

QS-Craft: Learning to Quantize, Scrabble and Craft for Conditional Human Motion Animation

Yuxin Hong¹, Xuelin Qian^{*1}, Simian Luo¹, Guodong Guo³, Xiangyang Xue^{1,2},
and Yanwei Fu¹

¹ School of Data Science, and MOE Frontiers Center for Brain Science, Shanghai
Key Lab of Intelligent Information Processing, Fudan University
{20210980140,xlqian,18300180157,yanweifu}@fudan.edu.cn

² School of Computer Science, Shanghai Key Lab of Intelligent Information
Processing, Fudan University
xyxue@fudan.edu.cn

³ Dept. of CSEE, West Virginia University
guodong.guo@mail.wvu.edu

Abstract. This paper studies the task of conditional Human Motion Animation (cHMA). Given a source image and a driving video, the model should animate the new frame sequence, in which the person in the source image should perform a similar motion as the pose sequence from the driving video. Despite the success of Generative Adversarial Network (GANs) methods in image and video synthesis, it is still very challenging to conduct cHMA due to the difficulty in efficiently utilizing the conditional guided information such as images or poses, and generating images of good visual quality. To this end, this paper proposes a novel model of learning to Quantize, Scrabble, and Craft (QS-Craft) for conditional human motion animation. The key novelties come from the newly introduced three key steps: quantize, scrabble and craft. Particularly, our QS-Craft employs transformer in its structure to utilize the attention architectures. The guided information is represented as a pose coordinate sequence extracted from the driving videos. Extensive experiments on human motion datasets validate the efficacy of our model.

1 Introduction

The task of conditional Human Motion Animation (cHMA) has attracted increasing attention in the vision community, as it can be utilized in many industrial applications such as computer games, advertisement, and animation industry [4, 17, 18]. In this work, we aim to solve the cHMA problem by transferring driving videos to animate the humans in the source. Particularly, given a source image and a driving video, the cHMA model should animate the new frame sequence, in which the person in the source image should perform similar motion as pose sequence from the driving video.

* corresponding author

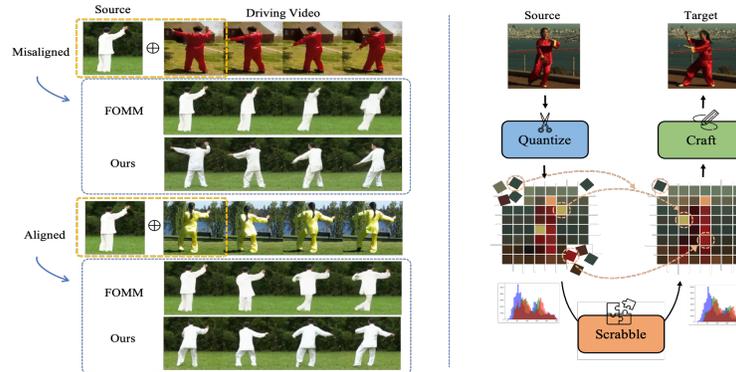


Fig. 1. Left: Examples of cHMA task. Right: idea illustration of our QS-Craft. Our method learns to quantize source image, scrabble quantized representation conditioned on driving video, and finally craft the synthesis of target images. Notice that when pose misalignment occurs (left-top), FOMM [27] fails to generate realistic motion transfer images, while our method accomplishes the task well and preserves finer details. Two histograms (right-bottom) indicate quite similar visual features of the source and target, producing intuition to play the *Scrabble* game.

To tackle this task, some classical approaches [3, 5, 30] are very domain-specific and rely on parametric models to represent humans. Generative models, such as Variational Auto-Encoders (VAE) [14, 29] and Generative Adversarial Networks (GANs) [2, 19, 21, 23], have shown good capacity on the image generation task, which could be the alternative ways to address our cHMA task. Unfortunately, these models typically demand expensive training from large collection of data in an *unconditioned* manner. On the other hand, motion flow based approaches [27, 26, 28] have received more and more attention due to its remarkable performance. The key idea of these methods is to learn motion flows between the source image and the driving frames, such that the target image features can be synthesized by wrapping. For cHMA, authors [26, 28, 27] proposed to use keypoints to produce flow maps, but these works additionally require modules to map the flow from sparse to dense. This is not only difficult to be learned but also may introduce lots of artifacts. For example, the misalignment of human motion in [27] in the source image and the first driving video frame may lead to an unreasonable animated result as shown in the left of Fig. 1.

In this paper, we advocate the way of motion flow yet from a different and novel perspective. As illustrated in Fig. 1, we propose to generate the target image by only rearranging the order of patches in the source image, on account of the observation that the source and the target images have very similar distribution. This is similar to a game of ‘Scrabble’ using the same letters to spell different words but following different rules. So, the key question is *how to learn the rule of Scrabble?* We argue that it should satisfy three requirements: (i) rules should be objective and independent of input; (ii) the order of patches

should be reasonable, including the texture of background and the identity of foreground; (iii) obviously it should be conditioned on the driving videos. To this end, we present a novel paradigm of learning to **Q**uantize, **S**crabble and **C**raft for conditional human motion animation (cHMA), thus dubbed QS-Craft.

Formally, given the input source image and the driving video, our QS-Craft model should have the following three key steps, as illustrated in the right of Fig. 1. (1) *Quantize*. Train an encoder to get the latent feature of the source image and then quantize it to produce a bag of discrete feature pixels by referring to the global codebook, which is trained in an end-to-end manner. (2) *Scrabble*. Conditioned on motion information from driving video, our QS-Craft trains a decoder-only transformer to rearrange the feature pixels computed in (1) to produce the corresponding feature pixels for the targets. By virtue of such a way, we can effectively exploit the distribution of visual features, to help synthesize the latent features of the target images. (3) *Craft*⁴. Finally, the reordered feature pixels will be passed to a decoder layer to produce the animated target images. Extensive experiments on benchmarks validate the efficacy of our model, and the quantitative and qualitative results show that our method outperforms the state-of-the-art competitors.

Contributions. The main contribution of this paper is to propose a novel paradigm for the conditional human motion animation: learning to Quantize, Scrabble and Craft (QS-Craft). The quantize step is first introduced to encode the input source image and driving videos into a discrete representation. Next, we employ transformer to learn to *scrabble* by mixing up pixel distributions and exploiting the known distribution of the source image to fit those in driving videos. The reordered feature pixels are decoded to *craft* the human motion animation. Based on this novel model framework, we can not only address motion transfer when pose misalignment occurs (in Fig. 1) but achieve better quantitative scores in video reconstruction compared to other state-of-the-art methods.

2 Related Works

Image Animation. Image animation and video re-targeting have drawn attention of computer vision researchers in recent years. Some of the previous approaches [5, 30, 3] are heavily domain-specific, which means they can only tackle animation of human faces, etc. When transferred to other situations, these methods might fail and are ineffective. Recycle-GAN [1] incorporates Spatio-temporal cues into a conditional GAN architecture to translate from one domain to another. However, it has to be trained with a strong prior of specific domains and cannot be generalized to even another individual. Similarly, the model in [33] aiming to address the motion transfer task is also exposed to the same problem. Compared to these works, our method does not need any prior during the training phase and can be domain-agnostic in the inference time.

X2Face [34] employs reference poses to warp a given image and obtain the final transferred video. It does not require priors but is better to face anima-

⁴ The name of Craft is inspired by the game of Minecraft.

tions than cHMA. Siarohin *et al.* introduced Monkey-Net [26]. This network can indeed animate arbitrary objects guided by detected key points, but it is only effective on lower resolution images. The authors in [27] and [28] proposed an affine transformation and a PCA-based motion estimation respectively to transfer driving motions into the target images with higher resolution. Note that both of them are based on warping motion flow in order to encode images in feature space then decoding to get the targets. It suggests that the latent feature in target images has to be encoded from scratch whilst our proposed model can rearrange the quantized feature sequences conditioned on driving videos for the motion transfer task. This discretization method ensures better quality of our generated images. Moreover, the main structure of our model is transformer, which is quite different from other methods mentioned above. And both the quantitative and qualitative results show the efficacy of this prominent block.

Visual Transformer. Since the introduction of the self-attention module [32], many transformer-based methods have achieved impressive performance in numerous Natural Language Processing (NLP) and vision tasks [24, 10, 11, 8, 16]. Autoregressively generating pixels by transformer can only be applied to low-resolution images due to costly computation and huge memory footprint [9, 15]. Many recent works, such as dVAE [25] and VQ-VAE [20], attempt to model discrete representation of images to reduce the total sequence length. Besides, VQ-GAN [12] added Generative Adversarial Networks (GANs) into VQ-VAE to improve training efficiency. Both VQ-VAE and VQ-GAN quantize an image according to a global codebook. In the second stage of [12], its transformer module has to look up in the whole codebook to generate latent features of targets, which may lead to accumulated errors and unreasonable generated results. Hence constructing a local dynamic codebook for each conditions-target pair can enhance the searching ability of the transformer; and thus the target images generated by the decoder will be more realistic. To achieve this goal, we propose to quantize one image twice to rearrange conditions for the target one. And this scrabbling also speed up convergence of the model.

3 Methodology

The purpose of this paper is to animate the human in a source image $x_s \in \mathbb{R}^{H \times W \times 3}$ conditioned on the motion in a driving video. For the target image $x_t \in \mathbb{R}^{H \times W \times 3}$, it should be semantically consistent with the source image, including the information of background and the content of the human texture. More importantly, it demands natural and realistic motion changes, which are guided by a series of driving frames (*i.e.*, the condition c). Following our inspiration, we assume that the target image x_t can be generated by reordering the latent features z_s of the source image in spatial dimensions, which can be expressed as,

$$x_t = G(\mathcal{T}(z_s, c)) = G(\mathcal{T}(E(x_s), c)) \quad (1)$$

where $E(\cdot)$ and $G(\cdot)$ denote the encoder and the decoder layer; $\mathcal{T}(\cdot)$ means the operation of ‘scrabble’, *i.e.*, putting latent features z_s together following a specific rule.

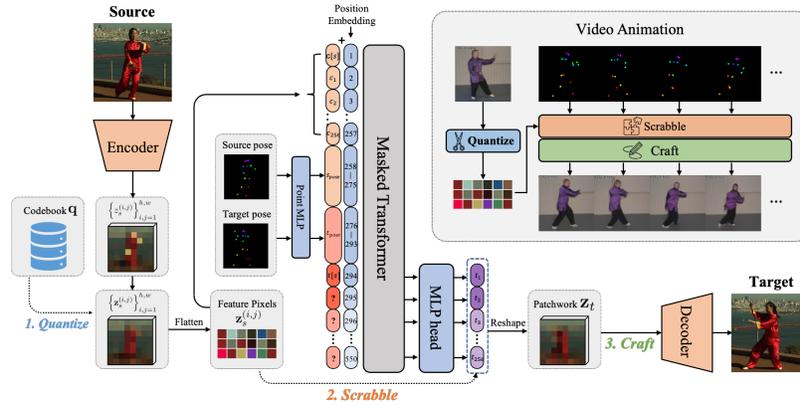


Fig. 2. The overall pipeline of QS-Craft. We first train an encoder-decoder based framework to quantize both source and target images. Then, a conditional transformer is designed to learn the scrabble game under the given driving videos. Finally, the decoder crafts the reordered features as output.

The overall of our framework is illustrated in Fig. 2, which is composed of three steps, *Quantize*, *Scrabble* and *Craft*. First, we perform the vector quantization mechanism [20] on the encoded latent features z_s of the source image, resulting in a bag of feature pixels $\{z_s^{(i,j)}\}_{i,j=1}^{h,w}$, where h and w indicate the height and width of the latent features, respectively. Then, we play a Scrabble game by selecting pixels from the bag and mixing them together, $\hat{z}_t = \mathcal{T}(z_s, c)$. This process must be meaningful and also be conditioned on the motions from driving videos. Finally, the patchwork \hat{z}_t is fed to the decoder to craft the image, which is realistic with natural motion changes.

Considering that $\mathcal{T}(\cdot)$ is operated on index-based representation, we divide the training procedure of our model into two stages, *image reconstruction with scrabble* and *learning scrabble rules with transformer*. For the rest of this section, we will follow these two training stages to elaborate our framework in Sec. 3.1 and Sec. 3.2, respectively. The objective functions will be elaborated in the supplementary material.

3.1 Image Reconstruction with Scrabble

In the first stage, we aim to teach the decoder G to generate a natural target image x_t from a known patchwork \hat{z}_t . We first utilize the encoder E to extract latent features from both source and target images, denoted as z_s and z_t . Considering the correlations between each feature pixel in z_s , we adopt one more step of vector quantization before fitting them together. Specifically, a learnable codebook $\mathbf{q} \in \mathbb{R}^{m \times c_q}$ is introduced, where m and c_q are the number and dimen-

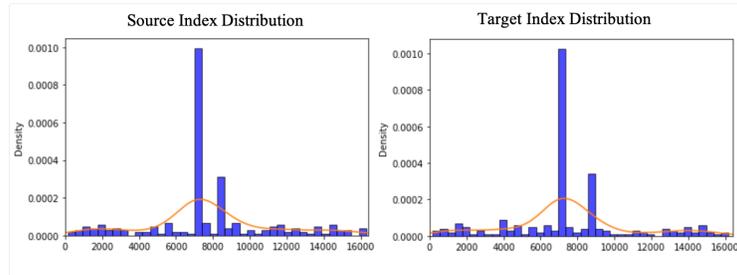


Fig. 3. Index distribution of a random sample pair \mathbf{z}_s and \mathbf{z}_t .

sion of codes, respectively. For each pixel $\left\{z_s^{(i,j)}\right\}_{i,j=1}^{h,w}$, we do quantization $\mathcal{Q}(\cdot)$ by replacing it with the closest codebook entry q_k in \mathbf{q} ⁵,

$$\mathbf{z}_s = \mathcal{Q}(z_s) := \arg \min_{q_k \in \mathbf{q}} \|z_s^{(i,j)} - q_k\| \quad (2)$$

where $\mathbf{z}_s \in \mathbb{R}^{h \times w \times c_q}$. After that, we obtain a bag of quantized feature pixels $\left\{\mathbf{z}_s^{(i,j)}\right\}_{i,j=1}^{h,w}$.

In order to reconstruct the target image x_t , we utilize the latent feature z_t as a reference to play scrabble with pixels (words) in \mathbf{z}_s . More concretely, we compute the distance between feature pixels in z_t and \mathbf{z}_s , and select the closest ones to fit a patchwork $\hat{z}_t \in \mathbb{R}^{h \times w \times c_q}$. The formulation can be written as,

$$x_t = G(\hat{z}_t), \quad \hat{z}_t := \arg \min_{z_k \in \mathbf{z}_s} \|z_t^{(i,j)} - z_k\| \quad (3)$$

To encourage the fidelity of the synthesized images, we reverse the generation flow in Eq. 3, that is, using a bag of quantized target feature pixels $\left\{\mathbf{z}_t^{(i,j)}\right\}_{i,j=1}^{h,w}$ to fit a patchwork of the source \hat{z}_s . Furthermore, we also introduce a perceptual loss [35, 31] and a Discriminator D [36, 17] to highly maintain the perceptual quality during training. Details of structures about the Encoder E , the Decoder G and the Discriminator D can be found in the supplementary.

Remark. Different from Eq. 2, Eq. 3 is much more important here. A straightforward explanation of Eq. 3 is to narrow down the search space in Eq. 2. Experiments in Sec. 4.2 empirically show that this constraint will not degrade the expressive power of the representation \hat{z}_s or \hat{z}_t . Even better, it is efficient and keeps more fidelity in background and foreground. Furthermore, to simply validate our explanation, we randomly choose a sample pair and plot their index distribution in Fig. 3. As we can see, the index distribution of \mathbf{z}_s and \mathbf{z}_t are quite similar; therefore, by scrabbling the indices in the source, we are able to reconstruct the target image.

⁵ We refer the readers to [12] for details of the codebook learning.

3.2 Learning Scrabble Rules with Transformer

Through the process described in Sec. 3.1, the Decoder G now can synthesize realistic images according to the given patchwork. Nevertheless, the patchwork is generated with the reference of target image, which is not available during inference. To this end, we propose learning the rule of Scrabble, such that the desired patchwork \hat{z}_t can be successfully produced *only with* a bag of source feature pixels \mathbf{z}_s and the condition motion information c .

Here, the condition c refers to the information of pose coordinates. We claim that directly utilizing RGB images as condition may cause the information leakage [6] when both the source and driving frames are from the same video. Empirical results in Sec. 4.3 suggest that pose coordinates can not only provide effective and straightforward condition, but also save lots of computational costs compared with using pose image [12]. More precisely, an off-the-shelf open pose detector is applied to detect and localize n keypoints of an object. Each point is represented by x - y coordinates. If a point is occluded or undetectable, we set $x = y = -1$ for indication. Then, we attach three fully-connected layers with ReLU activation to encode the 2D pose coordinates $p \in \mathbb{R}^{n \times 2}$, which can be expressed as,

$$c = \mathcal{F}(p) \in \mathbb{R}^{n \times n_c} \tag{4}$$

where $\mathcal{F}(\cdot)$ denotes the stacked fully-connected layers, n_c means the output dimension. We flatten the pose condition c as a sequence $\{c_i\}_{i=1}^n$, which is further incorporated with our transformer.

Speaking of the transformer, we adopt a decoder-only structure⁶ and first convert the feature representation of $\mathbf{z}_s \in \mathbb{R}^{h \times w \times c_q}$ and $\hat{z}_t \in \mathbb{R}^{h \times w \times c_q}$ to the sequence of index representation, *i.e.*, $\{s_i\}_{i=1}^l$ and $\{t_i\}_{i=1}^l$, where its value ranges from 0 to $m - 1$, indicating the index of each feature pixel in the codebook \mathbf{q} , and $l = h \times w$. Two embedding layers with the dimension of n_c are subsequently followed for $\{s_i\}$ and $\{t_i\}$, respectively⁷. During training phase, we concatenate the source embedding $\{s_i\}$ and the condition sequence $\{c_i\}$ as input. Note that the previous ground-truth target embedding $\{t_i\}_{i=1}^{j-1}$ are progressively attached to the input to predict the likelihood of the next target index t_j . The overall formulation can be written as,

$$p(t|s, c) = \prod_{j=1}^l p(t_j|s, c, t_{[s]}, t_{<j}) \tag{5}$$

where $t_{[s]} \in \mathbb{R}^{1 \times c_q}$ denotes a learnable start token, which is added after the condition sequence for the case of predicting the first target index t_1 , as shown in Fig. 2. For inference, we use the previously predicted index to replace the ground-truth ones. With regard to the multi-head self-attention in the transformer, we design a new attention mask \mathbf{M} with four sub-masks,

⁶ Please refer to the supplementary for the transformer structure.

⁷ We reuse the symbols of $\{s_i\}$ and $\{t_i\}$ after embedding for simplicity.

$$\mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix} = \begin{bmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{M}_{tril} \end{bmatrix} \quad (6)$$

where $\mathbf{B} \in \mathbb{R}^{(l+n) \times (l)}$ is a zero matrix, $\mathbf{A} \in \mathbb{R}^{(l+n) \times (l+n)}$ and $\mathbf{C} \in \mathbb{R}^{l \times (l+n)}$ are all-ones matrices, designed to learn the relationship among the source embedding $\{s_i\}$, the conditional information $\{c_i\}$ and the target embedding $\{t_i\}$; $\mathbf{M}_{tril} \in \mathbb{R}^{l \times l}$ means a standard lower triangular matrix filled with 1, so that the next target index can be deduced from the previous known information.

Recall the key idea of ‘Scrabble’ in the transformer, a patchwork \hat{z}_t is acquired *only with* feature pixels from \mathbf{z}_s , thus we apply a mask constraint to the output of index probabilities. Specifically, denote the output of the j -th target index from transformer as $v_j \in \mathbb{R}^m$, we mask out some elements if their indices are out of the bag of source index $\{s_i\}_{i=1}^l$,

$$v_j^k = -\text{inf} \quad \text{if} \quad k \notin \{s_i\}_{i=1}^l \quad (7)$$

where $k \in [0, m - 1]$ and v_j^k indicates the k -th element in v_j . We feed masked v_j into *Softmax* operation to obtain final probability of the j -th target index. Furthermore, in order to improve the fidelity and consistency of the synthesized images, especially for the foreground objects, we encourage the model to learn more correlations between target indices by re-weighting the loss of each predicted index. Since we intuitively strengthen the learning of the foreground area, we call this strategy as RoI (regions of interest) weight. Formulations are elaborated in the supplementary materials.

4 Experiments

Datasets. We evaluate our model on two widely-used benchmarks. (1) Tai-Chi-HD, collected from YouTube following [31], is a dataset of human bodies performing Tai Chi actions, consisting of 252 videos for training and 28 videos for testing. Following MRAA [28] preprocessing method, we finally obtain 3,049 and 285 video chunks for training and testing. All video frames are resized to 256×256 . (2) Penn Action (PA) dataset [37] contains 2,326 video sequences of 15 action classes. All video frames are resized to 256×256 after preprocessing.

Metrics. In our experiments, we use the following metrics to provide an in-depth comparisons with other competitors. (1) *Average Keypoint Distance (AKD)*, which means the average distance between the detected keypoints of the ground truth image sequences and the generated ones. For both two datasets, we employ the human-pose estimator in [7] (2) *Missing Keypoint Rate (MKR)*, Another metric evaluating the difference between the poses in the real images and the reconstructed, is the proportion of keypoints detected in ground truth but missing in the reconstruction. (3) *Fréchet Inception Distance (FID)* [13], which measures the quality of generated images. In this paper, we concentrate more on the foreground area (*i.e.*, human bodies) to evaluate the fidelity and consistency.

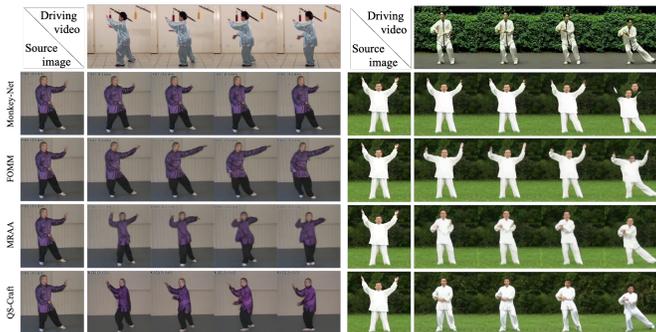


Fig. 4. Visualization of image animation on Tai-Chi-HD. Given the source image and driving videos, our animated images are better than those from competitors. Best viewed in color and zoom in.

Implementation Details. On all datasets, we train our model in two stages. (1) for the reconstruction training, our generative model is finetuned on ImageNet datasets [12]. In each training iteration, we randomly select 2 frames from the same video for training. Adam optimizer is applied with $\beta_1 = 0.5, \beta_2 = 0.9$, batch size 2 and the initial learning rate of $5e-5$. We drop the learning rate by half for every 70K steps. The total training iteration for the first stage is about 210K. (2) For the transformer training, we use Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.95$ and batch size 12. We do not change the initial learning rate but add the warmup strategy within the first 10K steps. The learning rate is linearly decayed to 0 gradually. We train the second stage 280K for Tai-Chi-HD and 270K for PA.

Competitors. Several related models are listed as competitors. For Tai-Chi-HD dataset, we compare our QS-Craft with three state-of-the-art models for animating, namely, Monkey-Net [26], FOMM [27] and MRVA and report both qualitative and quantitative results. For PA dataset, as for its complexity and variety of human motions, it is not suitable for the animation task whereas tends to be a dataset of pose-guiding. Hence, we compare with four pose-guided generative models, including PG^2 [17], PATN [38], PN-GAN [22] and MR-Net [35].

4.1 Qualitative Results

Note that the objective metrics used in our cMHA task can only reflect the general quality of synthesized images, while it is difficult for these metrics to directly evaluate whether the synthesized images are mimicking the human motion sequences from driving video. In that case, qualitative evaluations are much more important in cMHA task to directly reflect whether our model works well.

Human Animation. Figure 4 and 5 show animation results on Tai-Chi-HD dataset. It can be noticed that our method can generate more realistic images with accurate motion according to driving videos on the four randomly selected samples. As these competitors all highly rely on flow information and warping

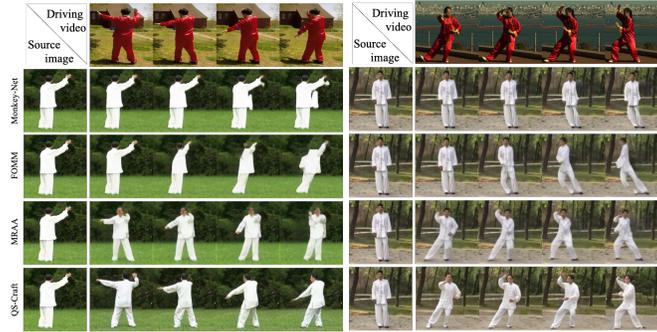


Fig. 5. More animation results on Tai-Chi-HD. Best viewed in color and zoom in.

Table 1. Video Reconstruction Results on Tai-Chi-HD. **Table 2.** Quantitative ablation study for pose-guided generation on PA dataset.

Methods	Tai-Chi-HD			Methods	Penn Action		
	AKD ↓	MKR ↓	FID ↓		AKD ↓	MKR ↓	FID ↓
Monkey-Net [26]	13.77	0.061	32.223	<i>w/o Scrabble</i>	31.878	0.326	42.721
FOMM [27]	7.02	0.031	25.910	<i>w/o RoI weight</i>	17.819	0.121	33.066
MRAA [28]	5.73	0.025	35.794	Full Model	16.358	0.121	30.136
Ours	4.61	0.017	25.064				

operations, once the source image is not aligned with the first frame of the given driving video, it will introduce either wrong poses or twisted human bodies in the following generated animation results. For example, in the left subfigure of Fig. 4, results of Monkey-Net in the first row indicate the failure of motion animation: it only copies the source image. In contrast, our QS-Craft introduces the *Scrabble* step to rearrange the discrete pixel representations from the source image conditioned on the driving motion, thus facilitating the large human poses between the source and driving videos. Moreover, FOMM can sometimes capture action information in the driving but collapses in most cases. For MRAA, it can indeed generate the animated human bodies with roughly correct poses but details such as “raising left foot” are missing. It should be emphasized that FOMM tends to fail in the case of misalignment between human motion in the source image and the first frame of driving video, which has been discussed in the limitation of its original paper. As for MRAA, it alleviates the misalignment problem to a certain extent but at the cost of losing a lot of details. Compared to these methods, our proposed QS-Craft can transfer source images accurately as well as retain most important details. More results about face animation can be found in the supplementary materials.

Pose Guidance. We further provide results on PA dataset, which is a more challenging and large dataset. The qualitative results are present in Fig. 7. We compare the competitors of PG², PATN, PN-GAN and MR-Net. Particularly, we show that PG² tends to generate blur results in most cases. Besides, PN-GAN

Table 3. Pose-guided Generation Results on PA dataset.

Methods	Penn Action		
	AKD ↓	MKR ↓	FID ↓
PG2[17]	20.577	0.279	78.615
PN-GAN[22]	19.637	0.167	47.096
PATN[38]	19.288	0.267	51.583
MR-Net[35]	13.663	0.169	58.796
Ours	16.358	0.121	30.136

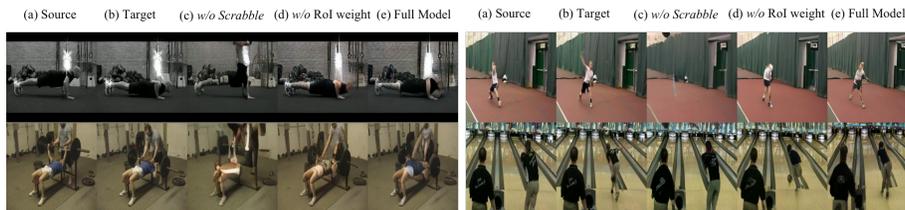


Fig. 6. Visualization of pose-guided generation on PA dataset. Our method is better than the competitors.

is inclined to copy the source image as target ones. Human body parts in results of PATN are twisted, resulting in unreasonable artifacts. For MR-Net, the area around the human is blurred, making the synthesized images unrealistic. Compared to these models, our QS-Craft demonstrates its efficiency in generating the accurate human poses guided by conditions.

4.2 Quantitative Results

We also give the general measurement of the image quality. Thus we compare the quantitative results on PA dataset of our proposed QS-Craft and other competitors in Tab. 3. As for Tai-Chi-HD dataset, since there is no ground truth according to the animation results to get quantitative scores, a sanity check by video reconstruction is conducted to demonstrate the effectiveness of our QS-Craft. Here we employ three metrics described above, Average Keypoint Distance (AKD), Missing Keypoint Rate (MKR) and Fréchet Inception Distance (FID) to measure the quality of generated results.

Note that for Monkey-Net, FOMM, and MRAA are not designed to tackle the complex images from PA dataset; and thus there is no available model for the direct comparison. Thus, the methods of PG², PATN, PN-GAN and MR-Net are compared here. We note that our QS-Craft outperforms all the baselines in AKD, MKR and FID. This intuitively reflects that our QS-Craft can deal with the complicated background with human motion, producing visually good synthesized results.

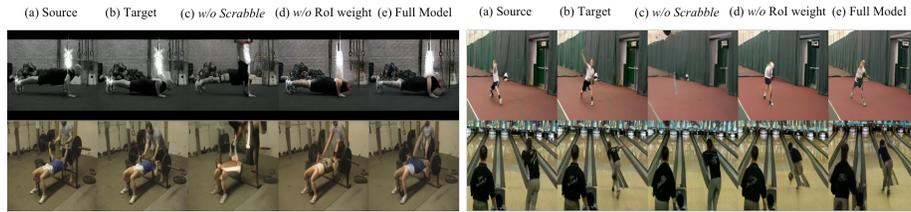
Sanity check by video reconstruction. We take this task as the sanity check of our model, as it is designed for animation in transferring settings. That

Table 4. Ablation study on different usages of pose conditions (report *w/o* RoI weight).

Methods	Tai-Chi-HD		
	AKD ↓	MKR ↓	FID ↓
Pose Flow	22.939	0.065	70.117
Pose Image	10.676	0.025	32.324
Pose VQ	19.194	0.052	56.285
Pose Coord	6.289	0.020	25.629

Table 5. User study results on Tai-Chi-HD: user preferences in favour of our approach.

Methods	Tai-Chi-HD	
	Test 1	Test 2
FOMM [27]	19.87%	25.29%
Ours	80.13%	74.71%

**Fig. 7.** Qualitative results in ablation study on PA dataset

is, source image and driving video are different. In particular, we give results of video reconstruction in Fig. 10. Image frames reconstructed by Monkey-Net suffer from the wrong motion compared against the ground truth. Furthermore, FOMM and MRAA can capture the general motion information whilst ignoring some body parts such as hand movements in the third column. In a word, our reconstructed images demonstrate the superiority to the quality and detailed texture, which highly supports the efficacy of our QS-Craft framework in addressing the cMHA task. The quantitative results on Tai-Chi-HD dataset are shown in Tab. 1. We observe that our QS-Craft outperforms other methods in all metrics, which indicated the superiority of our method.

User study. To complete our evaluation, a user study is conducted on Tai-Chi-HD dataset. We provide 32 users with two test sets of randomly selected animation results and ask them to select the most realistic and reasonable generation. Specifically, both test sets contain 25 different image animations, each of which involves one source image, one driving video and two animated generations. Results in Tab. 5 show that our QS-Craft is preferred over the competitor.

4.3 Ablation Study

Generalizability of Model. The first ablation study is to evaluate the generalizability of our method. We download some online videos to be the target driving. These online videos have complete novel driving poses, which are quite different from Tai-Chi-HD dataset. Qualitative results of this experiment demonstrate our method generalizes well in Fig. 8. Even for such challenging cases, Our method can still get good results, clearly beats MRAA.



Fig. 8. Online driving videos with source Tai-Chi-HD images.



Fig. 9. Qualitative comparisons of image reconstruction with different conditions.

Variants of Pose Conditions. In our framework, we particularly use pose coordinates, encoded by multilayer perceptrons, as conditional features. In order to verify the effectiveness of this design (termed *Pose Coord*), we conduct three variants: (1) *Pose Flow*: we follow [26] to build a motion net with five convolution layers to encode flow information. Differently, the flow features are further applied as a sequence in our transformer model, rather than warping source features [26]. (2) *Pose Image*: Similar to other studies [22, 35, 17], we use pose skeleton RGB images as condition information. Pose image features are extracted with five convolution layers. (3) *Pose VQ*: We use pose images to train a VQ-VAE model [20] and then the quantized pose embeddings are applied as conditions, which is the same as [12].

As illustrated in Fig. 9, the results achieved by our QS-Craft are apparently more natural than other variants. Particularly, our results are authentic with richer accurate motion information from the driving videos. Besides, quantitative results in Tab. 4 also validate the same conclusion. For all the metrics, our *Pose Coord* outperforms others apparently.



Fig. 10. Reconstruction results on Tai-Chi-HD. The synthesized results are compared among different methods.

Efficacy of the proposed modules. We further conduct studies to demonstrate the effectiveness of our proposed *Scrabble* step and RoI weight in Fig. 7 and Tab. 2. As we can find in Fig. 7, model trained without *Scrabble* step fails to generate reasonable target images with accurate pose information as it do not take into account the strong connection between the source and target and then accumulates prediction errors in inference. Additionally, model lack of RoI weight pays less attention to foregrounds compared to QS-Craft so it is inclined to generate blurred human motion or different clothes colors. And quantitative results of our QS-Craft in Tab. 2 outperform other variants, which indicates the superiority of our full model.

5 Conclusion

In this paper, we propose a novel method to animate objects conditioned on driving videos through three phases: *Quantize*, *Scrabble* and *Craft*. This transformer-based model can effectively generate semantically consistent and realistic results as we demonstrate above. Besides, especially compared to other methods, our QS-Craft do not need the source given to be aligned with the first frame of the driving video as the simple pose keypoints are enough for QS-Craft.

Limitation. Our QS-Craft also encounters some challenges, including smoothness in video animation, since no temporal information are involved. Besides, we utilize information in the source to generate the target which may cause failure in some extreme cases. For example, when the person’s body is frontal in the source and another person is back to camera in the driving, we cannot generate accurate faces according to the identity in the source. Hence involving more information about the source person is pretty important.

Acknowledgements. This work was supported by China Postdoctoral Science Foundation (2022M710746), the Science and Technology Major Project of Commission of Science and Technology of Shanghai (No.21XD1402500).

References

1. Bansal, A., Ma, S., Ramanan, D., Sheikh, Y.: Recycle-gan: Unsupervised video re-targeting. In: Proceedings of the European conference on computer vision (ECCV). pp. 119–135 (2018)
2. Berthelot, D., Schumm, T., Metz, L.: Began: Boundary equilibrium generative adversarial networks. arXiv preprint arXiv:1703.10717 (2017)
3. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: Proceedings of the 26th annual conference on Computer graphics and interactive techniques. pp. 187–194 (1999)
4. Cai, H., Bai, C., Tai, Y.W., Tang, C.K.: Deep video generation, prediction and completion of human action sequences. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 366–382 (2018)
5. Cao, C., Hou, Q., Zhou, K.: Displaced dynamic expression regression for real-time facial tracking and animation. ACM Transactions on graphics (TOG) **33**(4), 1–10 (2014)
6. Cao, C., Hong, Y., Li, X., Wang, C., Xu, C., Fu, Y., Xue, X.: The image local autoregressive transformer. In: NeuPIS (2021)
7. 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)
8. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European Conference on Computer Vision. pp. 213–229. Springer (2020)
9. Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: International Conference on Machine Learning. pp. 1691–1703. PMLR (2020)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
11. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
12. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12873–12883 (2021)
13. 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. arXiv preprint arXiv:1706.08500 (2017)
14. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
15. Kumar, M., Weissenborn, D., Kalchbrenner, N.: Colorization transformer. arXiv preprint arXiv:2102.04432 (2021)
16. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 (2021)
17. Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. arXiv preprint arXiv:1705.09368 (2017)

18. 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)
19. 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)
20. Oord, A.v.d., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning. arXiv preprint arXiv:1711.00937 (2017)
21. Qi, G.J.: Loss-sensitive generative adversarial networks on lipschitz densities. International Journal of Computer Vision **128**(5), 1118–1140 (2020)
22. Qian, X., Fu, Y., Xiang, T., Wang, W., Qiu, J., Wu, Y., Jiang, Y.G., Xue, X.: Pose-normalized image generation for person re-identification. In: Proceedings of the European conference on computer vision (ECCV). pp. 650–667 (2018)
23. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434 (2015)
24. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding by generative pre-training (2018)
25. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. arXiv preprint arXiv:2102.12092 (2021)
26. Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., Sebe, N.: Animating arbitrary objects via deep motion transfer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2377–2386 (2019)
27. Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., Sebe, N.: First order motion model for image animation. Advances in Neural Information Processing Systems **32**, 7137–7147 (2019)
28. Siarohin, A., Woodford, O.J., Ren, J., Chai, M., Tulyakov, S.: Motion representations for articulated animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13653–13662 (2021)
29. Sønderby, C.K., Raiko, T., Maaløe, L., Sønderby, S.K., Winther, O.: Ladder variational autoencoders. arXiv preprint arXiv:1602.02282 (2016)
30. Thies, J., Zollhofer, M., Stamminger, M., Theobalt, C., Nießner, M.: Face2face: Real-time face capture and reenactment of rgb videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2387–2395 (2016)
31. 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)
32. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. arXiv preprint arXiv:1706.03762 (2017)
33. Wang, T.C., Liu, M.Y., Zhu, J.Y., Liu, G., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. arXiv preprint arXiv:1808.06601 (2018)
34. Wiles, O., 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)
35. Xu, C., Fu, Y., Wen, C., Pan, Y., Jiang, Y.G., Xue, X.: Pose-guided person image synthesis in the non-iconic views. IEEE Transactions on Image Processing **29**, 9060–9072 (2020)
36. Yang, C., Wang, Z., Zhu, X., Huang, C., Shi, J., Lin, D.: Pose guided human video generation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 201–216 (2018)

37. Zhang, W., Zhu, M., Derpanis, K.G.: From actemes to action: A strongly-supervised representation for detailed action understanding. In: 2013 IEEE International Conference on Computer Vision. pp. 2248–2255 (2013). <https://doi.org/10.1109/ICCV.2013.280>
38. Zhu, Z., Huang, T., Shi, B., Yu, M., Wang, B., Bai, X.: Progressive pose attention transfer for person image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2347–2356 (2019)