125, Sept. 2017.

[31] M. Soleymani, D. Garcia, B. Jou, B. Schuller, S.-F. Chang, and M. Pantic, "A survey of multimodal sentiment analysis," *Image Vis. Comput.*, vol. 65, pp. 3–14, Sept. 2017.

[32] L. Zhang, S. Wang, and B. Liu, "Deep learning for sentiment analysis: A survey," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, Mar. 2018, 25 p.

[33] A. van den Oord et al., "WaveNet: A generative model for raw audio," arXiv Preprint, arXiv:1609.03499, Sept. 2016.

[34] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu, "Pixel recurrent neural networks," in *Proc. 33rd Int. Conf. Machine Learning*, New York City, NY, 2016, pp. 1747–1756.

[35] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," arXiv Preprint, arXiv:1312.6114, Dec. 2013.

[36] Y. Lee, A. Rabiee, and S.-Y. Lee, "Emotional end-to-end neural speech synthesizer," arXiv Preprint, arXiv:1711.05447, Nov. 2017.

[37] K. Akuzawa, Y. Iwasawa, and Y. Matsuo, "Expressive speech synthesis via modeling expressions with variational autoencoder," in *Proc. Annu. Conf. Int. Speech Communication Association*, Hyderabad, India.

[38] F. Qiao, N. Yao, Z. Jiao, Z. Li, H. Chen, and H. Wang, "Geometry-contrastive generative adversarial network for facial expression synthesis," arXiv Preprint, arXiv:1802.01822, Feb. 2018.

[39] H. Zhou, M. Huang, T. Zhang, X. Zhu, and B. Liu, "Emotional chatting machine: Emotional conversation generation with internal and external memory," in *Proc. 32nd Conf. Association Advancement Artificial Intelligence*, New Orleans, LA, 2018, pp. 730–738.

[40] B. Schuller et al., "Cross-corpus acoustic emotion recognition: Variances and strategies," *IEEE Trans. Affect. Comput.*, vol. 1, no. 2, pp. 119–131, July 2010.

[41] J. Han, Z. Zhang, M. Schmitt, M. Pantic, and B. Schuller, "From hard to soft: Towards more human-like emotion recognition by modelling the perception uncertainty," in *Proc. 25th ACM Int. Conf. Multimedia*, Mountain View, CA, 2017, pp. 890–897.

[42] Y. Kim, H. Lee, and E. M. Provost, "Deep learning for robust feature generation in audiovisual emotion recognition," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Vancouver, Canada, 2013, pp. 3687–3691.

[43] N. Cummins, S. Amiriparian, G. Hagerer, A. Batliner, S. Steidl, and B. Schuller, "An image-based deep spectrum feature representation for the recognition of emotional speech," in *Proc. 25th ACM Int. Conf. Multimedia*, Mountain View, CA, 2017, pp. 478–484.

[44] Z. Zhang, J. Han, J. Deng, X. Xu, F. Ringeval, and B. Schuller, "Leveraging unlabeled data for emotion recognition with enhanced collaborative semi-supervised learning," *IEEE Access*, vol. 6, pp. 22196–22209, Apr. 2018.

[45] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.

[46] X. Glorot, A. Bordes, and Y. Bengio, "Domain adaptation for large-scale sentiment classification: A deep learning approach," in *Proc. 28th Int. Conf. Machine Learning*, Bellevue, WA, 2011, pp. 513–520.

[47] J. Deng, R. Xia, Z. Zhang, Y. Liu, and B. Schuller, "Introducing shared-hidden-layer autoencoders for transfer learning and their application in acoustic emotion recognition," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Florence, Italy, 2014, pp. 4818–4822.

[48] Q. You, J. Luo, H. Jin, and J. Yang, "Robust image sentiment analysis using progressively trained and domain transferred deep networks," in *Proc. 29th Conf. Association Advancement Artificial Intelligence*, Austin, TX, 2015, pp. 381–388.

[49] J. Bao, D. Chen, F. Wen, H. Li, and G. Hua, "Towards open-set identity preserving face synthesis," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Salt Lake City, UT, 2018, pp. 6713–6722.

[50] B. Nojavanasghari, Y. Huang, and S. Khan, "Interactive generative adversarial networks for facial expression generation in dyadic interactions," arXiv Preprint, arXiv:1801.09092, Jan. 2018.

[51] Y. Huang and S. M. Khan, "DyadGAN: Generating facial expressions in dyadic interactions," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition Workshops*, Honolulu, HI, 2017, pp. 2259–2266.

[52] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: A method for automatic evaluation of machine translation," in *Proc. 40th Annu. Meeting Association Computational Linguistics*, Stroudsburg, PA, 2002, pp. 311–318.

- [53] C.-Y. Lin, "ROUGE: A package for automatic evaluation of summaries," in *Proc. Text Summarization Branches Out Workshop ACL*, Barcelona, Spain, 2004, 8 p.
- [54] J. Deng, N. Cummins, M. Schmitt, K. Qian, F. Ringeval, and B. Schuller, "Speech-based diagnosis of autism spectrum condition by generative adversarial network representations," in *Proc. Int. Conf. Digital Health*, London, UK, 2017, pp. 53–57.
- [55] M. Abdelwahab and C. Busso, "Domain adversarial for acoustic emotion recognition," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 26, no. 12, pp. 2423–2435, Dec. 2018.
- [56] H. Caesar, "Really-awesome-gan," 2017. [Online]. Available: <https://github.com/nightrome/really-awesome-gan>
- [57] A. Hindupur, "The-gan-zoo," 2018. [Online]. Available: <https://github.com/hindupuravinash/the-gan-zoo>
- [58] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," arXiv Preprint, arXiv:1701.07875, Mar. 2017.
- [59] J. Zhao, M. Mathieu, and Y. LeCun, "Energy-based generative adversarial network," in *Proc. 5th Int. Conf. Learning Representations*, Toulon, France, 2017.
- [60] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley, "Least squares generative adversarial networks," in *Proc. IEEE Int. Conf. Computer Vision*, Venice, Italy, 2017, pp. 2794–2802.
- [61] G.-J. Qi, "Loss-sensitive generative adversarial networks on Lipschitz densities," arXiv Preprint, arXiv:1701.06264, Jan. 2017.
- [62] S. Patel, A. Kakadiya, M. Mehta, R. Derasari, R. Patel, and R. Gandhi, "Correlated discrete data generation using adversarial training," arXiv Preprint, arXiv:1804.00925, 2018.
- [63] T. Che, Y. Li, A. P. Jacob, Y. Bengio, and W. Li, "Mode regularized generative adversarial networks," in *Proc. 5th Int. Conf. Learning Representations*, Toulon, France, 2017.
- [64] M. Mirza and S. Osindero, "Conditional generative adversarial nets," arXiv Preprint, arXiv:1411.1784, June 2014.
- [65] A. Odena, "Semi-supervised learning with generative adversarial networks," in *Proc. Data-Efficient Machine Learning Workshop in ICML*, New York, NY, 2016.
- [66] J. Donahue, P. Krähenbühl, and T. Darrell, "Adversarial feature learning," in *Proc. 5th Int. Conf. Learning Representations*, Toulon, France, 2017.
- [67] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Venice, Italy, 2017, pp. 2223–2232.
- [68] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim, "Learning to discover cross-domain relations with generative adversarial networks," in *Proc. 34th Int. Conf. Machine Learning*, Sydney, Australia, 2017, pp. 1857–1865.
- [69] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel, "InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets," in *Proc. Advances in Neural Information Processing Systems*, Barcelona, Spain, 2016, pp. 2172–2180.
- [70] L. Chongxuan, T. Xu, J. Zhu, and B. Zhang, "Triple generative adversarial nets," in *Proc. Advances in Neural Information Processing Systems*, Long Beach, CA, 2017, pp. 4091–4101.
- [71] J. Luo, Y. Xu, C. Tang, and J. Lv, "Learning inverse mapping by autoencoder based generative adversarial nets," in *Proc. Int. Conf. Neural Information Processing*, Guangzhou, China, 2017, pp. 207–216.
- [72] O. Mørgen, "C-RNN-GAN: Continuous recurrent neural networks with adversarial training," in *Proc. Constructive Machine Learning Workshop in NIPS*, Barcelona, Spain, 2016, 6 p.
- [73] T. Xu et al., "AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Salt Lake City, UT, 2018, pp. 1316–1324.
- [74] A. Jaiswal, W. AbdAlmageed, and P. Natarajan, "CapsuleGAN: Generative adversarial capsule network," arXiv Preprint, arXiv:1802.06167, Feb. 2018.
- [75] A. Creswell and A. A. Bharath, "Adversarial training for sketch retrieval," in *Proc. European Conf. Computer Vision*, Amsterdam, Netherlands, 2016, pp. 798–809.
- [76] W. R. Tan, C. S. Chan, H. E. Aguirre, and K. Tanaka, "ArtGAN: Artwork synthesis with conditional categorical GANs," in *Proc. IEEE Int. Conf. Image Processing*, Beijing, China, 2017, pp. 3760–3764.
- [77] S. Pascual, A. Bonafonte, and J. Serrà, "SEGAN: Speech enhancement generative adversarial network," in *Proc. 18th Annu. Conf. Int. Speech Communication Association*, Stockholm, Sweden, 2017, pp. 3642–3646.
- [78] C. Donahue, J. McAuley, and M. Puckette, "Synthesizing audio with generative adversarial networks," in *Proc. Workshop in 6th Int. Conf. Learning Representations*, Vancouver, Canada, 2018.
- [79] A. van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves, and K. Kavukcuoglu, "Conditional image generation with PixelCNN decoders," in *Proc. Advances in Neural Information Processing Systems*, Barcelona, Spain, 2016, pp. 4790–4798.
- [80] I. Caswell, O. Sen, and A. Nie, "Exploring adversarial learning on neural network models for text classification," 2015.
- [81] H. X. Pham, Y. Wang, and V. Pavlovic, "Generative adversarial talking head: Bringing portraits to life with a weakly supervised neural network," arXiv Preprint, arXiv:1803.07716, Mar. 2018.
- [82] Y. Song, J. Zhu, X. Wang, and H. Qi, "Talking face generation by conditional recurrent adversarial network," arXiv Preprint, arXiv:1804.04786, Apr. 2018.
- [83] D. Alvarez-Melis and J. Amores, "The emotional GAN: Priming adversarial generation of art with emotion," in *Proc. Advances in Neural Information Processing Systems*, Long Beach, CA, 2017, 4 p.
- [84] Q. Wang and T. Artières, "Motion capture synthesis with adversarial learning," in *Proc. 17th Int. Conf. Intelligent Virtual Agents*, Stockholm, Sweden, 2017, pp. 467–470.
- [85] B. Liu, J. Fu, M. P. Kato, and M. Yoshikawa, "Beyond narrative description: Generating poetry from images by multi-adversarial training," arXiv Preprint, arXiv:1804.08473, Apr. 2018.
- [86] L. Yu, W. Zhang, J. Wang, and Y. Yu, "SeqGAN: Sequence generative adversarial nets with policy gradient," in *Proc. 31st Conf. Association Advancement Artificial Intelligence*, San Francisco, CA, 2017, pp. 2852–2858.
- [87] J. Li, W. Monroe, T. Shi, S. Jean, A. Ritter, and D. Jurafsky, "Adversarial learning for neural dialogue generation," in *Proc. Empirical Methods in Natural Language Processing*, Copenhagen, Denmark, 2017, pp. 2157–2169.
- [88] Y. Zhou and B. E. Shi, "Photorealistic facial expression synthesis by the conditional difference adversarial autoencoder," in *Proc. 7th Int. Conf. Affective Computing and Intelligent Interaction*, San Antonio, TX, 2017, pp. 370–376.
- [89] L. Song, Z. Lu, R. He, Z. Sun, and T. Tan, "Geometry guided adversarial facial expression synthesis," arXiv Preprint, arXiv:1712.03474, Dec. 2017.
- [90] C. Hsu, H. Hwang, Y. Wu, Y. Tsao, and H. Wang, "Voice conversion from unaligned corpora using variational autoencoding Wasserstein generative adversarial networks," in *Proc. 18th Annu. Conf. Int. Speech Communication Association*, Stockholm, Sweden, 2017, pp. 3364–3368.
- [91] G. Perarnau, J. van de Weijer, B. Raducanu, and J. M. Álvarez, "Invertible conditional gans for image editing," in *Proc. Adversarial Training Workshop in NIPS*, Barcelona, Spain, 2016.
- [92] H. Ding, K. Sricharan, and R. Chellappa, "ExprGAN: Facial expression editing with controllable expression intensity," in *Proc. 32nd Conf. Association Advancement Artificial Intelligence*, New Orleans, LA, 2018, pp. 6781–6788.
- [93] C. Lee, Y. Wang, T. Hsu, K. Chen, H. Lee, and L. Lee, "Scalable sentiment for sequence-to-sequence chatbot response with performance analysis," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Calgary, Canada, 2018, pp. 6164–6168.
- [94] B. Schuller, Z. Zhang, F. Weninger, and F. Burkhardt, "Synthesized speech for model training in cross-corpus recognition of human emotion," *Int. J. Speech Technol.*, vol. 15, no. 3, pp. 313–323, Sept. 2012.
- [95] Y. Zhang, R. Barzilay, and T. S. Jaakkola, "Aspect-augmented adversarial networks for domain adaptation," *Trans. Assoc. Comput. Linguist.*, vol. 5, pp. 515–528, Dec. 2017.
- [96] X. Zhu, Y. Liu, J. Li, T. Wan, and Z. Qin, "Emotion classification with data augmentation using generative adversarial networks," in *Proc. Pacific-Asia Conf. Knowledge Discovery and Data Mining*, Melbourne, Australia, 2018, pp. 349–360.
- [97] S. Sahu, R. Gupta, G. Sivaraman, W. AbdAlmageed, and C. Y. Espy-Wilson, "Adversarial auto-encoders for speech based emotion recognition," in *Proc. 18th Annu. Conf. Int. Speech Communication Association*, Stockholm, Sweden, 2017, pp. 1243–1247.
- [98] Y. Ganin et al., "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2096–2130, Jan. 2016.
- [99] Z. Li, Y. Zhang, Y. Wei, Y. Wu, and Q. Yang, "End-to-end adversarial memory network for cross-domain sentiment classification," in *Proc. 26th Int. Joint Conf. Artificial Intelligence*, Melbourne, Australia, 2017, pp. 2237–2243.
- [100] J. Shen, Y. Qu, W. Zhang, and Y. Yu, "Wasserstein distance guided representation learning for domain adaptation," in *Proc. 32nd Conf. Association Advancement Artificial Intelligence*, New Orleans, LA, 2018, pp. 4058–4065.
- [101] H. Zhao, S. Zhang, G. Wu, J. Costeira, J. Moura, and G. Gordon, "Multiple source domain adaptation with adversarial learning," in *Proc. Workshop in 6th Int. Conf. Learning Representations*, Vancouver, Canada, 2018.
- [102] P. Liu, X. Qiu, and X. Huang, "Adversarial multi-task learning for text classification," in *Proc. 55th Annu. Meeting Association Computational Linguistics*, Vancouver, Canada, 2017, pp. 1–10.
- [103] X. Chen and C. Cardie, "Multinomial adversarial networks for multi-domain text classification," in *Proc. 16th Annu. Conf. North American Chapter Association Computational Linguistics*, New Orleans, LA, 2018, pp. 1226–1240.
- [104] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *Proc. 3rd Int. Conf. Learning Representations*, Vancouver, Canada, 2015.
- [105] J. Chang and S. Scherer, "Learning representations of emotional speech with deep convolutional generative adversarial networks," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, New Orleans, LA, 2017, pp. 2746–2750.
- [106] T. Miyato, S.-i. Maeda, M. Koyama, K. Nakae, and S. Ishii, "Distributional smoothing with virtual adversarial training," in *Proc. 4th Int. Conf. Learning Representations*, San Juan, PR, 2016.
- [107] T. Miyato, A. M. Dai, and I. Goodfellow, "Adversarial training methods for semi-supervised text classification," in *Proc. 5th Int. Conf. Learning Representations*, Toulon, France, 2017.
- [108] S. Sahu, R. Gupta, G. Sivaraman, and C. Espy-Wilson, "Smoothing model predictions using adversarial training procedures for speech based emotion recognition," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Calgary, Canada, 2018, pp. 4934–4938.
- [109] S. Arora, R. Ge, Y. Liang, T. Ma, and Y. Zhang, "Generalization and equilibrium in generative adversarial nets," in *Proc. 34th Int. Conf. Machine Learning*, Sydney, Australia, 2017, pp. 224–232.
- [110] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training GANs," in *Proc. Advances in Neural Information Processing Systems*, Barcelona, Spain, 2016, pp. 2226–2234.
- [111] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein, "Unrolled generative adversarial networks," in *Proc. 5th Int. Conf. Learning Representations*, Toulon, France, 2017.
- [112] I. O. Tolstikhin, S. Gely, O. Bousquet, C. Simon-Gabriel, and B. Schölkopf, "AdaGAN: Boosting generative models," in *Proc. Advances in Neural Information Processing Systems*, Long Beach, CA, 2017, pp. 5430–5439.
- [113] S. Yang et al., "Statistical parametric speech synthesis using generative adversarial networks under a multi-task learning framework," in *Proc. IEEE Automatic Speech Recognition and Understanding Workshop*, Okinawa, Japan, 2017, pp. 685–691.
- [114] R. I. Dunbar, A. Marriott, and N. D. Duncan, "Human conversational behavior," *Human Nature*, vol. 8, no. 3, pp. 231–246, Sept. 1997.