

Exploiting Images for Video Recognition with Hierarchical Generative Adversarial Networks

Feiwu Yu¹, Xinxiao Wu^{1*}, Yuchao Sun¹, Lixin Duan²

¹ Beijing Laboratory of Intelligent Information Technology, School of Computer Science, Beijing Institute of Technology

² Big Data Research Center, University of Electronic Science and Technology of China
 {yufeiwu,wuxinxiao,sunyuchao}@bit.edu.cn, lxduan@uestc.edu.cn

Abstract

Existing deep learning methods of video recognition usually require a large number of labeled videos for training. But for a new task, videos are often unlabeled and it is also time-consuming and labor-intensive to annotate them. Instead of human annotation, we try to make use of existing fully labeled images to help recognize those videos. However, due to the problem of domain shifts and heterogeneous feature representations, the performance of classifiers trained on images may be dramatically degraded for video recognition tasks. In this paper, we propose a novel method, called Hierarchical Generative Adversarial Networks (HiGAN), to enhance recognition in videos (i.e., target domain) by transferring knowledge from images (i.e., source domain). The HiGAN model consists of a *low-level* conditional GAN and a *high-level* conditional GAN. By taking advantage of these two-level adversarial learning, our method is capable of learning a domain-invariant feature representation of source images and target videos. Comprehensive experiments on two challenging video recognition datasets (i.e. UCF101 and HMDB51) demonstrate the effectiveness of the proposed method when compared with the existing state-of-the-art domain adaptation methods.

1 Introduction

Video recognition is an active research area because of its wide applications such as anomalous events detection, action retrieval, human behavior analysis and so forth. Thanks to the success of deep neural networks, the performance of video recognition has been dramatically improved. However, training deep video classifiers requires collecting and labeling a large number of videos, which is often time-consuming and labor-intensive. On the other hand, images are much easier and cheaper to collect and annotate, and there are also

many existing labeled image datasets which can be utilized. Furthermore, the computational cost of learning deep classifiers of images is much less than that of videos, so it would be beneficial a lot to transfer knowledge from source images to target videos [Duan *et al.*, 2012a; Zhang *et al.*, 2017; Gan *et al.*, 2016].

However, applying the images trained classifier to video data directly introduces the domain shift problem, where the variations in data between source and target significantly degrade the recognition performance at test time. To solve this problem, many domain adaptation methods have been well explored. A rich line of prior work has focused on learning shallow features by obtaining a symmetric transformation and minimizing a distance metric of domain discrepancy [Akaho, 2006; Wang and Mahadevan, 2011; Duan *et al.*, 2012a; Shi *et al.*, 2010; Duan *et al.*, 2012b; Tsai *et al.*, 2016; Zhang *et al.*, 2017]. Recently, studies have shown that deep neural networks can learn more transferable features for domain adaptation [Chen *et al.*, 2016; Long *et al.*, 2015a; 2016; 2017]. In order to represent videos more appropriately, spatiotemporal features that are totally different from image representations with different feature dimensions and physical meanings are usually extracted, which makes the problem worse.

To tackle these issues aforementioned, inspired by recent advances in generative adversarial networks (GANs) [Goodfellow *et al.*, 2014], we propose a new approach referred to as *Hierarchical Generative Adversarial Networks* (HiGAN) to transfer knowledge from images to videos by learning domain-invariant feature representations between them. As illustrated in Figure 1(a), our approach mainly consists of two components: a two-level hierarchical conditional GAN model and a domain adaptation model. The domain adaptation model is adopted to learn a common feature representation between source images and target video frames, called image-frame feature. The two-level HiGAN is designed to have a *low-level* conditional GAN and a *high-level* conditional GAN. The *low-level* conditional GAN is built to connect videos and their corresponding video frames by learning a mapping function from frame features to video features in the target domain. The *high-level* conditional GAN, on the other hand, is modeled to bridge the gap between source

*Corresponding author: Xinxiao Wu

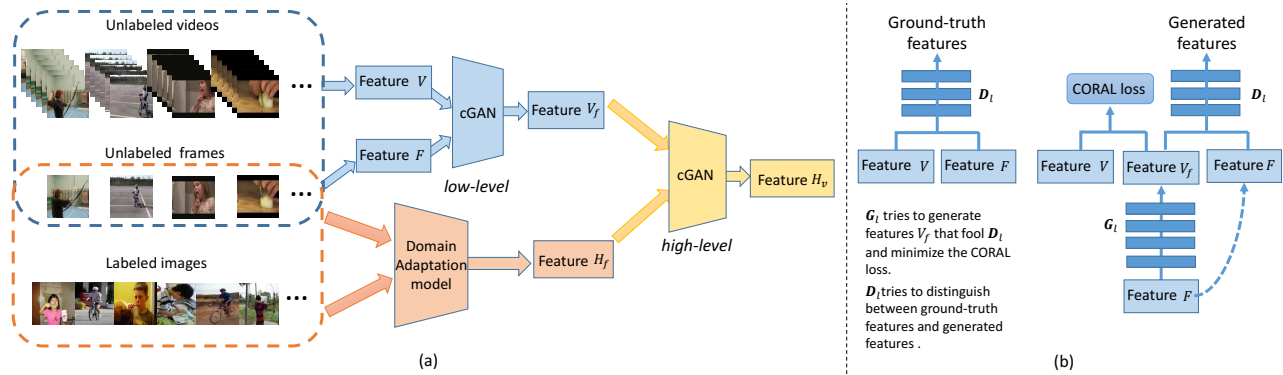


Figure 1: (a) The framework of the proposed HiGAN. We first train a Domain Adaptation model to learn features H_f . Then, we learn the *low-level* condition GAN with features V and F , and calculate the features V_f . After that, we train the *high-level* condition GAN with features V_f and H_f , and calculate the features H_v . (b) The training process of the *low-level* conditional GAN (the *high-level* conditional GAN has a similar structure).

images and target videos by formulating a mapping function from video features to image-frame features. Therefore, any classifiers trained on the source image domain can be effectively transferred to the target video domain for recognition, with the help of the transformable features between them. In HiGAN, the Correlation Alignment (CORAL) Loss [Sun and Saenko, 2016] is utilized to minimize the difference in second-order statistics between the generated features and the real features.

It is worth emphasizing that, by exploiting the correspondence between a video and its related frames, the projected video features, which keep the temporal motion of videos, can be learned with the absence of any paired image-video training examples in an unsupervised scenario. The experiments conducted on the UCF101 and HMDB51 datasets demonstrate that our method can achieve better results than the current state-of-the-art domain adaptation methods.

2 Related Work

To harness the information from large-scale image datasets, several works use images as auxiliary training data for video recognition [Duan *et al.*, 2012a; Zhang *et al.*, 2017; Gan *et al.*, 2016]. In [Duan *et al.*, 2012a], they proposed a new multiple domain adaptation method for event recognition by leveraging a large number of loosely labeled web images from different sources. In [Gan *et al.*, 2016], they jointly exploited source images and source videos for labeling-free video recognition, and proposed a mutually voting approach to filter noisy source images and video frames. In [Zhang *et al.*, 2017], they presented a classifier of image-to-video adaptation, which borrows the knowledge adapted from images, and utilizes the heterogeneous features of unlabeled videos to enhance the performance of action recognition. Different from these methods, our method takes the temporal information of videos into consideration without extracting keyframes from videos.

When transferring knowledge between different domains, the domain shift will cause the classifier learned from the source domain to perform poorly on the target domain. Numerous domain adaptation methods have been well explored

where the key focus is to learn domain-invariant feature representations. A common strategy is to find a mapping function that would align the source distribution with the target domain [Long *et al.*, 2014; 2015b; Sun *et al.*, 2016; Xu *et al.*, 2016; Busto and Gall, 2017; Akaho, 2006; Wang and Mahadevan, 2011; Shi *et al.*, 2010; Duan *et al.*, 2012b; Tsai *et al.*, 2016]. Recently, deep neural networks have been exploited for domain adaptation [Long *et al.*, 2015a; 2016; Chen *et al.*, 2016; Long *et al.*, 2017; Carlucci *et al.*, 2017]. In [Long *et al.*, 2015a; 2016; 2017], Long *et al.* embedded domain-adaptation modules into deep networks in which all the layers corresponding to task-specific features are adapted in a layerwise manner. These methods focus on the image-to-image adaptation. In [Chen *et al.*, 2016], the architecture of Transfer Neural Trees has been presented, which jointly solves the cross-domain feature adaptation and classification between heterogeneous domains in a semi-supervised scenario. In contrast, our method can transfer knowledge between images and videos in an unsupervised scenario.

Generative Adversarial Networks (GANs) [Goodfellow *et al.*, 2014] are a class of artificial intelligence algorithms that are implemented by a system of two competing models: a generative model and a discriminative model. These two networks compete with each other in a two-player min-max game: the generator is learning to produce as realistic as possible samples at confusing the discriminator, and the discriminator is learning to get as correct as possible results at distinguishing generated samples from real data. GANs have achieved impressive progress in domain adaptation [Bousmalis *et al.*, 2017; Sankaranarayanan *et al.*, 2017]. In [Sankaranarayanan *et al.*, 2017], they proposed an approach that brings the source and target distributions closer in a learned joint feature space. In [Bousmalis *et al.*, 2017], they presented a new approach that learns a transformation in the pixel space from one domain to another. Different from these homogeneous domain adaptation methods based on single GAN, our hierarchical GAN method can effectively learn a common feature representation shared by heterogeneous domains.

3 Hierarchical Generative Adversarial Networks

3.1 Model

Given a source domain with a set of labeled images and a target domain with a set of unlabeled videos that share the same categories, our goal is to learn a common feature representation shared by source and target domains, on which the classifier built from the source domain adapts well to the target domain. Since a video is composed of a sequence of frames, we assume that a short video clip has a relationship with any frames in it, which provides a natural correlation between frame features and video features. On the other hand, video frames are a collection of images, which could be easily adapted to the source image domain. We explore the image-to-video adaptation by leveraging the relationship among source images, target video clips and their corresponding frames.

Image-Frame Feature. Let $\mathcal{D}_s = \{\mathbf{x}_s^i, y_s^i\}_{i=1}^{n_s}$ represent a labeled dataset of n_s images from the source domain and $\mathcal{D}_t = \{\mathbf{x}_t^j\}_{j=1}^{n_t}$ denote an unlabeled dataset of n_t videos from the target domain. We divide each video into several clips with the same length, building up an unlabeled video clip domain $\mathcal{D}_v = \{\mathbf{x}_v^k\}_{k=1}^{n_v}$, where n_v is the number of video clips in total. For each clip, a frame is randomly selected and all the selected frames compose an unlabeled video frame domain $\mathcal{D}_f = \{\mathbf{x}_f^k\}_{k=1}^{n_v}$. Since both of \mathcal{D}_s and \mathcal{D}_f are image collections, we adopt a state-of-the-art deep domain adaptation model (i.e. JAN [Long *et al.*, 2017]) to learn the common feature (called image-frame feature) of source images and target frames. In the common image-frame feature space, the source images and the target frames are represented by $\mathbf{H}_s = [\mathbf{h}_s^1, \mathbf{h}_s^2, \dots, \mathbf{h}_s^{n_s}] \in \mathbb{R}^{n_s \times d_h}$ and $\mathbf{H}_f = [\mathbf{h}_f^1, \mathbf{h}_f^2, \dots, \mathbf{h}_f^{n_v}] \in \mathbb{R}^{n_v \times d_h}$, respectively, where d_h is the feature dimension.

Frame-to-Video Mapping. For the \mathcal{D}_v , we extract the C3D feature [Tran *et al.*, 2015] to describe the target video clips as $\mathbf{V} = [\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^{n_v}] \in \mathbb{R}^{n_v \times d_v}$, where d_v is the feature dimension. For the \mathcal{D}_f , we employ the ResNet [He *et al.*, 2016] to extract deep features of target frames as $\mathbf{F} = [\mathbf{f}^1, \mathbf{f}^2, \dots, \mathbf{f}^{n_v}] \in \mathbb{R}^{n_v \times d_f}$, where d_f is the feature dimension. Generally speaking, frame features and video features are from two heterogeneous feature spaces, which means $d_v \neq d_f$. Considering the correspondence between a video frame and its related video clip, we get a collection of n_v instances of video-frame pairs, denoted as $P = \{p_k\}_{k=1}^{n_v}$ with $p_k = (\mathbf{v}^k, \mathbf{f}^k)$. We aim at learning a mapping function from video frames to their related video clips. Specifically, the frame features \mathbf{F} are projected to video clip features \mathbf{V} as $\mathbf{V}_f = G_l(\mathbf{F}; \theta_{G_l})$, where $G_l(\cdot; \theta_{G_l})$ is the mapping function and $\mathbf{V}_f = [\mathbf{v}_f^1, \mathbf{v}_f^2, \dots, \mathbf{v}_f^{n_v}] \in \mathbb{R}^{n_v \times d_v}$ indicates the generated video feature from \mathbf{F} .

Video-to-Image Mapping. For each frame \mathbf{x}_f^k , it can be represented by two different features $\mathbf{v}_f^k \in \mathbf{V}_f$ and $\mathbf{h}_f^k \in$

\mathbf{H}_f , thus we have a collection of n_v two heterogeneous features pairs, expressed as $Q = \{q_k\}_{k=1}^{n_v}$ with $q_k = (\mathbf{h}_f^k, \mathbf{v}_f^k)$. Given Q , a projection $\mathbf{H}'_f = G_h(\mathbf{V}_f; \theta_{G_h})$ from \mathbf{V}_f to \mathbf{H}_f can be learned. Here, $G_h(\cdot; \theta_{G_h})$ is the mapping function and \mathbf{H}'_f is the generated image-frame feature from \mathbf{V}_f . Note that \mathbf{H}'_f and \mathbf{H}_s share the same feature space learned by JAN model. In fact, \mathbf{V}_f and \mathbf{V} also come from the same video feature space. Consequently, $G_h(\cdot; \theta_{G_h})$ can be actually considered as the projection from videos to images, and the video clip features \mathbf{V} can be projected to the image-frame space as $\mathbf{H}_v = G_h(\mathbf{V}; \theta_{G_h})$, where $\mathbf{H}_v = [\mathbf{h}_v^1, \mathbf{h}_v^2, \dots, \mathbf{h}_v^{n_v}] \in \mathbb{R}^{n_v \times d_h}$ indicates the generated image-frame feature from \mathbf{V} . Followed by averaging the clip features of each video, we obtain the final features for videos, denoted as $\mathbf{H}_t = [\mathbf{h}_t^1, \mathbf{h}_t^2, \dots, \mathbf{h}_t^{n_t}] \in \mathbb{R}^{n_t \times d_h}$, which can be directly compared against source image features \mathbf{H}_s .

Motivated by the success of GANs in multiple fields, we propose a HiGAN model where the *low-level* GAN and the *high-level* GAN are designed to learn the mapping functions G_l and G_h , respectively. Our whole objective includes two terms. One is the adversarial loss [Goodfellow *et al.*, 2014] for matching the distribution of generated features to the data distribution in the original domain. The other is the CORAL loss [Sun and Saenko, 2016] for minimizing the difference in second-order statistics between the synthesized features and the original features.

3.2 Loss

Adversarial Loss. For the mapping function G_l associated with the discriminator D_l , we denote the data distribution as $\mathbf{V} \sim P_{data}(\mathbf{V})$ and $\mathbf{F} \sim P_{data}(\mathbf{F})$, then the objective is formulated as:

$$\mathcal{L}_{GAN}(D_l, G_l, \mathbf{F}, \mathbf{V}) = E_{\mathbf{V} \sim P_{data}(\mathbf{V})}[\log D_l(\mathbf{V}|\mathbf{F})] + E_{\mathbf{F} \sim P_{data}(\mathbf{F})}[\log(1 - D(G_l(\mathbf{F})|\mathbf{F}))], \quad (1)$$

where G_l attempts to generate video features $G_l(\mathbf{F})$ that resemble the video features from \mathbf{V} , while D_l tries to distinguish between generated features $G_l(\mathbf{F})$, and ground-truth features \mathbf{V} . In other words, G_l aims at minimizing this objective against an adversary D_l that tries to maximize it, formulated by $\min_{G_l} \max_{D_l} \mathcal{L}_{GAN}(D_l, G_l, \mathbf{F}, \mathbf{V})$. We introduce a similar adversarial loss for G_h and D_h , given by $\min_{G_h} \max_{D_h} \mathcal{L}_{GAN}(D_h, G_h, \mathbf{V}_f, \mathbf{H}_f)$.

CORAL Loss. It has been proven that the GAN objective combined with another loss could produce significant good results [Zhu *et al.*, 2017]. In our method, we introduce the CORAL loss [Sun and Saenko, 2016] to minimize the difference in second-order statistics between the generated features and the real features. The CORAL Loss is simple, effective and can be easily integrated into a deep learning architecture.

For the generator G_l associated with D_l , \mathbf{V} is the real features and \mathbf{V}_f is the synthesized features. Suppose \mathbf{E}_v and \mathbf{E}_{v_f} denote the feature covariance matrices. The CORAL loss as the distance between the second-order statistics (covariance) of \mathbf{V} and \mathbf{V}_f is expressed as follows:

$$\mathcal{L}_{CORAL}(\mathbf{V}, \mathbf{V}_f) = \frac{1}{4d_v^2} \|\mathbf{E}_V - \mathbf{E}_{V_f}\|_F^2, \quad (2)$$

where $\|\cdot\|_F^2$ denotes the squared matrix Frobenius norm. The covariance matrices \mathbf{E}_V and \mathbf{E}_{V_f} are calculated by

$$\mathbf{E}_V = \frac{1}{n_v - 1} (\mathbf{V}^T \mathbf{V} - \frac{1}{n_v} (\mathbf{1}^T \mathbf{V})^T (\mathbf{1}^T \mathbf{V})), \quad (3)$$

$$\mathbf{E}_{V_f} = \frac{1}{n_v - 1} (\mathbf{V}_f^T \mathbf{V}_f - \frac{1}{n_v} (\mathbf{1}^T \mathbf{V}_f)^T (\mathbf{1}^T \mathbf{V}_f)), \quad (4)$$

where $\mathbf{1}$ is a column vector with all elements equal to 1.

For the generator G_h associated with D_h , we introduce a similar CORAL loss as well:

$$\mathcal{L}_{CORAL}(\mathbf{H}_f, \mathbf{H}'_f) = \frac{1}{4d_h^2} \|\mathbf{E}_{H_f} - \mathbf{E}_{H'_f}\|_F^2. \quad (5)$$

3.3 Objective

In order to prevent the learned parameters from overfitting, we introduce the regularization term:

$$\mathcal{L}_{reg}(D_l, G_l) = \sum_{l_D=1}^{L_D} \|W_{D_l}^{l_D}\|_F + \sum_{l_G=1}^{L_G} \|W_{G_l}^{l_G}\|_F, \quad (6)$$

$$\mathcal{L}_{reg}(D_h, G_h) = \sum_{l_D=1}^{L_D} \|W_{D_h}^{l_D}\|_F + \sum_{l_G=1}^{L_G} \|W_{G_h}^{l_G}\|_F, \quad (7)$$

where $W_{D_l}^{l_D}$, $W_{G_l}^{l_G}$, $W_{D_h}^{l_D}$ and $W_{G_h}^{l_G}$ represent the layer-wise parameters of networks. L_D and L_G denote layer numbers of discriminator and generator, respectively. Based above, the whole objective becomes

$$\begin{aligned} \mathcal{L}(D_l, G_l) &= \lambda_1 \mathcal{L}_{cGAN}(D_l, G_l, \mathbf{F}, \mathbf{V}) \\ &+ \lambda_2 \mathcal{L}_{CORAL}(\mathbf{V}, \mathbf{V}_f) \\ &+ \mathcal{L}_{reg}(D_l, G_l), \end{aligned} \quad (8)$$

$$\begin{aligned} \mathcal{L}(D_h, G_h) &= \lambda_3 \mathcal{L}_{cGAN}(D_h, G_h, \mathbf{V}_f, \mathbf{H}_f) \\ &+ \lambda_4 \mathcal{L}_{CORAL}(\mathbf{H}_f, \mathbf{H}'_f) \\ &+ \mathcal{L}_{reg}(D_h, G_h), \end{aligned} \quad (9)$$

where λ_1 , λ_2 , λ_3 and λ_4 are weight parameters that control the relative importance of adversarial loss and CORAL loss, respectively.

3.4 Algorithm

To learn the optimal feature representations, the generator G and the discriminator D compete with each other in a two-player minimax game, jointly minimizing the adversarial loss in Eq. (1) and the CORAL loss in Eq. (2). We aim to solve:

$$D_l^*, G_l^* = \arg \min_{G_l} \max_{D_l} \mathcal{L}(D_l, G_l), \quad (10)$$

$$D_h^*, G_h^* = \arg \min_{G_h} \max_{D_h} \mathcal{L}(D_h, G_h). \quad (11)$$

The algorithm of HiGAN is summarized in Algorithm 1.

Algorithm 1: Hierarchical Generative Adversarial Networks

Input: The source image domain \mathcal{D}_s , the target video domain \mathcal{D}_t , the video clip domain \mathcal{D}_v , the video frame domain \mathcal{D}_f , the deep features \mathbf{V} of \mathcal{D}_v , the deep features \mathbf{F} of \mathcal{D}_f

Output: The deep features \mathbf{H}_s of \mathcal{D}_s and \mathbf{H}_t of \mathcal{D}_t

- 1 Train the JAN model with \mathcal{D}_s and \mathcal{D}_f to learn features \mathbf{H}_s and \mathbf{H}_f of \mathcal{D}_s and \mathcal{D}_f , respectively.
 - 2 Learn the mapping function $G_l(\cdot; \theta_{G_l})$ with deep features \mathbf{F} and \mathbf{V} via Eq. (8), and calculate the new deep feature V_f of \mathcal{D}_f .
 - 3 Learn the mapping function $G_h(\cdot; \theta_{G_h})$ with deep features V_f and \mathbf{H}_f by Eq. (9), and calculate the new deep feature \mathbf{H}_v of \mathcal{D}_v .
 - 4 Average the clip features \mathbf{H}_v of each video to obtain the features \mathbf{H}_t of \mathcal{D}_t .
 - 5 Return \mathbf{H}_s and \mathbf{H}_t .
-

4 Experiments

4.1 Datasets

To evaluate the performance of our method, we conduct the experiments on two complex video datasets, i.e., UCF101 [Soomro *et al.*, 2012] and HMDB51 [Yao *et al.*, 2011]. For the UCF101 as the target video domain, the source images come from the Stanford40 dataset [Yao *et al.*, 2011]. For the HMDB51 as the target video domain, the source image domain consists of Stanford40 dataset and HII dataset [Tanisik *et al.*, 2016], denoted by EADs dataset.

Stanford40 and UCF101 (S→U): The UCF101 is a dataset of action videos collected from YouTube with 101 action categories. The Stanford40 dataset include diverse action images with 40 action categories. We choose 12 common categories between these two datasets.

EADs and HMDB51 (E→H): The HMDB51 dataset has 51 classes, containing 6766 video sequences. The EADs dataset consists of Stanford40 and HII datasets. The HII dataset has 10 action images and each class contains at least 150 images, forming a total of 1972 images. The 13 shared categories between the EDAs and the HMDB51 datasets are adopted in our experiment.

4.2 Experiment Setup

We split each target video into 16-frame clips without overlap, and all the clips from all the target videos construct the video-clip domain. The deep feature of each video clip is the 512D feature vector from the pool5 layer of 3D CveNets [Tran *et al.*, 2015] which are trained on a large-scale video dataset. For each video clip, we randomly sample one frame and all the frames from all the video clips compose the video-frame domain. The deep feature of frames is the 2048D vector extracted from the pool5 layer of ResNet [He *et al.*, 2016]. We utilize the JAN [Long *et al.*, 2017] method based on the ResNet to get the domain-invariant image-frame features between source images and target video frames from the pool5 layer of JAN with the dimension of 2048D.

Implementation details. To model the two generators in HiGAN, we deploy four-layered feed-forward neural networks activated by *relu* function, (i.e., $2048 \rightarrow 1024 \rightarrow 1024 \rightarrow 1024 \rightarrow 512$ for $G_l(f; \theta_{G_l})$ and $512 \rightarrow 1024 \rightarrow 1024 \rightarrow 2048 \rightarrow 2048$ for $G_h(v_f; \theta_{G_h})$). In terms of two discriminators, we both utilize three fully connected layers ($2560 \rightarrow 1280 \rightarrow 640 \rightarrow 1$) activated by *relu* function, except for the last layer.

In practice, we replace the negative log likelihood objective by a least square loss [Mao *et al.*, 2016], which performs more stably during training and generates higher quality results.

Since the adversarial losses and the CORAL losses have different orders of magnitude according to the experiments, and the *low-level* conditional GAN and the *high-level* one have very similar network structures, we set $\lambda_2 = \lambda_4 = 100$, $\lambda_1 = \lambda_3 = 1$ in Eq. (8) and Eq. (9) for all the experiments, to somehow balance the two types of losses. We employ the Adam solver [Kingma and Ba, 2014] with a batch size of 64. All the networks were trained from scratch with the learning rate of 0.00002 for the *low-level* conditional GAN and 0.000008 for *high-level* conditional GAN.

Related methods. We evaluate the effectiveness of our approach by investigating whether our generated features of source images and target videos are more transferable than other domain adaptation methods, that is, whether the recognition accuracy on the target videos could be improved. Specifically, the compared methods are: 1) MKL [Xu *et al.*, 2013]; TJM [Long *et al.*, 2014]; TKL [Long *et al.*, 2015b]; CORAL [Sun *et al.*, 2016]; LRSR [Xu *et al.*, 2016]; ATI [Busto and Gall, 2017]; KCCA [Akaho, 2006]; HEMAP [Shi *et al.*, 2010]; DAMA [Wang and Mahadevan, 2011]; HFA [Duan *et al.*, 2012b]; CDLS [Tsai *et al.*, 2016], which are traditional shallow domain adaptation methods, and 2) ResNet [He *et al.*, 2016]; DAN [Long *et al.*, 2015a]; RTN [Long *et al.*, 2016]; JAN [Long *et al.*, 2017]; DAL [Carlucci *et al.*, 2017], which are deep domain adaptation methods. The MKL is considered as a traditional baseline method which directly uses source classifiers in the target domain. Similarly, the ResNet is taken as a deep baseline method. As fair comparison with identical evaluation setting, the image-frame (i.e., JAN) features of source images are taken as source features in the traditional shallow domain adaptation methods. For all the methods, we use the accuracy for performance evaluation.

Note that MKL, TJM, TKL, CORAL, LRSR and ATI can only handle the homogeneous domain adaptation problem, when the data from the source and target domains are with the same type of feature. Therefore, for these methods, each target video is represented by the mean of the image-frame features of all the frames.

For the traditional heterogeneous domain adaptation methods of KCCA, HEMAP, DAMA, HFA and CDLS, the C3D feature is adopted to describe the target videos. Considering that all these methods require the labeled training data in the target domain, we randomly choose three videos per class as the labeled target data and take the rest videos as the test data. For each dataset, we repeat the sampling for 5 times and re-

Method	S→U	E→H
MKL [Xu <i>et al.</i> , 2013]	0.889	0.392
TJM [Long <i>et al.</i> , 2014]	0.898	0.314
TKL [Long <i>et al.</i> , 2015b]	0.902	0.391
CORAL [Sun <i>et al.</i> , 2016]	0.904	0.398
LRSR [Xu <i>et al.</i> , 2016]	0.893	0.380
ATI [Busto and Gall, 2017]	0.905	0.330
HiGAN (no labeled video)	0.954	0.446
KCCA [Akaho, 2006]	0.853	0.366
HEMAP [Shi <i>et al.</i> , 2010]	0.432	0.213
DAMA [Wang and Mahadevan, 2011]	0.860	0.438
HFA [Duan <i>et al.</i> , 2012b]	0.885	0.459
CDLS [Tsai <i>et al.</i> , 2016]	0.885	0.435
HiGAN (three labeled videos)	0.967	0.526

Table 1: Comparison between our method and the traditional shallow domain adaptation methods.

Method	S→U	E→H
ResNet [He <i>et al.</i> , 2016]	0.814	0.385
DAN [Long <i>et al.</i> , 2015a]	0.842	0.395
RTN [Long <i>et al.</i> , 2016]	0.838	0.402
JAN [Long <i>et al.</i> , 2017]	0.914	0.409
DAL [Carlucci <i>et al.</i> , 2017]	0.826	0.426
HiGAN	0.954	0.446

Table 2: Comparison between our method and the deep domain adaptation methods.

port the average results.

Regarding the deep methods of ResNet, DAN, RTN, JAN, DAL, since they only take images as input, the source images and target video frames are utilized to train the networks. The output scores (from the last *fc* layer) of all the frames are further averaged to determine the class label of the video. All the deep methods are implemented based on the Caffe framework [Jia *et al.*, 2014] and fine-tuned from Caffe-provided models of ResNet, which are pre-trained on the ImageNet 2012 dataset. We fine-tune all convolutional and pooling layers and train the classifier layer via back propagation.

4.3 Results

Comparison with existing domain adaptation methods.

Table 1 and Table 2 show the recognition accuracies of traditional shallow domain adaptation methods and deep domain adaptation methods, respectively. From them, we can notice that our method achieves the best performance compared with other state-of-the-art methods for both datasets, which explicitly demonstrates the effectiveness of transferring knowledge from images to videos for video recognition. The more detailed observations are as follows.

(1) When videos are represented by C3D features which are different from source image features, the traditional heterogeneous methods substantially perform worse than the traditional homogeneous methods on the S→U dataset, even though they have some labeled videos in the target domain. This shows that it is a quite challenging task to conduct domain adaptation between images and videos when they are respectively represented by different features. Compared with

Method	S→U	E→H
HiGAN_CORAL	0.726	0.335
HiGAN_Adversarial	0.921	0.439
HiGAN	0.954	0.446

Table 3: Ablation study: classification performance of S→U and E→H for different losses

them, our proposed HiGAN can achieve better performance owing to the strengths of taking video frames as a bridge to obtain the common feature shared by heterogeneous images and videos.

(2) When there are no labeled data in the target domain, both traditional methods and deep methods are better than the baseline methods (i.e., MKL and ResNet) on the S→U dataset, which indicates that these adaptation methods can improve action recognition performance, but the improvement is not significant. On the E→H dataset, traditional methods slightly underperform the baseline MKL. A possible explanation is that the huge difference between source images and target video frames on the E→H dataset leads to the negative transfer. On the other hand, deep transfer learning methods are better than the baseline ResNet, which verifies that deep networks are better at addressing negative transfer issue. Compared with them, our proposed HiGAN gains performance improvement on both datasets, which demonstrates that our method can attain respectable performances even though there is a huge difference between source images and target videos.

Analysis of the loss function. In Table 3, we assess the individual effect of the adversarial loss and the CORAL loss in HiGAN, with respect to the recognition accuracies in two datasets. We develop two variants of our HiGAN: one with the adversarial loss but without the CORAL loss, the other with the CORAL loss but without the adversarial loss. The two variants are referred to as HiGAN_Adversarial and HiGAN_CORAL, respectively. Removing the adversarial loss substantially degrades results, as does removing the CORAL loss. We therefore conclude that both terms are critical to our results.

Effectiveness of the HiGAN features for traditional domain adaptation. To validate the performance improvement of the proposed HiGAN features for video recognition, we compare it with the original features (i.e., JAN and C3D) used in the traditional shallow domain adaptation methods. Additional experiments are conducted with different features. Figure 2 (a) and (b) demonstrate the accuracy improvements on the S→U and E→H datasets, respectively. We can observe that, the recognition result could be improved with our proposed HiGAN features for both datasets.

Effectiveness of exploring images by HiGAN for video recognition. We also explore whether it is beneficial to transfer knowledge from images for video recognition by comparing our HiGAN features and the C3D features on the

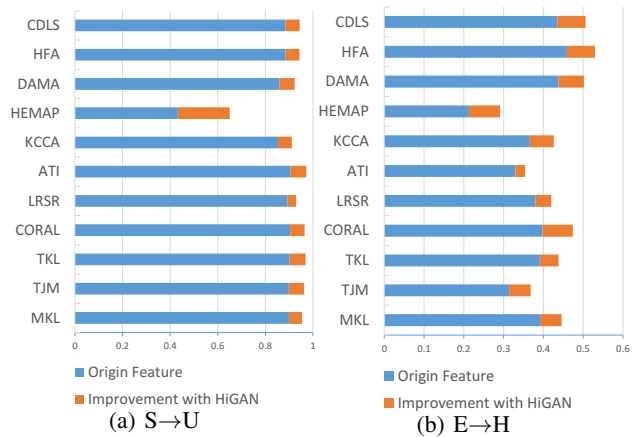


Figure 2: The accuracy improvement in the traditional shallow domain adaptation methods with our HiGAN features on both the S→U and E→H datasets. The blue bars show the accuracy of traditional methods trained on JAN and C3D features. The red bars show the absolute increase in accuracy of the traditional methods trained using our HiGAN features.

Method	UCF101	HMDB51
C3D [Tran <i>et al.</i> , 2015]	0.940	0.657
HiGAN	0.980	0.740

Table 4: Comparison between our method and C3D on the video recognition.

UCF101 and the HMDB51 datasets with the chosen categories. The experiments are separately done on the three train/test splits, and the result is averaged across three test splits. From table 4, it is interesting to notice that with the same training samples, our method can significantly improve the video recognition accuracy on both datasets, which validates the benefits of exploring related images for video recognition.

5 Conclusion

We have proposed a new *Hierarchical Generative Adversarial Networks* approach to transfer knowledge from images to videos. By taking advantage of these two-level adversarial learning with the CORAL loss, our method is capable of learning a domain-invariant feature representation between source images and target videos. Thus the image classifiers trained on the domain-invariant features can be effectively adapted to videos. We conduct experiments on two datasets and the results validate the effectiveness of our method.

The future work includes exploiting large-scale Web images for video recognition which will further improve the recognition accuracy. We also plan to extend our method to multiple source domain adaptation.

Acknowledgments

This work was supported in part by the Natural Science Foundation of China (NSFC) under grant No. 61673062 and No.

61472038.

References

- [Akaho, 2006] S Akaho. A kernel method for canonical correlation analysis. *arXiv preprint cs/0609071*, 2006.
- [Bousmalis *et al.*, 2017] K Bousmalis, N Silberman, D Dohan, D Erhan, and D Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*, 2017.
- [Busto and Gall, 2017] PP Busto and J Gall. Open set domain adaptation. In *ICCV*, 2017.
- [Carlucci *et al.*, 2017] FM Carlucci, L Porzi, B Caputo, E Ricci, and SR Bulò. Autodial: Automatic domain alignment layers. In *ICCV*, 2017.
- [Chen *et al.*, 2016] WY Chen, TM Hsu, YH Tsai, YC Wang, and MS Chen. Transfer neural trees for heterogeneous domain adaptation. In *ECCV*, pages 399–414, 2016.
- [Duan *et al.*, 2012a] L Duan, D Xu, and SF Chang. Exploiting web images for event recognition in consumer videos: A multiple source domain adaptation approach. In *CVPR*, pages 1338–1345, 2012.
- [Duan *et al.*, 2012b] L Duan, D Xu, and I Tsang. Learning with augmented features for heterogeneous domain adaptation. *arXiv preprint arXiv:1206.4660*, 2012.
- [Gan *et al.*, 2016] C Gan, C Sun, L Duan, and B Gong. Webly-supervised video recognition by mutually voting for relevant web images and web video frames. In *ECCV*, pages 849–866, 2016.
- [Goodfellow *et al.*, 2014] I Goodfellow, J Pouget-Abadie, M Mirza, B Xu, D Warde-Farley, S Ozair, A Courville, and Y Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.
- [He *et al.*, 2016] K He, X Zhang, S Ren, and J Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [Jia *et al.*, 2014] Y Jia, E Shelhamer, J Donahue, S Karayev, J Long, R Girshick, S Guadarrama, and T Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM MM*, pages 675–678, 2014.
- [Kingma and Ba, 2014] D Kingma and J Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Long *et al.*, 2014] M Long, J Wang, G Ding, J Sun, and PS Yu. Transfer joint matching for unsupervised domain adaptation. In *CVPR*, pages 1410–1417, 2014.
- [Long *et al.*, 2015a] M Long, Y Cao, J Wang, and M Jordan. Learning transferable features with deep adaptation networks. In *ICML*, pages 97–105, 2015.
- [Long *et al.*, 2015b] M Long, J Wang, J Sun, and S Yu Philip. Domain invariant transfer kernel learning. *TKDE*, 27(6):1519–1532, 2015.
- [Long *et al.*, 2016] M Long, H Zhu, J Wang, and MI Jordan. Unsupervised domain adaptation with residual transfer networks. In *NIPS*, pages 136–144, 2016.
- [Long *et al.*, 2017] M Long, H Zhu, J Wang, and MI Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, pages 2208–2217, 2017.
- [Mao *et al.*, 2016] X Mao, Q Li, H Xie, R YK Lau, and Z Wang. Multi-class generative adversarial networks with the l2 loss function. *arXiv preprint arXiv:1611.04076*, 2016.
- [Sankaranarayanan *et al.*, 2017] S Sankaranarayanan, Y Balaji, CD Castillo, and R Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. *arXiv preprint arXiv:1704.01705*, 2017.
- [Shi *et al.*, 2010] X Shi, Q Liu, W Fan, SY Philip, and R Zhu. Transfer learning on heterogeneous feature spaces via spectral transformation. In *ICDM*, pages 1049–1054. IEEE, 2010.
- [Soomro *et al.*, 2012] K Soomro, A Roshan Zamir, and M Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [Sun and Saenko, 2016] B Sun and K Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV 2016 Workshops*, 2016.
- [Sun *et al.*, 2016] B Sun, J Feng, and K Saenko. Return of frustratingly easy domain adaptation. In *AAAI*, volume 6, page 8, 2016.
- [Tanisik *et al.*, 2016] G Tanisik, C Zalluhoglu, and N Iklizler-Cinbis. Facial descriptors for human interaction recognition in still images. *PR*, 73:44–51, 2016.
- [Tran *et al.*, 2015] D Tran, L Bourdev, R Fergus, L Torresani, and M Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015.
- [Tsai *et al.*, 2016] YH Tsai, YR Yeh, and YC Wang. Learning cross-domain landmarks for heterogeneous domain adaptation. In *CVPR*, pages 5081–5090, 2016.
- [Wang and Mahadevan, 2011] C Wang and S Mahadevan. Heterogeneous domain adaptation using manifold alignment. In *IJCAI*, volume 22, page 1541, 2011.
- [Xu *et al.*, 2013] X Xu, IW Tsang, and D Xu. Soft margin multiple kernel learning. *TNNLS*, 24(5):749–761, 2013.
- [Xu *et al.*, 2016] Y Xu, X Fang, J Wu, X Li, and D Zhang. Discriminative transfer subspace learning via low-rank and sparse representation. *TIP*, 25(2):850–863, 2016.
- [Yao *et al.*, 2011] B Yao, X Jiang, A Khosla, AL Lin, L Guibas, and Li FF. Human action recognition by learning bases of action attributes and parts. In *ICCV*, pages 1331–1338, 2011.
- [Zhang *et al.*, 2017] J Zhang, Y Han, J Tang, Q Hu, and J Jiang. Semi-supervised image-to-video adaptation for video action recognition. *TCYB*, 47(4):960–973, 2017.
- [Zhu *et al.*, 2017] JY Zhu, T Park, P Isola, and AA. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.