

# Motion and Context-Aware Audio-Visual Conditioned Video Prediction

Yating Xu  
xu.yating@u.nus.edu

Conghui Hu  
conghui@nus.edu.sg

Gim Hee Lee  
gimhee.lee@nus.edu.sg

Department of Computer Science  
National University of Singapore  
Singapore

## Abstract

The existing state-of-the-art method for audio-visual conditioned video prediction uses the latent codes of the audio-visual frames from a multimodal stochastic network and a frame encoder to predict the next visual frame. However, a direct inference of per-pixel intensity for the next visual frame is extremely challenging because of the high-dimensional image space. To this end, we decouple the audio-visual conditioned video prediction into motion and appearance modeling. The multimodal motion estimation predicts future optical flow based on the audio-motion correlation. The visual branch recalls from the motion memory built from the audio features to enable better long-term prediction. We further propose context-aware refinement to address the diminishing of the global appearance context in the long-term continuous warping. The global appearance context is extracted by the context encoder and manipulated by motion-conditioned affine transformation before fusion with features of warped frames. Experimental results show that our method achieves competitive results on existing benchmarks.

## 1 Introduction

Humans perceive the world via multisensory processing. Particularly, audio-visual interactions are ubiquitous in our daily life. For instance, a person at a bus stop perceives an approaching bus through both audio and visual cues, yet his/her view is suddenly blocked by another pedestrian. Despite this visual barrier, the person’s cognitive system can still fill in the missing visual content using previously seen frames and unbroken audio information. It has also been proven in neurobiology studies that such audio-visual pairing has the advantage of increasing the activation of neurons [22] to enhance perceptual ability. Inspired




Figure 1: An illustration of audio-visual conditioned video prediction task. Given a full audio clip and a short sequence of past visual frames, it aims to predict corresponding future visual frames.

by these observations, our aim is to investigate the feasibility of equipping machines with the capacity to leverage audio-visual interaction, employing audio-visual conditioned video prediction as a testbed.

Fig. 1 illustrates the audio-visual conditioned video prediction task. Given a full audio clip and a short sequence of the past visual frames, the objective is to predict the missing future visual frames that are as close to the ground truth frames as possible. The given audio clips and visual frames act as guidance in predicting the motion and appearance of future visual frames.

Existing state-of-the-art work Sound2Sight [9] proposes to directly infer per-pixel intensity, which contains both motion and appearance of future images, from the audio-visual inputs. A multi-modal stochastic network is used to encode the audio-visual information into a latent space distribution. A random vector drawn from this latent space distribution is then concatenated with the embedding of the current visual frame to simultaneously predict the motion and appearance of objects for the next visual frame, which is extremely challenging due to the high-dimensional nature of image space. As a result, Sound2Sight struggles to preserve the visual content in the generated visual frames, *i.e.*, the appearance of the digits in Fig. 4. We postulate that this inability to retain the visual content is caused by the tightly entangled motion and appearance modeling in Sound2Sight.

To this end, we model the motion and appearance separately in the audio-visual conditioned video prediction. We propose multimodal motion estimation (MME) to predict future object motion in the form of optical flow by using the audio-motion correlation. The rhythm of the sound and the object motion are naturally aligned since the sound originates from object motion [8, 33]. Inspired by this natural phenomenon, we store the audio features in the motion memory to guide the visual motion prediction. In contrast to the internal memory in the ConvLSTMs that selectively remembers the past memory, the motion memory keeps all the past memory, which provides a good supplement in long-term motion prediction. To effectively bridge the audio-visual modality gap, we further design operator *Condense* to condense the large-size motion memory into a compact representation and operator *Recall* for the visual branch to use the audio motion context.

Although the next frame can now be generated by simply warping the current frame according to the predicted motion, the long-term continuous warping in the video prediction would result in the diminishing of the global appearance context. In order to effectively retain the appearance of objects, we creatively design a context-aware refinement (CAR) module. Other than just using a U-Net to refine the warped frames [4, 51], we design a context encoder to supplement the global appearance context by extracting the visual features from the given past frame. Since the appearance details of each future frame may vary due to object movement, we perform the motion-conditioned affine transformation on the global context feature before inserting it into the U-Net decoder for better adaptation of the motion variance at each time step.

Our contributions in this paper are as follows:

- We propose the multimodal motion estimation module to predict the optical flow of the next visual frame with the help of motion memory built from the audio features. We design the *Condense* and *Recall* operators to help the visual branch effectively utilize the audio motion memory.
- We introduce the context-aware refinement module to mitigate the loss of global appearance context in the long-term warpings by using the context encoder and motion-conditioned affine transformation.

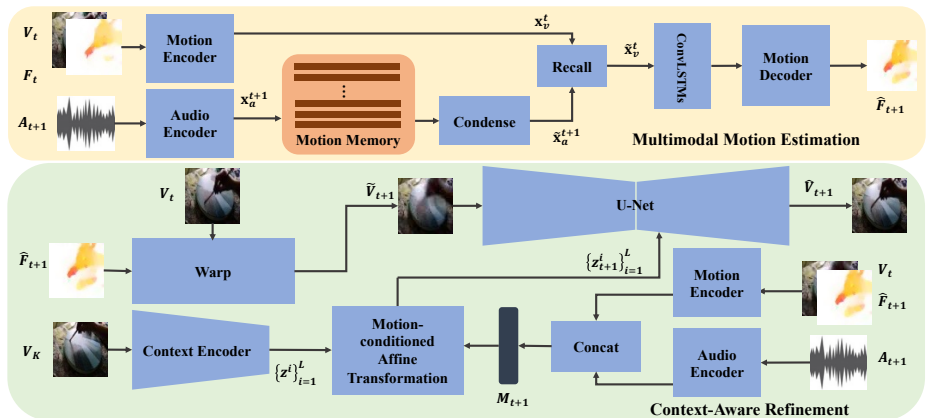


Figure 2: Our framework consists of the Multimodal Motion Estimation (MME) and Context-Aware Refinement (CAR) modules. MME predicts future optical flow by using audio information in the form of motion memory. CAR mitigates the loss of appearance context in the long-term recurrent warpings by using a context encoder and motion-conditioned affine transformation.

- We conduct extensive experiments on both synthetic and real-world datasets to verify the effectiveness of our method.

## 2 Related Work

**Video prediction.** Video prediction estimates future video frames close to ground truth future frames conditioned on a few past frames. It can be divided into two categories: prediction in pixel space [20, 24, 44, 45, 49], and prediction in low-dimensional space [25, 40, 41, 48]. Prediction in pixel space directly estimates the per-pixel intensity. For example, PredRNN [44] enables memory states to be updated across vertically stacked RNN layers and horizontally through all RNN states. Jin *et al.* [20] adopt multi-level wavelet transform to deal with motion blur and missing appearance details. Denton and Fergus [13] additionally consider stochasticity in the prediction by learning a prior distribution to capture the future dynamics.

Direct estimation of the per-pixel intensity of future visual frames is a very challenging task because of the high-dimensional image space. As a result, there are many works that seek to predict a low-dimensional representation, such as semantic map [4, 25, 39], human pose [39, 40] and optical flow [4, 4, 26, 27, 28]. Li *et al.* [26] generate consecutive multiple future frames from only one image by multiple time step flow prediction and flow-to-frame synthesis. Bei *et al.* [4] first infer future optical flow and semantic maps independently for each object in the scene and then use U-Net to inpaint the future images. MSPred [39] simultaneously predicts future images and abstract representations via the hierarchical predictor module.

Different from the visual-only video prediction methods, we additionally consider the audio information. The pioneer work Sound2Sight [9] learns a joint audio-visual latent embedding to capture future dynamics. However, the latent codes cause the tight entanglement of motion and appearance, thus making it difficult for the model to balance the capture of

correct motion and maintaining the shape of objects. Our solution is to model the motion and appearance separately through multimodal motion estimation and context-aware refinement, respectively.

**Audio-visual representation learning.** Audio-Visual representation learning learns a general representation in a self-supervised way by utilizing audio-visual correlation. There are mainly three types of correspondences: semantic similarity [1, 2, 3, 4, 5], temporal consistency [6, 7, 8] and spatial correspondence [9, 10]. Arandjelovic *et al.* [1] propose audio-visual correspondence learning by classifying whether the input audio and visual frame are from the same video or not. [5] design a pretext task to predict whether the audio stream and visual sequence are temporally aligned. [8] leverage audio-visual spatial correspondence by predicting whether the two channels of audio are flipped. In addition to learning a general representation, there are also many downstream audio-visual tasks, such as sound source localization [11, 12, 13], audio-visual source separation [14, 15, 16], audio-visual video parsing [17, 18] and talking face generation [19, 20, 21]. In this paper, we study the task of audio-visual conditioned video prediction and leverage audio-motion correlation.

### 3 Method

**Problem definition.** Let us denote a video with  $T$  frames as  $\mathcal{V} = \{\{V_1, A_1\}, \dots, \{V_T, A_T\}\}$ , where  $\{V_t, A_t\}$  represents a visual frame  $V_t$  and its corresponding audio clip  $A_t$ . A visual frame  $V_t$  is an image of size  $C \times W \times H$ , where  $C$  is the number of channels,  $W$  and  $H$  are the width and height of the image. An audio clip  $A_t$  is a  $\mathcal{F} \times \mathcal{T}$  two-dimensional spectrogram, where  $\mathcal{F}$  denotes the frequency range and  $\mathcal{T}$  represents the time duration of the audio clip. Given the first  $K$  visual frames  $\{V_1, \dots, V_K\}$  and the whole audio  $\{A_1, \dots, A_K, \dots, A_T\}$ , the goal of the audio-visual conditioned video prediction task is to predict the missing  $T - K$  visual frames  $\{V_{K+1}, \dots, V_T\}$  to be as similar to the ground truth visual frames as possible. For brevity, we drop the frame index  $K$  and denote the indices of the previous, current and next predicted frames as  $t - 1$ ,  $t$  and  $t + 1$ , respectively.

**Method overview.** Fig. 2 shows our proposed framework which consists of two modules: 1) multimodal motion estimation (MME) and 2) context-aware refinement (CAR). Our MME and CAR modules are designed to model motion and appearance, respectively. MME predicts future optical flow by taking visual and audio inputs. It stores the past and current audio information in the motion memory and the visual branch recalls from the condensed motion memory to better infer long-term future motions. CAR renders future frames by taking global appearance context into account. The warped images are supplemented with global appearance features extracted by the context encoder, and modulated by motion-conditioned affine transformation to get the final predicted future frames.

#### 3.1 Multimodal Motion Estimation

The MME predicts optical flow in a recurrent way. For ease of warping, the MME operates on backward optical flow, *i.e.* both input  $F_t$  and output  $\hat{F}_{t+1}$  are backward optical flow. We use both audio and visual modalities in this module because they both contain clear motion information. To predict the  $t + 1$  frame, the motion encoder  $E_m$  takes the concatenation of the current optical flow  $F_t$  and visual frame  $V_t$  as inputs, and outputs the motion feature  $x'_t \in \mathbb{R}^{w \times h \times c}$ . Note that  $V_t$  is a given visual frame if  $t \leq K$ , otherwise it is generated by warping  $V_{t-1}$  according to the predicted optical flow  $F_t$ . We do not refine the warped image

in MME. Concurrently, the audio encoder  $E_a$  extracts the audio feature  $x_a^{t+1} \in \mathbb{R}^{1 \times c}$  from input audio signal  $A_{t+1}$ .

For our recurrent MME to effectively facilitate motion prediction, it is crucial to guarantee the presence of long-term memory for the audio information. However, the straightforward choice of LSTM is notorious for only having a short memory, where the memory cell is quickly saturated to remember only the latest input. As a result, it is difficult to infer long-term future motions. We thus design an external motion memory MM to store the audio features. In contrast to the memory cell inside LSTM that selectively remembers past memory, our motion memory MM keeps all the past memory without abandon. Our advantage is that it does not lose any past memory, and thus efficiently keeps all long-term information. The MM stores audio features up to time  $t + 1$ ,  $\text{MM} = \{x_a^n\}_{n=1}^{t+1} \in \mathbb{R}^{(t+1) \times c}$ . Then, visual features *Recall* from MM via cross-attention as the query (**Q**), and MM is used to generate the key (**K**) and value (**V**). During the long-term prediction, the size of MM increases linearly with the prediction length, making it difficult for the visual branch to catch useful information from the large motion memory. We address this motion memory complexity problem by *Condensing* the full MM into a compact memory  $\tilde{x}_a^{t+1} \in \mathbb{R}^{1 \times c}$  as follows:

$$\text{Atten}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}(\mathbf{Q}\mathbf{K}^\top) \mathbf{V}, \quad (1)$$

$$\text{Condense}(\text{MM}) = \text{Atten}(\text{MM}, \text{MM}, \text{MM})[-1], \quad \tilde{x}_a^{t+1} = x_a^{t+1} + \text{Condense}(\text{MM}),$$

where  $[-1]$  denotes the last slice of the memory. The Atten operation follows [E3]. The visual branch *Recalls* from the condensed audio memory  $\tilde{x}_a^{t+1}$  and obtain the enhanced motion feature  $\tilde{x}_v^t \in \mathbb{R}^{w \times h \times c}$  as follow:

$$\text{Recall}(x_v^t, \tilde{x}_a^{t+1}) = \text{Atten}(x_v^t, \tilde{x}_a^{t+1}, \tilde{x}_a^{t+1}), \quad \tilde{x}_v^t = x_v^t + \text{Recall}(x_v^t, \tilde{x}_a^{t+1}). \quad (2)$$

We tile  $x_v^t$  into the feature vector of size  $l \times c$ , where  $l = w \times h$  before computing Eqn. 2.  $\tilde{x}_v^t$  is then fed into ConvLSTMs and motion decoder  $D_m$  to predict optical flow  $\hat{F}_{t+1}$  at next time step as follows:

$$\{h_{t+1}, o_{t+1}\} = \text{ConvLSTM}(\tilde{x}_v^t, h_t), \quad \hat{F}_{t+1} = D_m(h_{t+1}), \quad (3)$$

where  $h_{t+1}$  is the hidden state and  $o_{t+1}$  is the cell state of the ConvLSTMs at time step  $t + 1$ .

## 3.2 Context-Aware Refinement

The Context-Aware Refinement (CAR) module is then designed to generate the next visual frame according to the optical flow predicted by MME. Since audio cannot convey dense pixel information, we rely on the visual modality for appearance modeling. We start by warping the current visual frame  $V_t$  using the predicted optical flow  $\hat{F}_{t+1}$  as follow:

$$\hat{V}_{t+1}(p) = V_t(p + \hat{F}_{t+1}(p)), \quad (4)$$

where  $p$  is the pixel coordinates. Similar to MME,  $V_t$  is the given visual frame if  $t \leq K$ , otherwise it is the output of CAR.

The warping procedure is performed recurrently by taking the current prediction as input and warping it into the next frame. It causes the warped image to become increasingly blurry and consequently leads to the loss of appearance context as shown in Fig. 3. The warped images gradually lose the shape of the brush.

To alleviate the blurry visual frame problem, we first use a U-Net [62] to refine the local details. Furthermore, we address the loss of appearance context by providing external appearance information. The context encoder  $E_c$  extracts global appearance context  $Z = \{z^i\}_{i=1}^L$  from the last given visual frame  $V_K$ .  $z^i \in \mathbb{R}^{c^i \times h^i \times w^i}$  is the feature at the  $i$ -th layer of  $E_c$  and  $L$  is the total number of layers. The context encoder resembles the U-Net encoder in its design, albeit shallower structure. Given the varying appearance of each frame caused by object motion, it is necessary to adjust the global context accordingly for effective prediction. Inspired by the success of text-image fusion using affine transformation [66], we propose motion-conditioned affine transformation to refine the global context for each frame. We first obtain the motion feature  $M_{t+1} \in \mathbb{R}^{2c}$  by concatenating the output of the motion encoder  $E_m$  and audio encoder  $E_a$  as follow:

$$M_{t+1} = E_m(V_t, \hat{F}_{t+1}) || E_a(A_t), \quad (5)$$

where  $||$  denotes the concatenation operation. We also employ audio to compute  $M_{t+1}$  as it contains motion-related information. The transformation parameters  $\gamma_{t+1}^i \in \mathbb{R}^c$  and  $\beta_{t+1}^i \in \mathbb{R}^c$  for the global context feature at the  $i$ -th layer are obtained from two separate MLPs, which are both conditioned on motion feature  $M_{t+1}$  so that the transformed feature is able to adapt to motion variance:

$$\gamma_{t+1}^i = \text{MLP}_1^i(M_{t+1}), \quad \beta_{t+1}^i = \text{MLP}_2^i(M_{t+1}). \quad (6)$$

Then, channel-wise scaling and shifting is performed on  $z^i$  to obtain the adapted context  $z_{t+1}^i \in \mathbb{R}^{c^i \times h^i \times w^i}$ :

$$z_{t+1}^i = \gamma_{t+1}^i \cdot z^i + \beta_{t+1}^i. \quad (7)$$

Finally, we insert  $z_{t+1}^i$  into its corresponding layer in the U-Net decoder.

### 3.3 Optimization

We train our model in two stages due to the different convergence speeds of MME and CAR. We optimize MME in the first stage, and then we fix MME and optimize CAR in the second stage. We train MME by optimizing an optical flow reconstruction loss  $\mathcal{L}_{\text{flow}}$  and a smoothness regularization term  $\mathcal{L}_{\text{smooth}}$ :

$$\mathcal{L}_{\text{MME}} = \mathcal{L}_{\text{flow}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}}, \quad (8)$$

where  $\lambda_{\text{smooth}}$  is a weighting hyperparameter for  $\mathcal{L}_{\text{smooth}}$ .  $\mathcal{L}_{\text{flow}}$  computes a mean-squared error between the predicted optical flow and ground truth optical flow as follows:

$$\mathcal{L}_{\text{flow}} = \sum_{t=K+1}^T \|F_t - \hat{F}_t\|_2^2. \quad (9)$$

Following [48, 61], we have a smoothness regularization term  $\mathcal{L}_{\text{smooth}}$  for our optical flow estimation:

$$\mathcal{L}_{\text{smooth}} = \sum_{t=K+1}^T \|\nabla \hat{F}_t\|_1 e^{-\|\nabla V_t\|}, \quad (10)$$




Figure 3: Example of the warped images. Resultant images gradually lose the shape of the brush.

Method	Type	SSIM $\uparrow$				PSNR $\uparrow$			
		Fr 6	Fr 15	Fr 25	Mean	Fr 6	Fr 15	Fr 25	Mean
Denton and Fergus [10]	V	0.9265	0.8300	0.7999	—	18.59	14.65	13.98	—
MSPred [34]	V	0.9400	0.8903	0.8060	0.8846	21.57	19.02	17.04	19.08
Vougioukas et al. [11]	M	0.8600	0.8571	0.8573	—	15.17	14.99	15.01	—
Sound2Sight [1]	M	0.9505	0.8780	0.8749	0.8910	22.22	18.16	17.84	18.70
Our Method	M	<b>0.9608</b>	<b>0.9158</b>	<b>0.8990</b>	<b>0.9195</b>	<b>23.61</b>	<b>19.88</b>	<b>18.99</b>	<b>20.23</b>

Table 1: Comparison on Multimodal MovingMNIST with 5 seen frames. Higher values are better for SSIM and PSNR. The best results are marked in bold. Type M denotes method using audio-visual modality and Type V denotes only using visual modality.




Figure 4: Qualitative results on Multimodal MovingMNIST.

where  $\nabla$  is the gradient operator. The supervision for CAR is an image reconstruction loss  $\mathcal{L}_V$ , which computes the mean-squared error between the predicted and ground truth frames:

$$\mathcal{L}_V = \sum_{t=K+1}^T \|\mathbf{V}_t - \hat{\mathbf{V}}_t\|_2^2. \quad (11)$$

## 4 Experiments

**Datasets.** We conduct experiments on both synthetic and real-world datasets taken from [9]: 1) **Multimodal MovingMNIST.** This dataset is an extension of stochastic MovingMNIST [34] by adding artificial sound. Although MovingMNIST is synthetic, it is a standard benchmark in video prediction [13, 24, 45]. In Multimodal MovingMNIST, each digit is equipped with a unique tone and the amplitude of the tone is inversely proportional to the distance of the digit from the origin. The sound changes momentarily whenever the digit hits the boundary. This dataset consists of 8,000 training and 1,000 test videos. 2) **YouTube Painting.** It contains painting videos from YouTube. In each video, a painter is painting on a canvas in an indoor environment and there is a clear sound of the brush strokes. It contains 4.8 K training, and 500 test videos. Each video is of  $64 \times 64$  resolution, 30fps, and 3 seconds long. 3) **AudioSet-Drums.** This dataset contains selected videos from the drums class of AudioSet [16]. The video clips are selected such that the drum player is visible when the drum beat is audible. It contains 6 K training, and 1 K test videos. Each video is of  $64 \times 64$  resolution, 30fps, and 3 seconds long.

**Implementation.** The video frames are resized to  $48 \times 48$  in Multimodal MovingMNIST and  $64 \times 64$  in YouTube Painting and AudioSet-Drums. The intensity values of the video frames are normalized into  $[0, 1]$ . We directly use the precomputed spectrogram by [9] as audio input. We use the Adam optimizer [21] and the learning rate is set to  $1e-4$  and  $1e-3$  for MME and CAR, respectively. The coefficient  $\lambda_{\text{smooth}}$  is set to 0.01. We use



Method	Type	SSIM $\uparrow$				PSNR $\uparrow$			
		Fr 16	Fr 30	Fr 45	Mean	Fr 16	Fr 30	Fr 45	Mean
Denton and Fergus [12]	V	0.9779	0.6654	0.4193	—	32.52	16.05	11.84	—
MSPred [39]	V	0.9648	0.8991	0.8617	0.8965	33.42	26.42	24.25	26.65
Vougioukas <i>et al.</i> [9]	M	0.9281	0.9126	0.9027	—	26.97	25.58	24.78	—
Sound2Sight [9]	M	0.9716	0.9261	0.9074	0.9264	31.91	26.73	25.17	26.95
Our Method	M	<b>0.9848</b>	<b>0.9284</b>	<b>0.9104</b>	<b>0.9313</b>	<b>35.12</b>	<b>27.19</b>	<b>25.53</b>	<b>27.70</b>

Table 2: Comparison on YouTube Painting with 15 seen frames.

Method	Type	SSIM $\uparrow$				PSNR $\uparrow$			
		Fr 16	Fr 30	Fr 45	Mean	Fr 16	Fr 30	Fr 45	Mean
Denton and Fergus [12]	V	0.9706	0.6606	0.5097	—	30.01	16.57	13.49	—
MSPred [39]	V	0.9799	0.9382	0.9214	0.9389	33.55	27.30	25.91	27.60
Vougioukas <i>et al.</i> [9]	M	0.8986	0.8905	0.8866	—	23.62	23.14	22.91	—
Sound2Sight [9]	M	0.9875	0.9524	0.9434	0.9544	34.23	27.73	<b>26.71</b>	28.13
Our Method	M	<b>0.9896</b>	<b>0.9533</b>	<b>0.9437</b>	<b>0.9558</b>	<b>35.00</b>	<b>27.76</b>	26.68	<b>28.22</b>

Table 3: Comparison on AudioSet-Drums with 15 seen frames.

4-layer ConvLSTMs in MME and insert  $L = 4$  layers of context feature in CAR. Ground truth optical flow is precomputed using the TV-L1 algorithm from [52]. We implement our framework with Pytorch on two Nvidia GeForce GTX1080Ti GPUs.

**Evaluation setup.** On Multimodal MovingMNIST, we show 5 initial frames and full audio, and we predict the next 15 frames during training and 20 frames during evaluation. On YouTube Painting and AudioSet-Drums, we show 15 initial visual frames and full audio, and predict the next 15 visual frames during training and 30 visual frames during evaluation. We evaluate our method by the structure similarity (SSIM) [46] and Peak Signal to Noise Ratio (PSNR). We report SSIM and PSNR on selected visual frames. The mean SSIM and PSNR of the full predicted sequence are also reported to represent average performance.

## 4.1 Results

We compare our results with Sound2Sight [9], Vougioukas *et al.* [42], MSPred [39] and Denton and Fergus [12]. We re-run the source code of Sound2Sight [9] to compute the average score over 100 times of sampling instead of using the best result of the samples drawn from sampling reported by Sound2sight. In practice, it is statistically more meaningful to use the mean instead of the best result since the stochastic approach cannot always achieve the highest result in every sampling. Moreover, we do not know the ground truths of future visual frames. Therefore, it is unreasonable for [9] to search for the best samples drawn from their stochastic model using the ground truths.

**Multimodal MovingMNIST.** Tab. 1 shows quantitative results on Multimodal MovingM-




Figure 5: Qualitative results on YouTube Painting.



Method	YouTube	MNIST	AudioSet
V	11.25	—	—
V + Recall	10.43	6.40	4.63
MME	<b>9.85</b>	<b>4.82</b>	<b>3.97</b>

Table 4: Analysis of audio in MME. AEPE results are presented in  $10^{-2}$  scale.

Method	SSIM $\uparrow$			
	Fr 16	Fr 30	Fr 45	Mean
MME+Unet	<b>0.9873</b>	0.9105	0.8761	0.9143
MME+Unet+ContextEnc	0.9821	0.9266	0.9091	0.9291
MME+CAR	0.9848	<b>0.9284</b>	<b>0.9104</b>	<b>0.9313</b>

Table 5: Analysis of CAR on YouTube Painting.

NIST. Our method outperforms others by a large margin across all time steps. Fig. 4 shows the quantitative results. The first four rows show the comparison of the predicted visual frames sequence. The last two rows show the corresponding optical flow results. We follow the flow color coding of [9]. It can be seen that our method is able to correctly predict the trajectory and also maintain the complex shape of digit ‘3’ in long-term prediction. On the contrary, the predicted sequence from Sound2Sight and MSPred gradually lose the shape of the digit or even deviate from the correct trajectory.

**YouTube Painting.** Tab. 2 shows quantitative results on YouTube Painting. Our model again surpasses other methods. Fig. 5 shows the qualitative comparison with Sound2Sight. Due to entangled modeling of motion and appearance, the result of Sound2Sight gradually loses the shape of the hand. Furthermore, the hand becomes stationary at the top right corner of the image. In contrast, our model maintains a relatively good appearance and generally follows the movement of the hand in ground truth frames. We can also see that our predicted optical flows are close to the ground truths.

**AudioSet-Drums.** Tab. 3 shows quantitative results on AudioSet-Drums. Our model surpasses others at almost every time step. We notice that our superiority is not as obvious on AudioSet-Drums as on the other two datasets. The reason may be that motion only exists in a small region of frames in AudioSet-Drums, and is hard for our model to fully capture it. Nevertheless, our model still achieves the best overall result.

## 4.2 Ablation Study

**Effectiveness of audio-motion correlation in MME.** As audio is directly influencing optical flow prediction, we analyze the effectiveness of audio in MME. We compare the quality of optical flow using average endpoint error (AEPE) [19, 35] in Tab. 4. The baseline ‘V’ predicts optical flow without audio. ‘V+Recall’ *Recall* from audio motion memory, but does not condense it. ‘MME’ is our full model in stage 1 with both *condense* and *recall*. MME has the lowest error in the flow prediction, suggesting the audio is effective in helping the visual motion prediction.

**Effectiveness of Context-Aware Refinement.** We analyze context-aware refinement (CAR) in Tab. 5. ‘MME+Unet’ only uses U-Net for refinement in stage 2. ‘MME+Unet+ContextEnc’ adds a context encoder but does not perform affine transformation over the context feature. ‘MME+CAR’ is our full model with motion-conditioned affine transformation. We can see from the first and second rows that a single U-Net is insufficient to model image appearance in long-term prediction, and thus it is very important to supply additional context information. We can see from the second and last rows that our motion-conditioned affine transformation is able to adjust the global context feature for better performance in long-term predictions.

**Input video with higher resolution.** We use the video frames of original resolution  $96 \times 96$  as input. The results are shown in Tab. 6. Our model consistently outperforms [9], demonstrating its effectiveness even when provided with higher-resolution videos.

## 5 Conclusions

In this paper, we address the task of audio-visual conditioned video prediction by decoupling it into motion and appearance. We design the MME and CAR by considering the influence of audio and visual modalities in each component. MME leverages audio-motion correlation by using audio as an external motion memory in predicting future optical flow. CAR addresses the loss of appearance context in long-term prediction by adding a context encoder. To make the global context feature compatible with motion variance, we perform the motion-conditioned affine transformation to adjust the feature. The effectiveness of our proposed method is analyzed both quantitatively and qualitatively.

Method	Fr 16	Fr 30	Fr 45	Mean
Sound2Sight [11]	0.9556	0.9107	0.8936	0.9117
Ours	<b>0.9827</b>	<b>0.9252</b>	<b>0.9069</b>	<b>0.9280</b>

Table 6: SSIM results on YouTube Painting with  $96 \times 96$  videos.

**Acknowledgement.** This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2021-024), and the Tier 2 grant MOE-T2EP20120-0011 from the Singapore Ministry of Education.

## References

- [1] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 208–224. Springer, 2020.
- [2] Adil Kaan Akan, Erkut Erdem, Aykut Erdem, and Fatma Güney. Slamp: Stochastic latent appearance and motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14728–14737, 2021.
- [3] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017.
- [4] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European conference on computer vision (ECCV)*, pages 435–451, 2018.
- [5] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. *Advances in neural information processing systems*, 29:892–900, 2016.
- [6] Simon Baker, Daniel Scharstein, JP Lewis, Stefan Roth, Michael J Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *International journal of computer vision*, 92(1):1–31, 2011.
- [7] Xinzhu Bei, Yanchao Yang, and Stefano Soatto. Learning semantic-aware dynamics for video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 902–912, 2021.
- [8] D Blokhintzev. The propagation of sound in an inhomogeneous and moving medium i. *The Journal of the Acoustical Society of America*, 18(2):322–328, 1946.

- [9] Moitreya Chatterjee and Anoop Cherian. Sound2sight: Generating visual dynamics from sound and context. In *European Conference on Computer Vision*, pages 701–719. Springer, 2020.
- [10] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Localizing visual sounds the hard way. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16867–16876, 2021.
- [11] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7832–7841, 2019.
- [12] Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. You said that? *arXiv preprint arXiv:1705.02966*, 2017.
- [13] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. In *International Conference on Machine Learning*, pages 1174–1183. PMLR, 2018.
- [14] Ruohan Gao and Kristen Grauman. Co-separating sounds of visual objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3879–3888, 2019.
- [15] Ruohan Gao, Changan Chen, Ziad Al-Halah, Carl Schissler, and Kristen Grauman. Visualechoes: Spatial image representation learning through echolocation. In *European Conference on Computer Vision*, pages 658–676. Springer, 2020.
- [16] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 776–780. IEEE, 2017.
- [17] Di Hu, Feiping Nie, and Xuelong Li. Deep multimodal clustering for unsupervised audiovisual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9248–9257, 2019.
- [18] Di Hu, Rui Qian, Minyue Jiang, Xiao Tan, Shilei Wen, Errui Ding, Weiyao Lin, and Dejing Dou. Discriminative sounding objects localization via self-supervised audiovisual matching. *Advances in Neural Information Processing Systems*, 33, 2020.
- [19] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017.
- [20] Beibei Jin, Yu Hu, Qiankun Tang, Jingyu Niu, Zhiping Shi, Yinhe Han, and Xiaowei Li. Exploring spatial-temporal multi-frequency analysis for high-fidelity and temporal-consistency video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4554–4563, 2020.

- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [22] Thomas Knöpfel, Yann Sweeney, Carola I Radulescu, Nawal Zabouri, Nazanin Doostdar, Claudia Clopath, and Samuel J Barnes. Audio-visual experience strengthens multisensory assemblies in adult mouse visual cortex. *Nature communications*, 10(1):1–15, 2019.
- [23] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. *arXiv preprint arXiv:1807.00230*, 2018.
- [24] Sangmin Lee, Hak Gu Kim, Dae Hwi Choi, Hyung-II Kim, and Yong Man Ro. Video prediction recalling long-term motion context via memory alignment learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3054–3063, 2021.
- [25] Wonkwang Lee, Whie Jung, Han Zhang, Ting Chen, Jing Yu Koh, Thomas Huang, Hyungsuk Yoon, Honglak Lee, and Seunghoon Hong. Revisiting hierarchical approach for persistent long-term video prediction. *arXiv preprint arXiv:2104.06697*, 2021.
- [26] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Flow-grounded spatial-temporal video prediction from still images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 600–615, 2018.
- [27] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P Xing. Dual motion gan for future-flow embedded video prediction. In *proceedings of the IEEE international conference on computer vision*, pages 1744–1752, 2017.
- [28] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *Proceedings of the IEEE international conference on computer vision*, pages 4463–4471, 2017.
- [29] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12475–12486, 2021.
- [30] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018.
- [31] Junting Pan, Chengyu Wang, Xu Jia, Jing Shao, Lu Sheng, Junjie Yan, and Xiaogang Wang. Video generation from single semantic label map. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2019.
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [33] Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, and In So Kweon. Learning to localize sound source in visual scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4358–4366, 2018.

- [34] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852. PMLR, 2015.
- [35] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018.
- [36] Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis. *arXiv preprint arXiv:2008.05865*, 2020.
- [37] Yapeng Tian, Dingzeyu Li, and Chenliang Xu. Unified multisensory perception: Weakly-supervised audio-visual video parsing. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 436–454. Springer, 2020.
- [38] H Tijdeman. On the propagation of sound waves in cylindrical tubes. *Journal of Sound and Vibration*, 39(1):1–33, 1975.
- [39] Angel Villar-Corrales, Ani Karapetyan, Andreas Boltres, and Sven Behnke. Mspreed: Video prediction at multiple spatio-temporal scales with hierarchical recurrent networks. In *British Machine Vision Conference (BMVC)*, 2022.
- [40] Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, and Honglak Lee. Learning to generate long-term future via hierarchical prediction. In *international conference on machine learning*, pages 3560–3569. PMLR, 2017.
- [41] Ruben Villegas, Dumitru Erhan, Honglak Lee, et al. Hierarchical long-term video prediction without supervision. In *International Conference on Machine Learning*, pages 6038–6046. PMLR, 2018.
- [42] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. End-to-end speech-driven facial animation with temporal gans. *arXiv preprint arXiv:1805.09313*, 2018.
- [43] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.
- [44] Yunbo Wang, Mingsheng Long, Jianmin Wang, Zhifeng Gao, and Philip S Yu. Predrnn: Recurrent neural networks for predictive learning using spatiotemporal lstms. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 879–888, 2017.
- [45] Yunbo Wang, Lu Jiang, Ming-Hsuan Yang, Li-Jia Li, Mingsheng Long, and Li Fei-Fei. Eidetic 3d lstm: A model for video prediction and beyond. In *International conference on learning representations*, 2018.
- [46] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

- [47] Yu Wu and Yi Yang. Exploring heterogeneous clues for weakly-supervised audio-visual video parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1326–1335, 2021.
- [48] Yue Wu, Rongrong Gao, Jaesik Park, and Qifeng Chen. Future video synthesis with object motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5539–5548, 2020.
- [49] SHI Xingjian, Zhoung Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015.
- [50] Karren Yang, Bryan Russell, and Justin Salamon. Telling left from right: Learning spatial correspondence of sight and sound. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9932–9941, 2020.
- [51] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1983–1992, 2018.
- [52] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-l 1 optical flow. In *Joint pattern recognition symposium*, pages 214–223. Springer, 2007.
- [53] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *Proceedings of the European conference on computer vision (ECCV)*, pages 570–586, 2018.
- [54] Hang Zhao, Chuang Gan, Wei-Chiu Ma, and Antonio Torralba. The sound of motions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1735–1744, 2019.