

Learning Fine-Grained Motion Embedding for Landscape Animation

Hongwei Xue*
gh051120@mail.ustc.edu.cn
University of Science and Technology
of China
Hefei, China

Bei Liu
bei.liu@microsoft.com
Microsoft Research
Beijing, China

Huan Yang
huayan@microsoft.com
Microsoft Research
Beijing, China

Jianlong Fu
jianf@microsoft.com
Microsoft Research
Beijing, China

Houqiang Li
lihq@ustc.edu.cn
University of Science and Technology
of China
Hefei, China

Jiebo Luo
jluo@cs.rochester.edu
University of Rochester
Rochester, NY

ABSTRACT

In this paper we focus on landscape animation, which aims to generate time-lapse videos from a single landscape image. Motion is crucial for landscape animation as it determines how objects move in videos. Existing methods are able to generate appealing videos by learning motion from real time-lapse videos. However, current methods suffer from inaccurate motion generation, which leads to unrealistic video results. To tackle this problem, we propose a model named **FGLA** to generate high-quality and realistic videos by learning **Fine-Grained** motion embedding for **Landscape Animation**. Our model consists of two parts: (1) a motion encoder which embeds time-lapse motion in a fine-grained way. (2) a motion generator which generates realistic motion to animate input images. To train and evaluate on diverse time-lapse videos, we build the largest high-resolution **Time-lapse** video dataset with **Diverse** scenes, namely **Time-lapse-D**, which includes 16,874 video clips with over 10 million frames. Quantitative and qualitative experimental results demonstrate the superiority of our method. In particular, our method achieves relative improvements by 19% on LIPIS and 5.6% on FVD compared with state-of-the-art methods on our dataset. A user study carried out with 700 human subjects shows that our approach visually outperforms existing methods by a large margin.

CCS CONCEPTS

• **Applied computing** → **Media arts**; • **Computing methodologies** → **Image representations**.

*This work was performed when Hongwei Xue was visiting Microsoft Research as a research intern.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '21, October 20–24, 2021, Virtual Event, China

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00

<https://doi.org/10.1145/3474085.3475421>

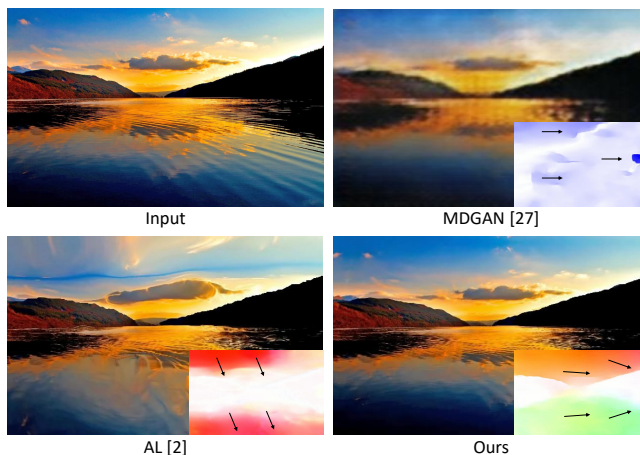


Figure 1: Methods implicitly learning motion (e.g., MDGAN [27]) tend to generate videos with subtle and irrational motion. Explicit motion learning by direct encoding optical flows (e.g., AL [2]) can alleviate this but sometimes generates inaccurate motion. We visualize optical flows to show the motion. Different colors of visualized optical flows represent different motion patterns. The motion direction is roughly marked by black arrows. [Best viewed in color.]

KEYWORDS

image animation, optical flow, datasets

ACM Reference Format:

Hongwei Xue, Bei Liu, Huan Yang, Jianlong Fu, Houqiang Li, and Jiebo Luo. 2021. Learning Fine-Grained Motion Embedding for Landscape Animation. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21)*, October 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3474085.3475421>

1 INTRODUCTION

Image animation aims to generate a time-lapse video from a still image in an automatic way. This topic has drawn great attention in recent years [6, 34, 36], as video representation has become an evolving trend in content consumption that makes people more engaged. More specifically, landscape animation [2, 27, 32] can benefit a broad range of applications, such as video footage production, social engagement boosting, virtual background animation for online

video conferences. Thus, we mainly focus on landscape animation in this paper.

For landscape videos, motion is crucial as it determines how objects move in the video (e.g., moving clouds, flowing water). However, motion learning is quite difficult due to motion’s uncertainty and complexity. Uncertainty means it is difficult to predict the future from one image. For example, clouds in the sky might move left, right, forward, or backward in next frames. Thus, the model requires capabilities of learning diverse motion. Meanwhile, motion in a landscape video is complex as objects usually have different motion patterns and each pattern should be well aligned to the object. Therefore, although motion learning is crucial, it’s challenging to learn diverse and realistic motion.

Existing works on landscape animation model motion in an implicit or explicit way. GAN-based methods use 3D-CNNs to implicitly learn motion from real videos [27]. This implicit learning usually results in subtle and irrational motion, which does not correspond to scene’s layout, such as the result of MDGAN [27] in Figure 1 (right-top). This is because the spatial-temporal domain of videos is too large for generative models to learn. Besides, the lack of motion embedding leads to a lack of diverse video generation capabilities, which limits the practicality of these methods. To solve these problems, some works use an encoder-decoder paradigm to explicitly learn motion by using a convolutional encoder to embed estimated optical flows into a latent space [2, 32]. Relying on optical flows as the representation of motion, these works are able to well learn motion then generate vivid videos. Also, the capacity of generating diverse videos can be supported by sampling from motion embedding latent space.

Although current motion embedding can well handle the uncertainty of future frames, it still suffers from motion’s complexity, especially when an image involves complex scenes. Compared to images, optical flows lack of semantics that can be learned and distinguished by convolutional networks as it is only the derivative of pixel coordinates. Thus, the generated videos often have inaccurate motion, including inaccurate alignment with scene’s layout and unified motion among global frame (see Figure 1, AL [2]). Inaccurate motion may cause the objects in videos to move unrealistically and even distorted.

To better learn motion for a more realistic landscape animation, we propose a novel fine-grained motion embedding learning approach. Our approach is able to embed motion information well then generate accurate and realistic motion. To achieve this, our model consists of a fine-grained motion encoder and a motion generator. The fine-grained motion encoder decomposes the motion space into a fine-grained latent space guided by semantics. To be more specific, our fine-grained motion encoder is based on partial convolution by which motion space can be well decomposed without interference with each other. Moreover, the scaling factor of each partial convolution ensures that the motion information of small objects or textures in the frame will not be ignored. Then a latent map and the input image are spatially aligned then fed into the motion generator to generate motion for animation. The output video is generated frame-by-frame by warping the input image according to the generated motion.

Existing works mostly focus on sky time-lapse videos. To train and evaluate on diverse videos, we introduce a new dataset, namely

Time-lapse-D, which consists of 16,874 high-resolution **Time-lapse** videos of **Diverse** scene with category annotations. The videos in our dataset have over 10 million high-resolution frames in total, of which over 40% are 1280×720 and over 30% are 1920×1080 resolution. To the best of our knowledge, this is the largest time-lapse video dataset so far, which can provide large-scale and high-quality data support for image animation tasks.

Quantitative and qualitative experiments on our dataset demonstrate the superiority of our proposed method. We also conduct a user study to evaluate the comparison of our method with existing works and the confusion ratio with the real video. These experiments demonstrate that our method outperforms existing works by a large margin. Besides, our fine-grained learning makes it possible to control different objects in frames with different motion patterns.

Our key contributions are as follows:

- 1 We propose a novel fine-grained motion embedding method for the landscape animation task to generate high-quality video with realistic motion.
- 2 We build the largest time-lapse video dataset (Time-lapse-D) which includes diverse natural scenes with over 10 million frames.
- 3 We conduct extensive experiments, including subjective and objective comparison experiments and qualitative experiments. Our method achieves relative improvements by 19% on LIPIS and 5.6% on FVD compared with state-of-the-art methods on our dataset.

2 RELATED WORK

In this section, we briefly review related works of general video generation from a still image. Then we specifically summarize related works on landscape animation.

Video Generation from Still Image. Traditional non-learning-based approaches using examples can reproduce realistic motion and content without complex parameters by directly transferring reference videos [16–18, 20, 21]. These approaches heavily depend on large datasets and huge computational resources. Also, existing techniques often impose heavy manual processes for specifying, e.g., alpha mattes, flow fields and regions for fluid.

The past few years have witnessed dramatic advances in learning-based approaches. Some Generative Adversarial Networks (GAN) [3] based models [23, 34, 36] and conditional variational auto-encoders (CVAE) based models [12, 19] automatically generate videos conditioned on the content of input images. Among them, explicitly motion learning is achieved by using sparse trajectories [4], human poses [34], or optical flows [12]. In this paper, we adapt optical flows as motion representation as flow estimation approach has been well developed.

Landscape Animation. Landscape animation targets to generate a time-lapse video from a landscape image. This task can be decomposed into two subtasks: color animation and motion animation [2]. Color animation is animating the color of a landscape at different hours in one day (e.g., time-varying colors in the sky) [1, 15]. Motion animation is animating the scene at a smaller time scale (e.g., moving clouds) [2, 14, 27, 32]. As the data required by these two subtasks are at different time scales, existing works mostly focus on one subtask. We focus on and refer to motion animation in the rest of this paper. It is worth noting that we only refer to motion




Figure 2: The overall pipeline of FGLA. During motion embedding phase, (a) we first extract semantics and motion features by a semantic segmentation module and an optical flow estimation module respectively; (b) a fine-grained motion encoder takes the optical flow and segmentation masks of the input image then generates a latent set. Different colors in latent set Z_t represent different semantics. For frame generation, (c) the semantic segmentation masks and the latent set are reorganized into a latent map which is fed into a generator with the input frame; (d) the generator generates an optical flow for the motion to warp the input frame to obtain the generated frame. Different colors of visualized optical flow represent different directions and the higher saturation means higher speed and vice versa. [Best viewed in color.]

animation part when mention AL [2] as it independently trains two models on different datasets to solve these two subtasks.

MDGAN [27] learns the whole spatial-temporal space of videos and directly generates all video frames. MDGAN presents a 3D-CNN based two-stage approach which generates videos of realistic content for each frame in the first stage and then refines the generated video from the first stage. Although MDGAN enforces generated videos to be closer to real videos with regard to motion dynamics in the second stage, it’s too hard for 3D-CNN to model the whole spatial-temporal domain, resulting in subtle and irrational motion. Besides, color distortion and blur are observed on their results. DTVNet [32] and AL [2] explicitly learn the motion by motion embedding. DTVNet introduces optical flow into GAN and consists of two submodules. An optical flow encoder embeds a sequence of optical flow maps to a normalized motion vector. A dynamic video generator generates all frames from the motion vector and the input image. DTVNet alleviates color distortion but still suffers from the blur problem, and also tends to generate videos with subtle motion. AL [2] uses optical flows in both embedding and generating stages. An encoder-decoder paradigm is used to encode motion into a latent space and then generates optical flow for warping. A video is generated frame-by-frame by warping the input image.

Results generated by AL have high-quality frames without blur or color distortion. However, the insufficiency of motion modeling often leads to inaccurate and global monotonic motion thus fails in complex scenes. In this paper, we improve motion modeling by embedding fine-grained motion guided by semantics. The inaccurate and global motion problem is well alleviated by our approach.

Some methods model motion by sampling from pre-defined transformations [14]. DeepLandscape [14] extends StyleGAN model [10] by augmenting it with learning pairwise consistency of frames in videos. Their generated video frames are sometimes discontinuous due to their method does not learn full temporal dynamics of videos. As a result, They can generate appealing videos of time-lapse clouds but often fails on other scenes.

3 APPROACH

In this section, we will introduce the proposed FGLA model from aspects of model structure, network training and fine-grained controllable generation. Figure 2 shows the whole pipeline of our approach. We first extract semantics by image segmentation and motion features in optical flow. During motion embedding, the fine-grained motion encoder takes a backward optical flow F_t and semantic segmentation masks M_t of the input frame I_t and then generates a

latent set Z_t . Each variable in the set is corresponding to a unique semantic. Then the latent set is organized as a latent map L_t fed into a motion generator which takes the input frame and generates an optical flow \hat{F}_t to warp the input frame. In this section we will introduce the details.

3.1 Fine-grained Motion Embedding

To extract semantic and motion features of an image, we use ADE20k-pretrained [35] segmentation network to obtain a set of segmentation binary masks $M_t^1, M_t^2, \dots, M_t^N$ and use FlowNet 2.0 [8] to estimate optical flow F_t of an image pair (I_t, I_{t+1}) .

One landscape video usually have unique motion patterns for different semantics. To independently encode the motion information of each semantic part, we adopt partial convolution [13] as our extractor. In each layer of partial convolution, the filter weights \mathbf{W} and the corresponding bias b are learnable parameters. \mathbf{X} and \mathbf{M} denote the input feature maps and the corresponding binary mask in each partial convolution layer. The partial convolution is expressed as:

$$x' = \begin{cases} \mathbf{W}^T(\mathbf{X} \odot \mathbf{M}) \frac{\text{sum}(\mathbf{1})}{\text{sum}(\mathbf{M})} + b, & \text{if } \text{sum}(\mathbf{M}) > 0, \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where \odot denotes element-wise multiplication and sum denotes element summation. We use the same way of updating the binary mask as in [13]:

$$m' = \begin{cases} 1, & \text{if } \text{sum}(\mathbf{M}) > 0, \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

Through partial convolutional computing, the output semantic latent variables depend only on their corresponding semantics. The scaling factor $\text{sum}(\mathbf{1})/\text{sum}(\mathbf{M})$ adjusts for the differences in proportion of different semantics.

We denote the fully partial convolution network \mathcal{E} in the encoder. With each binary mask M_t^i in t -th frame of i -th semantics, we feed the flow F_t into the encoder \mathcal{E} to generate latent variable z_t^i for each semantics separately as follows:

$$z_t^i = \mathcal{E}(F_t, M_t^i). \quad (3)$$

As a result, we get a set of latent variables and each variable z_t^i models the motion information of a unique semantic:

$$Z_t = [z_t^1, z_t^2, \dots, z_t^N]. \quad (4)$$

3.2 Latent Mapping and Motion Generation

The main challenge of landscape animation is to drive the given image to transform realistically over time and the content will not change too much. There are two advantages of the flow-warping strategy: realistic and high-fidelity. First, every pixel of generated frames is transformed from pixels of original images, making results realistic. Second, an optical flow is smooth in a small neighborhood. As a result, an optical flow can be interpolated to high resolution. Therefore our model can directly generate high-fidelity landscape videos, without using interpolation or super-resolution techniques to increase frames' size. For above reasons, our generator generates frames based on flow-warping manner: we first generate optical flow for the motion, and then warp the input image with optical flow for next frame.

Algorithm 1 Fine-grained Controllable Landscape Animation

Input: An input image I_0 , A set of reference videos R_1, R_2, \dots, R_L .

Output: A generated video $V = \{\hat{I}_0, \hat{I}_1, \dots, \hat{I}_T\}$.

- 1: Load a motion encoder \mathcal{E} , a motion generator \mathcal{P} ;
 - 2: Let $\hat{I}_0 = I_0$, $t = 0$;
 - 3: **while** $t < T$ **do**
 - 4: Compute segmentation masks $M_t^1, M_t^2, \dots, M_t^N$ for \hat{I}_t ;
 - 5: Compute optical flows F_t^k and segmentation masks $M_t^{1,k}, M_t^{2,k}, \dots, M_t^{N,k}$ for each reference video R_k ;
 - 6: Compute $z_t^{i,k} = \mathcal{E}(F_t^k, M_t^{i,k})$, $i \in [1, N]$ is the i th semantic class. Then get $Z_t^k = [z_t^{1,k}, z_t^{2,k}, \dots, z_t^{N,k}]$ for each reference video R_k ;
 - 7: Controllably get $Z_t = [z_t^{1,r_1}, z_t^{2,r_2}, \dots, z_t^{N,r_N}]$, where z_t^{i,r_i} is the i th semantic latent variable of the r_i th reference video, $r_i \in [1, L]$.
 - 8: Fill a latent map $L_t(x, y) = Z_t[i]$ when $M_t^i(x, y) = 1$ for each pixel (x, y) in latent map;
 - 9: Compute generated optical flow $\hat{F}_t = \mathcal{P}(L_t \oplus \hat{I}_t)$, where \oplus is concatenation;
 - 10: Accumulate flow $\hat{F}_{0 \rightarrow t} = \hat{F}_t + \hat{F}_{0 \rightarrow t-1}$ if $t > 0$ else \hat{F}_0
 - 11: Warp the input image I_0 to generate frame $\hat{I}_{t+1} = \text{Warp}(I_0, \hat{F}_{0 \rightarrow t})$
 - 12: $t = t + 1$;
 - 13: **end while**
-

The learned semantic latent variables are supposed to represent motion of each semantics without the spatial information of videos. Thus, we propose a latent mapping module to combine the image content and its corresponding latent variables into one latent map. Given the given landscape image I_t and the obtained latent Z_t , we reorganize the latent set as the latent map L_t by filling the semantic mask M_t with the corresponding latent as follows:

$$L_t(x, y) = z_k \quad \text{when} \quad M_t^k(x, y) = 1, \quad (5)$$

where (x, y) is coordinate. Then we concatenate L_t with I_t and feed them into the motion generator to generate optical flow:

$$\hat{F}_t = \mathcal{P}(L_t \oplus I_t). \quad (6)$$

Similar to [2, 19], we use skip-connections in motion generator. This U-Net [9] like architecture generates an optical flow that spatially corresponds well to the image.

Using \tanh , our generator infers an optical flow \hat{F}_t which value ranges from -1 to 1. Similar to [2], we divide the flow by a constant $c > 1$ since the objects do not move significantly in a single timestep. While we choose a larger c to generate optical flow with a larger amplitude range. In this paper, we set c to 32.

With the generated optical flow \hat{F}_t , the input image I_t is warped to generate next frame \hat{I}_{t+1} as follows:

$$\hat{I}_{t+1} = \text{Warp}(I_t, \hat{F}_t). \quad (7)$$

3.3 Network Training

For supervised training of our landscape animation model, we require a ground-truth video of each input image. However, this kind of ground-truth videos are very difficult to be obtained. For this reason, we utilize a self-supervised training mechanism by




Figure 3: Comparisons of frames sampled from videos generated by MDGAN [27], DTVNet [32], AL [2] and ours. To view details, we zoom in the last frames. We highlight some defects with red rectangles. [Best viewed in color.]

using the horizontally flipped version as the ground-truth video similar to [1]. Flipping can prevent our model from learning the spatial correlation between input images and ground-truth videos. To keep the value of the optical flow constant in each pixel, we conduct the flip operation after estimating the optical flow on the original image.

The training objectives of the landscape animation framework consists of three parts: frame loss, flow loss, and edge-preserving smoothing loss. We adopt a L1 loss to compute distance between the generated output frame \hat{I}_{t+1} and the ground-truth image I_{t+1} :

$$\mathcal{L}_{frame} = \|\hat{I}_{t+1} - I_{t+1}\|_1. \quad (8)$$

As the warping operation is sampling and interpolating pixels on the original image, the frame loss has very limited effect in areas with small color changes or some color patterns. We add an L1 loss to measure the difference of generated optical flow \hat{F} and the ground-truth optical flow F in each timestep:

$$\mathcal{L}_{flow} = \|\hat{F}_t - F_t\|_1. \quad (9)$$

Similar to [2], we adopt a weighted total variation loss applied to the output flow field for the purpose of edge-preserving smoothing:

$$\mathcal{L}_{tv} = \sum_{\mathbf{p}, \mathbf{q} \in N(\mathbf{p})} w \left(\hat{I}_{t+1}(\mathbf{p}), \hat{I}_{t+1}(\mathbf{q}) \right) \|\hat{F}_t(\mathbf{p}) - \hat{F}_t(\mathbf{q})\|_1, \quad (10)$$

$$w(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|_1}{\sigma}\right), \quad (11)$$

where $N(\mathbf{p})$ indicates the right and above neighbors of \mathbf{p} . This loss function makes the model tend to output close optical flows in areas where the pixel value changes little. The final objective of our model is the combination of the above three losses:

$$\mathcal{L}_{total} = \mathcal{L}_{frame} + \alpha * \mathcal{L}_{flow} + \beta * \mathcal{L}_{tv}, \quad (12)$$

where α and β are the weights of \mathcal{L}_{flow} and \mathcal{L}_{tv} respectively.

3.4 Fine-grained Controllable Animation

Instead of randomly sampling from a latent codebook or latent space, our approach makes it possible to generate videos in a fine-grained controllable manner. Each variable in the latent set represents one unique semantics. This kind of latent set makes it possible to transfer semantic motion patterns from one or multiple reference

Table 1: Comparison between Sky Time-lapse [27] and our proposed dataset.

	Sky Time-lapse [27]	Time-lapse-D
Clips	2049	16874
Frames	1.2×10^6	1.0×10^7
Resolution	640×360	40% 1280×720 30% 1920×1080
Scene Categories	2	13

videos. We can simply get latent Z_t with different semantics from multiple references R_1, R_2, \dots, R_L :

$$Z_t = [z_t^{1,r_1}, z_t^{2,r_2}, \dots, z_t^{N,r_N}], \quad (13)$$

where z_t^{i,r_i} is the i th semantic latent variable of the r_i th reference video, $r_i \in [1, L]$. The whole inference procedure is shown in Algorithm 1.

During inference, similar to [2], we recurrently use the generated frame \hat{I}_t as the input of next motion generation while we warp optical flows sequentially to accumulate flows so that we can reconstruct each output frame from the first frame. Through this design, we can avoid generating gradually blurry frames due to error accumulation in multiple sampling.

4 EXPERIMENTS

4.1 Dataset

Time-lapse videos are usually used for the experiments of landscape animation. One recent dataset that is used by [2, 32] is Sky Time-lapse dataset [27], which contains dynamic sky scenes, such as the cloudy sky with moving clouds and the starry sky with moving stars. There are 1,825 training video clips and 224 testing video clips. Note that by video clips we refer to real video clips instead of divided fixed-length clips. The original size of each frame is $3 \times 640 \times 360$. This dataset contains a small number of video clips and is limited to sky scenes. To better model the motion of time-lapse videos, we build a new large dataset by collecting more than 10,000 high-resolution time-lapse videos from Flickr. By PySceneDetect¹, these videos are automatically cut into video clips based on the scene

¹<https://pyscenedetect.readthedocs.io>

Table 2: Comparison with existing works and ablation study on Time-lapse-D. The \uparrow indicates that the larger the value, the better the model performance, and vice versa for \downarrow . For PSNR and SSIM, we mask out dynamic region pixels.

Method	Masked PSNR \uparrow	Masked SSIM \uparrow	FID \downarrow	LPIPS \downarrow	FVD \downarrow
MDGAN [27]	27.85	0.841	60.38	0.225	1177.75
DTVNet [32]	31.41	0.904	59.37	0.165	1007.81
AL [2]	42.77	0.981	59.08	0.026	831.46
Ours w/o LM	41.38	0.981	59.27	0.028	901.62
Ours w/o flip	44.23	0.980	58.78	0.022	803.26
Ours w/ \mathcal{L}_{flow}	42.01	0.979	58.79	0.025	799.71
Ours w/ \mathcal{L}_{frame}	45.61	0.984	59.03	0.025	802.42
Ours w/o \mathcal{L}_{tv}	44.99	0.980	58.95	0.024	787.33
Ours	45.53	0.982	58.77	0.021	784.99

Table 3: User study results. Our approach outperforms other approaches by a large margin.

Comparison Methods	User Preference
MDGAN vs. Ours	4 vs. 96
DTVNet vs. Ours	6 vs. 94
AL vs. Ours	22 vs. 78
MDGAN vs. GT	2 vs. 98
DTVNet vs. GT	4 vs. 96
AL vs. GT	13 vs. 87
Ours vs. GT	31 vs. 69

variety. We utilize crowdsourcing to manually filter out clips that have obvious camera movement or quick zoom-in/zoom-out. We then have the remaining video clips annotated with categories of different dynamic scenes, such as sky, water and hayfield.

The dataset is named as **Time-lapse-D**, in which “D” indicates diverse scenes. The statistics of Time-lapse-D are shown in Table 1. Time-lapse-D contains 16,874 video clips with diverse scene categories in total, includes 14,418 clips for training and 2,456 clips for testing. It is about 8 times larger than Sky Time-lapse dataset [27]. The videos in the test set are selected from different scene categories in the same proportion, to ensure that the test set contains all categories of scenes. Each video clip contains 607 frames on average. Most of the clips are with high resolution, with over 40% are 1280×720 and over 30% are 1920×1080 . The dynamic patterns in our dataset are much more diverse, including common categories (i.e., cloud, star, water, waterfall, hayfield, traffic, smoke, tree, sun, moon, aurora, fire, fog).

This is the largest high-resolution time-lapse video dataset with the most diverse categories. In this paper, we mainly focus on five moving scenes that are feasible for flow-warping (i.e., sky, tree, grass, water and waterfall). Even though some categories (e.g., traffic and smoke) are not our targets in this work, this dataset provides high-quality and diverse video data to facilitate the research of video prediction, classification, generation and dynamic texture synthesis.

4.2 Implementation Details

In this work, we use FlowNet 2.0 [8] as our optical flow estimator. The FlowNet 2.0 model takes two $3 \times 256 \times 256$ frames as input and estimates a $2 \times 256 \times 256$ backward optical flow as output. We use

HRNetV2 [22] as our segmentation model and fix the resolution of all outputs to 256×256 . As we focus on five different types of scenes, we aggregate the categories of the data into six groups: sky, tree, grass, water, waterfall and others. We horizontally flip and downsample the flow and semantic masks, then feed them into the fine-grained motion encoder. The encoder takes a $2 \times 128 \times 128$ flow and a $6 \times 128 \times 128$ mask as input and outputs a 6×2 latent matrix. As described in Section 3, we get a $2 \times 256 \times 256$ latent map and feed it into the generator. The motion generator takes the latent map and a $3 \times 256 \times 256$ frame as input to generate a $2 \times 256 \times 256$ optical flow. During training stage, we use Adam [11] optimizer for all modules with learning rate of $1e^{-4}$. Similar to [2], Instance Normalization [24] followed by Leaky ReLU [28] with slope of 0.1 is used in all the convolution layers in both encoder and generator. We empirically set $\sigma = 0.1$ for edge-preserving smoothing loss. The weights of the frame reconstruction loss, flow reconstruction loss and smoothing loss α and β are set to 0.5 and 5 respectively.

To train our model, we first sample one from every four frames for each clip to avoid learning motion that is too subtle. Then we resize the frames into square images of size 256×256 . Before feeding the frames to the model, we normalize the color values to $[-1, 1]$. No additional pre-processing is required. To avoid training biases by long video clips, we train each pair of frames sampled randomly for each video clip in each epoch. Under the above settings, our model converges in 500 epochs.

4.3 Comparison with State-of-the-art methods

We conduct a series of quantitative, qualitative experiments and user study to verify the quality of generated videos by our model comparing with existing methods. Three recent DL-based models with released codes are used for comparison: MDGAN [27], DTVNet [32] and Animating-Landscape (AL) [2]. We also try to train Deeplandscape [14] on our dataset, but the model is hard to converge and shows unsatisfactory quality. They also reported in paper that their model fails on their new video dataset due to the dataset’s scale and motion’s distribution. Although our dataset is large, the diversity of scenes and motion make it more difficult for a StyleGAN to converge. Besides, lacking of full temporal dynamics damages the visual effect of scenes other than clouds. Thus, we only do partial experiments on Sky-Timelapse dataset.

All methods use the first frame of videos in the test set as input and output 32 frames as a sequence at 256×256 resolution. For

Table 4: Comparison with existing works and ablation study on Sky Time-lapse [27]. The \uparrow indicates that the larger the value, the better the model performance, and vice versa for \downarrow . For PSNR and SSIM, we mask out dynamic region pixels.

Method	Masked PSNR \uparrow	Masked SSIM \uparrow	FID \downarrow	LPIPS \downarrow	FVD \downarrow
MDGAN [27]	31.43	0.911	67.32	0.186	1086.66
DTVNet [32]	35.80	0.950	67.34	0.149	774.85
AL [2]	38.65	0.961	62.59	0.058	441.35
Deeplandscape [14]	43.32	0.981	61.91	0.038	411.08
Ours w/o LM	36.38	0.943	63.68	0.053	830.09
Ours w/o flip	45.99	0.981	61.82	0.025	367.77
Ours w/ \mathcal{L}_{flow}	41.97	0.958	61.73	0.034	365.62
Ours w/ \mathcal{L}_{frame}	46.94	0.989	61.98	0.029	370.56
Ours w/o \mathcal{L}_{tv}	46.02	0.983	61.75	0.028	361.57
Ours	46.85	0.986	61.73	0.025	354.60

MDGAN, we directly fine-tune their released pre-trained model on our new dataset since the two-stage generative model is difficult to train. MDGAN lacks of diverse video generation capabilities and only one output is generated for one input. Thus we can only generate one instance for each test video.

For DTVNet, although the motion vector is normalized, the model cannot guarantee that the flow encoder’s output following a certain distribution during training. Thus we randomly sample 5 videos to calculate optical flows using FlowNet 2.0 before feeding them into the flow encoder.

For AL, we use the same way of saving codebook as motion latent in the paper. During evaluation, five latent codes are randomly sampled from the codebook to generate 5 sequences and we take the average evaluation results of these 5 sequences.

For our model, we randomly sample 5 videos’ embeddings and use them separately to generate 5 videos. Evaluation of each metric is the average of these 5 results.

Quantitative Evaluation. We choose PSNR, SSIM [26], FID (Fréchet Inception Distance) [5] and LPIPS [33] to evaluate both quality and consistency of generated frames in output videos. These metrics are calculated between first frame and generated video frames. FID and LPIPS are respectively used to measure the quality of generated frames and perceptual dissimilarity between each generated frame and the corresponding first frame. For PSNR and SSIM, we measure the static consistency, which means we need a mask to mask out dynamic region pixels. Given L continuous frames, we calculate the mask as follows:

$$mask(i, j) = \frac{1}{L} \sum_{i \in [0, L)} \|I_{i+1}^{GT}(i, j) - I_i^{GT}(i, j)\|_1 > t, \quad (14)$$

where threshold t is to 2.5 in our experiments. However, perfect image quality and static consistency can be achieved by simply not animating anything at all and they cannot reflect the quality of generated videos. Thus we use Fréchet Video Distance [25] to measure the quality of generated videos, especially on their motion. User studies are further conducted to evaluate their quality from human sense.

We conduct quantitative experiments on two datasets: our Time-lapse-D and Sky Time-lapse [27]. For evaluation on Sky Time-lapse, we directly use official released models of the three methods. The comparison results are shown in Table 2 and 4. Our model gains big improvements on LPIPS and FVD compared with other methods.

These two metrics are more close to human perception measures. LPIPS is insensitive to pixel displacement and FVD measures the difference between the generated videos and real videos. Our huge improvement in these two metrics demonstrates that our model can generate more realistic videos. By the better performance of AL and our approach, we found that optical flow can generate better frames compared with 3D-CNN-based methods. This is because it is difficult for the 3D-CNN-based methods to directly generate high-quality frames. Overall, evaluation results indicate that our approach outperforms the other three baselines in all metrics, which illustrates that our model can generate more high-quality and realistic videos than other methods.

Qualitative Evaluation. We give two examples to visually show the result comparison between ours and other approaches in Figure 3. As in red rectangles that show explicit defects, MDGAN suffers from the distortion in color. DTVNet tends to generate blur frames gradually, which can also be observed from the results shown in their paper. Although AL can generate frames with good details, it often generates distorted motion when multiple semantics with different motion are included in the same frame, which results in unrealistic content distortion in generated frames. Results generated by our method preserve details in high-fidelity while having more obvious motion (as shown by the position of clouds). These results show that our model can generate results with good quality of both content and motion compared with other approaches.

User Study. To better evaluate the video quality, we further perform a user study. we randomly choose 100 start images from our test set and generate corresponding videos by different methods using the same settings as described in quantitative evaluation. Then the generated videos by different methods are simultaneously shown to annotators and we can get a result indicating “which one is better”. Totally, the aforementioned experiment with 700 human subjects is conducted by 14 real workers.

Comparison results shown in Table 3 demonstrate that our approach outperforms the other three methods by a large margin and our superiority is prominent when comparing with MDGAN and DTVNet. Comparison results with ground-truth videos show that although it is still challenging to generate results that are completely comparable to real videos, our method has the ability to generate some results that can be confused as real videos.




Figure 4: Visualizations of frames sampled from videos generated with references for water and sky semantics. Different reference videos are used to independently animate different semantics (i.e., sky and water in this figure) of the input image. We sample frames to show the content and visualize optical flows to show the motion. Different colors of visualized optical flow represent different directions and the higher saturation means higher speed and vice versa. [Best viewed in color.]




Figure 5: Sampling from generated videos using reference at different speeds: 0.25x, 1x, 4x. We sample frames to show the content and visualize optical flows to show the motion. Different colors of visualized optical flow represent different directions and the higher saturation means higher speed and vice versa. [Best viewed in color.]

4.4 Ablation Study

We design ablation studies to show how each proposal contributes to the final performance. In Table 2 and 4, w/o LM denotes feeding flow concatenated with segmentation map into a CNN encoder and thus no latent mapping. w/o flip indicates no flipping during training. w/ \mathcal{L}_{flow} and w/ \mathcal{L}_{frame} respectively indicate only use \mathcal{L}_{flow} and \mathcal{L}_{frame} as loss function. w/o \mathcal{L}_{tv} removes the \mathcal{L}_{tv} item compared with our full model. From the results, our full model gains the best performance on most metrics.

4.5 Fine-grained Control of Generation

In this subsection we use some qualitative results to show the capability of our method to generate videos under fine-grained control. The results are shown in Figure 4 and Figure 5. To facilitate the presentation in the paper, we visualize the optical flow to show the motion of the video. Different colors of visualized optical flow represent different directions and the higher saturation means higher speed and vice versa. As our method can generate videos of any resolution, we unify the resolution of all results to 640×360 .

In Figure 4, we choose sky and water as two representative semantics in this experiment. From top to down are three animating results from the same input image using different references. First and second row refer to different video references for sky and the same reference for water. According to visualized optical flows,

the generated results have similar motion patterns to reference videos on each semantics. First and third row shows that we can controllably fix motion of one semantics by using reference without that semantics. These results demonstrate that our method avoids the global motion problem on multiple semantics. In Figure 5, we show that our method can also control the speed of motion.

As our approach encodes motion in a fine-grained manner, two benefits are achieved. First, we can controllably animate each semantics in input images to achieve the desired dynamic effect. Secondly, each reference video can be used to extract motion information of one or a few semantics. The constraint of reference videos is greatly reduced, which is obvious while animating an image with many semantics.

5 CONCLUSIONS

In this paper, we propose a fine-grained motion embedding approach to animate single landscape images to time-lapse videos with accurate and realistic motion. Our model consists of two parts: a fine-grained motion encoder independently encodes the motion patterns of different semantics. Then a motion generator generates motion for animation. The result video is generated by warping frames with the generated motion in an accumulated way. We collect the largest high-resolution time-lapse video dataset (Time-lapse-D) with multiple scenes. Experimental results show that our approach can generate more realistic and high quality videos compared to state-of-the-art methods by both quantitative and qualitative evaluations. In the future, we will study on robust approach to deal with more diverse scenes, such as flower blossoming and traffic flow from the view of style transfer [7]. And methods used in video inpainting [30, 31] and super resolution [29] can also be applied to overcome the occlusion.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (NSFC) under Grants 61836011 and 62021001. The authors would like to thank all users for voluntarily participating our user study.

REFERENCES

- [1] Chia-Chi Cheng, Hung-Yu Chen, and Wei-Chen Chiu. 2020. Time Flies: Animating a Still Image With Time-Lapse Video As Reference. In *CVPR*. 5641–5650.
- [2] Yuki Endo, Yoshihiro Kanamori, and Shigeru Kuriyama. 2019. Animating landscape: self-supervised learning of decoupled motion and appearance for single-image video synthesis. *TOG* 38, 6 (2019), 1–19.
- [3] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *NeurIPS*. 2672–2680.
- [4] Zekun Hao, Xun Huang, and Serge Belongie. 2018. Controllable video generation with sparse trajectories. In *CVPR*. 7854–7863.
- [5] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*. 6626–6637.
- [6] Qiyang Hu, Adrian Waelchli, Tiziano Portenier, Matthias Zwicker, and Paolo Favaro. 2018. Video synthesis from a single image and motion stroke. *arXiv preprint arXiv:1812.01874* (2018).
- [7] Zhiyuan Hu, Jia Jia, Bei Liu, Yaohua Bu, and Jianlong Fu. 2020. Aesthetic-aware image style transfer. In *Proceedings of the 28th ACM International Conference on Multimedia*. 3320–3329.
- [8] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. 2017. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*. 2462–2470.
- [9] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *CVPR*. 1125–1134.
- [10] Tero Karras, Samuli Laine, and Timo Aila. 2019. A style-based generator architecture for generative adversarial networks. In *CVPR*. 4401–4410.
- [11] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [12] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. 2018. Flow-grounded spatial-temporal video prediction from still images. In *ECCV*. 600–615.
- [13] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. 2018. Image inpainting for irregular holes using partial convolutions. In *ECCV*. 85–100.
- [14] Elizaveta Logacheva, Roman Suvorov, Oleg Khomenko, Anton Mashikhin, and Victor Lempitsky. 2020. DeepLandscape: adversarial modeling of landscape video. In *ECCV*. 256–272.
- [15] Seonghyeon Nam, Chongyang Ma, Menglei Chai, William Brendel, Ning Xu, and Seon Joo Kim. 2019. End-to-end time-lapse video synthesis from a single outdoor image. In *CVPR*. 1409–1418.
- [16] Makoto Okabe, Ken Anjyo, Takeo Igarashi, and Hans-Peter Seidel. 2009. Animating pictures of fluid using video examples. In *CGF*, Vol. 28. 677–686.
- [17] Makoto Okabe, Ken Anjyo, and Rikio Onai. 2011. Creating fluid animation from a single image using video database. In *CGF*, Vol. 30. 1973–1982.
- [18] Makoto Okabe, Yoshinori Dobashi, and Ken Anjyo. 2018. Animating pictures of water scenes using video retrieval. *The Visual Computer* 34, 3 (2018), 347–358.
- [19] Junting Pan, Chengyu Wang, Xu Jia, Jing Shao, Lu Sheng, Junjie Yan, and Xiaogang Wang. 2019. Video generation from single semantic label map. In *CVPR*. 3733–3742.
- [20] Ekta Prashnani, Maneli Noorkami, Daniel Vaquero, and Pradeep Sen. 2017. A phase-based approach for animating images using video examples. In *CGF*, Vol. 36. 303–311.
- [21] Yichang Shih, Sylvain Paris, Frédo Durand, and William T Freeman. 2013. Data-driven hallucination of different times of day from a single outdoor photo. *TOG* 32, 6 (2013), 1–11.
- [22] Ke Sun, Yang Zhao, Borui Jiang, Tianheng Cheng, Bin Xiao, Dong Liu, Yadong Mu, Xinggang Wang, Wenyu Liu, and Jingdong Wang. 2019. High-resolution representations for labeling pixels and regions. *arXiv preprint arXiv:1904.04514* (2019).
- [23] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. 2018. Mocogan: Decomposing motion and content for video generation. In *CVPR*. 1526–1535.
- [24] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. 2016. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022* (2016).
- [25] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. 2018. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717* (2018).
- [26] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *TIP* 13, 4 (2004), 600–612.
- [27] Wei Xiong, Wenhan Luo, Lin Ma, Wei Liu, and Jiebo Luo. 2018. Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks. In *CVPR*. 2364–2373.
- [28] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. 2015. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853* (2015).
- [29] Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Baining Guo. 2020. Learning Texture Transformer Network for Image Super-Resolution. In *CVPR*.
- [30] Yanhong Zeng, Jianlong Fu, and Hongyang Chao. 2020. Learning Joint Spatial-Temporal Transformations for Video Inpainting. In *ECCV*.
- [31] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Baining Guo. 2019. Learning Pyramid-Context Encoder Network for High-Quality Image Inpainting. In *CVPR*. 1486–1494.
- [32] Jiangning Zhang, Chao Xu, Liang Liu, Mengmeng Wang, Xia Wu, Yong Liu, and Yunliang Jiang. 2020. DTVNet: Dynamic Time-lapse Video Generation via Single Still Image. In *ECCV*. 300–315.
- [33] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*. 586–595.
- [34] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris Metaxas. 2018. Learning to forecast and refine residual motion for image-to-video generation. In *ECCV*. 387–403.
- [35] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. 2019. Semantic understanding of scenes through the ADE20K dataset. *IJCV* 127, 3 (2019), 302–321.
- [36] Yipin Zhou, Yale Song, and Tamara L Berg. 2018. Image2GIF: Generating cinemagraphs using recurrent deep Q-Networks. In *WACV*. 170–178.