



VideoMamba: Spatio-Temporal Selective State Space Model

Jinyoung Park^{*}, Hee-Seon Kim^{*}, Kangwook Ko^{*}, Minbeom Kim, and Changick Kim

Korea Advanced Institute of Science and Technology
{jinyoungpark, hskim98, kw.ko, alsqja1754, changick}@kaist.ac.kr
<https://github.com/jinyjelly/VideoMamba>

Abstract. We introduce VideoMamba, a novel adaptation of the pure Mamba architecture, specifically designed for video recognition. Unlike transformers that rely on self-attention mechanisms leading to high computational costs by quadratic complexity, VideoMamba leverages Mamba’s linear complexity and selective SSM mechanism for more efficient processing. The proposed Spatio-Temporal Forward and Backward SSM allows the model to effectively capture the complex relationship between non-sequential spatial and sequential temporal information in video. Consequently, VideoMamba is not only resource-efficient but also effective in capturing long-range dependency in videos, demonstrated by competitive performance and outstanding efficiency on a variety of video understanding benchmarks. Our work highlights the potential of VideoMamba as a powerful tool for video understanding, offering a simple yet effective baseline for future research in video analysis.

Keywords: Efficient Video Recognition · State Space Models · Mamba

1 Introduction

In the field of natural language processing, Transformer [43] has demonstrated remarkable performance. After the success of Vision Transformer [9], transformers began to be utilized across various computer vision problems, surpassing the performance of previous CNN-based methods [2, 5, 7, 46, 54]. However, the core operator of Transformers, self-attention, presents a challenge due to its quadratic complexity. This becomes particularly problematic in video recognition tasks that require the processing of multiple frames, limiting the applicability in resource-constrained environments.

Recently, Mamba [15] is introduced, offering an appealing alternative to this problem. Based on structured Selective State Space Models (SSMs), Mamba employs a selective scan mechanism that dynamically adjusts parameters based on

^{*} Equal contribution.




Fig. 1: Performance and efficiency comparisons among ImageNet-1K pre-trained video models. VideoMamba shows superior or comparable performance to VideoSwin-T [33], while having clear advantage in terms of reduced GFLOPs and memory consumption compared to VideoSwin-T [33] and VideoMAE-S [45].

the input. With the selective scan mechanism and hardware-aware algorithm, Mamba effectively captures long-range dependency with contextual awareness, while maintaining linear complexity. Consequentially, Mamba achieves superior performance compared to transformers in various 1D sequence modeling problems, such as language modeling, audio processing, and genomics. Efforts to apply this innovative mechanism to a range of vision tasks, such as image classification and segmentation, are currently underway [30, 32, 56].

Inspired by the success of Mamba, we propose **VideoMamba**, the first video model to offer a comprehensive analysis of adapting the pure Mamba architecture for video tasks. Our video model not only demands lower computational resources but also delivers competitive performance compared to its transformer counterparts of similar size. While the Mamba model’s linear complexity and ability to capture long-range dependency are well-suited for video applications, representing the spatio-temporal information of videos into a 1D sequence presents a significant challenge. To effectively address the issue of handling non-sequential spatial information, recently introduced variants of the mamba models in vision [29, 32, 56] adopt bidirectional scanning methods. Building on this foundation, VideoMamba also adopts the bidirectional scanning strategy for video data, presenting a pioneering study on expanding vision models to video applications and making the pre-trained model available. However, video data introduces a more challenging scenario due to the interlinked relationship between non-sequential spatial information (e.g., the position and posture of a person in specific frames) and sequential temporal information (e.g., changes in a person’s actions over frames). Our carefully designed experiments aim to tackle the inherent complexity of processing spatio-temporal information. To extend the bidirectional SSM module effectively, we develop it into *Spatio-Temporal Forward and Backward SSM* modules. Focusing on defining the backward scanning direction, we explore the effects of spatio-temporal scanning among spatial, temporal, and spatio-temporal reversal. Through extensive experiments and ablation studies, we provide insights and examine various design choices and training recipes. Additionally, we analyze how the VideoMamba model responds to temporal consistency, investigating whether the model merely treats video as a

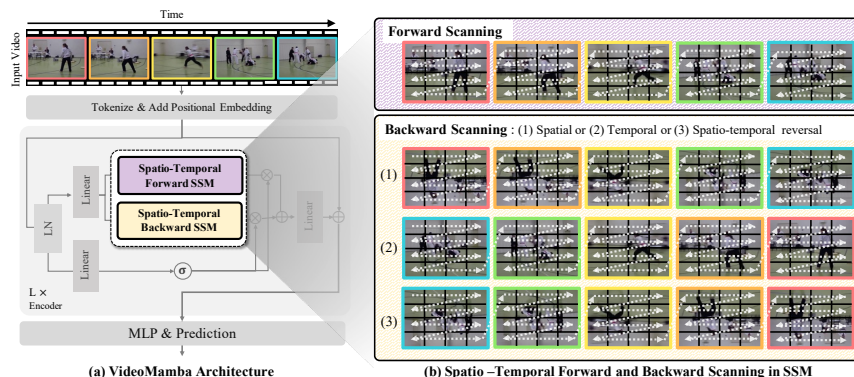


Fig. 2: Comprehensive View of VideoMamba’s Framework. (a) Architecture of VideoMamba. This includes the initial tokenization of video frames, addition of positional embeddings, and processing through encoder blocks that utilize proposed Spatio-Temporal Forward and Backward SSMs for thorough video analysis. (b) Process of Spatio-Temporal Forward and Backward Scanning within the SSMs, with white dashed arrows indicating the scanning direction of video tokens.

bundle of images or not. Moreover, an analysis of one of Mamba’s key components, Delta, offers valuable insights into how the model captures and processes spatio-temporal information.

VideoMamba has demonstrated its competitiveness across various video benchmarks. It exhibits either superior performance in HMDB 51 [28] and Something-Something V2 [14], or comparable performance in Kinetics-400 [25] compared to other video models. Its efficiency stands out, as VideoMamba’s outstanding reduction in GFLOPs and memory usage is highlighted in Fig. 1.

The contributions of this paper are:

- We introduce VideoMamba, a pioneering study of a simple yet efficient pure-Mamba-based model, highlighting its potential for future advancements in video recognition.
- Through Spatio-Temporal Forward and Backward SSM, we tackle the unique challenge of integrating non-sequential spatial with sequential temporal information in video processing.
- We demonstrate VideoMamba’s competitive performance and its reduced computational demands compared to conventional transformer models.
- Our extensive experiments and analysis underscore the Mamba model’s strength as an effective tool for video processing.

2 Related work

2.1 Video Understanding

Traditionally, CNNs [11, 20, 23, 26, 31, 37, 41] have been the standard backbone architectures in video modeling. These CNNs leverage 3D convolutions [3, 6, 11,

36,41] or factorize spatial and temporal convolutions [19,37,42,51] for efficiency. Before the rise of pure Transformer-based models, the integration of attention mechanisms [13,27,48], renowned for their efficiency in capturing long-range dependencies within data, into CNN frameworks achieved promising results. The success of non-local networks [48] and hybrid models [49,53] led to the development of pure Transformer architectures specifically designed for video recognition [1,2,4,10,33,34,47,50]. However, a significant challenge with pure Transformer architectures lies in their computational complexity for lengthy videos. This stems from the quadratic complexity of the attention mechanism as the sequence length increases [43]. To address this issue, several approaches have been proposed, including factorization techniques [2,4], windowing mechanisms [33], and masked token reconstruction [12,40,45]. These methods achieve promising results in handling long sequences while capturing spatial and temporal information concurrently. While pure Transformers demonstrate remarkable capabilities, the inherent quadratic complexity with sequence length remains an open challenge that necessitates further exploration for efficient video understanding.

2.2 State Space Models

Derived from the classical state space models (SSMs) [24], structured SSMs (S4) [16,17] have emerged as promising frameworks in modeling long-distance sequences with a linear increase in computational cost. S4’s success sparked a wave of research, resulting in diverse S4-inspired models that capture long-range dependencies in sequential data [18,38]. These variants achieve competitive performance on various tasks [22,35,44]. S4’s strength lies in its adherence to Linear Time Invariance (LTI), ensuring consistent output for identical inputs irrespective of their temporal application. While LTI systems present several advantages, they also introduce limitations, particularly in handling time-variant dynamics. The constraint that the internal state transition matrix remains constant across the sequence restricts the model’s ability to adapt to evolving content, limiting its application in scenarios requiring content-based reasoning.

Recently, Mamba [15] addressed the limitations by introducing a selective state-space model that dynamically adjusts its parameters based on the input sequence. This flexibility allows Mamba to perform context-dependent reasoning, significantly enhancing its applicability across various domains from language and speech [15] to complex visual data [21,29,30,32,56]. While concurrent work [29] shows the potential for extension into multi-dimensions, its particular application in video recognition tasks remains unexplored.

In this work, we propose VideoMamba, an extension of the Mamba model specifically designed for video understanding tasks. It leverages Mamba’s capabilities to enhance long-range modeling for video data. This approach builds upon the recent advancements in SSMs for video tasks. Unlike prior work, S4ND [35], which struggled to capture input-dependent information, VideoMamba incorporates a selective scan mechanism to address this limitation.

3 Preliminary

3.1 State Space Model

State space models are linear time-invariant systems that map 1-D input sequence to 1-D output sequence through a latent state. Mathematically, these models can be expressed as simple ordinary differential equations (ODE) as follows:

$$\begin{aligned} h'(t) &= Ah(t) + Bx(t), \\ y(t) &= Ch(t) + Dx(t), \end{aligned} \tag{1}$$

where $x(t) \in \mathbb{R}$ is a continuous input sequence, $y(t) \in \mathbb{R}$ is an output sequence, and $h(t) \in \mathbb{R}^N$ denotes a latent state.

To enable the processing of discrete signals in Equations 1, a discretization process is necessary. The most commonly used method for discretization is the zero-order hold (ZOH) technique, which transforms parameters A and B for continuous signals into parameters \bar{A} and \bar{B} for discrete signals through the step size parameter Δ . The parameter discretization process is formulated as follows:

$$\begin{aligned} \bar{A} &= \exp(\Delta A), \\ \bar{B} &= (\exp(\Delta A) - I)(\Delta A)^{-1}B, \\ \bar{C} &= C. \end{aligned} \tag{2}$$

After the discretization, the continuous system in Equations 1 can be written in discrete form as

$$\begin{aligned} h_k &= \bar{A}h_{k-1} + \bar{B}x_k, \\ y_k &= \bar{C}h_k, \end{aligned} \tag{3}$$

where x_k and y_k are discrete input and output signals.

3.2 Selective SSM

While SSMs have demonstrated outstanding performance in various tasks with sequential data, as previously mentioned, they suffer from the inherent limitation of being an LTI system. In other words, with parameters A , B , C , and Δ remaining constant across all time steps, the model's computations are independent of the input, making it challenging to address problems requiring contextual awareness. Recently, to overcome these limitations, Mamba [15] is proposed, leveraging a *selective scan* mechanism. In Mamba, model parameters such as B , C , and Δ are dynamically determined based on the input, enabling the model to understand the context of input sequences. We adopt this selective SSM as the core operator in our proposed model.

The key parameter of selective SSM, Δ , acts as a gating mechanism that controls the influence of specific elements within the state transition matrices




Fig. 3: Details of proposed Spatio-Temporal SSM. The figure illustrates the external and internal operations of the Spatio-Temporal Forward & Backward SSM. Here, the backward scanning method represents Spatio-temporal reversal.

(A, B, and C), and these matrices determine how the model’s hidden state evolves over time. More specifically, as stated in the original paper [15], large Δ represents that the hidden state is ignored and current input is highlighted, and small Δ represents that current input is ignored. In the context of video understanding, Δ enables VideoMamba to selectively focus on relevant aspects of the hidden state for updates, aiding in context-dependent reasoning. More analysis and visualization of Δ is given in Sec. 5.6.

4 VideoMamba

The overall architecture of the proposed VideoMamba is presented in Fig. 2. First, the sampled video clip is transformed into video tokens via a video tokenizer (Sec. 4.1). These video tokens are then added with positional embeddings (Sec. 4.2), which contain positional information. These tokens, along with the class token, form the input for the model. The input tokens pass through L layers of the VideoMamba encoder (Sec. 4.3), where spatio-temporal forward and backward SSM (Sec. 4.4) is applied. Finally, after passing through the last layer, the class token is normalized and fed to a video classification head, which consists of a single MLP layer, to generate the final prediction of the model. To overcome the difficulties in model training due to the relatively small size of video datasets, we initialize our video models using pre-trained image models, capitalizing on the strengths offered by large-scale image datasets. The detailed explanations for each component are provided in the following subsections.

4.1 Video Tokenizer

Video tokenizer maps the sampled video clip $V \in \mathbb{R}^{T \times H \times W \times C}$ into a sequence of video tokens $z = [z_1, z_2, \dots, z_{n_t \times n_h \times n_w}]$, where T , H , W and C represent the temporal length, height, width and channel of the video, respectively and each video token $z_i \in \mathbb{R}^d$ is a d dimensional feature embedding. First, we divide the video clip into non-overlapping tubelets with size of $s_t \times s_h \times s_w$, resulting in $n_t \cdot n_h \cdot n_w$ number of tubelets, where $n_t = \lfloor \frac{T}{s_t} \rfloor$, $n_h = \lfloor \frac{H}{s_h} \rfloor$, $n_w = \lfloor \frac{W}{s_w} \rfloor$.

Then, we utilize 3D convolutional layer to extract video tokens from each tubelet. The video tokenizer is initialized from pretrained image model, by inflating the 2D convolutional layer to the 3D convolutional layer by expanding the weight tensor to the temporal axis and averaging.

4.2 Positional Embedding

Positional embedding was first introduced in Transformer [43], and is widely used in transformer-based vision models [2, 9]. However, positional embedding is not employed in SSMs [15, 16, 35], since the inherent recurrent nature of SSM negates the necessity for positional information. Even though, considering the characteristics of video, the use of positional embedding offers the advantage of supplementing the model with additional spatio-temporal information for each token. Therefore, as an SSM-based video understanding model, to explore the effect of using positional embedding, we consider several options for positional embedding: **(1) not using** any positional embedding, **(2) Sinusoidal** positional embedding and **(3) learnable** positional embedding. For options (2) and (3), the positional embedding $P \in \mathbb{R}^{n_t \cdot n_h \cdot n_w \times d}$ is added to input tokens z .

When using learnable positional embedding, we can initialize the learnable positional embedding P using learned positional embedding from the pretrained image model $P_{image} \in \mathbb{R}^{n_h \cdot n_w \times d}$, which can be regarded as case when $n_t = 1$. We consider several initialization methods: (3-1) expanding the learned positional embedding P_{image} to the temporal axis by replicating it n_t times, (3-2) interpolating in spatial dimension, which is interpolating the learned positional embedding $P \in \mathbb{R}^{n_h \cdot n_w \times d}$ to $\mathbb{R}^{(n_h \cdot n_w \times n_t) \times d}$, (3-3) interpolating in embedding dimension, which is interpolating the learned positional embedding P_{image} to $P \in \mathbb{R}^{n_h \cdot n_w \times (d \times n_t)}$ and reshaping, and (3-4) random initialization.

4.3 VideoMamba Encoder Block

After the pre-processing steps mentioned in Sec. 4.1 and Sec. 4.2, we obtain $n_t \cdot n_h \cdot n_w$ video tokens. The video tokens are prepended with a class token, and fed into the encoder of our proposed VideoMamba, which comprises L encoder blocks. The architectural design of VideoMamba’s encoder block adopts the design of previous works [15, 32, 56], incorporating layer normalization, 1D convolution, and SSM block. For the SSM block, we propose spatio-temporal forward and backward SSM, which is described in the next section.

4.4 Spatio-Temporal Forward and Backward SSM

To effectively handle the non-sequential spatial information in VideoMamba encoder block, we process $n_t \cdot n_h \cdot n_w$ tokens through both forward and backward directions. Since temporal information is inherently ordered, The forward direction can be straightforwardly defined as flattening $n_t \cdot n_h \cdot n_w$ tokens. On the other hand, in the backward direction, the question arises: should we reverse temporal information as we do with spatial information, or not? The combination of

sequential and non-sequential spatio-temporal information complicates the determination of the appropriate reversing approach. With this consideration, we explore three distinct methods for defining the backward scanning direction of the spatio-temporal token sequence, as outlined in Fig. 2-(b).

Spatio-temporal reversal. The first approach is a spatio-temporal reversal, which is equivalent to reversing the order of all flattened $n_t \cdot n_h \cdot n_w$ tokens, compared to $n_t \cdot n_h \cdot n_w$ tokens in the forward direction. Figure 3 illustrates this process. This approach has the advantage of preserving the overall order of input between forward and backward in a similar way to images, akin to handling the video by concatenating n_t frames in a columnar direction, creating a vertically long image.

Spatial reversal. The second method is a spatial reversal, which does not flip all tokens but only reverses each $n_h \cdot n_w$ set of tokens, maintaining the order along the temporal axis. This approach preserves the temporal order of data across both forward and backward paths, giving the model a clear temporal flow in the given video.

Temporal reversal. The last method is a temporal reversal, maintaining the order among the $n_h \cdot n_w$ tokens while reversing their temporal sequence only. This approach can enrich the understanding of the model for the temporal dynamics, by providing inverse event progression without altering the spatial integrity of the frames.

5 Experiments

5.1 Experimental Setup

Datasets. For action recognition tasks, we adopt Kinetics-400 (K400) [25], Something-Something V2 (SSV2) [14], and HMDB51 (HMDB) [28] datasets. The Kinetics-400 dataset contains ~ 240 k training videos and 20k validation videos from 400 human action classes. Something-Something V2 dataset consists of 168.9k training videos and 24.7k validation videos over 174 classes. HMDB51 dataset is relatively smaller than the Kinetics and the SSV2 datasets, containing around 9.5k training videos and 3.5k validation videos from 51 classes. In all experiments, the models are trained on the training videos and then evaluated on the validation videos from the above datasets.

Implementation details. By default, we set the number of blocks L to 24. To align with the model sizes of VideoSwin series, we set the hidden state dimension d to 384. For video tokenizer, tubelet size is set to $s_t = 2$, $s_h = s_w = 16$.

We employed the AdamW optimizer for training. A linear warm-up and a cosine decay learning rate scheduler were both used to train the models for a fixed number of epochs, with a batch size of 64. We initialized the backbone of the network with weights pre-trained on ImageNet1K, while randomly initializing the head layers. The initial learning rate is set to $3e-4$. Stronger data

augmentation techniques were employed, including label smoothing [39], RandAugment [8], and random erasing [55]. For inference, we followed the approach described in [40] by using multiple views (crops) of the video. The final prediction score was computed by averaging the scores from each view.

Kinetics. We sampled 16 frames from each video with a temporal stride of 2 and resized them to 224×224 . We used an AdamW optimizer for 30 epochs with a cosine decay learning rate scheduler and 1 epoch of linear warm-up. A batch size of 64 was utilized.

SSV2. Similar to Kinetics, we used 16 frames with a temporal stride of 2 resized to 224×224 . AdamW optimizer was used for training for 35 epochs with a learning rate scheduler and warm-up. Batch size and data augmentation, except for reverse augmentation, were consistent with Kinetics.

HMDB51. We used the same training strategy as Kinetics for HMDB, with 50 epochs for training.

5.2 Analysis of Model’s Dependency on Temporal Consistency

In this experiment, we aim to assess the video model’s understanding of temporal arrangements by reordering the input frames. This investigation is crucial for the initial expansion and analysis of image models into video understanding. For instance, if the model merely treats video as a bundle of images, the results would likely be similar regardless of how the video is reordered. To do this, the original video frames (ex. indexed 1 through 8) were subjected to various reordering strategies to test their effects. Our reordering strategies for this experiment include:

(1) Interleaved reordering: Implementing a pattern that consistently alternates between distant frames, this method was designed to test the model’s robustness to extreme disruptions in narrative flow. For instance, the sequence could follow the order $[1 \rightarrow 8 \rightarrow 2 \rightarrow 7 \rightarrow 3 \rightarrow 6 \rightarrow 4 \rightarrow 5]$. **(2) Pairwise reordering:** In this approach, the frames were reordered in pairs, specifically creating the sequence $[1 \rightarrow 2 \rightarrow 7 \rightarrow 8 \rightarrow 5 \rightarrow 6 \rightarrow 3 \rightarrow 4]$ by grouping frames in twos. This was done to ensure that, in scenarios where the tubelet size is set to 2, feature generation would not encounter significant issues. **(3) Block-wise reordering:** This strategy involved reorganizing the video frames into blocks, each comprising half of consecutive frames. For example, the sequence $[5 \rightarrow 6 \rightarrow 7 \rightarrow 8 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 4]$ is created.

The experimental result on the HMDB dataset is presented in Table 1. The result demonstrates a clear relationship between the complexity of temporal rearrangement and the performance of our model. The model achieves the best performance when processing videos in their original sequential order, showing its reliance on the inherent temporal flow for optimal understanding. As the temporal disruption became more severe, from blockwise reordering to pairwise reordering and then to the most distracting interleaved reordering, the performance of the model is progressively declined. These results show that our model interprets the action in video by reflecting their temporal order.

Table 1: Experimental result of VideoMamba’s dependency on temporal consistency. We report Top-1 accuracy of ImageNet pretrained model on HMDB dataset.

Reordering strategy	HMDB
Interleaved	51.3
Pairwise	53.5
Block-wise	56.5
Sequential	58.9

Table 2: Different backward scanning methods in spatio-temporal forward and backward SSM. We report the Top-1 accuracy on SSV2 and HMDB, using ImageNet pretrained model.

Backward scanning	SSV2	HMDB
Spatial reversal	61.9	43.3
Temporal reversal	63.3	52.9
Spatio-Temporal reversal	64.7	55.2

Table 3: Positional embedding and initialization method from ImageNet pretrained model on SSV2 and HMDB, as stated in Section 4.2.

Pos. Emb.	Initialize	SSV2	HMDB
None	-	63.2	48.7
Sinusoid	-	63.3	47.5
Learned	Random	63.4	47.9
	Interpolate-s	63.6	49.4
	Interpolate-d	63.6	51.5
	Expanding	63.7	58.9

Table 4: Adding regularization on HMDB dataset. Random augmentation [8] and label smoothing [39] are progressively added on K400 initialization.

	HMDB
ImageNet init.	53.6
K400 init.	66.6
+ Random aug. [8]	67.6
+ Label smoothing [39]	68.6

5.3 Spatio-Temporal Forward and Backward SSM

As mentioned in Section 4.4, we compare three different methods for spatio-temporal backward scanning: (1) spatial reversal, (2) temporal reversal, and (3) spatio-temporal reversal. Table 2 shows the experimental result on SSV2 and HMDB datasets. We did not use an additional class token in this experiment, since the position of the class token can affect the performance. The results show that spatio-temporal reversal is the most effective method for backward SSM, emphasizing the importance of the complementary relationship of token orders in forward and backward SSM. In contrast, spatial reversal is shown to be less advantageous, as it keeps the relative positions of most tokens the same in both the forward and backward paths, impeding the model from fully benefiting from the advantages of bidirectional scanning. Therefore, we select spatio-temporal reversal as our backward scanning method.

5.4 Ablational Study

Positional embedding. In this section, we investigate the impact of different positional embedding strategies (Sec. 4.2) and their initialization methods on the SSV2 and HMDB datasets. As shown in Tab. 3, Omitting positional embeddings leads to lower performance, highlighting their importance. Using sinusoidal

Table 5: Ablation study on the number of input frames on SSV2 and HMDB dataset, fine-tuning from ImageNet pretrained model.

Frames	GFLOPs	SSV2	HMDB
8	17.1	61.0	52.7
16	34.2	63.7	58.9
32	68.5	64.2	59.3

Table 6: Ablation study on embedding dimension of model on K400, SSV2, HMDB dataset. We used ImageNet pretrained model, and Top-1 accuracy is reported.

Dim.	GFLOPs	K400	SSV2	HMDB
192	8.8	69.5	54.6	56.5
384	34.2	76.2	63.7	68.6

embeddings provides a slight improvement, but using learnable positional embedding with appropriate initialization yields better performance. Notably, the initialization by temporal expanding method stands out, achieving the highest accuracy across both datasets, as shown in our table. This shows the effectiveness of appropriately initialized positional embeddings from the image model in enhancing the model’s ability to process video content, giving the model additional spatiotemporal information.

Adding regularization. As a pioneering work in video recognition model based on Mamba, we conducted experiments to discover training recipes for efficient learning of the model. We sequentially evaluated the impact of various regularization techniques on training on HMDB in Tab. 4. Using Kinetics-400 initialization considerably increased the accuracy, illustrating the benefits of domain-specific pretraining. Using random augmentation and label smoothing also lead to a slight improvement, each resulting in an improvement of 1%. This progression demonstrates regularization methods as random augment and label smoothing are effective in training the VideoMamba model on small datasets as well.

Number of frames. We conducted a comparative analysis of the model’s performance across varying numbers of frames, focusing on the impact of frame count on accuracy in Tab. 5. It was observed that the model delivered its most robust performance with the longest input sequence of 32 frames. This outcome underscores the model’s ability to effectively process long-range data, achieving superior performance with linear computational complexity. This efficiency showcases the model’s adeptness in handling extensive temporal information, a significant advantage over the quadratic complexity typically associated with self-attention mechanisms.

Embedding dim We perform an ablation study on the embedding dimension of the model on K400, SSV2, and HMDB datasets. We compared two different embedding dimension sizes, 192 and 384, and GFLOPs and Top-1 accuracy are reported in Tab. 6. The results clearly show that even in very lightweight scenarios, the model performs well, while a larger embedding dimension can further enhance the understanding ability of VideoMamba.

Table 7: Comparison with previous work on HMDB51. ‡ denotes results from [29] and † is reproduced number for fair comparison. Magnitudes are Mega (10^6) for Param. The subscript denotes the trained epoch of the model. “N/A” indicates the numbers are not available for us.

Method	Backbone	Pretrain	Frames	Param	Top-1
I3D [6]	Inception	IN-1K	30	25.0	49.8
SpeedNet [3]	S3D-G	K400	64	9.0	48.8
VTHCL [52]	SlowOnly-R50	K400	32	32.0	67.9
MemDPC [19]	R-2D3D	K400	40	32.0	54.5
CVRL [36]	SlowOnly-R50	K400	32	32.0	49.2
VideoSwin-T† [33]	Swin-T	IN-1K	32	27.9	54.4
VideoSwin-T† [33]	Swin-T	K400	32	27.9	69.9
VideoSwin-S‡ [33]	Swin-T	IN-1K	32	54.0	58.1
VideoMAE _{4800e} [40]	ViT-B	-	16	87.0	62.6
VideoMAE _{4800e} [40]	ViT-B	K400	16	87.0	73.3
S4ND-ConvNeXt-3D‡ [35]	ConvNeXt	IN-1K	30	29.0	55.2
Mamba-ND‡ [29]	Mamba	IN-1K	32	36.0	59.0
VideoMamba	Mamba	IN-1K	16	26.3	58.9
VideoMamba	Mamba	IN-1K	32	26.8	59.3
VideoMamba	Mamba	K400	16	26.3	68.6
VideoMamba	Mamba	K400	32	26.8	75.7

5.5 Comparison to Various Video Models

HMDB51 The performance of VideoMamba on the HMDB51 dataset, provided in Table 7, demonstrates its superiority not only against traditional transformer models like VideoSwin but also against SSM models such as Mamba-ND and S4ND-ConvNeXt-3D. When pretrained on ImageNet-1K, VideoMamba exhibits decent performance with a Top-1 accuracy of 58.9% for 16 frames and 59.3% for 32 frames. Furthermore, when we pretrain VideoMamba on K400, a dataset more aligned with video content, its performance notably increases to a Top-1 accuracy of 68.6% for 16 frames, outperforming all of the compared models.

Compared to VideoSwin-T, a conventional transformer model, VideoMamba utilizes significantly fewer FLOPs and fewer parameters, surpassing VideoSwin-T by 4.9%. Furthermore, we also outperform concurrent mamba-based architecture [29] and the previous SSM-based method [35], still with fewer parameters.

Something-Something V2 Table 8 shows that our VideoMamba achieves high performance with reduced computational needs on Something-Something V2. It outperforms TimeSformer, which has a significantly larger number of parameters. While VideoSwin-T requires 88 GFLOPs, VideoMamba efficiently runs on just 34 GFLOPs for 16 frames and 68 GFLOPs for 32 frames, also having fewer parameters (26.4M and 26.8M, respectively). It achieves Top-1 accuracies of 63.7% and 64.2% for 16 and 32 frames, respectively, outperforming VideoSwin-T by 7% in terms of Top-1 accuracy. This highlights VideoMamba’s superior efficiency and its ability to handle datasets that require detailed interpretation of spatial and temporal dynamics with fewer resources.

Table 8: Comparison with previous work on Something-Something V2. “Views” indicates temporal clip \times spatial crop and \dagger is reproduced number. Magnitudes are Mega (10^6) for Param. The subscript denotes the trained epoch of the model.

Method	Backbone	Pretrain	Frames	View	GFLOPs	Param	Top-1	Top-5
SlowFast [11]	ResNet101	K400	8+32	1 \times 3	106	53.3	63.1	87.6
TSM-RGB [31]	ResNet50	K400	16	2 \times 3	62	42.9	63.4	88.2
VideoSwin-T \dagger [33]	Swin-T	IN-1K	32	1 \times 3	88	27.8	52.3	81.9
VideoSwin-T \dagger [33]	Swin-T	K400	32	1 \times 3	88	27.8	57.2	85.7
TimeSformer [4]	ViT-B	IN-21K	8	1 \times 3	196	121.4	59.5	N/A
TimeSformer [4]	ViT-L	IN-21K	64	1 \times 3	5549	430.0	62.4	N/A
ViViT FE [2]	ViT-L	IN-21K+K400	32	4 \times 3	995	N/A	65.9	89.9
VideoMAE _{2400e} [40]	ViT-S	-	16	2 \times 3	57	22.0	66.8	90.3
VideoMamba	Mamba	IN-1K	16	2 \times 3	34	26.3	63.7	87.8
VideoMamba	Mamba	IN-1K	32	2 \times 3	68	26.8	64.2	88.0

Table 9: Comparison with previous work on Kinetics-400. The subscript denotes the trained epoch of the model.

Method	Backbone	Pretrain	Frames	View	GFLOPs	Param	Top-1	Top-5
R(2+1)D [42]	R(2+1)D	-	32	10 \times 1	75	61.8	72.0	90.0
I3D [6]	Inception	IN-1K	32	10 \times 3	108	25.0	72.1	90.3
NL I3D [48]	ResNet101	IN-1K	32	10 \times 3	359	62.0	77.7	93.3
SlowFast [11]	R101+NL	-	16+64	10 \times 3	234	59.9	79.8	93.9
MViT-S [10]	MViT-S	-	16	5 \times 1	33	26.1	76.0	92.1
MViT-B, 16 \times 4 [10]	MViT-B	-	16	5 \times 1	71	36.6	78.4	93.5
VideoSwin-T [33]	Swin-T	IN-1K	32	4 \times 3	88	28.2	78.8	93.6
VideoMAE _{1600e} [40]	ViT-S	-	16	5 \times 3	57	22.0	79.0	93.8
VideoMamba	Mamba	IN-1K	16	5 \times 3	34	26.4	76.1	92.5
VideoMamba	Mamba	IN-1K	32	5 \times 3	68	26.8	77.7	93.3

Kinetics-400 Table 9 shows that VideoMamba also demonstrates remarkable efficiency on the Kinetics-400 dataset, offering competitive performance with a smaller computational resource. In addition to these efficiencies, VideoMamba achieves Top-1 accuracies of 76.1% and 77.7% for 16 and 32 frames, respectively, demonstrating its capability to process video data effectively with less resource consumption compared to VideoSwin-T and other models of similar size.

5.6 Analysis of Delta

In this section, we explore the significance of Δ in VideoMamba, highlighting their role in emphasizing the informative features within videos. We specifically examine how Δ evolves to focus on essential spatio-temporal details over less relevant background noise. Our visual analysis, presented in Fig. 4, demonstrates the selective nature of Δ in the Spatio-Temporal SSM of VideoMamba. Initially, the model focuses on the overall scene with high Δ values, using the hidden state’s context to distinguish important features. As the layer gets deeper, VideoMamba shifts its focus towards dynamic elements, such as motion, by adjusting Δ values

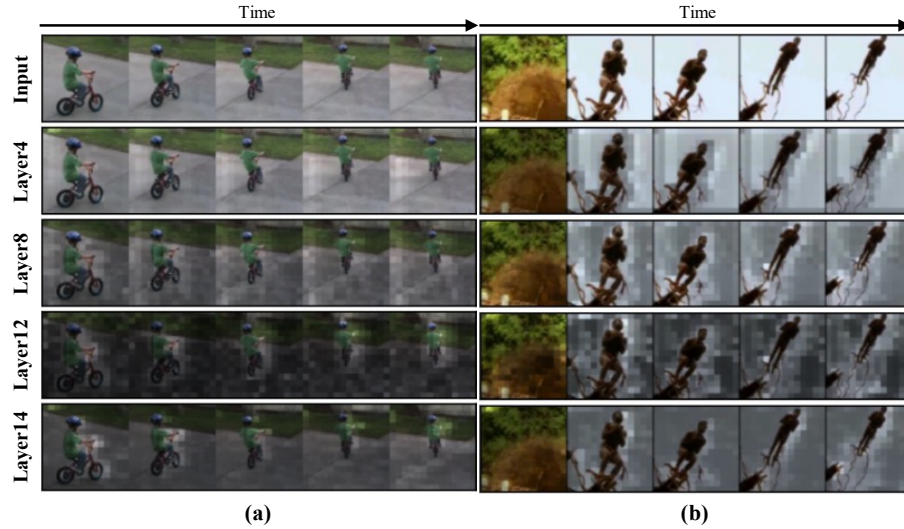


Fig. 4: Delta visualization. Delta plays a crucial role in VideoMamba’s context-dependent reasoning by enabling selective emphasis on important aspects when updating the hidden state. The GT label of (a) is "ride bike," and label of (b) is "dive".

to highlight areas of significant change. For example, increased Δ values around a moving child in Fig. 4-(a) display the model’s ability to focus on motion and complex features, rather than static background elements. Similarly, in Fig. 4-(b), Δ also prioritizes the diving individual while disregarding the background of the initial frame. This shows how VideoMamba effectively captures long-range dependency with contextual awareness.

6 Conclusion

This study has introduced VideoMamba, an innovative model that marks a significant advancement in video analysis by adapting the pure Mamba architecture to address the complex requirements of video content. Through its utilization of a Spatio-Temporal Selective SSM mechanism, VideoMamba efficiently processes the intricate interplay of spatial and temporal information, achieving a notable balance between computational efficiency and analytical precision. Our extensive evaluations demonstrate VideoMamba’s superior performance across various datasets, showcasing its ability to outperform existing models with its adept handling of long-range dependencies and complex video dynamics. VideoMamba not only sets new standards in the field but also lays the groundwork for future research, promising to drive significant progress in video recognition and analysis.

Acknowledgements This work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No.2020-0-00153, Penetration Security Testing of ML Model Vulnerabilities and Defense).

References

1. Akbari, H., Yuan, L., Qian, R., Chuang, W.H., Chang, S.F., Cui, Y., Gong, B.: Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems* **34**, 24206–24221 (2021)
2. Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6836–6846 (2021)
3. Benaim, S., Ephrat, A., Lang, O., Mosseri, I., Freeman, W.T., Rubinstein, M., Irani, M., Dekel, T.: Speednet: Learning the speediness in videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9922–9931 (2020)
4. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: *ICML*. vol. 2, p. 4 (2021)
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *European conference on computer vision*. pp. 213–229. Springer (2020)
6. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 6299–6308 (2017)
7. Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12299–12310 (2021)
8. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*. pp. 702–703 (2020)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations* (2020)
10. Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., Feichtenhofer, C.: Multiscale vision transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6824–6835 (2021)
11. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6202–6211 (2019)
12. Feichtenhofer, C., Li, Y., He, K., et al.: Masked autoencoders as spatiotemporal learners. *Advances in neural information processing systems* **35**, 35946–35958 (2022)
13. Gavriluyk, K., Sanford, R., Javan, M., Snoek, C.G.: Actor-transformers for group activity recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 839–848 (2020)

14. Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., et al.: The "something something" video database for learning and evaluating visual common sense. In: Proceedings of the IEEE international conference on computer vision. pp. 5842–5850 (2017)
15. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752 (2023)
16. Gu, A., Goel, K., Re, C.: Efficiently modeling long sequences with structured state spaces. In: International Conference on Learning Representations (2021)
17. Gu, A., Johnson, I., Goel, K., Saab, K., Dao, T., Rudra, A., Ré, C.: Combining recurrent, convolutional, and continuous-time models with linear state space layers. *Advances in neural information processing systems* **34**, 572–585 (2021)
18. Gupta, A., Gu, A., Berant, J.: Diagonal state spaces are as effective as structured state spaces. *Advances in Neural Information Processing Systems* **35**, 22982–22994 (2022)
19. Han, T., Xie, W., Zisserman, A.: Memory-augmented dense predictive coding for video representation learning. In: European conference on computer vision. pp. 312–329. Springer (2020)
20. Hara, K., Kataoka, H., Satoh, Y.: Learning spatio-temporal features with 3d residual networks for action recognition. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 3154–3160 (2017)
21. Hu, V.T., Baumann, S.A., Gui, M., Grebenkova, O., Ma, P., Fischer, J.S., Ommer, B.: Zigma: A dit-style zigzag mamba diffusion model. arXiv preprint arXiv:2403.13802 (2024)
22. Islam, M.M., Bertasius, G.: Long movie clip classification with state-space video models. In: European Conference on Computer Vision. pp. 87–104. Springer (2022)
23. Ji, S., Xu, W., Yang, M., Yu, K.: 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence* **35**(1), 221–231 (2012)
24. Kalman, R.E.: A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* (1960)
25. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)
26. Kondratyuk, D., Yuan, L., Li, Y., Zhang, L., Tan, M., Brown, M., Gong, B.: Movinets: Mobile video networks for efficient video recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16020–16030 (2021)
27. Kozlov, A., Andronov, V., Gritsenko, Y.: Lightweight network architecture for real-time action recognition. In: Proceedings of the 35th Annual ACM Symposium on Applied Computing. pp. 2074–2080 (2020)
28. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: Hmdb: a large video database for human motion recognition. In: 2011 International conference on computer vision. pp. 2556–2563. IEEE (2011)
29. Li, S., Singh, H., Grover, A.: Mamba-nd: Selective state space modeling for multi-dimensional data. arXiv preprint arXiv:2402.05892 (2024)
30. Liang, D., Zhou, X., Wang, X., Zhu, X., Xu, W., Zou, Z., Ye, X., Bai, X.: Point-mamba: A simple state space model for point cloud analysis. arXiv preprint arXiv:2402.10739 (2024)

31. Lin, J., Gan, C., Han, S.: Tsm: Temporal shift module for efficient video understanding. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7083–7093 (2019)
32. Liu, Y., Tian, Y., Zhao, Y., Yu, H., Xie, L., Wang, Y., Ye, Q., Liu, Y.: Vmamba: Visual state space model. arXiv preprint arXiv:2401.10166 (2024)
33. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3202–3211 (2022)
34. Neimark, D., Bar, O., Zohar, M., Asselmann, D.: Video transformer network. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3163–3172 (2021)
35. Nguyen, E., Goel, K., Gu, A., Downs, G., Shah, P., Dao, T., Baccus, S., Ré, C.: S4nd: Modeling images and videos as multidimensional signals with state spaces. *Advances in neural information processing systems* **35**, 2846–2861 (2022)
36. Qian, R., Meng, T., Gong, B., Yang, M.H., Wang, H., Belongie, S., Cui, Y.: Spatiotemporal contrastive video representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6964–6974 (2021)
37. Qiu, Z., Yao, T., Mei, T.: Learning spatio-temporal representation with pseudo-3d residual networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5533–5541 (2017)
38. Smith, J.T., Warrington, A., Linderman, S.W.: Simplified state space layers for sequence modeling. arXiv preprint arXiv:2208.04933 (2022)
39. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2818–2826 (2016)
40. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems* **35**, 10078–10093 (2022)
41. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 4489–4497 (2015)
42. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 6450–6459 (2018)
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
44. Wang, J., Zhu, W., Wang, P., Yu, X., Liu, L., Omar, M., Hamid, R.: Selective structured state-spaces for long-form video understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6387–6397 (2023)
45. Wang, L., Huang, B., Zhao, Z., Tong, Z., He, Y., Wang, Y., Wang, Y., Qiao, Y.: Videomae v2: Scaling video masked autoencoders with dual masking (2023)
46. Wang, N., Zhou, W., Wang, J., Li, H.: Transformer meets tracker: Exploiting temporal context for robust visual tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1571–1580 (2021)
47. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 568–578 (2021)

48. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7794–7803 (2018)
49. Wu, C.Y., Feichtenhofer, C., Fan, H., He, K., Krahenbuhl, P., Girshick, R.: Long-term feature banks for detailed video understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 284–293 (2019)
50. Wu, C.Y., Krahenbuhl, P.: Towards long-form video understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1884–1894 (2021)
51. Xie, S., Sun, C., Huang, J., Tu, Z., Murphy, K.: Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In: Proceedings of the European conference on computer vision (ECCV). pp. 305–321 (2018)
52. Yang, C., Xu, Y., Dai, B., Zhou, B.: Video representation learning with visual tempo consistency. arXiv preprint arXiv:2006.15489 (2020)
53. Yin, M., Yao, Z., Cao, Y., Li, X., Zhang, Z., Lin, S., Hu, H.: Disentangled non-local neural networks. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16. pp. 191–207. Springer (2020)
54. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., et al.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6881–6890 (2021)
55. Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y.: Random erasing data augmentation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 13001–13008 (2020)
56. Zhu, L., Liao, B., Zhang, Q., Wang, X., Liu, W., Wang, X.: Vision mamba: Efficient visual representation learning with bidirectional state space model. arXiv preprint arXiv:2401.09417 (2024)