

Generative Image Dynamics

Zhengqi Li

Richard Tucker

Noah Snavely

Aleksander Holynski

Google Research

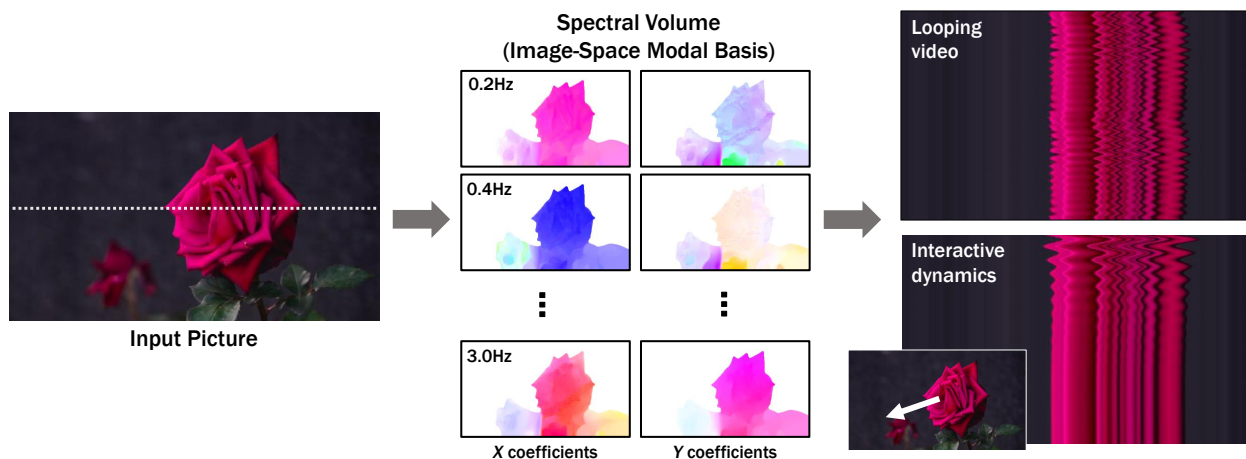


Figure 1. We model a generative image-space prior on scene motion: from a single RGB image, our method generates a *spectral volume* [23], a motion representation that models dense, long-term pixel trajectories in the Fourier domain. Our learned motion priors can be used to turn a single picture into a seamlessly looping video, or into an interactive simulation of dynamics that responds to user inputs like dragging and releasing points. On the right, we visualize output videos as space-time X - t slices (along the input scanline shown on the left).

Abstract

We present an approach to modeling an image-space prior on scene motion. Our prior is learned from a collection of motion trajectories extracted from real video sequences depicting natural, oscillatory dynamics of objects such as trees, flowers, candles, and clothes swaying in the wind. We model dense, long-term motion in the Fourier domain as spectral volumes, which we find are well-suited to prediction with diffusion models. Given a single image, our trained model uses a frequency-coordinated diffusion sampling process to predict a spectral volume, which can be converted into a motion texture that spans an entire video. Along with an image-based rendering module, the predicted motion representation can be used for a number of downstream applications, such as turning still images into seamlessly looping videos, or allowing users to interact with objects in real images, producing realistic simulated dynamics (by interpreting the spectral volumes as image-space modal bases). See our project page for more results: generative-dynamics.github.io.

1. Introduction

The natural world is always in motion, with even seemingly static scenes containing subtle oscillations as a result of wind, water currents, respiration, or other natural rhythms. Emulating this motion is crucial in visual content synthesis—human sensitivity to motion can cause imagery without motion (or with slightly unrealistic motion) to seem uncanny or unreal.

While it is easy for humans to interpret or imagine motion in scenes, training a model to learn or produce realistic scene motion is far from trivial. The motion we observe in the world is the result of a scene’s underlying physical dynamics, i.e., forces applied to objects that respond according to their unique physical properties—their mass, elasticity, etc—quantities that are hard to measure and capture at scale. Fortunately, measuring them is not necessary for certain applications: e.g., one can simulate plausible dynamics in a scene by simply analyzing some observed 2D motion [23].

This same observed motion can also serve as a supervisory signal in learning dynamics across scenes—because although observed motion is multi-modal and grounded in complex physical effects, it is nevertheless often predictable:

candles will flicker in certain ways, trees will sway, and their leaves will rustle. For humans, this predictability is ingrained in our systems of perception: by viewing a still image, we can imagine plausible motions— or, since there might have been many possible such motions, a *distribution* of natural motions conditioned on that image. Given the facility with which humans are able to model these distributions, a natural research problem is to model them computationally.

Recent advances in generative models, in particular conditional diffusion models [44, 84, 86], have enabled us to model rich distributions, including distributions of real images conditioned on text [72–74]. This capability has enabled several new applications, such as text-conditioned generation of diverse and realistic image content. Following the success of these image models, recent work has extended these models to other domains, such as videos [7, 43] and 3D geometry [76, 99, 100, 102].

In this paper, we model a generative prior for *image-space scene motion*, i.e., the motion of all pixels in a single image. This model is trained on motion trajectories automatically extracted from a large collection of real video sequences. In particular, from each training video we compute motion in the form of a *spectral volume* [22, 23], a frequency-domain representation of dense, long-range pixel trajectories. Spectral volumes are well-suited to scenes that exhibit oscillatory dynamics, e.g., trees and flowers moving in the wind. We find that this representation is also highly effective as an output of a diffusion model for modeling scene motion. We train a generative model that, conditioned on a single image, can sample spectral volumes from its learned distribution. A predicted spectral volume can then be directly transformed into a motion texture—a set of long-range, per-pixel motion trajectories—that can be used to animate the image. The spectral volume can also be interpreted as an *image-space modal basis* for use in simulating interactive dynamics [22].

We predict spectral volumes from input images using a diffusion model that generates coefficients one frequency at a time, but coordinates these predictions across frequency bands through a shared attention module. The predicted motions can be used to synthesize future frames (via an image-based rendering model)—turning still images into realistic animations, as illustrated in Fig. 1.

Compared with priors over raw RGB pixels, priors over motion capture more fundamental, lower-dimensional structure that efficiently explains long-range variations in pixel values. Hence, generating intermediate motion leads to more coherent long-term generation and more fine-grained control over animations. We demonstrate the use of our trained model in several downstream applications, such as creating seamless looping videos, editing the generated motions, and enabling interactive dynamic images via image-space modal bases, i.e., simulating the response of object dynamics to user-applied forces [22].

2. Related Work

Generative synthesis. Recent advances in generative models have enabled photorealistic synthesis of images conditioned on text prompts [16, 17, 24, 72–74]. These text-to-image models can be augmented to synthesize video sequences by extending the generated image tensors along a time dimension [7, 9, 43, 62, 83, 105, 105, 110]. While these methods can produce video sequences that capture the spatiotemporal statistics of real footage, these videos often suffer from artifacts like incoherent motion, unrealistic temporal variation in textures, and violations of physical constraints like preservation of mass.

Animating images. Instead of generating videos entirely from text, other techniques take as input a still picture and animate it. Many recent deep learning methods adopt a 3D-Unet architecture to produce video volumes directly [27, 36, 40, 47, 53, 92]. These models are effectively the same video generation models (but conditioned on image information instead of text), and exhibit similar artifacts to those mentioned above. One way to overcome these limitations is to not directly generate the video content itself, but instead animate an input source image through image-based rendering, i.e., moving the image content around according to motion derived from external sources such as a driving video [51, 79–81, 98], motion or 3D geometry priors [8, 29, 46, 63, 64, 66, 89, 96, 100, 101, 103, 108], or user annotations [6, 18, 20, 33, 38, 97, 104, 107]. Animating images according to motion fields yields greater temporal coherence and realism, but these prior methods either require additional guidance signals or user input, or utilize limited motion representations.

Motion models and motion priors. In computer graphics, natural, oscillatory 3D motion (e.g., water rippling or trees waving in the wind) can be modeled with noise that is shaped in the Fourier domain and then converted to time-domain motion fields [78, 87]. Some of these methods rely on a modal analysis of the underlying dynamics of the system being simulated [22, 25, 88]. These spectral techniques were adapted to animate plants, water, and clouds from single 2D pictures by Chuang *et al.* [20], given user annotations. Our work is especially inspired by Davis [23], who connected modal analysis of a scene with the motions observed in a video of that scene, and used this analysis to simulate interactive dynamics from a video. We adopt the frequency-space *spectral volume* motion representation from Davis *et al.*, extract this representation from a large set of training videos, and show that spectral volumes are suitable for predicting motion from single images with diffusion models.

Other methods have used various motion representations in *prediction* tasks, where an image or video is used to inform a deterministic future motion estimate [34, 70], or a more rich *distribution* of possible motions [93, 95, 103]. However,

many of these methods predict an optical flow motion estimate (i.e., the instantaneous motion of each pixel), not full motion trajectories. In addition, much of this prior work is focused on tasks like activity recognition, not on synthesis tasks. More recent work has demonstrated the advantages of modeling and predicting motion using generative models in a number of closed-domain settings such as humans and animals [2, 19, 28, 71, 90, 106].

Videos as textures. Certain moving scenes can be thought of as a kind of texture—termed *dynamic textures* [26]—that model videos as space-time samples of a stochastic process. Dynamic textures can represent smooth, natural motions like waves, flames, or moving trees, and have been widely used for video classification, segmentation or encoding [12–15, 75]. A related kind of texture, called a *video texture*, represents a moving scene as a set of input video frames along with transition probabilities between any pair of frames [65, 77]. A number of methods estimate dynamic or video textures through analysis of scene motion and pixel statistics, with the aim of generating seamlessly looping or infinitely varying output videos [1, 21, 32, 58, 59, 77]. In contrast to much of this work, our method learns priors in advance that can then be applied to single images.

3. Overview

Given a single picture I_0 , our goal is to generate a video $\{\hat{I}_1, \hat{I}_2, \dots, \hat{I}_T\}$ featuring oscillatory motions such as those of trees, flowers, or candle flames swaying in the breeze. Our system consists of two modules: a motion prediction module and an image-based rendering module. Our pipeline begins by using a latent diffusion model (LDM) to predict a spectral volume $\mathcal{S} = (S_{f_0}, S_{f_1}, \dots, S_{f_{K-1}})$ for the input I_0 . The predicted spectral volume is then transformed to a motion texture $\mathcal{F} = (F_1, F_2, \dots, F_T)$ through an inverse discrete Fourier transform. This motion determines the position of each input pixel at every future time step.

Given a predicted motion texture, we then animate the input RGB image using a neural image-based rendering technique (Sec. 5). We explore applications of this method, including producing seamless looping animations and simulating interactive dynamics, in Sec. 6.

4. Predicting motion

4.1. Motion representation

Formally, a motion texture is a sequence of time-varying 2D displacement maps $\mathcal{F} = \{F_t | t = 1, \dots, T\}$, where the 2D displacement vector $F_t(\mathbf{p})$ at each pixel coordinate \mathbf{p} from input image I_0 defines the position of that pixel at a future time t [20]. To generate a future frame at time t , one can splat pixels from I_0 using the corresponding displacement




Figure 2. **Left:** We visualize the average power spectrum for the X and Y motion components extracted from real videos, shown as the blue and green curves. Natural oscillation motions are composed primarily of low-frequency components, and so we use the first $K = 16$ terms, marked with red dots. **Right:** we show a histogram of the amplitude of Fourier terms at 3.0 Hz after (1) scaling amplitude by image width and height (blue), or (2) frequency adaptive normalization (red). Our adaptive normalization prevents the coefficients from concentrating at extreme values.

map D_t , resulting in a forward-warped image I'_t :

$$I'_t(\mathbf{p} + F_t(\mathbf{p})) = I_0(\mathbf{p}). \quad (1)$$

If our goal is to produce a video via a motion texture, then one choice would be to predict a time-domain motion texture directly from an input image. However, the size of the motion texture would need to scale with the length of the video: generating T output frames implies predicting T displacement fields. To avoid predicting such a large output representation for long videos, many prior animation methods either generate video frames autoregressively [7, 29, 57, 60, 92], or predict each future output frame independently via an extra time embedding [4]. However, neither strategy ensures long-term temporal consistency of generated videos.

Fortunately, many natural motions can be described as a superposition of a small number of harmonic oscillators represented with different frequencies, amplitude and phases [20, 23, 25, 50, 68]. Because these underlying motions are quasi-periodic, it is natural to model them in the frequency domain. Hence, we adopt from Davis *et al.* [23] an efficient frequency space representation of motion in a video called a *spectral volume*, visualized in Fig. 3. A spectral volume is the temporal Fourier transform of per-pixel trajectories extracted from a video.

Given this motion representation, we formulate the motion prediction problem as a multi-modal image-to-image translation task: from an input image to an output motion spectral volume. We adopt latent diffusion models (LDMs) to generate spectral volumes comprised of a $4K$ -channel 2D motion spectrum map, where $K \ll T$ is the number of frequencies modeled, and where at each frequency we need four scalars to represent the complex Fourier coefficients for the x - and y -dimensions. Note that the motion trajectory of a pixel at future time steps $\mathcal{F}(\mathbf{p}) = \{F_t(\mathbf{p}) | t = 1, 2, \dots, T\}$ and its representation as a spectral volume $\mathcal{S}(\mathbf{p}) = \{S_{f_k}(\mathbf{p}) | k = 0, 1, \dots, \frac{T}{2} - 1\}$ are related by the Fast Fourier transform (FFT):

$$\mathcal{S}(\mathbf{p}) = \text{FFT}(\mathcal{F}(\mathbf{p})). \quad (2)$$




Figure 3. **Motion prediction module.** We predict a spectral volume S through a frequency-coordinated denoising model. Each block of the diffusion network ϵ_θ interleaves 2D spatial layers with attention layers (red box, right), and iteratively denoises latent features z^n . The denoised features are fed to a decoder \mathcal{D} to produce S . During training, we concatenate the downsampled input I_0 with noisy latent features encoded from a real motion texture via an encoder \mathcal{E} , and replace the noisy features with Gaussian noise z^N during inference (left).

How should we select the K output frequencies? Prior work in real-time animation has observed that most natural oscillation motions are composed primarily of low-frequency components [25, 68]. To validate this observation, we computed the average power spectrum of the motion extracted from 1,000 randomly sampled 5-second real video clips. As shown in the left plot of Fig. 2, the power spectrum of the motion decreases exponentially with increasing frequency. This suggests that most natural oscillation motions can indeed be well represented by low-frequency terms. In practice, we found that the first $K = 16$ Fourier coefficients are sufficient to realistically reproduce the original natural motion in a range of real videos and scenes.

4.2. Predicting motion with a diffusion model

We choose a latent diffusion model (LDM) [73] as the backbone for our motion prediction module, as LDMs are more computationally efficient than pixel-space diffusion models, while preserving synthesis quality. A standard LDM consists of two main modules: (1) a variational autoencoder (VAE) that compresses the input image to a latent space through an encoder $z = E(I)$, then reconstructs the input from the latent features via a decoder $I = D(z)$, and (2) a U-Net based diffusion model that learns to iteratively denoise features starting from Gaussian noise. Our training applies this process not to RGB images but to spectral volumes from real video sequences, which are encoded and then diffused for n steps with a pre-defined variance schedule to produce noisy latents z^n . The 2D U-Nets are trained to denoise the noisy latents by iteratively estimating the noise $\epsilon_\theta(z^n; n, c)$ used to update the latent feature at each step $n \in (1, 2, \dots, N)$. The training loss for the LDM is written as

$$\mathcal{L}_{\text{LDM}} = \mathbb{E}_{n \in \mathcal{U}[1, N], \epsilon^n \in \mathcal{N}(0, 1)} [\|\epsilon^n - \epsilon_\theta(z^n; n, c)\|^2] \quad (3)$$

where c is the embedding of any conditional signal, such as text, or, in our case, the first frame of the training video sequence, I_0 . The clean latent features z^0 are then passed through the decoder to recover the spectral volume.

Frequency adaptive normalization. One issue we observed is that motion textures have particular distribution characteristics across frequencies. As visualized in the left plot of Fig. 2, the amplitude of the spectral volumes spans a range of 0 to 100 and decays approximately exponentially with increasing frequency. As diffusion models require that the absolute values of the output are between -1 and 1 for stable training and denoising [44], we must normalize the coefficients of S extracted from real videos before using them for training. If we scale the magnitudes of these coefficients to $[0, 1]$ based on the image dimensions as in prior work [29, 76], nearly all the coefficients at higher frequencies will end up close to zero, as shown in the right plot of Fig. 2. Models trained on such data can produce inaccurate motions, since during inference, even small prediction errors can cause large relative errors after denormalization.

To address this issue, we employ a simple but effective frequency adaptive normalization method: First, we independently normalize Fourier coefficients at each frequency based on statistics computed from the training set. Namely, for each individual frequency f_j , we compute the 95th percentile of Fourier coefficient magnitudes over all input samples and use that value as a per-frequency scaling factor s_{f_j} . We then apply a power transformation to each scaled Fourier coefficient to pull it away from extreme values. In practice, we observe that a square root performs better than other nonlinear transformations such as log or reciprocal. In summary, the final coefficient values of spectral volume $S(\mathbf{p})$ at

frequency f_j (used for training our LDM) are computed as

$$S'_{f_j}(\mathbf{p}) = \text{sign}(S_{f_j}) \sqrt{\left| \frac{S_{f_j}(\mathbf{p})}{s_{f_j}} \right|}. \quad (4)$$

As shown on the right plot of Fig. 2, after applying frequency adaptive normalization, the spectral volume coefficients distribute more evenly.

Frequency-coordinated denoising. The straightforward way to predict a spectral volume \mathcal{S} with K frequency bands is to output a tensor of $4K$ channels from a single diffusion U-Net. However, as in prior work [7], we observe that training a model to produce a large number of channels can yield over-smoothed, inaccurate outputs. An alternative would be to independently predict each individual frequency slice by injecting an extra frequency embedding into the LDM [4], but this design choice would result in uncorrelated predictions in the frequency domain, leading to unrealistic motion.

Therefore, inspired by recent video diffusion work [7], we propose a frequency-coordinated denoising strategy, illustrated in Fig. 3. In particular, given an input image I_0 , we first train an LDM ϵ_θ to predict a single 4-channel frequency slice of spectral volume S_{f_j} , where we inject an extra frequency embedding along with the time-step embedding into the LDM. We then freeze the parameters of this LDM ϵ_θ , introduce attention layers interleaved with the 2D spatial layers of ϵ_θ across the K frequency bands, and fine-tune. Specifically, for a batch size B , the 2D spatial layers of ϵ_θ treat the corresponding $B \cdot K$ noisy latent features of channel size C as independent samples with shape $\mathcal{R}^{(B \cdot K) \times C \times H \times W}$. The attention layer then interprets these as consecutive features spanning the frequency axis, and we reshape the latent features from previous 2D spatial layers to $\mathcal{R}^{B \times K \times C \times H \times W}$ before feeding them to the attention layers. In other words, the frequency attention layers are fine-tuned to coordinate all frequency slices so as to produce coherent spectral volumes. In our experiments, we see that the average VAE reconstruction error improves from 0.024 to 0.018 when we switch from a single 2D U-Net to a frequency-coordinated denoising module, suggesting an improved upper bound on LDM prediction accuracy; in Sec. 7.3, we also show that this design choice improves video generation quality.

5. Image-based rendering

We now describe how we take a spectral volume \mathcal{S} predicted for a given input image I_0 and render a future frame \hat{I}_t at time t . We first derive a motion texture in the time domain using the inverse temporal FFT applied at each pixel $\mathcal{F}(\mathbf{p}) = \text{FFT}^{-1}(\mathcal{S}(\mathbf{p}))$. To produce a future frame \hat{I}_t , we adopt a deep image-based rendering technique and perform splatting with the predicted motion field F_t to forward-warp the encoded I_0 , as shown in Fig. 4. Since forward warping can lead to holes, and multiple source pixels can map to the same output

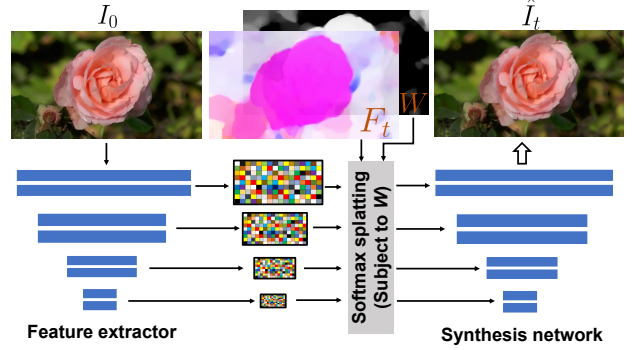


Figure 4. **Rendering module.** We fill in missing content and refine the warped input image using a deep image-based rendering module, where multi-scale features are extracted from the input image I_0 . Softmax splatting is then applied over the features with a motion field F_t from time 0 to t (subject to the weights W). The warped features are fed to an image synthesis network to produce the rendered image \hat{I}_t .

2D location, we adopt the feature pyramid softmax splatting strategy proposed in prior work on frame interpolation [67].

Specifically, we encode I_0 through a feature extractor network to produce a multi-scale feature map. For each individual feature map at scale j , we resize and scale the predicted 2D motion field F_t according to the resolution. As in Davis *et al.* [22], we use predicted flow magnitude as a proxy for depth to determine the contributing weight of each source pixel mapped to its destination location. In particular, we compute a per-pixel weight, $W(\mathbf{p}) = \frac{1}{T} \sum_t \|F_t(\mathbf{p})\|_2$ as the average magnitude of the predicted motion texture. In other words, we assume large motions correspond to moving foreground objects, and small or zero motions correspond to background. We use motion-derived weights instead of learnable ones as [46] because we observe that in the single-view case, learnable weights are not effective for addressing disocclusion ambiguities.

With the motion field F_t and weights W , we apply softmax splatting to warp the feature map at each scale to produce a warped feature. The warped features are then injected into the corresponding blocks of an image synthesis decoder to produce a final rendered image \hat{I}_t .

We jointly train the feature extractor and synthesis networks with start and target frames (I_0, I_t) randomly sampled from real videos, using the estimated flow field from I_0 to I_t to warp encoded features from I_0 , and supervising predictions \hat{I}_t against I_t with a VGG perceptual loss [49].

6. Applications

Image-to-video. Our system enables the animation of a single still picture by first predicting a motion spectral volume from the input image and generating an animation by applying our image-based rendering module to the motion

texture transformed from the spectral volume. Since we explicitly model scene motion, this allows us to produce slow-motion videos by linearly interpolating the motion texture, or to magnify (or minify) animated motions by adjusting the amplitude of the predicted spectral volume coefficients.

Seamless looping. Many applications require videos that loop seamlessly, where there is no discontinuity between the start and end of the video. Unfortunately, it is hard to find a large collection of seamlessly looping videos for training. Instead, we devise a method to use our motion diffusion model, trained on regular non-looping video clips, to produce seamless looping video. Inspired by recent work on guidance for image editing [3, 30], our method is a *motion self-guidance* technique that guides the motion denoising sampling processing using explicit looping constraints. In particular, at each iterative denoising step during inference, we incorporate an additional motion guidance signal alongside standard classifier-free guidance [45], where we enforce each pixel’s position and velocity at the start and end frames to be as similar as possible:

$$\begin{aligned} \tilde{e}^n &= (1 + w)\epsilon_\theta(z^n; n, c) - w\epsilon_\theta(z^n; n, \emptyset) + u\sigma^n \nabla_{z^n} \mathcal{L}_g^n \\ \mathcal{L}_g^n &= \|F_T^n - F_1^n\|_1 + \|\nabla F_T^n - \nabla F_1^n\|_1 \end{aligned} \quad (5)$$

where F_t^n is the predicted 2D displacement field at time t and denoising step n . w is the classifier-free guidance weight, and u is the motion self-guidance weight. In the supplemental video, we apply a baseline appearance-based looping algorithm [58] to generate a looping video from our non-looping output, and show that our motion self-guidance technique produces seamless looping videos with less distortion and fewer artifacts.

Interactive dynamics from a single image. Davis *et al.* [22] show that the spectral volume, evaluated at certain resonant frequencies, can approximate an *image-space modal basis* that is a projection of the vibration modes of the underlying scene (or, more generally, captures spatial and temporal correlations in oscillatory dynamics), and can be used to simulate the object’s response to a user-defined force. We adopt this modal analysis method [22, 69], allowing us to write the image-space 2D motion displacement field for the object’s physical response as a weighted sum of motion spectrum coefficients S_{f_j} modulated by the state of complex modal coordinates $\mathbf{q}_{f_j}(t)$ at each simulated time step t :

$$F_t(\mathbf{p}) = \sum_{f_j} S_{f_j}(\mathbf{p}) \mathbf{q}_{f_j}(t) \quad (6)$$

We simulate the state of the modal coordinates $\mathbf{q}_{f_j}(t)$ via an explicit Euler method applied to the equations of motion for a decoupled mass-spring-damper system represented in modal space [22, 23, 69]. We refer readers to supplementary material and original work for a full derivation. Note that our method produces an interactive scene from a *single picture*, whereas these prior methods required a video as input.

Method	Image Synthesis		Video Synthesis			
	FID	KID	FVD	FVD ₃₂	DTFVD	DTFVD ₃₂
TATS [35]	65.8	1.67	265.6	419.6	22.6	40.7
Stochastic I2V [27]	68.3	3.12	253.5	320.9	16.7	41.7
MCVD [92]	63.4	2.97	208.6	270.4	19.5	53.9
LFDM [66]	47.6	1.70	187.5	254.3	13.0	45.6
DMVFN [48]	37.9	1.09	206.5	316.3	11.2	54.5
Endo <i>et al.</i> [29]	10.4	0.19	166.0	231.6	5.35	65.1
Holynski <i>et al.</i> [46]	11.2	0.20	179.0	253.7	7.23	46.8
Ours	4.03	0.08	47.1	62.9	2.53	6.75

Table 1. **Quantitative comparisons on the test set.** We report both image synthesis and video synthesis quality. Here, KID is scaled by 100. Lower is better for all error. See Sec. 7.1 for descriptions of baselines and error metrics.

7. Experiments

Implementation details. We use an LDM [73] as the backbone for predicting spectral volumes, for which we use a VAE with a continuous latent space of dimension 4. We train the VAE with an L_1 reconstruction loss, a multi-scale gradient consistency loss [54–56], and a KL-divergence loss with respective weights of 1, 0.2, 10^{-6} . We train the same 2D U-Net used in the original LDM work to perform iterative denoising with a simple MSE loss [44], and adopt the attention layers from [41] for frequency-coordinated denoising. For quantitative evaluation, we train both VAE and LDM on images of size 256×160 from scratch for fair comparisons, and it takes around 6 days to converge using 16 Nvidia A100 GPUs. For our main quantitative and qualitative results, we run the motion diffusion model with DDIM [85] for 250 steps. We also show generated videos of up to a resolution of 512×288 , created by fine-tuning a pre-trained image inpainting LDM model [73] on our dataset.

We adopt ResNet-34 [39] for the feature extractor in our IBR module. Our image synthesis network is based on an architecture for conditional image inpainting [57, 109]. Our rendering module runs in real-time at 25FPS on a Nvidia V100 GPU during inference. We adopt universal guidance [3] to produce seamless looping videos, where we set weights $w = 1.75$, $u = 200$, and use 500 DDIM steps with 2 self-recurrence iterations.

Data. We collect and process a set of 3,015 videos of natural scenes exhibiting oscillatory motions from online sources our own captures. We withhold 10% of videos for testing and use the remainder for training. To extract ground truth motion trajectories, we apply a coarse-to-fine flow method [10, 61] between each selected starting image and every future frame of the video. As training data, we take every 10th video frame as input images and derive the corresponding ground truth spectral volumes using the computed motion trajectories across the following 149 frames. In total, our data consists of over 150K image-motion pairs.




Figure 5. $X-t$ slices of videos generated by different approaches. From left to right: input image and corresponding $X-t$ video slices from the ground truth video, from videos generated by three baselines [27, 29, 46, 92], and finally videos generated by our approach.




Figure 6. **Sliding window FID and DTFVD.** We show sliding window FID with window size 30 frames, and DTFVD with size 16 frames, for videos generated by different methods.

Baselines. We compare our approach to recent single-image animation and video prediction methods. Endo *et al.* [29] and DMVFN [48] predict instantaneous 2D motion fields and render future frames auto-regressively. Holynski *et al.* [46] instead simulates motion through a single static Eulerian motion description. Other recent work such as Stochastic Image-to-Video (Stochastic-I2V) [27], TATS [35], and MCVD [92] adopt VAEs, transformers, or diffusion models to directly predict raw video frames; LFDN [66] generates future frames by predicting flow volumes and warping latents in a diffusion model. We train all the above methods on our data using their respective open-source implementations.¹

We evaluate the quality of the videos generated by our approach and by prior baselines in two ways. First, we evaluate the quality of individual synthesized frames using metrics designed for image synthesis tasks. We adopt the Fréchet Inception Distance (FID) [42] and Kernel Inception Distance (KID) [5] to measure the average distance between the distributions of generated frames and ground truth frames.

Second, to evaluate the quality and temporal coherence of synthesized videos, we adopt the Fréchet Video Dis-

¹We use the open-source implementation of [46] from Fan *et al.* [82].

Method	Image Synthesis		Video Synthesis			
	FID	KID	FVD	FVD ₃₂	DTFVD	DTFVD ₃₂
Repeat I_0	-	-	237.5	316.7	5.30	45.6
$K = 4$	3.92	0.07	60.3	78.4	3.12	8.59
$K = 8$	3.95	0.07	52.1	68.7	2.71	7.37
$K = 24$	4.09	0.08	48.2	65.1	2.50	6.94
w/o adaptive norm.	4.53	0.09	62.7	80.1	3.16	8.19
Independent pred.	4.00	0.08	52.5	71.3	2.70	7.40
Volume pred.	4.74	0.09	53.7	71.1	2.83	7.79
Baseline splat [46]	4.25	0.09	49.5	66.8	2.83	7.27
Full ($K = 16$)	4.03	0.08	47.1	62.9	2.53	6.75

Table 2. **Ablation study.** Sec. 7.3 describes each configuration.

tance [91] with window size 16 (FVD) and 32 (FVD₃₂), based on an I3D model [11] trained on the Human Kinetics datasets [52]. To more faithfully reflect synthesis quality for the natural oscillation motions we seek to generate, we also adopt the Dynamic Texture Fréchet Video Distance [27], which measures distance from videos with window size 16 (DTFVD) and size 32 (DTFVD₃₂), using a I3D model trained on the Dynamic Textures Database [37], a dataset consisting primarily of natural motion textures.

We further use a sliding window FID of window size 30 frames, and a sliding window DTFVD with window size 16 frames, as in [57, 60], to measure how generated video quality degrades over time. For all methods, we evaluate metrics at 256×128 resolution by center-cropping.

7.1. Quantitative results

Table 1 shows quantitative comparisons between our approach and baselines on our test set. Our approach signifi-



Figure 7. We show generated future frames from three recent large video diffusion models [31, 36, 97].

cantly outperforms prior single-image animation baselines in terms of both image and video synthesis quality. Specifically, our much lower FVD and DT-FVD distances suggest that the videos generated by our approach are more realistic and more temporally coherent. Further, Fig. 6 shows the sliding window FID and sliding window DT-FVD distances of generated videos from different methods. Thanks to the global spectral volume representation, videos generated by our approach do not suffer from degradation over time.

7.2. Qualitative results

We visualize qualitative comparisons between videos as spatio-temporal $X-t$ slices of the generated videos, a standard way of visualizing small motions in a video [94]. As shown in Fig. 5, our generated video dynamics more strongly resemble the motion patterns observed in the corresponding real reference videos (second column), compared to other methods. Baselines such as Stochastic I2V [27] and MCVF [92] fail to model both appearance and motion realistically over time. Endo *et al.* [29] and Holynski *et al.* [46] produce video frames with fewer artifacts but exhibits over-smooth or non-oscillatory motion over time. We refer readers to supplementary material to assess the quality of generated video frames and estimated motion across different methods.

7.3. Ablation study

We conduct an ablation study to validate the major design choices in our motion prediction and rendering modules, comparing our full configuration with different variants. Specifically, we evaluate results using different numbers of frequency bands $K = 4, 8, 16, 24$. We observe that increasing the number of frequency bands improves video prediction quality, but the improvement is marginal with more than 16 frequencies. Next, we remove adaptive frequency normalization from the ground truth spectral volumes, and instead just scale them based on input image width and height (*w/o adaptive norm.*). Additionally, we remove the frequency coordinated-denoising module (*Independent pred.*), or replace it with a simpler DM where a tensor volume of $4K$ channel spectral volumes are predicted jointly via a single 2D U-net diffusion model (*Volume pred.*). Finally, we compare results where we use a baseline rendering method that applies softmax splatting over single-scale features subject to learnable weights as used by [46] (*Baseline splat*). We also add a baseline where the generated video is a volume by



Figure 8. **Limitations.** We show examples of rendered future frames (even), and overlay of input and rendered images (odd). Our method can produce artifacts in regions of thin objects or large motions, and regions requiring filling large amount of new contents.

repeating input image N times (Repeat I_0). From Table 2, we observe that all simpler or alternative configurations lead to worse performance compared with our full model.

7.4. Comparing to large video models

We further perform a user study and compare our generated animations with ones from recent large video diffusion models: AnimateDiff [36], ModelScope [97] and Gen-2 [31], which predict video volumes directly. On a randomly selected 30 videos from the test set, we ask users “which video is more realistic?”. Users report a 80.9% preference for our approach over others. Moreover, as shown in Fig. 7, we observe that the generated videos from these baselines are either unable to adhere to the input image content, or exhibit gradual color drift and distortion over time. We refer readers to the supplementary material for full comparisons.

8. Discussion and conclusion

Limitations. Since our approach only predicts lower frequencies of a spectral volume, it can fail to model non-oscillating motions or high-frequency vibrations—this may be resolved by using learned motion bases. Furthermore, the quality of generated videos relies on the quality of underlying motion trajectories, which may degrade in scenes with thin moving objects or objects with large displacements. Even if correct, motions that require generating large amounts of new unseen content may also cause degradations (Fig. 8).

Conclusion. We present a new approach for modeling natural oscillation dynamics from a single still picture. Our image-space motion prior is represented with spectral volumes, a frequency representation of per-pixel motion trajectories, which we find to be efficient and effective for prediction with diffusion models, and which we learn from collections of real world videos. Spectral volumes are predicted using frequency-coordinated latent diffusion model and are used to animate future video frames through an image-based rendering module. We show that our approach produces photo-realistic animations from a single picture and significantly outperforms prior baselines, and that it can enable several downstream applications such as creating seamlessly looping or interactive image dynamics.

Acknowledgements. We thank Abe Davis, Rick Szeliski, Andrew Liu, Boyang Deng, Qianqian Wang, Xuan Luo, and Lucy Chai for fruitful discussions and helpful comments.

References

- [1] Aseem Agarwala, Ke Colin Zheng, Chris Pal, Maneesh Agrawala, Michael Cohen, Brian Curless, David Salesin, and Richard Szeliski. Panoramic video textures. In *ACM SIGGRAPH 2005 Papers*, pages 821–827. 2005.
- [2] Hyemin Ahn, Esteve Valls Mascaro, and Dongheui Lee. Can we use diffusion probabilistic models for 3d motion prediction? *arXiv preprint arXiv:2302.14503*, 2023.
- [3] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 843–852, 2023.
- [4] Hugo Bertiche, Niloy J Mitra, Kuldeep Kulkarni, Chun-Hao P Huang, Tuanfeng Y Wang, Meysam Madadi, Sergio Escalera, and Duygu Ceylan. Blowing in the wind: Cyclenet for human cinemagraphs from still images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 459–468, 2023.
- [5] Mikolaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. *arXiv preprint arXiv:1801.01401*, 2018.
- [6] Andreas Blattmann, Timo Milbich, Michael Dorcenwald, and Björn Ommer. ipoke: Poking a still image for controlled stochastic video synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14707–14717, 2021.
- [7] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22563–22575, 2023.
- [8] Richard Strong Bowen, Richard Tucker, Ramin Zabih, and Noah Snavely. Dimensions of motion: Monocular prediction through flow subspaces. In *2022 International Conference on 3D Vision (3DV)*, pages 454–464. IEEE, 2022.
- [9] Tim Brooks, Janne Hellsten, Miika Aittala, Ting-Chun Wang, Timo Aila, Jaakko Lehtinen, Ming-Yu Liu, Alexei Efros, and Tero Karras. Generating long videos of dynamic scenes. *Advances in Neural Information Processing Systems*, 35:31769–31781, 2022.
- [10] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 25–36. Springer, 2004.
- [11] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [12] Dan Casas, Marco Volino, John Collomosse, and Adrian Hilton. 4d video textures for interactive character appearance. In *Computer Graphics Forum*, volume 33, pages 371–380. Wiley Online Library, 2014.
- [13] Antoni B Chan and Nuno Vasconcelos. Mixtures of dynamic textures. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 1, pages 641–647. IEEE, 2005.
- [14] Antoni B Chan and Nuno Vasconcelos. Classifying video with kernel dynamic textures. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–6. IEEE, 2007.
- [15] Antoni B Chan and Nuno Vasconcelos. Modeling, clustering, and segmenting video with mixtures of dynamic textures. *IEEE transactions on pattern analysis and machine intelligence*, 30(5):909–926, 2008.
- [16] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023.
- [17] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11315–11325, 2022.
- [18] Tsai-Shien Chen, Chieh Hubert Lin, Hung-Yu Tseng, Tsung-Yi Lin, and Ming-Hsuan Yang. Motion-conditioned diffusion model for controllable video synthesis. *arXiv preprint arXiv:2304.14404*, 2023.
- [19] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18000–18010, 2023.
- [20] Yung-Yu Chuang, Dan B Goldman, Ke Colin Zheng, Brian Curless, David H Salesin, and Richard Szeliski. Animating pictures with stochastic motion textures. In *ACM SIGGRAPH 2005 Papers*, pages 853–860. 2005.
- [21] Vincent C Couture, Michael S Langer, and Sebastien Roy. Omnistereo video textures without ghosting. In *2013 International Conference on 3D Vision-3DV 2013*, pages 64–70. IEEE, 2013.
- [22] Abe Davis, Justin G Chen, and Frédo Durand. Image-space modal bases for plausible manipulation of objects in video. *ACM Transactions on Graphics (TOG)*, 34(6):1–7, 2015.
- [23] Myers Abraham Davis. *Visual vibration analysis*. PhD thesis, Massachusetts Institute of Technology, 2016.
- [24] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [25] Julien Diener, Mathieu Rodriguez, Lionel Baboud, and Lionel Reveret. Wind projection basis for real-time animation of trees. In *Computer graphics forum*, volume 28, pages 533–540. Wiley Online Library, 2009.
- [26] Gianfranco Doretto, Alessandro Chiuso, Ying Nian Wu, and Stefano Soatto. Dynamic textures. *International journal of computer vision*, 51:91–109, 2003.
- [27] Michael Dorcenwald, Timo Milbich, Andreas Blattmann, Robin Rombach, Konstantinos G. Derpanis, and Bjorn Ommer. Stochastic image-to-video synthesis using cinns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3742–3753, June 2021.
- [28] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali Thabet, and Arsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition*, pages 481–490, 2023.
- [29] Yuki Endo, Yoshihiro Kanamori, and Shigeru Kuriyama. Animating landscape: Self-supervised learning of decoupled motion and appearance for single-image video synthesis. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia 2019)*, 38(6):175:1–175:19, 2019.
- [30] Dave Epstein, Allan Jabri, Ben Poole, Alexei A Efros, and Aleksander Holynski. Diffusion self-guidance for controllable image generation. *arXiv preprint arXiv:2306.00986*, 2023.
- [31] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7346–7356, 2023.
- [32] Matthew Flagg, Atsushi Nakazawa, Qiushuang Zhang, Sing Bing Kang, Young Kee Ryu, Irfan Essa, and James M Rehg. Human video textures. In *Proceedings of the 2009 symposium on Interactive 3D graphics and games*, pages 199–206, 2009.
- [33] Jean-Yves Franceschi, Edouard Delasalles, Mickaël Chen, Sylvain Lamprier, and Patrick Gallinari. Stochastic latent residual video prediction. In *International Conference on Machine Learning*, pages 3233–3246. PMLR, 2020.
- [34] Ruohan Gao, Bo Xiong, and Kristen Grauman. Im2Flow: Motion hallucination from static images for action recognition. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [35] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. *arXiv preprint arXiv:2204.03638*, 2022.
- [36] Yuwei Guo, Ceyuan Yang, Anyi Rao, Yaohui Wang, Yu Qiao, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023.
- [37] Isma Hadji and Richard P Wildes. A new large scale dynamic texture dataset with application to convnet understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 320–335, 2018.
- [38] Zekun Hao, Xun Huang, and Serge Belongie. Controllable video generation with sparse trajectories. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7854–7863, 2018.
- [39] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [40] Yingqing He, Menghan Xia, Haoxin Chen, Xiaodong Cun, Yuan Gong, Jinbo Xing, Yong Zhang, Xintao Wang, Chao Weng, Ying Shan, et al. Animate-a-story: Storytelling with retrieval-augmented video generation. *arXiv preprint arXiv:2307.06940*, 2023.
- [41] Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion models for high-fidelity video generation with arbitrary lengths. *arXiv preprint arXiv:2211.13221*, 2022.
- [42] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [43] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- [44] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [45] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [46] Aleksander Holynski, Brian L Curless, Steven M Seitz, and Richard Szeliski. Animating pictures with Eulerian motion fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5810–5819, 2021.
- [47] Tobias Hoppe, Arash Mehrjou, Stefan Bauer, Didrik Nielsen, and Andrea Dittadi. Diffusion models for video prediction and infilling. *Trans. Mach. Learn. Res.*, 2022, 2022.
- [48] Xiaotao Hu, Zhewei Huang, Ailin Huang, Jun Xu, and Shuchang Zhou. A dynamic multi-scale voxel flow network for video prediction. *ArXiv*, abs/2303.09875, 2023.
- [49] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 694–711. Springer, 2016.
- [50] Hitoshi Kanda and Jun Ohya. Efficient, realistic method for animating dynamic behaviors of 3d botanical trees. In *2003 International Conference on Multimedia and Expo. ICME'03. Proceedings (Cat. No. 03TH8698)*, volume 2, pages II–89. IEEE, 2003.
- [51] Johanna Karras, Aleksander Holynski, Ting-Chun Wang, and Ira Kemelmacher-Shlizerman. Dreampose: Fashion image-to-video synthesis via stable diffusion. *arXiv preprint arXiv:2304.06025*, 2023.
- [52] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [53] Alex X Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic adversarial video prediction. *arXiv preprint arXiv:1804.01523*, 2018.
- [54] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snively, Ce Liu, and William T Freeman. Learning the depths of moving people by watching frozen people. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4521–4530, 2019.
- [55] Zhengqi Li and Noah Snively. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2041–2050, 2018.
- [56] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snively. Dynibar: Neural dynamic image-based

- rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4273–4284, 2023.
- [57] Zhengqi Li, Qianqian Wang, Noah Snavely, and Angjoo Kanazawa. Infinitenature-zero: Learning perpetual view generation of natural scenes from single images. In *European Conference on Computer Vision*, pages 515–534. Springer, 2022.
- [58] Jing Liao, Mark Finch, and Hugues Hoppe. Fast computation of seamless video loops. *ACM Transactions on Graphics (TOG)*, 34(6):1–10, 2015.
- [59] Zicheng Liao, Neel Joshi, and Hugues Hoppe. Automated video looping with progressive dynamism. *ACM Transactions on Graphics (TOG)*, 32(4):1–10, 2013.
- [60] Andrew Liu, Richard Tucker, Varun Jampani, Ameesh Makadia, Noah Snavely, and Angjoo Kanazawa. Infinite nature: Perpetual view generation of natural scenes from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14458–14467, 2021.
- [61] Ce Liu. *Beyond pixels: exploring new representations and applications for motion analysis*. PhD thesis, Massachusetts Institute of Technology, 2009.
- [62] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10209–10218, 2023.
- [63] Aniruddha Mahapatra and Kuldeep Kulkarni. Controllable animation of fluid elements in still images. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [64] Arun Mallya, Ting-Chun Wang, and Ming-Yu Liu. Implicit warping for animation with image sets. *Advances in Neural Information Processing Systems*, 35:22438–22450, 2022.
- [65] Medhini Narasimhan, Shiry Ginosar, Andrew Owens, Alexei A Efros, and Trevor Darrell. Strumming to the beat: Audio-conditioned contrastive video textures. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3761–3770, 2022.
- [66] Haomiao Ni, Changhao Shi, Kai Li, Sharon X Huang, and Martin Renqiang Min. Conditional image-to-video generation with latent flow diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18444–18455, 2023.
- [67] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5437–5446, 2020.
- [68] Shin Ota, Machiko Tamura, Kunihiko Fujita, T Fujimoto, K Muraoka, and Norishige Chiba. $1/f$ /sup/spl beta/noise-based real-time animation of trees swaying in wind fields. In *Proceedings Computer Graphics International 2003*, pages 52–59. IEEE, 2003.
- [69] Automne Petitjean, Yohan Poirier-Ginter, Ayush Tewari, Guillaume Cordonnier, and George Drettakis. Modalnerf: Neural modal analysis and synthesis for free-viewpoint navigation in dynamically vibrating scenes. In *Computer Graphics Forum*, volume 42, 2023.
- [70] Silvia L. Pinteá, Jan C. van Gemert, and Arnold W. M. Smeulders. Déjà vu: Motion prediction in static images. In *Proc. European Conf. on Computer Vision (ECCV)*, 2014.
- [71] Sigal Raab, Inbal Leibovitch, Guy Tevet, Moab Arar, Amit H Bermano, and Daniel Cohen-Or. Single motion diffusion. *arXiv preprint arXiv:2302.05905*, 2023.
- [72] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- [73] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [74] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- [75] Payam Saisan, Gianfranco Doretto, Ying Nian Wu, and Stefano Soatto. Dynamic texture recognition. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 2, pages II–II. IEEE, 2001.
- [76] Saurabh Saxena, Charles Herrmann, Junhwa Hur, Abhishek Kar, Mohammad Norouzi, Deqing Sun, and David J. Fleet. The surprising effectiveness of diffusion models for optical flow and monocular depth estimation, 2023.
- [77] Arno Schödl, Richard Szeliski, David H Salesin, and Irfan Essa. Video textures. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 489–498, 2000.
- [78] Mikio Shinya and Alain Fournier. Stochastic motion—motion under the influence of wind. *Computer Graphics Forum*, 11(3), 1992.
- [79] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2377–2386, 2019.
- [80] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in neural information processing systems*, 32, 2019.
- [81] Aliaksandr Siarohin, Oliver J Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. Motion representations for articulated animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13653–13662, 2021.
- [82] Chen Qian Kwan-Yee Lin Hongsheng Li Siming Fan, Jingtian Piao. Simulating fluids in real-world still images. *arXiv preprint, arXiv:2204.11335*, 2022.
- [83] Ivan Skorokhodov, Sergey Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3626–3636, 2022.
- [84] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan,

- and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [85] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv:2010.02502*, October 2020.
- [86] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [87] Jos Stam. *Multi-scale stochastic modelling of complex natural phenomena*. PhD thesis, 1995.
- [88] Jos Stam. Stochastic dynamics: Simulating the effects of turbulence on flexible structures. *Computer Graphics Forum*, 16(3), 1997.
- [89] Ryusuke Sugimoto, Mingming He, Jing Liao, and Pedro V Sander. Water simulation and rendering from a still photograph. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022.
- [90] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022.
- [91] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018.
- [92] Vikram Voleti, Alexia Jolicoeur-Martineau, and Christopher Pal. Mcvd: Masked conditional video diffusion for prediction, generation, and interpolation. In *(NeurIPS) Advances in Neural Information Processing Systems*, 2022.
- [93] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *Neural Information Processing Systems*, 2016.
- [94] Neal Wadhwa, Michael Rubinstein, Frédo Durand, and William T Freeman. Phase-based video motion processing. *ACM Transactions on Graphics (ToG)*, 32(4):1–10, 2013.
- [95] Jacob Walker, Carl Doersch, Abhinav Gupta, and Martial Hebert. An uncertain future: Forecasting from static images using variational autoencoders. In *Proc. European Conf. on Computer Vision (ECCV)*, 2016.
- [96] Jacob Walker, Abhinav Gupta, and Martial Hebert. Dense optical flow prediction from a static image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2443–2451, 2015.
- [97] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiuniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Videocomposer: Compositional video synthesis with motion controllability. *arXiv preprint arXiv:2306.02018*, 2023.
- [98] Yaohui Wang, Di Yang, Francois Bremond, and Antitza Dantcheva. Latent image animator: Learning to animate images via latent space navigation. *arXiv preprint arXiv:2203.09043*, 2022.
- [99] Frederik Warburg, Ethan Weber, Matthew Tancik, Aleksander Holynski, and Angjoo Kanazawa. Nerfbusters: Removing ghostly artifacts from casually captured nerfs. *arXiv preprint arXiv:2304.10532*, 2023.
- [100] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. *arXiv preprint arXiv:2210.04628*, 2022.
- [101] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Photo wake-up: 3d character animation from a single photo. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5908–5917, 2019.
- [102] Jamie Wynn and Daniyar Turmukhambetov. DiffusioNeRF: Regularizing Neural Radiance Fields with Denoising Diffusion Models. In *CVPR*, 2023.
- [103] Tianfan Xue, Jiajun Wu, Katherine L Bouman, and William T Freeman. Visual dynamics: Stochastic future generation via layered cross convolutional networks. *Trans. Pattern Analysis and Machine Intelligence*, 41(9):2236–2250, 2019.
- [104] Shengming Yin, Chenfei Wu, Jian Liang, Jie Shi, Houqiang Li, Gong Ming, and Nan Duan. Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. *arXiv preprint arXiv:2308.08089*, 2023.
- [105] Sihyun Yu, Kihyuk Sohn, Subin Kim, and Jinwoo Shin. Video probabilistic diffusion models in projected latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18456–18466, 2023.
- [106] Mingyuan Zhang, Xinying Guo, Liang Pan, Zhongang Cai, Fangzhou Hong, Huirong Li, Lei Yang, and Ziwei Liu. Remodiffuse: Retrieval-augmented motion diffusion model. *arXiv preprint arXiv:2304.01116*, 2023.
- [107] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. *arXiv preprint arXiv:2305.13077*, 2023.
- [108] Jian Zhao and Hui Zhang. Thin-plate spline motion model for image animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3657–3666, 2022.
- [109] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I Chang, and Yan Xu. Large scale image completion via co-modulated generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2021.
- [110] Daquan Zhou, Weimin Wang, Hanshu Yan, Weiwei Lv, Yizhe Zhu, and Jiashi Feng. Magicvideo: Efficient video generation with latent diffusion models. *arXiv preprint arXiv:2211.11018*, 2022.