

# MIMO Is All You Need : A Strong Multi-in-Multi-Out Baseline for Video Prediction

Shuliang Ning<sup>1,2\*</sup>, Mengcheng Lan<sup>2\*</sup>, Yanran Li<sup>3</sup>, Chaofeng Chen<sup>4</sup>, Qian Chen<sup>5</sup>,  
Xunlai Chen<sup>5†</sup>, Xiaoguang Han<sup>2,1†</sup>, Shuguang Cui<sup>2,1</sup>

<sup>1</sup> FNii, CUHKSZ

<sup>2</sup> SSE, CUHKSZ

<sup>3</sup> The University of Edinburgh

<sup>4</sup> Nanyang Technological University

<sup>5</sup> Shenzhen Meteorological Bureau

shuliangning@link.cuhk.edu.cn, {lanmengchengds, chaofenghust}@gmail.com, yli19@ed.ac.uk,  
{chenqian, chenxunlai}@weather.sz.gov.cn, {hanxiaoguang, shuguangcui}@cuhk.edu.cn

## Abstract

The mainstream of the existing approaches for video prediction builds up their models based on a Single-In-Single-Out (SISO) architecture, which takes the current frame as input to predict the next frame in a recursive manner. This way often leads to severe performance degradation when they try to extrapolate a longer period of future, thus limiting the practical use of the prediction model. Alternatively, a Multi-In-Multi-Out (MIMO) architecture that outputs all the future frames at one shot naturally breaks the recursive manner and therefore prevents error accumulation. However, only a few MIMO models for video prediction are proposed and they only achieve inferior performance due to the date. The real strength of the MIMO model in this area is not well noticed and is largely under-explored. Motivated by that, we conduct a comprehensive investigation in this paper to thoroughly exploit how far a simple MIMO architecture can go. Surprisingly, our empirical studies reveal that a simple MIMO model can outperform the state-of-the-art work with a large margin much more than expected, especially in dealing with long-term error accumulation. After exploring a number of ways and designs, we propose a new MIMO architecture based on extending the pure Transformer with local spatio-temporal blocks and a new multi-output decoder, namely MIMO-VP, to establish a new standard in video prediction. We evaluate our model in four highly competitive benchmarks. Extensive experiments show that our model wins 1st place on all the benchmarks with remarkable performance gains and surpasses the best SISO model in all aspects including efficiency, quantity, and quality. A dramatic error reduction is achieved when predicting 10 frames on Moving MNIST and Weather datasets respectively. We believe our model can serve as a new baseline to facilitate the future research of video prediction tasks. The code will be released.

## Introduction

Video prediction aims to forecast future frames of a video sequence conditioned on previous frames. Since anticipat-

\*These authors contributed equally.

†Both are Corresponding Authors.

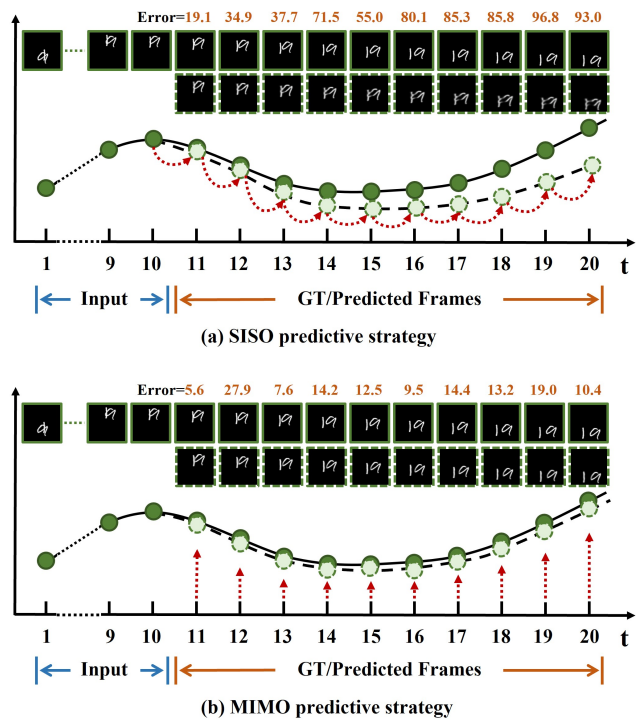


Figure 1: Illustration of MIMO and SISO on Moving MNIST: The Errors are from (a) PhyDNet (Guen and Thome 2020) (b) Our model. Figures with solid and dashed frames denote the ground-truth and predicted images respectively. The red arrows indicate the different error propagation way.

ing the future is such a fundamental capability for intelligent agents, video prediction has received increasing attention and benefited many applications such as precipitation forecasting (Xingjian et al. 2015), autonomous driving (Kwon and Park 2019), and action recognition (Liu et al. 2017). In the practical scenario of these real-life applications, an accurate prediction of the longer-term future will be greatly

beneficial.

However, the current approaches are quite limited on the prediction horizon. Most of them develop sophisticated models based on Single-In-Single-Out (SISO) architecture such as Recurrent Neural Networks (Wang et al. 2017, 2018a, 2019; Guen and Thome 2020; Yu et al. 2020; Lin et al. 2020; Wu et al. 2021; Lee et al. 2021a), Convolutional Long Short-Term Memory (ConvLSTM) (Xingjian et al. 2015) and their variants. These recurrent models learn a hidden state for the history information and generate the next frame conditioned on the current predicted frame. As a result, their quality and accuracy will degrade quickly when the generated frames go longer due to the Butterfly Effects of the small error produced in the earlier frames. Especially, video data is high-dimensional and their models have great complexity, the small errors are easy to be amplified into serious compound errors over time. Some hierarchical models predict the low-dimensional structures of the frames firstly and transfer them back into pixel space to reduce the error accumulation. They either extract the human pose keypoints (Kim et al. 2019; Walker et al. 2017; Villegas et al. 2017), face landmarks (Yang et al. 2018; Yan et al. 2018) or semantic label maps (Lee et al. 2021b). With the constraints of the high level structures, their methods can generate semantic reasonable motions in a very long time period. However, these methods are infeasible to be applied to unstructured video data such as weather and radar videos. More importantly, all of these methods only alleviate the problem by reducing the error generated at once but did not solve it essentially by breaking the error propagation chain.

In contrast, we notice that some approaches for prediction tasks build up their models in a Multi-In-Multi-Out (MIMO) manner which takes all the inputs into the encoder and output all the frames from the decoder at once, such as DVF (Liu et al. 2017) and SimVP (Gao et al. 2022). Compared with SISO, their future predictions do not rely on the previous ones and errors are not accumulated. Nevertheless, only few attempts for video prediction take MIMO architecture and their performances remain inferior to the best SISO models (Chang et al. 2021). *Have the ability of MIMO architecture been underestimated for video prediction?*

We are curious to figure out how far the MIMO architecture can achieve in video prediction. We implement MIMO architecture by a pure Transformer (Vaswani et al. 2017) model which demonstrates great success in many vision tasks recently as a general backbone since it is superior to capturing the global dependencies of input frames and output frames. However, it is worth noting that simply applying the trending Transformer models or their variants (Swin Transformer (Liu et al. 2021a), etc) is not able to deal with this problem successfully. Although Vision Transformer-based models (Arnab et al. 2021; Bertasius, Wang, and Torresani 2021; Liu et al. 2022) have achieved promising performance in video recognition, their extensions (Rakhimov et al. 2020; Weissenborn, Täckström, and Uszkoreit 2019) for video prediction are very limited and only output a single frame at once. Our experiments show Vanilla Transformer model only leads to inferior model performance since they struggle to exploit two very important clues for videos – *the*

*spatial dependencies of each frame and the complex spatio-temporal dynamics*. Notably, it is non-trivial to extend a basic Transformer from time series prediction problem for a spatio-temporal series prediction task.

To fill this gap, our objective of this paper is to develop an effective MIMO architecture for video prediction by carrying out a series of operations to tailor the basic architecture. Due to our best knowledge, this perspective has never been investigated in-depth. Various MIMO architectures for video prediction are empirically compared and studied in this work.

For the global spatio-temporal dependency clues, we incorporate the 2D convolutional structure into the multi-head attention to simultaneously capture the long term temporal dependencies of sequence and preserve the spatial information of frames. For the most important spatio-temporal clues, we design a Spatio-temporal block which replaces the simple forward layers both in the encoder and decoder. This block will capture the local Spatio-temporal context to benefit the prediction. While many concrete implementations of this block are possible to capture the local context, we reveal that a simple design of 3D convolutions layers is already capable to yield considerably performance gains than individual art methods. On Moving MNIST dataset, this block leads to a decrease of 16% MSE error.

Moreover, another problem is the decoder of the typical Transformer cannot achieve MIMO straightforwardly. To deal with that, we design a new Multi-Output decoder which takes all the placeholder of the output frames as input features. Then we rectify the last layer of the decoder to generate the output frames in one shot. This decoder captures more dependencies of the subsequent frames than recurrent models. Towards this end, a new video prediction framework based on the MIMO manner is proposed in the field. We illustrate the error growth curve is remarkably reduced compared with existing work.

Our contribution can be summarized in the following: (1) We reveal that the Multi-In-Multi-Out strategy plays a crucial role in addressing the long-standing error accumulation problem in video prediction for the first time. (2) We advocate a new MIMO architecture which leverages the Transformer-based architecture by exploiting the spatial context and local Spatio-temporal correlations of video data. A new decoder is designed to realize our MIMO idea. (3) Comprehensive experiments on four benchmarks are carried out to validate the effectiveness and fruitful insights are provided. Experimental results show that the proposed MIMO model achieves superior performance over the state-of-the-art SISO models in terms of accuracy and fidelity.

## Related Work

### Single-in-Single-Out Models.

Mainstream video prediction models (Xingjian et al. 2015; Shi et al. 2017; Wang et al. 2017, 2018a; Oliu, Selva, and Escalera 2018; Wang et al. 2018b, 2019; Guen and Thome 2020; Yu et al. 2020; Lin et al. 2020; Su et al. 2020; Wu et al. 2021; Lee et al. 2021a) follow the typical recurrent strategy to synthesis the future frames, which takes single

frame in the sequence as input and output the next one by the recurrent unit. As a prior work, Shi *et al.* (Xingjian et al. 2015) proposed ConvLSTM to address the problem. After that, its variants PredRNN (Wang et al. 2017) and MIM (Wang et al. 2019) are proposed which capture both spatial and temporal representations. MotionRNN (Wu et al. 2021) and PhyDNet (Guen and Thome 2020) explored different frameworks like GRU or two-branch methods to achieve promising results. However, it is still an inherent problem for them to capture long-term dependency and extrapolate longer sequences. To alleviate the pixel-level error propagation, a lot of work tried to learn the high-level semantic structures first, such as human pose (Villegas et al. 2017), segmentation maps (Lee et al. 2021b), then translate these to pixels space. LMC-Memory (Lee et al. 2021a) proposed to predict long-term future frames by storing and recalling long-term motion contexts of video sequences. However, all of their error still accumulates at a compound rate. In contrast, our MIMO model fundamentally changes the error propagation way rather than alleviating the error per frame.

### Multi-in-Multi-Out Models.

Instead of recurrent ways, MIMO models generate all future frames at once. It has been proved that these models (Bontempi 2008; Taieb, Sorjamaa, and Bontempi 2010) for multi-step-ahead time series prediction show superiority over the others. Compared with SISO, MIMO is able to capture the relationships of frames not only happened before but also happened afterwards. Few attempts proposed MIMO-based models to anticipate videos. For example, Liu *et al.* (Liu et al. 2017) proposed DVF, which used a 3D convolution autoencoder to learn voxel flows and synthesize video frames by flowing pixel values from existing ones. FutureGAN (Aigner and Körner 2018) is also an encoder-decoder GAN model by stacking Spatio-temporal 3d convolutions, it could predict multiple future frames at once. SimVP (Gao et al. 2022) is a simple model that is completely built upon CNN. The MIMO predictive strategy of SimVP makes it achieve significant performance only with basic CNNs. However, they did not exploit the real strength of MIMO architectures. Thus we aim to develop a top-performing MIMO model to rise the attention for them in video prediction.

### Transformer Models for Videos.

The advantage of modeling long-term dependency of videos enables Transformers (Vaswani et al. 2017; Dosovitskiy et al. 2020; Liu et al. 2021a) to be widely applied in the video domain (Zhou et al. 2018; Gabeur et al. 2020; Girdhar et al. 2019; Liu et al. 2020; Arnab et al. 2021; Dzabaraev et al. 2021; Liu et al. 2021b; Wang et al. 2021; Zhang, Hao, and Ngo 2021). For example, Zhou *et al.* (Zhou et al. 2018) proposed an end-to-end video captioning model based on Transformer. Kim *et al.* (Kim et al. 2018) proposed a novel Spatio-temporal Transformer Network (STTN) for Video Restoration. The works in (Gabeur et al. 2020; Dzabaraev et al. 2021)

investigated the potential of Transformers on video retrieval tasks. Besides, recent advances (Girdhar et al. 2019; Zhang, Hao, and Ngo 2021; Arnab et al. 2021; Liu et al. 2021b) focused on designing the T Transformer architectures for video classification. VisTR (Wang et al. 2021) is a video instance segmentation framework which outputs the masks for each instance in parallel. However, the Transformer aiming at video prediction is still underdeveloped and our work fills this gap.

## Proposed Method

### Problem Formulation

Suppose we observe a video sequence  $\mathcal{S}_{t-m+1:t} := \{\mathcal{X}_{t-m+1}, \mathcal{X}_{t-m+2}, \dots, \mathcal{X}_t\}$ , where  $\mathcal{X}_t \in \mathbb{R}^{C_0 \times H_0 \times W_0}$  represents the frame (channel  $C_0$ , height  $H_0$ , width  $W_0$ ) at the current time step  $t$ . The goal of the Spatio-temporal sequence prediction is to generate the most likely length- $n$  sequence in the future given the previous length- $m$  sequence:

$$\mathcal{S}_{t+1:t+n}^* = \arg \max_{\mathcal{S}_{t+1:t+n}} p(\mathcal{S}_{t+1:t+n} | \mathcal{S}_{t-m+1:t}) \quad (1)$$

**SISO predictive strategy.** To solve the problem, the Single-In-Single-Out architectures are widely used. They can naturally preserve the time order logic of the generated frames. Their prediction paradigm can be written as:

$$\mathcal{X}_{t'+1} = F(\mathcal{X}_{t'}, M_{t'}) + \sigma_{t'}, \quad (2)$$

where  $F$  denotes the model,  $M_{t'}$  is the memory cell that remembers the past historical information.  $\sigma_{t'}$  is a Tensor noise term.  $\mathcal{X}_{t'}$  is the observed frame ( $t' \leq t$ ) or the generated frame ( $t' > t$ ). Note that SISO predictive strategy is susceptible to the error accumulation problem. For simplicity, consider an Auto Regressive with  $K=1$  and without exogenous inputs model:

$$\mathcal{X}_t = A\mathcal{X}_{t-1} + \sigma_t \quad (3)$$

where  $A$  is a learnable Tensor. It is easy to have:

$$\mathcal{X}_{t+n} = A^{t+n-1}\mathcal{X}_t + \sum_{k=t}^{t+n-1} A^k \sigma_{t+n-k} + \sigma_{t+n} \quad (4)$$

It is clear that the error  $R_{t+n} = \sum_{k=t}^{t+n-1} A^k \sigma_{t+n-k} + \sigma_{t+n}$  will accumulate as time step  $t+n$  increases, which makes accurate prediction challenging. Besides, the long-term dependency problem inherited from RNNs would further limit the performance of the SISO strategy on video prediction.

**MIMO predictive strategy.** There is an alternative way to deploy MIMO predictive strategy:

$$\mathcal{S}_{t+1:t+n} = F(\mathcal{S}_{t-m+1:t}) + \Sigma_{t+1:t+n}, \quad (5)$$

where  $\Sigma_{t+1:t+n} := \{\sigma_{t+1}, \dots, \sigma_{t+n}\}$  is the normal prediction error. Therefore, the MIMO strategy breaks the error accumulation by synthesising all the frames in parallel. However, it is not sufficient to deduce MIMO models always outperform SISO models. Actually, the single error  $\sigma_{t+i}$  of one frame can be extremely high since the MIMO models may miss some important relationships which SISO models naturally capture. Notably, it is simple to present MIMO architectures but it is non-trivial to equip them with a strong ability to model long-term and short-term Spatio-temporal dependencies which are important for complex video data.

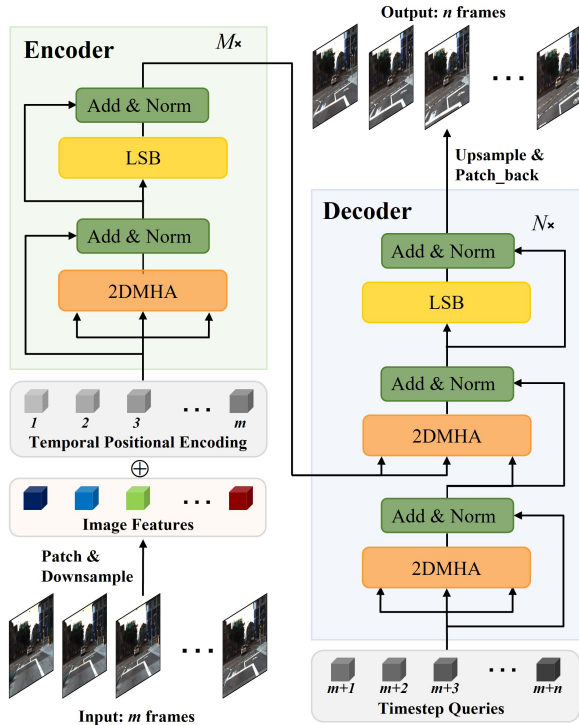


Figure 2: Illustration of MIMO-VP architecture.

## Our MIMO-VP Architecture

Encouraging by the great success of Transformer (Vaswani et al. 2017) models, we build up our MIMO architecture specific for video prediction based on the Transformer backbone. Their placeholder embedding method also helps to capture the global dependency of the frames. However, a pure Transformer is a Multi-In-Single-Out architecture and achieves only moderate performance in video prediction since they lack of ability to model complex Spatio-temporal dynamics. There is still a big gap in our goal to tailor it for video data. To tackle the issue, we propose a series of surgeries to enhance the ability of the model to capture spatial structures and local Spatio-temporal dependencies. Our final model is presented in Figure 2.

The formulation of our full model can be written as:

$$\mathcal{S}_{t+1:t+n} = \Phi_d(\Phi_e(\mathcal{S}_{t-m+1:t}), \mathcal{T}_{t+1:t+n}) \quad (6)$$

where  $\Phi_e$  and  $\Phi_d$  are the encoder and decoder modules respectively.  $\mathcal{T}_{t+1:t+n}$  is a length- $n$  timestep queries. For an input video  $\mathcal{S}_{t-m+1:t} \in \mathbb{R}^{m \times C_0 \times H_0 \times W_0}$  with  $m$  frames, we first decompose each frame into non-overlapping patches and then obtain the sequence-level feature maps  $h \in \mathbb{R}^{m \times C \times H \times W}$  after several convolution layers.

Note that, we operate 2D Convolution Multi-Head Attention (MHA) first to generate the query map and paired key-value map of each frame. It benefits the model to simultaneously learn the temporal correlation and preserve the spatial information of sequence. Moreover, different from the existing auto-regressive prediction paradigm of Transformers, we do not apply the mask operation on the multi-head

self-attention in the decoder block. Because our decoder inputs are time step embedding that is already known. In our model, each generated frame is conditioned not only on the past sequence but also on its future frames, which preserves the dependency of future frames. After exploration, we figure out two operations can boost the performance drastically. The details are provided in the following.

**Local Spatio-Temporal Block** We figure out one reason that the vanilla Transformer is limited in performance since they lack the ability to capture the short-term variation of sequences which is crucial for video generation. Therefore, we are motivated to install a local Spatio-temporal block which aims to learn the high order relationship between time step queries  $\mathcal{T}$  and their corresponding output sequence  $\mathcal{S}$ . We find out a simple way to implement this idea by installing this block in the place of the FCN layers of the encoder and decoder. Theoretically, this block can be any neural network which has the ability to capture Spatio-temporal information for video data. While a more complex design potentially will result in more gains, we hope to design this block as simple as possible to validate the importance of local Spatio-temporal clues for video prediction. Inspired by most MIMO frameworks in video domains, we choose a simple 3D Convolution layer which is widely used to capture these clues as our block. Specifically, the block is consisted of two embedding layers with each is a 3DConv, LayerNorm, SiLU block. The kernel size of 3DConv is  $3 \times 3 \times 3$ .

**Multi-Out Decoder** To achieve our goal, we need to rectify the original decoder to enable it to produce several future frames parallelly. The challenge here is how to maintain the order information of in video sequences. Considering the permutation-invariant character of the Transformer architecture, we here focus on the temporal positional encoding of the sequence. Specifically, we define a time step vector  $T = [1, 2, \dots, m+n]$ . Then we embed  $T \in \mathbb{R}^{1 \times (m+n)}$  to  $\hat{T} \in \mathbb{R}^{(m+n) \times C}$  by an embedding layer. Finally, we expant  $\hat{T}$  to  $\mathcal{T} \in \mathbb{R}^{(m+n) \times C \times H \times W}$  along the height and width dimensions. Accordingly, the temporal positional encoding of input sequence is  $\mathcal{T}_{1:m} \in \mathbb{R}^{m \times C \times H \times W}$ . After that, we regard  $\mathcal{T}_{m+1:m+n} \in \mathbb{R}^{n \times C \times H \times W}$  as time step queries and input them to the decoder module. In this way, our new decoder is capable of dealing with multi-output scenarios. SISO strategy is somehow a specific case of MIMO strategy when the input length is one. The details of this Multi-Out decoder are illustrated in *supplementary materials*.

## Experimental Results

### Dataset Descriptions

**Moving MNIST** (Srivastava, Mansimov, and Salakhudinov 2015) is a standard synthetic dataset for video prediction. We generate training set following (Guen and Thome 2020), and adopt the widely used testing set of 10,000 sequences provided by (Srivastava, Mansimov, and Salakhudinov 2015).

**Human3.6M** (Ionescu et al. 2013) dataset is a real-world human pose dataset comprising 17 kinds of human actions and 3.6 million poses. We use subjects S1, S5, S6, S7 and S8 for training, and subjects S9, S11 for testing. We predict 4 frames in the future conditioned on the 4 observed frames.

Method	10 - 10			10 - 30		
	MSE ↓	MAE ↓	SSIM ↑	MSE ↓	MAE ↓	SSIM ↑
ConvLSTM (Shi et al. 2017)	103.3	182.9	0.707	-	-	-
MIM (Wang et al. 2019)	44.2	101.1	0.910	-	-	-
CrevNet(Yu et al. 2019)	22.3	-	0.949	-	-	-
ConvTTLSTM (Su et al. 2020)	56.0	92.9	0.913	105.7	-	0.840
PhyDNet (Guen and Thome 2020)	24.2	70.3	0.947	58.7	144.1	0.852
LMC-Memory (Lee et al. 2021a)	41.5	-	0.924	73.2	-	0.879
MotionRNN (Wu et al. 2021)	29.2	87.7	0.933	64.2	156.1	0.844
SimVP (Gao et al. 2022)	23.8	68.9	0.948	-	-	-
<b>MIMO-VP</b>	<b>17.7</b>	<b>51.6</b>	<b>0.964</b>	<b>31.6</b>	<b>79.1</b>	<b>0.934</b>

Table 1: Quantitative comparison of different models on Moving MNIST dataset.

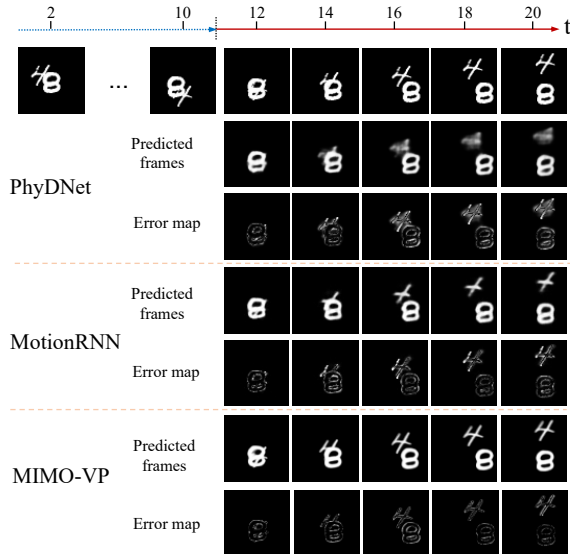


Figure 3: Prediction examples on Moving MNIST dataset.




Figure 4: Prediction examples on Human3.6M dataset.

**Weather** dataset is a real radar echo dataset for precipitation nowcasting, which is obtained from the local weather bureau. Our training set and testing set contain 9,600 and 2400 radar sequences (resolution  $128 \times 128 \times 1$ ) respectively.

**KITTI** dataset is one of the most popular datasets for autonomous driving and also a benchmark for computer vision algorithms. We predict 5 frames in the future when taking 5 observed frames as input.

**Implementation Details** We use Adam (Kingma and Ba 2014) optimizer with 0.0005 learning rate and L1+L2 loss to train our model. We implemented the model with PyTorch and conducted experiments on NVIDIA V100 GPUs. Mean Squared Error (MSE), Mean Absolute Error (MAE), the Structural Similarity(SSIM), LPIPS (Zhang et al. 2018) and Peak Signal to Noise Ratio (PSNR) are evaluated in our work. All these metrics are averaged for all predicted frames. Additionally, we also take Critical Success Index (CSI) into account for Weather dataset. Details are contained in *supplementary materials*.

## Quantitative and Qualitative Comparison

**Results on Moving MNIST.** Firstly, we report the results which use 10 frames to predict 10 future frames. As shown in Table 1 and Figure 3, our model significantly outperforms all the other models. Current RNN-based models get blurry and wrong predictions in this case, and their prediction error grows larger over time. On the contrary, our model can predict high quality frames. To further explore our performance on long-term prediction, we conduct an experiment by predicting 30 future frames conditioned on 10 input frames. As shown in Table 1, our method far surpasses the others in predicting 30 frames in terms of the MSE metric. Compared with PhyDNet, our MIMO-VP improves by 46.2% in MSE, which is much larger than 26.9% in 10-10 prediction task.

**Results on Human3.6M.** Table 2 shows that MIMO-VP achieves the best performance on Human3.6M. Compared with PhyDNet, MIMO-VP improves by 35.8% in MAE. In Figure 4, complicated background and human motion give much more challenge for predicting accurate and reasonable results. While our MIMO-VP can still product relatively precise future frames comparing with other models.

Method	MAE/100 ↓	SSIM ↑	PSNR ↑
MIM	17.8	0.790	20.57
PhyDNet	16.2	0.901	-
LMC-Memory	17.0	0.820	21.05
MotionRNN	14.8	0.846	22.16
<b>MIMO-VP</b>	<b>10.4</b>	<b>0.941</b>	<b>24.52</b>

Table 2: Quantitative comparisons on Human3.6M dataset.

Method	MAE/100 ↓	SSIM ↑	LPIPS ↓	PSNR ↑
CrevNet	56.3	0.587	0.454	17.046
PhyDNet	44.9	0.674	0.403	19.159
LMC-Memory	44.6	0.660	0.410	18.692
MotionRNN	45.0	0.652	0.384	18.931
<b>MIMO-VP</b>	<b>40.9</b>	<b>0.703</b>	<b>0.308</b>	<b>19.616</b>

Table 3: Quantitative comparisons on KITTI dataset.

Method	MAE ↓	SSIM ↑	CSI-20 ↑	CSI-30 ↑
MIM	402.6	0.725	0.402	0.164
PhyDNet	280.1	0.859	0.557	0.368
MotionRNN	336.9	0.809	0.519	0.290
<b>MIMO-VP</b>	<b>235.6</b>	<b>0.865</b>	<b>0.595</b>	<b>0.400</b>

Table 4: Quantitative comparisons on Weather dataset.

**Results on KITTI.** As shown in Table 3, our proposed MIMO-VP achieves consistent improvement in all evaluation metrics. Compared with LMC-Memory, our model improves by 8.3% in MAE. Quantitative visualization results in Figure 5 show that our model can preserve much more details of both background and foreground. For instance, from the right part of last predicted frames, we can find that the results of our MIMO-VP is much similar to GroundTruth while other methods produce obviously blurry frames.

**Results on Weather.** From Table 4, we can find that our MIMO-VP has significantly improvement in CSI scores within different thresholds, which indicates our approach can predict more credible radar echoes for precipitation nowcasting. The visualization radar echo maps are shown in Figure 6. Different color represents different radar echo intensity (dBZs), *i.e.* yellow indicates higher intensity than green. Additionally, the center region of the radar echo is zoomed in and placed below each predicted frame. We find that our MIMO-VP predicts more correct yellow regions than other methods, which means MIMO-VP is a more suitable method for precipitation nowcasting.

**The comparison of different MIMOs** Quantitative comparisons are detailed in Table 5, where 3D convolution is a fully convolutional network. For ViT (Dosovitskiy et al. 2020) and Swin-Transformer (Liu et al. 2021a), we replace the classification layer with convolution operation to be consistent with original image size. We observe that all other MIMO methods can not achieve state-of-the-art performance when compared with the best SISO method—CrevNet (Yu et al. 2019) except ours. To verify the effectiveness of our MIMO-VP, we carry out an ablation study to analyse each key component.

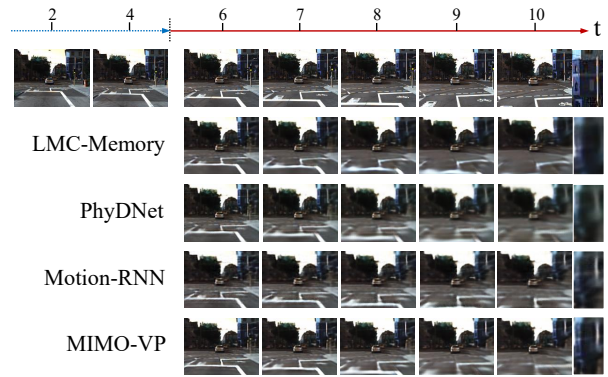


Figure 5: Prediction examples on KITTI dataset.

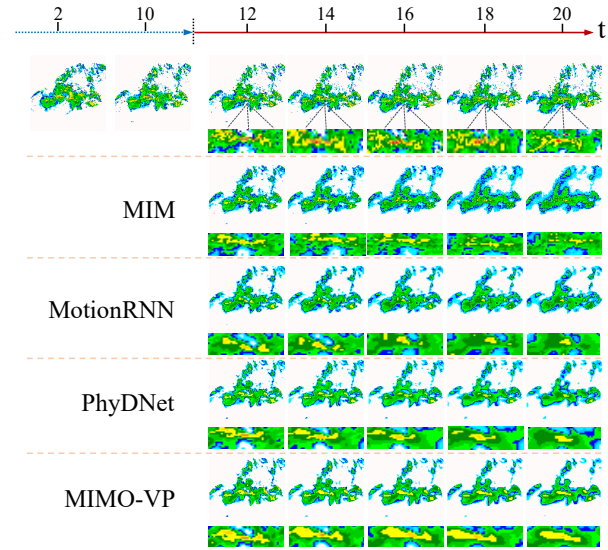


Figure 6: Prediction examples on Weather dataset.

### Ablation Study of Each Component

Quantitative results on Moving MNIST dataset are shown in Table 6. (1) **The Spatial Dependency of Frames.** We can see that MIMO-VP without 2DMHA results in a large performance degradation in comparison with MIMO-VP, with an increase of 15.1 and 32.8 in terms of MSE and MAE respectively. This phenomenon indicates that 2DMHA mechanism, which is able to capture the global long-term spatiotemporal relationships of sequences, makes an important contribution to the accurate video prediction. (2) **The Local Spatio-temporal Block (LSB).** The results of MIMO-VP without LSB also show inferior to those of MIMO-VP, which verifies that the local spatial-temporal variations of sequence are helpful for prediction. (3) **The Multi-Out Decoder.** MISO-VP denotes a variant of MIMO-VP that adopts the MISO predictive strategy similar to vanilla Transformer. It infers the next one frame using the predicted frames, in a recursive manner. Experimental results show that MIMO-VP beats down the MISO-VP with performance improves by 18.0% and 14.0% in MSE and MAE respectively. The

Method	MSE ↓	MAE ↓	SSIM ↑	#Params (M)
CrevNet	22.3	-	0.949	-
3D convolution	32.8	84.4	-	19.8
ViT	50.2	126.5	0.887	17.5
Swin-Transformer	28.8	77.2	0.940	21.8
SimVP	23.8	68.9	0.948	15.8
<b>MIMO-VP</b>	<b>17.7</b>	<b>51.6</b>	<b>0.964</b>	20.2

Table 5: Results of MIMO models on Moving MNIST.

Method	MSE ↓	MAE ↓
Without 2DMHA	32.8	84.4
Without LSB	21.2	59.2
MISO-VP	22.8	62.8
<b>MIMO-VP</b>	<b>17.7</b>	<b>51.6</b>

Table 6: Ablation Study for Each Components.

Method	MSE ↓	MAE ↓	SSIM ↑
MIMO-VP_F	21.6	59.4	0.955
<b>MIMO-VP</b>	<b>17.7</b>	<b>51.6</b>	<b>0.964</b>

Table 7: Comparison of MIMO-VP and its variant on Moving MNIST dataset.

reason may be that recursive single-out predictive strategy will trigger the error accumulation problem and overlook the stochastic dependency among the future frames.

## Comparison of MIMO and SISO

**Significance of ‘MI’: long-term dependency** Comparison results between MIMO-VP and PhyDNet of predicting 10 future frames on Moving MNIST with different input lengths are shown in Figure 7a. It is observed that the MSE metric of MIMO-VP drops quickly with the input length increasing from 5 to 50. PhyDNet improves significantly when the input length increases from 5 to 10, while improves slightly from 10 to 50. Precisely, the MSE of MIMO-VP and PhyDNet decreases by 64.3% and 24.4% respectively from 10 input frames to 50 input frames. Those results demonstrate that our MIMO-VP has a stronger ability to capture the long-term spatiotemporal dependency of video sequences for accurate prediction, while PhyDNet, an RNN-based method, suffers from the long-term dependency problem because of the limited capacity of memory.

**Significance of ‘MO’: free of error accumulation** Figure 7b illustrates the frame-wise MSE of MIMO-VP and PhyDNet, corresponding to the sample in Figure 3. We can observe that the curve of PhyDNet increase dramatically over time because PhyDNet suffers from the error accumulation problem. By contrast, the curve of MIMO-VP increase slightly because the normal predictive error, which is related to the digits overlapping, increases over time. To further verify that the error accumulation problem is even worse in long-term case, we compare the results of predicting 10, 30, 50 and 90 frames on Moving MNIST dataset. In training phase, all the models are trained by using 10 frames to predict next 10 frames, while in testing phase, 10, 30, 50, and 90




Figure 7: (a) Comparison of MIMO-VP and PhyDNet by predicting 10 future frames conditioned on different length of input sequence; (b) Frame-wise MSE of MIMO-VP and PhyDNet on Moving MNIST dataset.

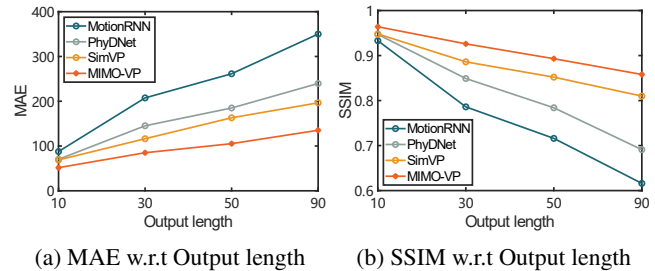


Figure 8: Results of long-term prediction on Moving MNIST dataset.

frames are predicted in a recurrent manner. Experimental results are shown in Figure 8a and Figure 8b, from which we can observe that the error accumulation problems of SISO models (MotionRNN, PhyDNet) is much severe than MIMO models (SimVP, MIMO-VP). Thus, ‘MO’ models can overcome the error accumulation problem in video prediction.

**Future frame dependency** We compare our MIMO-VP with its variant *MIMO-F*, a baseline which removes the self-attention module in the decoder so that it can not capture all the dependency of frames. Predictive results are shown in Table 7, from which we can see MIMO-VP achieves much better performance than *MIMO-F*. This phenomenon verifies the benefit of MIMO-VP in preserving the dependency among the future frames for accurate prediction.

## Conclusion

In this paper, we propose a new MIMO architecture for video prediction, which overcomes the long-term dependencies and error accumulation problems inherently introduced in the Single-In-Single-Out (SISO) models. The results show that our model wins 1st place in many evaluations and challenge datasets and surpasses the state-of-the-arts SISO models in a large gap, especially over a longer period. By this work, we encourage the researchers in prediction area to shift more attention on MIMO models. We believe our model could serve as a new backbone and facilitate the future research of video prediction.

## Acknowledgments

This work is supported by National Key Research and Development Program of China for Intergovernmental Cooperation with grant No. 2019YFE0110100, the Basic Research Project No. HZQB-KCZYZ-2021067 of Hetao Shenzhen-HK S&T Cooperation Zone, NSFC with Grant No. 62293482, the National Key R&D Program of China with grant No. 2018YFB1800800, the Shenzhen Outstanding Talents Training Fund 202002, the Guangdong Research Projects No. 2017ZT07X152 and No. 2019CX01X104, the Guangdong Provincial Key Laboratory of Future Networks of Intelligence (Grant No. 2022B1212010001), the Shenzhen Key Laboratory of Big Data and Artificial Intelligence (Grant No. ZDSYS201707251409055), Science and technology innovation team project of Guangdong Meteorological Bureau No.GRMCTD202104, and Innovation and Development Project of China Meteorological Administration (Grant No. CXFZ2022J002). Thanks Oisín Mac Aodha for his advices on our work.

## References

- Aigner, S.; and Körner, M. 2018. Futuregan: Anticipating the future frames of video sequences using spatio-temporal 3d convolutions in progressively growing gans. *arXiv preprint arXiv:1810.01325*.
- Arnab, A.; Deghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6836–6846.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding? In *ICML*, volume 2, 4.
- Bontempi, G. 2008. Long term time series prediction with multi-input multi-output local learning. *Proc. 2nd ESTSP*, 145–154.
- Chang, Z.; Zhang, X.; Wang, S.; Ma, S.; Ye, Y.; Xinguang, X.; and Gao, W. 2021. MAU: A Motion-Aware Unit for Video Prediction and Beyond. *Advances in Neural Information Processing Systems*, 34: 26950–26962.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Deghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Dzabraev, M.; Kalashnikov, M.; Komkov, S.; and Petiushko, A. 2021. Mdmmt: Multimodal multimodal transformer for video retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3354–3363.
- Gabeur, V.; Sun, C.; Alahari, K.; and Schmid, C. 2020. Multi-modal transformer for video retrieval. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, 214–229. Springer.
- Gao, Z.; Tan, C.; Wu, L.; and Li, S. Z. 2022. SimVP: Simpler Yet Better Video Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3170–3180.
- Girdhar, R.; Carreira, J.; Doersch, C.; and Zisserman, A. 2019. Video action transformer network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 244–253.
- Guen, V. L.; and Thome, N. 2020. Disentangling physical dynamics from unknown factors for unsupervised video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11474–11484.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2013. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7): 1325–1339.
- Kim, T. H.; Sajjadi, M. S.; Hirsch, M.; and Scholkopf, B. 2018. Spatio-temporal transformer network for video restoration. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 106–122.
- Kim, Y.; Nam, S.; Cho, I.; and Kim, S. J. 2019. Unsupervised Keypoint Learning for Guiding Class-Conditional Video Prediction. In *Advances in Neural Information Processing Systems*, 3814–3824.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kwon, Y.-H.; and Park, M.-G. 2019. Predicting future frames using retrospective cycle gan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1811–1820.
- Lee, S.; Kim, H. G.; Choi, D. H.; Kim, H.-I.; and Ro, Y. M. 2021a. Video Prediction Recalling Long-term Motion Context via Memory Alignment Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3054–3063.
- Lee, W.; Jung, W.; Zhang, H.; Chen, T.; Koh, J. Y.; Huang, T.; Yoon, H.; Lee, H.; and Hong, S. 2021b. Revisiting Hierarchical Approach for Persistent Long-Term Video Prediction. In *International Conference on Learning Representations, ICLR 2021. International Conference on Learning Representations, ICLR*.
- Lin, Z.; Li, M.; Zheng, Z.; Cheng, Y.; and Yuan, C. 2020. Self-attention convlstm for spatiotemporal prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 11531–11538.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021a. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*.
- Liu, Z.; Luo, S.; Li, W.; Lu, J.; Wu, Y.; Sun, S.; Li, C.; and Yang, L. 2020. Convtransformer: A convolutional transformer network for video frame synthesis. *arXiv preprint arXiv:2011.10185*.
- Liu, Z.; Ning, J.; Cao, Y.; Wei, Y.; Zhang, Z.; Lin, S.; and Hu, H. 2021b. Video swin transformer. *arXiv preprint arXiv:2106.13230*.
- Liu, Z.; Ning, J.; Cao, Y.; Wei, Y.; Zhang, Z.; Lin, S.; and Hu, H. 2022. Video swin transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3202–3211.



- Liu, Z.; Yeh, R. A.; Tang, X.; Liu, Y.; and Agarwala, A. 2017. Video frame synthesis using deep voxel flow. In *Proceedings of the IEEE International Conference on Computer Vision*, 4463–4471.
- Oliu, M.; Selva, J.; and Escalera, S. 2018. Folded recurrent neural networks for future video prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 716–731.
- Rakhimov, R.; Volkhonskiy, D.; Artemov, A.; Zorin, D.; and Burnaev, E. 2020. Latent video transformer. *arXiv preprint arXiv:2006.10704*.
- Shi, X.; Gao, Z.; Lausen, L.; Wang, H.; Yeung, D.-Y.; Wong, W.-k.; and Woo, W.-c. 2017. Deep learning for precipitation nowcasting: A benchmark and a new model. In *Advances in neural information processing systems*, 5617–5627.
- Srivastava, N.; Mansimov, E.; and Salakhudinov, R. 2015. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, 843–852.
- Su, J.; Byeon, W.; Kossaifi, J.; Huang, F.; Kautz, J.; and Anandkumar, A. 2020. Convolutional Tensor-Train LSTM for Spatio-Temporal Learning. *Advances in Neural Information Processing Systems*, 33.
- Taieb, S. B.; Sorjamaa, A.; and Bontempi, G. 2010. Multiple-output modeling for multi-step-ahead time series forecasting. *Neurocomputing*, 73(10-12): 1950–1957.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Villegas, R.; Yang, J.; Zou, Y.; Sohn, S.; Lin, X.; and Lee, H. 2017. Learning to generate long-term future via hierarchical prediction. In *international conference on machine learning*, 3560–3569. PMLR.
- Walker, J.; Marino, K.; Gupta, A.; and Hebert, M. 2017. The pose knows: Video forecasting by generating pose futures. In *Proceedings of the IEEE international conference on computer vision*, 3332–3341.
- Wang, Y.; Gao, Z.; Long, M.; Wang, J.; and Philip, S. Y. 2018a. Predrnn++: Towards a resolution of the deep-in-time dilemma in spatiotemporal predictive learning. In *International Conference on Machine Learning*, 5123–5132. PMLR.
- Wang, Y.; Jiang, L.; Yang, M.-H.; Li, L.-J.; Long, M.; and Fei-Fei, L. 2018b. Eidetic 3d lstm: A model for video prediction and beyond. In *International Conference on Learning Representations*.
- Wang, Y.; Long, M.; Wang, J.; Gao, Z.; and Yu, P. S. 2017. Predrnn: Recurrent neural networks for predictive learning using spatiotemporal lstms. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 879–888.
- Wang, Y.; Xu, Z.; Wang, X.; Shen, C.; Cheng, B.; Shen, H.; and Xia, H. 2021. End-to-end video instance segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8741–8750.
- Wang, Y.; Zhang, J.; Zhu, H.; Long, M.; Wang, J.; and Yu, P. S. 2019. Memory in memory: A predictive neural network for learning higher-order non-stationarity from spatiotemporal dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9154–9162.
- Weissenborn, D.; Täckström, O.; and Uszkoreit, J. 2019. Scaling Autoregressive Video Models. In *International Conference on Learning Representations*.
- Wu, H.; Yao, Z.; Wang, J.; and Long, M. 2021. Motion-RNN: A flexible model for video prediction with spacetime-varying motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15435–15444.
- Xingjian, S.; Chen, Z.; Wang, H.; Yeung, D.-Y.; Wong, W.-K.; and Woo, W.-c. 2015. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, 802–810.
- Yan, X.; Rastogi, A.; Villegas, R.; Sunkavalli, K.; Shechtman, E.; Hadap, S.; Yumer, E.; and Lee, H. 2018. Mt-vae: Learning motion transformations to generate multimodal human dynamics. In *Proceedings of the European conference on computer vision (ECCV)*, 265–281.
- Yang, C.; Wang, Z.; Zhu, X.; Huang, C.; Shi, J.; and Lin, D. 2018. Pose guided human video generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 201–216.
- Yu, W.; Lu, Y.; Easterbrook, S.; and Fidler, S. 2019. CrevNet: Conditionally Reversible Video Prediction. *arXiv preprint arXiv:1910.11577*.
- Yu, W.; Lu, Y.; Easterbrook, S. M.; and Fidler, S. 2020. Efficient and Information-Preserving Future Frame Prediction and Beyond. In *ICLR*.
- Zhang, H.; Hao, Y.; and Ngo, C.-W. 2021. Token Shift Transformer for Video Classification. In *Proceedings of the 29th ACM International Conference on Multimedia*, 917–925.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhou, L.; Zhou, Y.; Corso, J. J.; Socher, R.; and Xiong, C. 2018. End-to-end dense video captioning with masked transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8739–8748.