

Cross-Identity Motion Transfer for Arbitrary Objects through Pose-Attentive Video Reassembling

Subin Jeon¹[0000-0003-1651-2249], Seonghyeon Nam¹[0000-0003-1266-3568],
Seoung Wug Oh¹[0000-0002-8498-0864], and
Seon Joo Kim^{1,2}[0000-0001-8512-216X]

¹ Yonsei University

{subinjeon, shnam, sw.oh, seonjookim}@yonsei.ac.kr

² Facebook

Abstract. We propose an attention-based networks for transferring motions between arbitrary objects. Given a source image(s) and a driving video, our networks animate the subject in the source images according to the motion in the driving video. In our attention mechanism, dense similarities between the learned keypoints in the source and the driving images are computed in order to retrieve the appearance information from the source images. Taking a different approach from the well-studied warping based models, our attention-based model has several advantages. By reassembling non-locally searched pieces from the source contents, our approach can produce more realistic outputs. Furthermore, our system can make use of multiple observations of the source appearance (e.g. front and sides of faces) to make the results more accurate. To reduce the training-testing discrepancy of the self-supervised learning, a novel cross-identity training scheme is additionally introduced. With the training scheme, our networks is trained to transfer motions between different subjects, as in the real testing scenario. Experimental results validate that our method produces visually pleasing results in various object domains, showing better performances compared to previous works.

Keywords: Motion transfer, Generative Adversarial Network, Video to video translation

1 Introduction

Motion transfer is a task of transferring motion between different subjects. In other words, it generates a video conditioned on the appearance extracted from source image(s) and the motion patterns from a driving video. Here, the appearance means an identity of a subject observed in source image(s) and the motion patterns refer to a sequence of poses that change continuously. Motion transfer has long been studied for its practical applications such as video editing and

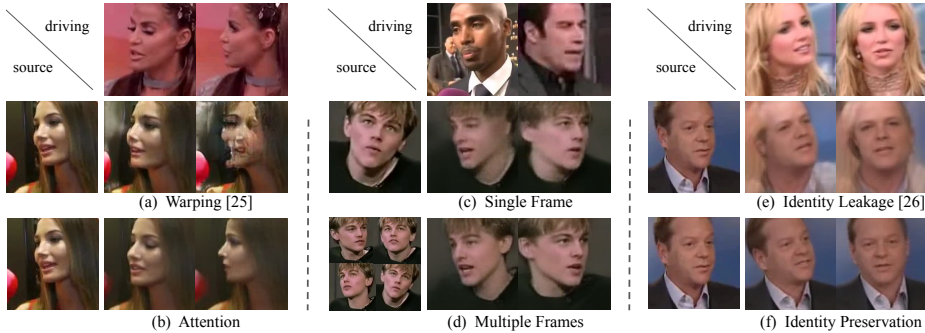


Fig. 1. (a) Previous motion transfer methods is based on spatial warping of single image. (b) We propose a non-local search based approach that handles a wide range of motion. (c) A single source image contains insufficient appearance data. (d) Our method can utilize multiple sources with various views. (e) Self-supervision causes identity-leakage in motion transfer. (f) We propose a cross-identity training which is effective in preserving identities of a source image.

virtual/augmented reality. Recent progress in deep generative model has shown the capability of generating highly realistic images, and it has led to vigorous attempts to transfer motions in learning basis [34, 3, 6, 35, 25, 32].

The majority of previous learning-based approaches relies on pre-computed (or given) landmark annotations to represent poses [34, 6, 19, 28]. With an accurate pose representation, those methods have achieved impressive results especially in the human face and body domains where the keypoint extraction is well-studied. However, the dependency on pre-computed keypoints limits their application. For example, [19, 28] are only applicable to specific domains (e.g. face and body), and [34, 6] work only for specific subjects.

To overcome the above issue, several attempts have been made to transfer object motions without using pre-defined pose representations [25, 35, 26] recently. They loosened the constraint on requiring the pre-defined keypoints by employing the self-supervised learning scheme for the pose representation [9], enabling applications beyond specific object categories. To learn to extract representations and transfer motion simultaneously, previous methods commonly take warping-based approach to synthesize images in various poses from a source image. However, we argue that the warping operation is not the best option for this task due to several limitations. In this paper, we highlight those limitations and propose a novel attention-based approach for arbitrary motion transfer.

While the warping operation works well for handling small motions, it has some challenges in achieving realistic synthesis. First, it is difficult to model large and complex motions. For example, it is more difficult to create a side view than to raise head slightly from the source image with a front view, as illustrated in Fig. 1 (a). Second, it is hard to make use of more than one source image. Most of previous works cannot use more than one source image by its design [19, 28,

25], and some methods rely on heuristics to combine results obtained with several source images [35]. Observations from different poses and views are essential in completely recovering occluded parts in various views. As shown in Fig. 1 (d), using multiple sources is beneficial in synthesizing various views.

We propose a novel attention-based method to transfer a variety of motion from arbitrary object categories. The main idea of our method is to reassemble the source image by finding proper visual appearance in the source image for each part of the driving pose. Compared with the warping-based approaches, our method produces more realistic output as it non-locally searches the best source content to synthesize every parts of an image.

Furthermore, our framework is flexible in the number of source images. While our approach already shows better results than the previous approaches with a single source image, we can maximize the performance by providing additional observations. Taking the advantage of multiple source images, we can potentially resolve the occlusion issue in the single observation case. Previous warping-based methods are designed only to use single source image and they focus on synthesizing unseen parts or different views with the limited appearance information. We let our network leverage multiple source images by broadening the comparison region to multiple source images and it shows the capability to utilize appearance in various views and poses (Fig. 1(c),(d)).

In addition, we propose a cross-identity training scheme that enables realistic motion transfer between subjects with significantly different appearances. Previously, motion transfer networks are usually trained by transferring poses in the same subject even though the actual task is to transfer poses between different subjects. While this training trick is required to learn the networks in an unsupervised way, it is not appropriate for the real testing scenario. For example, models trained in this way often fails to maintain the personal shape that is unique for each subject independent from its poses (Fig. 1(e)). With the proposed training scheme, the network can learn to perform the motion transfer between different subjects, as in the real testing, without additional supervision. Furthermore, our cross-identity reconstruction loss encourages the network to preserve the identity-specific property (e.g. the shape of a human face) after the motion transfer.

To sum up, our contribution can be summarized as follows:

- We propose a pose-attentive video reassembling network for motion transfer of arbitrary objects. Our network implicitly learns the pose representation specialized for transferring motion without annotation. Through the non-local search based reassembling of visual pieces, we overcome the limitation of warping based motion transfers.
- Our framework can naturally process and take advantage of multiple images or video as the source, which helps to synthesize more realistic outputs.
- We propose a cross-identity training method to preserve identity-relevant properties such as object shape while changing the pose. This training strategy matches the testing scenarios, therefore, is more realistic compared to previous works that train on same subjects.

2 Related work

Remarkable achievements in the generative modelling [12] have given rise to interests in video generation, leading to extensive development of video generation methods [31, 24, 29, 30, 37, 8]. As a type of video generation, motion transfer refers to the task of synthesizing motion extracted from a driving video on different subjects. In this section, we address major approaches of motion transfer.

Video-to-video Translation Motion transfer has been addressed from the perspective of video-to-video translation. This class of works [34, 3, 6] proposed methods to learn a mapping between two videos in order to convert a domain of videos into another. In those frameworks, the generator memorizes diverse aspects of the source appearance using a sequence of frames in the source video and learns to transfer the appearance on the driving video. These methods have the advantage of utilizing the rich appearance information of diverse pose and view in the source video, making realistic results. However, they require time-consuming training per each subject using a large amount of source images related to the subject to train the subject-specific generator.

Pose Guided Image Generation This class refers to the image generation conditioned on a source image and a driving pose. In order to represent the pose to condition on, they use keypoint annotations or off-the-shelf keypoint detector. Using the keypoint annotations and a single source image, many works [19, 20, 27, 2, 17, 28, 10] attempted to generate the source subject in a novel pose by warping a single source image. This approach has strength in transferring motion on new subjects at the test time, but experiences difficulties caused by insufficient appearance information in the single source image. Recently, multi-source based methods [36, 32, 18, 13] have emerged to overcome the drawback of a single source image. These works are similar to our work in that they can transfer motion on new subjects with multiple source images. However, we tackle more challenging case because we need to extract pose from images directly and generalize to various object categories, whereas this stream of works is restricted to a specific object category (human body or face).

Motion transfer for arbitrary objects This class of works learns to extract pose representation implicitly from driving frames to transfer motion, extending object scopes to arbitrary objects. X2Face [35] and Monkey-net [25] transfer motion by learning to warp a source image to pose of a driving image in unsupervised way. These warping-based method showed superiority in transferring local motion, but experienced difficulties in synthesizing motion when the movement becomes larger and more complex. To solve the large motion problem, Siarohin [26] et al. enhanced Monkey-net by taking the local movement into consideration using local affine transformation of keypoints. Our method also falls in to this category of transferring motion on arbitrary objects. Different from previous works, we propose a non-local motion matching mechanism to capture drastic motion changes in order to overcome the limitation of previous warping-based approaches.

3 Method

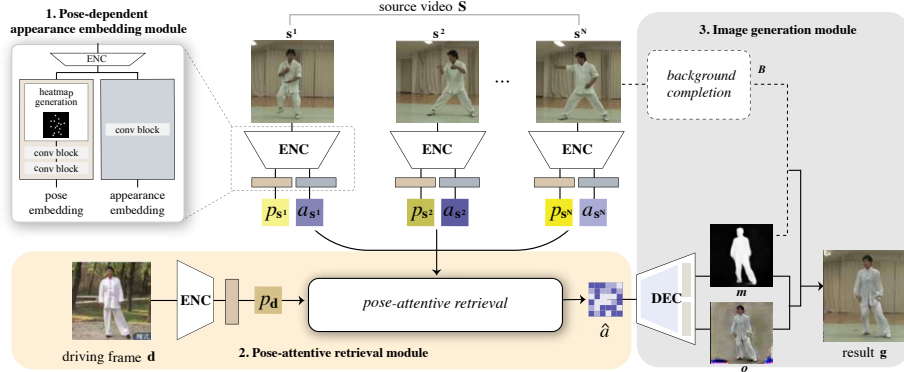


Fig. 2. Overview of our work. It consists of (1) pose-dependent appearance embedding module, (2) pose-attentive retrieval module, and (3) image generation module. The pose-dependent appearance embedding module takes frames in the source video to embed (pose, appearance) pairs. Pose attention block retrieves appearance correspondent to the driving pose. The driving pose is extracted using the shared pose embedding network. Finally, image generation module generates the results.

3.1 Overview

Given a source video $\mathbf{S} = \{\mathbf{s}^n\}_{n=1,2,\dots,N}$ and a driving video $\mathbf{D} = \{\mathbf{d}^t\}_{t=1,2,\dots,T}$, our task is to synthesize a video where the motion of the foreground object is similar to that of \mathbf{D} while the appearance is same as that of \mathbf{S} . Note that we deal with both cases of the single source frame $N = 1$ and the multiple source frames $N > 1$.

The overall architecture is illustrated in Fig. 2. Our model consists of following modules: (1) pose-dependent appearance embedding module, (2) pose-attentive retrieval module, and (3) image generation module. The pose-dependent appearance embedding module takes S as input, and extracts the pose and the appearance representations for each frame in S independently, using a shared embedding network. The pose attentive retrieval module extracts the pose from the driving frame, and searches for proper appearance features in the source frames using our spatio-temporal attention mechanism. Finally, our decoder synthesizes outputs using the appearance features and one of the source images for the background.

Our network is trained using a large number of videos containing subjects in the same object category. At each iteration, the network is trained using two videos with different identities by minimizing the self-reconstruction loss and the cross-identity reconstruction loss. After training, it generates output video

frame conditioned on S and each frame d^t in the driving video. In the following, we describe our method in detail.

3.2 Network Architecture

Pose-dependent appearance embedding Given a source video S , we extract poses P_S and their corresponding appearance feature A_S . Each s^n , n -th frame of S , is processed independently through the shared embedding network, as illustrated in the upper side of Fig. 2. It starts by putting a single image s^n into the shared encoder. The features extracted from the encoder, then, goes through two parallel streams, for the pose and the appearance representation, respectively. In the pose stream, it first reduces the number of channels to K with a convolutional layer. Here, K implicitly determines the number of keypoints to be extracted by the network. Each channel is condensed to represent a key point in a spatial domain by using the method proposed in [16]. In order to capture a local relationship between keypoints, the extracted keypoints are processed by a couple of convolutional layers. In case of the appearance stream, the appearance representation is embedded through a convolutional layer from the encoder feature. Each s^n gets a (p_{s^n}, a_{s^n}) pair, resulting in $P_S = \{p_{s^n}\}_{n=1,2,\dots,N}$ and $A_S = \{a_{s^n}\}_{n=1,2,\dots,N}$, which are concatenated in time-dimension respectively.

Pose-attentive retrieval block Given P^S , A^S , and a driving frame d , the goal is to synthesize a figure with the same identity as in the source video but in the pose of the subject in the driving frame. To this end, a computational block that matches poses between frames to retrieve appearance information for the synthesis is designed. We take insights from an attention mechanism in the previous memory networks [22, 23]. However, different from [22, 23] that matches visual features, we explicitly design our module to match pose keypoints with a different motivation.

To match between poses, the driving pose p_d is first extracted through the same pose embedding network used for the source video. Then, similarities between the source poses P_S and the driving pose p_d are computed to determine where to retrieve relevant appearance information from the source appearance A_S . Specifically, soft weights are computed in a non-local manner by comparing every pixel locations in the source pose embedding P_S with every pixels of the driving pose embedding p_d . Finally, the appearance \hat{a} correspondent to p_d is reassembled by taking a weighted summation of the source appearance features A_S with regard to the soft weights. Our appearance retrieval operation can be summarized as:

$$\hat{a}^i = \frac{1}{Z} \sum_{\forall j} f(p_d^i, P_S^j) A_S^j, \quad (1)$$

where $Z = \sum_{\forall j} f(p_d^i, P_S^j)$ is a normalization factor, and i and j indicate the pixel location of the driving and the source pose features, respectively. At first, it computes the attention map using the similarity function f . The similarity function f is defined as follow:

$$f(p_d^i, P_S^j) = \exp(p_d^i \cdot P_S^j), \quad (2)$$




Fig. 3. Self-reconstruction and cross-identity reconstruction training scheme.

where \cdot is dot-product.

Image generation Contrary to the object parts that change according to the driving pose, the background should remain still regardless of the variation of poses. Thus, the image generation module should focus on generating the object parts while the background are directly obtained from the source images. To this end, the decoder produces two outputs by dividing the output branches; one to create the object parts o and the other to create a mask for the object m . m is a single channel probability map obtained by applying the sigmoid function. We compute the background image B by feeding one of source images to a simple encoder-decoder network to inpaint the occluded part. Formally, we get the result by

$$g = m \odot o + (1 - m) \odot B, \quad (3)$$

where \odot represents element-wise multiplication.

3.3 Training and Losses

Our generator is trained using two main losses: self-reconstruction loss and cross-identity reconstruction loss. The distinction between them depends on whether they use the same identity of objects on the source and the driving images, as illustrated in Fig. 3.

Self-reconstruction loss If source images s_A and the driving image d_A have the same identity A , the generated image $g_A^{d_A}$ has the identity of A and the pose from d_A , which is equal to d_A . Therefore the network can be trained by using d_A as the ground-truth image of $g_A^{d_A}$ in a self-supervised way. Note that we split a single video of an identity into source images $s_{identity}$ and driving image $d_{identity}$, and there is no overlap between them.

We impose the reconstruction loss on the pair of $(g_A^{d_A}, d_A)$ to train the overall network. Additionally, we add adversarial loss for the quality of generated image. Thus the self-reconstruction loss is formulated as follows:

$$L_{self} = L_{GAN}(G, D) + \lambda_{FM} L_{FM}(G, D) + \lambda_{VGG} L_{VGG}(g_A^{d_A}, d_A). \quad (4)$$

Here, we give adversarial loss $L_{GAN}(G, D)$ using the least-square GAN formulation [21]. $L_{FM}(G, D)$ represents the discriminator feature-matching loss [33] and $L_{VGG}(g_A^{d_A}, d_A)$ is the perceptual loss using the pretrained VGG19 network.

This objective is commonly used to train previous motion transfer networks in a self-supervised way [35, 25]. However, using this self-reconstruction alone can possibly cause severe issues in the real testing, where subjects in the source and the driving image are different. For example, only working with samples from the same video, the model can simply learn to assume the source and the driving subjects to have the same identity. Thus the model may not be prepared for dealing with subjects that show a significant appearance difference. We address this issue by additionally imposing a novel cross-identity reconstruction loss that simulates the real testing cases, where the identity of subjects in the source and the driving are different.

Cross-identity reconstruction loss Consider transferring motion between different identities, which is our real task. To train a network in this scenario, it requires a lot of pairs of different subjects in the same posture. However, obtaining such pairs of images is not only time-consuming, but also results in annotations that involve subjective judgments. We overcome the challenge by two consecutive generation procedures between different subjects. The key idea behind our method is that we synthesize an image that has a different identity but the same pose as the one in the source video, then we use it as a pseudo-driving frame to transform an arbitrary frame in the source. In this case, the output image should be the same as the original source frame with the same pose since the poses of both the driving and the source image are identical. By minimizing the pixel distance, the network is enforced to preserve the identity-relevant content in the output. Although our method shares a similar concept with existing self-supervision methods [3, 38], our method presents a novel approach for tackling the cross-identity motion transfer in that it provides a way of supervision using different identities through the pseudo-driving image.

Fig. 3 illustrates the procedure of the cross-identity reconstruction scheme. First, we synthesize a pseudo-driving frame using two different videos. Specifically, we use d_A and s_B as the driving and the source image, respectively. The generator synthesizes the output $r = g_B^{d_A}$, where the object B mimics the pose in d_A . Then, we perform another generation using the output r as the pseudo-driving image and s_A as the source images. In this case, the generator outputs g_A^r that has the identity of A and the pose of r . Because the pose of r is d_A , g_A^r and d_A can be expected to be identical. We minimize a reconstruction loss between this pair of images with an adversarial loss similar to the self-reconstruction:

$$L_{cross} = L_{GAN}(G, D) + \lambda_{FM}L_{FM}(G, D) + \lambda_{VGG}L_{VGG}(g_A^r, d_A). \quad (5)$$

Full objective Our final training objective is described as

$$L_{total} = L_{self} + L_{cross}, \quad (6)$$

which is optimized by solving a minimax optimization problem on an adversarial learning framework.

3.4 Inference

At the inference time, our network takes a driving video and source images as input. Note that there is no limitation to the number of source images. At first, it performs the pose-dependent appearance embedding on all of the source images. Each source image is processed individually using the shared network. Regardless of the number of source frames used at the training time, our network can fully leverage diverse observations from all the source images. The resulting image is generated based on a frame in the driving video using the pose-attentive retrieval module. To transfer motions throughout the driving video, every frame in the driving video is forwarded one-by-one.

4 Experiments

4.1 Experimental setup

Datasets We conduct experiments on diverse object categories using following datasets. **VoxCeleb2** [7] dataset is a large dataset of human speaking videos downloaded from YouTube. VoxCeleb2 contains 150,480 videos of 6,112 celebrities with 145,569/ 4,911 train/test split. **Thai-Chi-HD** [26] contains 3,334 videos of Thai-chi performers. Among them, we use 3,049 videos for training and 285 videos for testing, following the official split. **BAIR robot pushing dataset** [11] is a collection of videos about robot arms moving and pushing various objects on the table. It provides 42,880 training and 128 test videos. Note that the motion contained in the 3 datasets is different from each other, which is face, body, and machine, respectively. By using them, we show the effectiveness of our method on a variety of objects.

Baselines We compare our network with unsupervised motion transfer methods, which are able to expand on arbitrary objects. Particularly for VoxCeleb2, we additionally compare our methods with multi-source motion transfer methods that use facial landmarks detected from a external landmark detector. To the best of our knowledge, we are not aware of any multi-source motion transfer methods designed for arbitrary objects. All of our baselines are listed below:

- X2Face [35]: It transfers motion by warping objects (especially face) in the direction of x and y axis. Also, they utilize multiple source images by averaging all of the appearance features.
- Monkey-net [25]: It uses keypoints learned in self-supervised way to extract motion from general object categories. It transfers motion by estimating optical flows between keypoints.
- First-order [26]: To express more complex and large motion, it uses keypoints and affine transformation for each keypoint.




Fig. 4. Qualitative results of single source animation. We conduct experiments on VoxCeleb2 [7], Thai-Chi-HD [26], and BAIR [11] datasets. Source images and driving images represent appearance and pose to condition on, respectively.

- NeuralHead [36]: It transfers the appearance of source frames into a set of face landmarks. It fine-tunes subject-specific generator parameters using few shots of source images at test time.

Metrics We evaluate motion transfer methods with the following criteria. We follow metrics proposed in Monkey-net [25] and extend the evaluation scope by applying them on self reconstruction results and the motion transfer on different identity.

- Fréchet Inception Distance (FID) [15]: This score indicates the quality of generated images. It is computed by comparing the feature statistics of generated and real images to estimate how realistic the generated images are.
- Average Keypoint Distance (AKD): It represents the average distance between keypoints of generated and driving videos, and evaluates whether the

Table 1. Quantitative comparison on various datasets. For all metrics, smaller values are better. The red and blue color indicate the first and the second ranks, respectively.

	VoxCeleb2				Thai-Chi-HD				BAIR	
	self		cross		self		cross		self	cross
	FID	AKD	FID	AED	FID	AKD	FID	AED	FID	FID
X2Face [35]	48.63	22.23	55.46	0.69	94.99	4.88	101.74	1.42	375.75	379.60
Monkey-net [25]	27.98	18.26	52.49	0.66	30.43	2.85	101.64	2.67	24.28	75.57
First-order [26]	35.39	2.52	78.21	0.78	18.55	1.47	37.38	2.15	31.80	52.39
Ours	17.24	1.90	36.39	0.62	35.60	2.74	48.33	2.02	20.75	34.81

motion of driving video is transferred or not. We extract face and body keypoints using external keypoint detection networks [4, 5].

- Average Euclidean Distance (AED): It shows the degree of identity preservation of generated images. In case of face and human body, we exploit external face recognition network [1, 14] to embed identity-related features and compute average euclidean distance between them.

4.2 Experimental results

Single image animation of arbitrary objects As existing methods for arbitrary objects are not capable of taking multiple source images, we first compare those methods in the setting of single source image. Table 1 shows the quantitative comparison on all datasets. In VoxCeleb2 and BAIR, our network outperforms all baselines methods in all metrics. In Thai-Chi-HD, however, our results are worse than those of First-order [26] especially for self-reconstruction. Although First-order improves the warping-based approach using a two-stage warping with hallucinating occluded background, its generalization is limited to only Thai-Chi-HD with a self-supervision setting. On the other hand, our method performs well on all datasets in general. Fig. 4 shows qualitative results. As can be seen, our method shows the effectiveness in transferring various types of motion. Warping-based methods work for the small movement change, e.g. leg regions of the third row in Fig. 4. However, they have difficulties in synthesizing complex scenes with challenging viewpoint changes.

Comparison with multi-source methods We compare results of using multiple frames. For fair comparison with NeuralHead [36], we use the test set and the results provided by the authors of NeuralHead [36]. The test set is composed of 3 subsets with 1, 8, and 32 source frames in a self-reconstruction setting.

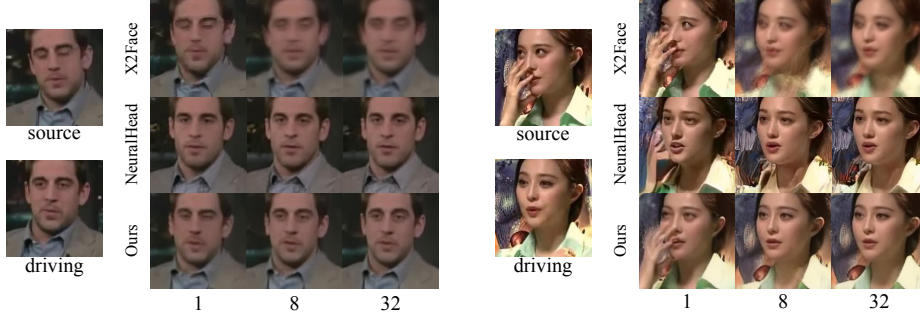
Table 2 represents the quantitative comparison. As can be seen, our method significantly outperforms X2Face and is competitive to NeuralHead. This is because NeuralHead uses keypoints extracted from the pre-trained landmark detector and produces sharper images, while our method learns to estimate the pose without supervision. However, NeuralHead is specifically designed for facial videos with external landmarks, therefore the applicability to arbitrary objects

Table 2. Quantitative results of few-shot motion transfer methods on VoxCeleb2.

	X2Face [35]			NeuralHead [36]			Ours		
	1	8	32	1	8	32	1	8	32
FID	97.37	110.29	127.32	50.10	45.88	44.85	56.19	60.17	60.80
AKD	113.79	96.18	142.09	4.48	3.71	3.45	7.47	4.26	4.09
AED	0.67	0.80	0.79	0.54	0.45	0.43	0.46	0.33	0.29

Table 3. Ablation study on the cross-identity training scheme.

	self			cross	
	FID	AKD	AED	FID	AED
w/o cross	18.49	5.79	0.39	55.11	0.76
w/ cross	17.24	1.90	0.38	36.39	0.62

**Fig. 5.** Comparison with face few-shot motion transfer methods. We compare results with varying number of source images.

is limited. Note that our method outperforms both baselines on AED. It indicates that our method better preserves the identity of the source due to our cross-identity training scheme, which is also observed in Fig. 5.

4.3 Human evaluation

We conducted a user study using Amazon Mechanical Turk (AMT) to compare our results with previous warping based motion transfer methods [35, 25, 26] subjectively. We randomly selected 50 results of VoxCeleb2 [7], Thai-chi-HD [26] and BAIR robot pushing [11] datasets. We used a single frame for source, and uniformly sampled 32 frames from the driving video. Source images and corresponding driving videos are matched randomly to make pairs of distinct identities. In the user study, users ranked the result videos according to the following criteria: (1) whether the identity of the source image is preserved or not, and (2) whether the motion of the driving video is transferred or not. Each question was rated by 10 people and we averaged the total ranks. The results are shown in Table 4. As can be seen, our results are preferred over previous warping-based methods on VoxCeleb2 and BAIR with a large margin. Even though our method achieves lower quantitative scores than First-order [26] on Thai-Chi-HD in Table 1, the user study shows that our results are preferred by users as much as those of First-order.

Table 4. Human evaluation on AMT. The scores indicate average ranking values among 4 methods.

	X2Face [35]	Monkey-net [25]	First-order [26]	Ours
VoxCeleb2	3.116	2.884	2.414	1.586
Thai-Chi-HD	3.438	3.048	1.792	1.722
BAIR	3.58	2.104	2.48	1.836



Fig. 6. Ablation on the cross-identity reconstruction training scheme.

4.4 Analysis

Efficacy of cross-identity reconstruction loss We conduct ablation study on the cross-identity training scheme using VoxCeleb2 dataset by transferring motion between people of different genders and ages. The quantitative and qualitative results are shown in Table 3 and Fig. 6. As represented in the first row of each example in Fig. 6, the self-reconstruction training can disentangle shape and appearance. However, it is not able to further disentangle identity-specific shape such as face shape or skeleton. For example, results without the cross-identity reconstruction training show the tendency that the identity of source images is not preserved for all the driving videos. On the other hand, the identity is well preserved across diverse driving videos with the cross-identity training scheme.

Moreover, the results using only the self-reconstruction training scheme show implausible output in regions surrounding the face such as hair, neck, and shoulder. We attribute it to the fact that the self-reconstruction training scheme focuses on the facial region, which is the most different part between frames in the same video. With cross-identity training, our network is enforced to consider all parts of a person, thereby further improves the quality. The efficacy of our cross-identity reconstruction training scheme is also demonstrated in Table 3. When adding our cross-identity reconstruction, both image quality and identity preservation are improved.

Efficacy of pose-attentive retrieval module For a deeper analysis on pose-attentive retrieval module, we replace the module with following baselines. **ConcatAppearance** concatenates pose and appearance features. It can take only a

Table 5. Ablation study on pose-attentive retrieval module.

Method	self		cross	
	FID	AKD	FID	AED
ConcatAppearance	17.57	4.21	63.52	0.92
ConcatNearest	23.69	4.27	52.13	0.87
ConcatAverage	51.35	103.78	52.67	1.35
Ours	17.24	1.90	36.39	0.62




Fig. 7. Ablation on using multiple source frames on Thai-Chi-HD.



Fig. 8. Visualization of keypoints and attention. We visualize learned keypoints on Thai-Chi-HD and BAIR datasets in the first row. The second row represents activated attention maps according to driving images.

single source image as input. **ConcatNearest** takes multiple source images as input and then selects a single image that has the most similar pose with the driving pose using Euclidean distance of the keypoints. **ConcatAverage** takes the average of feature maps of multiple images. As can be seen in Table 5, **ConcatAppearance** lowers the quality, especially in AKD. **ConcatNearest** could not improve the quality, as well. Moreover, **ConcatAverage** deteriorates performance.

Using multiple source images at inference time Fig. 7 demonstrates the effectiveness of utilizing multiple source images. The driving images are shown on the first row, and the second and the third row show the output images when 1 or 4 source images are given, respectively. With a single source image, reconstructing occluded parts is difficult as shown in the head region in the second row. On the other hand, exploiting multiple source images alleviates the issue by copying and pasting the occluded region from the source images. The performance gain is also shown quantitatively in Table 2.

Visualization To understand the learned representation of our method, we visualize the learned keypoints and attention in Fig. 8. The first row illustrates 10 keypoints extracted from the given frames. In the second row, we visualize the attention map in the source images corresponding to the query points in the driving pose.

5 Conclusion

In this paper, we presented a pose-attentive video reassembling method to tackle the problem of cross-identity motion transfer method for arbitrary objects. Unlike existing warping-based approaches, our pose-attentive reassembling method better handles a wide range of motion and is also able to exploit multiple sources images to robustly reconstruct occluded parts. Furthermore, our cross-identity training scheme disentangles identity-relevant content from object pose to preserve identity in the setting of cross-identity motion transfer.

Acknowledgement This work was conducted by Center for Applied Research in Artificial Intelligence (CARAI) grant funded by DAPA and ADD (UD190031RD).

References

1. Amos, B., Ludwiczuk, B., Satyanarayanan, M., et al.: Openface: A general-purpose face recognition library with mobile applications. *CMU School of Computer Science* **6** (2016)
2. Balakrishnan, G., Zhao, A., Dalca, A.V., Durand, F., Guttag, J.: Synthesizing images of humans in unseen poses. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 8340–8348 (2018)
3. Bansal, A., Ma, S., Ramanan, D., Sheikh, Y.: Recycle-gan: Unsupervised video retargeting. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 119–135 (2018)
4. Bulat, A., Tzimiropoulos, G.: How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks). In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1021–1030 (2017)
5. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7291–7299 (2017)
6. Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 5933–5942 (2019)
7. Chung, J.S., Nagrani, A., Zisserman, A.: Voxceleb2: Deep speaker recognition. *Proc. Interspeech 2018* pp. 1086–1090 (2018)
8. Denton, E., Fergus, R.: Stochastic video generation with a learned prior. *arXiv preprint arXiv:1802.07687* (2018)
9. Denton, E.L., et al.: Unsupervised learning of disentangled representations from video. In: *Advances in neural information processing systems*. pp. 4414–4423 (2017)
10. Ding, H., Sricharan, K., Chellappa, R.: Exprgan: Facial expression editing with controllable expression intensity. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
11. Ebert, F., Finn, C., Lee, A.X., Levine, S.: Self-supervised visual planning with temporal skip connections. *arXiv preprint arXiv:1710.05268* (2017)
12. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in neural information processing systems*. pp. 2672–2680 (2014)
13. Ha, S., Kersner, M., Kim, B., Seo, S., Kim, D.: Marionette: Few-shot face reenactment preserving identity of unseen targets. In: *Proceedings of the AAAI Conference on Artificial Intelligence* (2020)
14. Hermans, A., Beyer, L., Leibe, B.: In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737* (2017)
15. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: *Advances in Neural Information Processing Systems*. pp. 6626–6637 (2017)
16. Jakab, T., Gupta, A., Bilen, H., Vedaldi, A.: Unsupervised learning of object landmarks through conditional image generation. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 31*, pp. 4016–4027. Curran Associates, Inc. (2018), <http://papers.nips.cc/paper/7657-unsupervised-learning-of-object-landmarks-through-conditional-image-generation.pdf>
17. Kulkarni, T.D., Whitney, W.F., Kohli, P., Tenenbaum, J.: Deep convolutional inverse graphics network. In: *Advances in neural information processing systems*. pp. 2539–2547 (2015)

18. Lathuilière, S., Sangineto, E., Siarohin, A., Sebe, N.: Attention-based fusion for multi-source human image generation. In: *The IEEE Winter Conference on Applications of Computer Vision*. pp. 439–448 (2020)
19. Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. In: *Advances in Neural Information Processing Systems*. pp. 406–416 (2017)
20. Ma, L., Sun, Q., Georgoulis, S., Van Gool, L., Schiele, B., Fritz, M.: Disentangled person image generation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 99–108 (2018)
21. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2794–2802 (2017)
22. Oh, S.W., Lee, J.Y., Xu, N., Kim, S.J.: Video object segmentation using space-time memory networks. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 9226–9235 (2019)
23. Oh, S.W., Lee, S., Lee, J.Y., Kim, S.J.: Onion-peel networks for deep video completion. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4403–4412 (2019)
24. Saito, M., Matsumoto, E., Saito, S.: Temporal generative adversarial nets with singular value clipping. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2830–2839 (2017)
25. Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., Sebe, N.: Animating arbitrary objects via deep motion transfer. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2377–2386 (2019)
26. Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., Sebe, N.: First order motion model for image animation. In: *Advances in Neural Information Processing Systems* 32, pp. 7137–7147. Curran Associates, Inc. (2019), <http://papers.nips.cc/paper/8935-first-order-motion-model-for-image-animation.pdf>
27. Siarohin, A., Sangineto, E., Lathuilière, S., Sebe, N.: Deformable gans for pose-based human image generation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3408–3416 (2018)
28. Tran, L., Yin, X., Liu, X.: Disentangled representation learning gan for pose-invariant face recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1415–1424 (2017)
29. Tulyakov, S., Liu, M.Y., Yang, X., Kautz, J.: Mocogan: Decomposing motion and content for video generation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1526–1535 (2018)
30. Villegas, R., Yang, J., Zou, Y., Sohn, S., Lin, X., Lee, H.: Learning to generate long-term future via hierarchical prediction. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. pp. 3560–3569. JMLR. org (2017)
31. Vondrick, C., Pirsivash, H., Torralba, A.: Generating videos with scene dynamics. In: *Advances In Neural Information Processing Systems*. pp. 613–621 (2016)
32. Wang, T.C., Liu, M.Y., Tao, A., Liu, G., Catanzaro, B., Kautz, J.: Few-shot video-to-video synthesis. In: *Advances in Neural Information Processing Systems*. pp. 5014–5025 (2019)
33. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8798–8807 (2018)

34. Wang, T.C., Liu, M.Y., Zhu, J.Y., Yakovenko, N., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. In: *Advances in Neural Information Processing Systems*. pp. 1152–1164 (2018)
35. Wiles, O., Sophia Koepke, A., Zisserman, A.: X2face: A network for controlling face generation using images, audio, and pose codes. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 670–686 (2018)
36. Zakharov, E., Shysheya, A., Burkov, E., Lempitsky, V.: Few-shot adversarial learning of realistic neural talking head models. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 9459–9468 (2019)
37. Zhao, L., Peng, X., Tian, Y., Kapadia, M., Metaxas, D.: Learning to forecast and refine residual motion for image-to-video generation. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 387–403 (2018)
38. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2223–2232 (2017)