

# H-MBA: Hierarchical MamBa Adaptation for Multi-Modal Video Understanding in Autonomous Driving

Siran Chen<sup>1,2</sup>, Yuxiao Luo<sup>1,5</sup>, Yue Ma<sup>4</sup>, Yu Qiao<sup>1,3</sup>, Yali Wang<sup>1,3,\*</sup>

<sup>1</sup> Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Science, Beijing, China

<sup>3</sup> Shanghai Artificial Intelligence Laboratory, Shanghai, China

<sup>4</sup> The Hong Kong University of Science and Technology, Hong Kong, China

<sup>5</sup> The Hong Kong Polytechnic University, Hong Kong, China

chensiran17@mails.ucas.ac.cn, luoyuxiao@polyu.connect.hk, mayuefighting@gmail.com,  
{yu.qiao, yl.wang}@siat.ac.cn

## Abstract

With the prevalence of Multimodal Large Language Models (MLLMs), autonomous driving has encountered new opportunities and challenges. In particular, multi-modal video understanding is critical to interactively analyze what will happen in the procedure of autonomous driving. However, videos in such a dynamical scene that often contains complex spatial-temporal movements, which restricts the generalization capacity of the existing MLLMs in this field. To bridge the gap, we propose a novel **Hierarchical MamBa Adaptation (H-MBA)** framework to fit the complicated motion changes in autonomous driving videos. Specifically, our H-MBA consists of two distinct modules, including Context Mamba (C-Mamba) and Query Mamba (Q-Mamba). First, C-Mamba contains various types of structure state space models, which can effectively capture multi-granularity video context for different temporal resolutions. Second, Q-Mamba flexibly transforms the current frame as the learnable query, and attentively selects multi-granularity video context into query. Consequently, it can adaptively integrate all the video contexts of multi-scale temporal resolutions to enhance video understanding. Via a plug-and-play paradigm in MLLMs, our H-MBA shows the remarkable performance on multi-modal video tasks in autonomous driving, e.g., for risk object detection, it outperforms the previous SOTA method with 5.5% mIoU improvement.

## 1 Introduction

With the rapid advancement of modern artificial intelligence technologies, autonomous driving (AD) has made significant progress. Traditional end-to-end autonomous driving relies on precise environmental perception to make safe predictions and planning. Although it can directly generate planning routes or control signals from raw sensor data, this process resembles a black box that excludes human drivers, making them difficult to understand the driving process and interact with the driving system. To address the challenges, numerous studies have introduced datasets and methods (Kim et al. 2018; Deruyttere et al. 2019; Kim et al. 2019; Atakishiyev et al. 2021; Malla et al. 2023; Jin et al. 2023) to

incorporate natural language interaction and interpretability. Equipped with language models (i.e. BERT (Devlin et al. 2018)), driving sensor data could be translated into natural language prompts to reveal how the models understand driving scenes. However, these methods can only respond to predefined questions due to the limitations of language models, failing to work when faced with open-domain questions. Multi-modal large language models (MLLMs) with extensive general knowledge and reasoning capabilities have been regarded as the future direction for such problems (Zhou et al. 2023), offering the potential as AI agent for AD tasks.

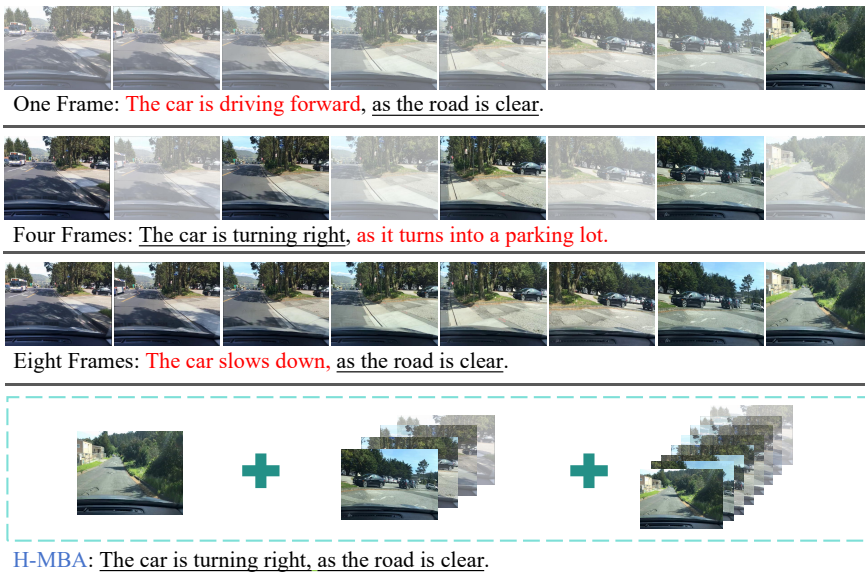
However, the existing MLLM approaches (Ding et al. 2023; Xu et al. 2023; Fu et al. 2024) are limited to model complex temporal characteristics in the first perspective driving videos, e.g., background changes, motions of vehicles and pedestrians, and traffic signals. Hence, they usually fail in multi-modal video understanding in AD tasks, for example, video MLLMs often get low performance in risk object localization (Ding et al. 2023). To solve the above limitation, we introduce a novel Hierarchical MamBa Adaptation (H-MBA) paradigm, which can effectively and efficiently adapt MLLMs with hierarchical mamba modeling for video understanding in complex driving scenarios.

Specifically, H-MBA consists of two distinct modules, Context Mamba (C-Mamba) and Query Mamba (Q-Mamba). C-Mamba can flexibly learn video contexts by developing multi-granularity mamba models in multi-scale temporal branches. In particular, we introduce the low and high temporal resolution branches to model appearance and motion contexts in the driving videos, as shown in Fig 1(a), the low temporal resolution branch captures more obvious motion changes, while the high resolution branch provides additional details. To master the diversified details in these ego-like videos, With both of them, we get more comprehensive understanding of the video. we further leverage three mamba structures for each temporal branch, namely T-Mamba, DST-Mamba and JST-Mamba, to respectively learn Time-only, Divided Space-Time and Joint Space-Time granularity contexts. After obtaining all the contexts, Q-Mamba generates a learnable query from the current frame, and attentively integrates multi-granularity contexts into this query, via performing latent mamba in both temporal

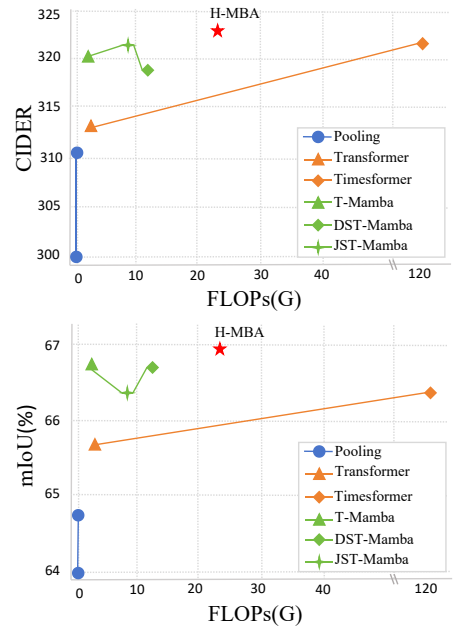
\*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

GT: The car turns right, because there was no oncoming cars.



(a) Multi-Scale Design



(b) Performance vs. GFLOPs

Figure 1: **Motivation.** (a) Previous models fail to give the correct description and justification with only single scale video input, while H-MBA combines multi-scale features and gets the most appropriate answer. (b) Performance of risk localization on Drama (Up/Down: Caption/Detection). Mamba blocks show advantage in the trade-off of performance and computation compared with attention modules, and our H-MBA achieves the best balance.

branches. In this case, the query could adaptively leverage rich video contexts within MLLMs to enhance the temporal understanding of driving videos.

To sum up, there are three contributions in our paper. First, we are the pioneer to develop a mamba paradigm for hierarchical video understanding in the driving scenarios. Compared to attention of video modeling (Bertasius, Wang, and Torresani 2021), our distinct mamba design exhibits the preferable performance with high computation efficiency in AD tasks, e.g., our H-MBA achieves higher performance in the caption and detection tasks, but only costs about one-fifth of computation FLOPs compared to Timesformer, as shown in Fig 1(b). Second, our H-MBA is a plug-and-play video adaption module for MLLMs. Taking advantage of novel C-Mamba and Q-Mamba designs, it can flexibly guide MLLMs to learn driving video representations, by adaptive extraction and integration of multi-scale multi-granularity contexts. Finally, we conduct our H-MBA on multi-modal video understanding benchmarks in autonomous driving, including DRAMA (Malla et al. 2023) and BDD-X (Kim et al. 2018). The extensive results show that, our H-MBA achieves the state-of-the-art performance, e.g., it gets 66.9% mIoU on risk localization, with 5.5% improvement compared with the previous SOTA approach (Malla et al. 2023).

## 2 Related Work

**Multimodal Large Language Models** With the significant success of Large Language Models (LLMs) (Brown et al. 2020; Chowdhery et al. 2023; Devlin et al. 2018; Wei

et al. 2021; Lin et al. 2024), there are increasing research interest to expand the ability of LLMs to deal with multi-modal inputs and tasks. Flamingo (Alayrac et al. 2022) and PaLM-E (Driess et al. 2023) seamlessly fuse image and text on massive image-text pairs, achieving great performance breakthrough among various multi-modal tasks; BLIP-2 (Li et al. 2023a) uses a lightweight Q-Former to bridge the gap of image and text modalities with little parameters and paves the way for LLaVA (Liu et al. 2024a), MiniGPT-4 (Zhu et al. 2023), InstructBLIP (Dai et al. 2024); Then the exploration extends into video domain, VideoChat (Li et al. 2023b), VideoChatGPT (Maaz et al. 2023) and Video-LLaMA (Zhang et al. 2023) further train their temporal module with video instruction-tuning data. These models take multi-modal inputs and have conversations with users in multiple rounds, showing some basic logical reasoning ability. Moreover, several works, e.g. Shikra (Chen et al. 2023), ContextDET (Zang et al. 2023) endow MLLMs with perceptual capabilities, enabling them to output bounding boxes, but the detection is just on image level. While our H-MBA enhances Shikra with the proposed H-Mamba, could not only understand driving video inputs with basic captions, but also provide bounding boxes of the risky object.

**State Space Model Series** The State Space Models (SSMs) have attracted widespread attention due to their high efficiency on modeling the dynamics and dependencies of long-term language sequences, S4 (Gu, Goel, and Ré 2021) designs a structured state-space sequence model especially for long-range dependencies, with only liner com-

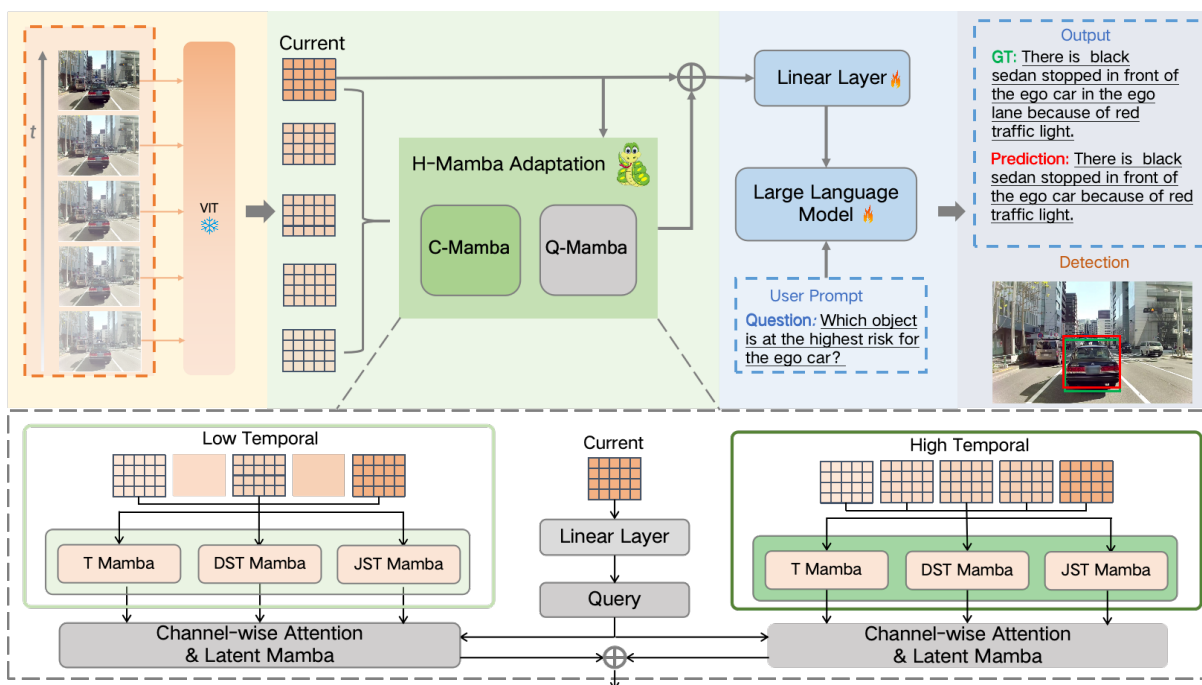


Figure 2: **Pipeline of H-MBA framework.** We design an extra H-Mamba Adaptation block to process video input, the hierarchical refers to high and low temporal resolution and different Mamba-style modules here. After the fuse of Q-Mamba adapter, the multi-scale features are aligned with text query prompt and sent to the LLM to get the final answer.

plexity, and followed by various variants (Smith, Warrington, and Linderman 2022; Dao et al. 2023). Furthermore, Mamba (Gu and Dao 2023) with data-dependent SSM layer and parallel scan selection mechanism (S6) surpasses transformers in long sequence tasks. For the vision domain, Vision Mamba (Zhu et al. 2024) and VMamba (Liu et al. 2024b) utilize different scan mechanisms for 2D image processing, ViVim (Yang, Xing, and Zhu 2024) and VideoMamba (Li et al. 2024b) expand the scan for 3D videos, these Mamba blocks show advantage in effectiveness and efficiency compared to transformer-based modules.

**MLLMs for Autonomous Driving.** Traditional end-to-end autonomous driving exclude human drivers from involvement in the driving process like black boxes, making it hard for drivers to understand and interact with the system, though some works (Wang et al. 2021; Kim et al. 2018; Deruyttere et al. 2019; Jin et al. 2023; Malla et al. 2023) aim to interpret vehicle status from sensor inputs in natural language, they can only solve predefined questions with poor generality. MLLMs provide new directions to address such problems, in fact, they have been widely applied in various downstream tasks as AI agent (Luo et al. 2023; Liang et al. 2023; Li et al. 2024a; Karabacak and Margetis 2023). DriveLikeHuman (Fu et al. 2024) conducts simulation to prove the potential of LLM to make human-like decisions. DriveGPT4 (Xu et al. 2023) is fine-tuned on large GPT-generated conversations to enable models to respond to user queries. HiLM-D (Ding et al. 2023) leverages high and low resolution branches to let models put more attention on the location of small risky objects. However, they

merely utilize existing temporal processing methods while ignoring the characteristics of driving videos, for instance, HiLM-D uses ST-Adapter (Pan et al. 2022) to process video inputs, though getting higher caption performance, the detection results show a decline than the image inputs. Our method not only allows user queries as prompts to perform various tasks, but also incorporates multi-scale/granularity paradigm, which enables more comprehensive understanding for driving videos and boosts performance on the tasks.

### 3 Method

#### Overview

The framework of our model is presented in Fig 2. We choose Shikra (Chen et al. 2023) pre-trained with large location-aware data as the baseline model, which is able to give the coordinates of the output object, in addition to the basic QA ability. The structure is straightforward and simple, consisting of pre-trained CLIP (Radford et al. 2021) ViT-L/14 as the frozen visual encoder, Vicuna-7/13B as the LLM and a fully connected layer for feature transition and alignment. The original Shikra can only handle image input, it's not enough for the driving scenarios where the understanding of temporal information is very important. While preserving the baseline structure, we expand the input phase to take multi-frame video inputs, then leverage the H-MBA to model the spatio-temporal representation. H-MBA consists of Context Mamba (C-Mamba) to capture multi-granularity video context and Query Mamba (Q-Mamba) to adaptively learn and fuse these features, they will be introduced in detail later. The processed features are added to the

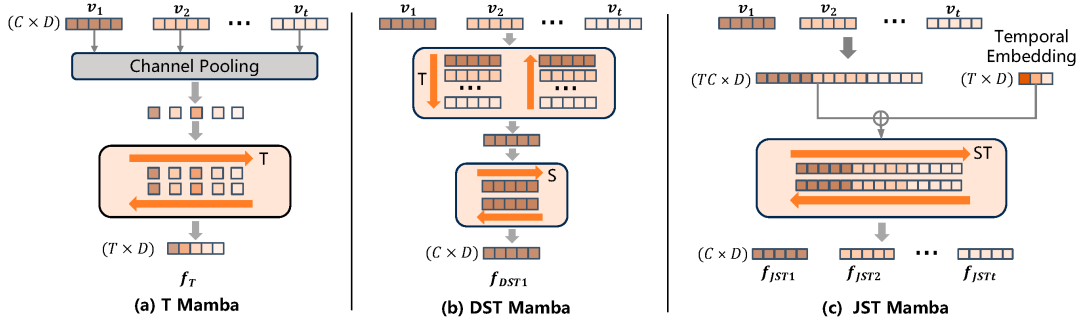


Figure 3: **Illustration of three different space-time Mamba sequences in our C-Mamba.** So we could get multi-granularity video features to fit for various tasks.

original current frame in a residual manner to preserve the pre-training performance. Users could flexibly give query prompts for different tasks, e.g., “What’s the ego car’s action in the video and explain the reasons.”, “Which object is at the highest risk?”. Then the pre-trained LLM receives the multi-modal inputs and gives the corresponding answer.

### C-Mamba: Context Mamba for Multi-Granularity Video Modeling

To better capture the object motion and model background changes in complex driving scenarios, we use a hierarchical C-Mamba to process multi-frame video context. The “hierarchical” has two meanings here. Firstly, considering that speed serves as the “bridge” connecting spatio-temporal changes, the driving video context is highly correlated to the speed of the ego car, for example, the ego car may stop, drive quickly or turn to another lane, the video contexts vary a lot correspondingly. Thus we introduce the paradigm that incorporates both high and low temporal resolutions to fit the diverse speed dynamics, the low branch observes more obvious motion changes, while the high branch involves more detailed information. Secondly, for each temporal resolution branch, we utilize different Mamba state space architectures to learn multi-granularity spatio-temporal video features.

To be more specific, we design three Mamba-based modules, i.e., temporal only sequence Mamba, divided space-time sequence Mamba and joint space-time sequence Mamba, each block learns specific ST patterns. Note that we leverage bidirectional Mamba (Bi-Mamba) to model vision-specific tasks here, which adapts another flattened visual sequences through simultaneous forward and backward SSMs. We label the frame features from frozen visual encoder as  $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_t \in \mathbb{R}^{C \times D}$ ,  $t$  is the frame number,  $C$  is the number of patch channel per frame, and  $D$  is the feature dimension. The three distinct structures could be seen in Fig 3, and introduced below.

**T Mamba.** For temporal only sequence Mamba, it first pools the patch dimension of each frame feature to get the overall frame representation, then the  $t$  frame features are sent to the bidirectional mamba block to learn the temporal relation, this module is suitable for modeling the overall changes of the background:

$$\mathbf{f}_T = \text{Mamba}(\text{pooling}(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_t)) \quad (1)$$

**DST Mamba.** For the divided space-time sequence Mamba (DST), we first separately process the temporal sequence for each patch channel  $\mathbf{v}_i^c$  with Mamba block, so the patches could involve motion change information of the local part, labeled as  $\mathbf{v}_T$ . Then the processed patch sequences of the  $i_{th}$  frame  $\mathbf{v}_{T_i}$  are sent to spatial Mamba block and get the divided space-time features:

$$\mathbf{f}_{DST} = \text{Mamba}(\mathbf{v}_T) \quad (2)$$

$$\mathbf{v}_T^c = \text{Mamba}(\mathbf{v}_1^c, \mathbf{v}_2^c, \dots, \mathbf{v}_t^c) \quad (3)$$

**JST Mamba.** For the joint space-time sequence Mamba (JST), thanks to the  $O(n)$  computational complexity of Mamba algorithm, we are able to jointly process all the patch sequences, which is computationally extravagant for transformer-based method. The  $t$  frame patch feature sequences are first concatenated, we add a learnable temporal embedding (TE) for each frame to distinguish temporal differences. The  $t \times C$  sequences could be sent to Mamba to jointly learn the global spatio-temporal relation.

$$\mathbf{f}_{JST} = \text{Mamba}(\text{Concat}(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_t) + \text{TE}) \quad (4)$$

### Q-Mamba: Query Mamba for Multi-Scale Video Context Adaptation

After obtaining the multi-granularity features, another important question is how to fuse these extra knowledge to the pre-trained MLLM. We have three kinds of designs here. The first straightforward and simple way is to directly add all the weighted learned features to the original visual features, we name it directly add (DA) adapter as the baseline method. Then, motivated by the inception (Szegedy et al. 2015) network which processes features with different convolution kernels and concatenate them in parallel, we also use such paradigm and concat all the features in the feature dimension  $D$  and get  $\mathbf{f}_{cat} \in \mathbb{R}^{C \times (n * D)}$ ,  $n$  is the number of different granularity features. After that, we use a linear-layer to transform the features form  $n * D$  to  $D$ , and this format is called inception concat (IC) adapter.

$$\mathbf{F}_{IC} = \text{FC}(\text{Concat}(\mathbf{f}_1, \dots, \mathbf{f}_n)) \quad (5)$$

Furthermore, we design a Query Mamba (Q-Mamba) adapter. Considering that the visual encoder is frozen during fine-tuning, the original frame visual feature from the

Method	MLLM	B4 ↑	M ↑	R ↑	C ↑	S ↑	mIoU ↑
SAT (Xu et al. 2015)	✗	0.531	0.386	0.694	3.583	0.544	-
ResNet-101 (He et al. 2016)	✗	-	-	-	-	-	60.0
LCP (Malla et al. 2023)	✗	0.520	0.379	0.688	3.501	0.526	59.7
LCP(with OF) (Malla et al. 2023)	✗	0.547	0.391	0.700	3.724	0.560	61.4
BLIP-2 (Li et al. 2023a)	✓	0.506	0.386	0.651	2.585	0.537	46.3
LLaVA (Li et al. 2024a)	✓	-	-	-	-	-	47.2
Shikra (Chen et al. 2023)	✓	0.513	0.379	0.677	3.002	0.539	64.0
Video-Chat (Li et al. 2023b)	✓	0.368	0.290	0.482	1.530	0.435	-
HiLM-D(Video) (Ding et al. 2023)	✓	-	-	-	-	-	59.2
Video-LLaMA (Zhang et al. 2023)	✓	0.481	0.384	0.640	2.135	0.507	42.8
H-MBA	✓	<b>0.536</b>	<b>0.394</b>	<b>0.697</b>	<b>3.227</b>	<b>0.558</b>	<b>66.9</b>

Table 1: **Performance comparison on DRAMA.** Note that we use the detection results from HiLM-D (Ding et al. 2023), but its caption results are conducted on their extra private labeled data, we report the caption results of the original data. Our H-MBA gets the highest risk detection result, achieving 66.9% mIoU, with 5.5% improvement than previous SOTA, LCP(with OF), OF denotes optical flow.

encoder is aligned with LLM and plays a vital role in the pre-training and inference, in order to preserve the performance of the original model as much as possible, we flexibly transform the current frame feature as query  $Q^c$  for each patch channel, and the multi-scale video context features are treated as keys and values, so each channel-wise query  $Q^c$  could adaptively learn what’s valuable from the multi-level features of the patch. This ensures that the model could learn the most appropriate features for different tasks through a unified framework. Then a latent mamba is used to perform channel interaction and learn an overall understanding. We call this Q-Mamba, the process is perceiver (Jaegle et al. 2021) style, and can be formulated as:

$$F_{Q^c} = \sum_{i=1}^n \text{Mamba}(\text{CrossAttn}(Q^c, f_i^c, f_i^c)) \quad (6)$$

## 4 Experiments

**Datasets and Evaluation metrics.** DRAMA (Malla et al. 2023) is a benchmark evaluating the visual reasoning of driving risks, meanwhile, it provides important object bounding boxes with captions describing their risk from the ego-car perspective. The whole dataset contains 17,785 two-second interactive driving scenarios, while considering spatio-temporal relationships from videos. BDD-X (Kim et al. 2018) is a driving-domain caption dataset, consisting of nearly 7000 videos that are collected from BDD100K, and the videos are manually captioned with vehicle behaviors, such as accelerating, and also accompanied with text justification for the behavior. The videos are divided into around 29,000 clips, and the clips’ length duration ranges from 1s to 40s, so the processing of temporal information is critical. We perform two tasks for the driving videos:(1) language caption, including the DRAMA and BDD-X datasets, the captioning performance is evaluated with standard metrics, such as BLEU-4(B4) , METEORP(M), ROGUE(R), CIDER(C) and SPICE(S). (2) Detection for DRAMA, the performance is evaluated with mean intersection over union (mIoU) for

the prediction of bounding box.

**Implementations details.** Following previous methods (Chen et al. 2023; Li et al. 2023a,b), we use the pre-trained Shikra (Chen et al. 2023) as the baseline, and finetune the model with task specific data. We uniformly sample  $L$  frames for a video, the frames are resized and cropped to  $224 \times 224$  resolution. For DRAMA dataset,  $L = 5$ , and we use the Spot-Captioning training format, which allows a bounding box along with the output caption. And for BDD-X, the videos are longer, we set  $L = 8$  and use the Instruct-tuning format. For other MLLM based methods, unless specifically mentioned, we follow the official training code with the same experimental settings. All the experiments are done with 4 A6000 GPUs, we train the model for 5 epochs with  $2e^{-5}$  learning rate in cosine annealing schedule (Loshchilov and Hutter 2016).

### SOTA Comparison

We conduct experiments on DRAMA and BDD-X datasets and compare the results with both MLLM and non-MLLM based methods. Note that, the non-MLLM models are designed for specific tasks, limited to singular functionalities, the results are gray in the tables. In contrast, the MLLM-based approaches offer greater flexibility, enabling tailored responses based on user questions, all MLLMs are finetuned in same setting with official scripts, and as a unified framework, we only use visual inputs for all the tasks.

The risk localization task on DRAMA requires the output of the object with highest risk for the ego car along with the bounding box, different from traditional vehicle object detection, the model should not only have detection and localization ability, but also have a general understanding of the traffic environment and vehicle motion, which is helpful for risk warning and improving driving safety. As shown in Table 1, for the caption scores of risk objects, our H-MBA obtains the optimal results among MLLM models, moreover, we also achieve the highest detection mIoU score, i.e., 66.9%, with 5.5% improvement than previous SOTA

Method	MLLM	Signal	Description					Justification				
			B4↑	M↑	R↑	C↑	S↑	B4↑	M↑	R↑	C↑	S↑
S2VT (Venugopalan et al. 2015)	✗	✗	0.302	0.275	–	1.798	–	0.063	0.112	–	0.534	–
ADAPT-32 (Jin et al. 2023)	✗	✓	0.346	0.306	0.628	2.475	0.597	0.114	0.152	0.320	1.026	0.256
ADAPT-8 (Jin et al. 2023)	✗	✓	0.329	0.290	0.609	2.257	0.579	0.084	0.134	0.306	0.837	0.235
BLIP-2 (Li et al. 2023a)	✓	✗	0.165	0.231	0.497	1.076	0.481	0.043	0.088	0.218	0.455	0.087
Video-Chat (Li et al. 2023b)	✓	✗	0.145	0.211	0.473	0.947	0.451	0.038	0.080	0.166	0.402	0.073
Video-LLaMA (Zhang et al. 2023)	✓	✗	0.168	0.238	0.509	1.045	0.486	0.045	0.086	0.221	0.464	0.101
Shikra (Chen et al. 2023)	✓	✗	0.244	0.256	0.549	1.349	0.521	0.072	0.125	0.288	0.709	0.209
H-MBA	✓	✗	<b>0.301</b>	<b>0.285</b>	<b>0.601</b>	<b>1.996</b>	<b>0.562</b>	<b>0.088</b>	<b>0.137</b>	<b>0.308</b>	<b>0.949</b>	<b>0.238</b>
ADAPT-32 (Jin et al. 2023)	✗	✓	0.355	0.310	–	2.589	–	0.123	0.163	–	1.161	–
DriveGPT-4 (Xu et al. 2023)	✓	✓	0.300	0.298	–	2.140	–	0.094	0.146	–	1.027	–
H-MBA	✓	✗	0.310	0.296	–	2.188	–	0.098	0.148	–	1.088	–

Table 2: **Performance comparison on BDD-X.** The top part of the table is on full test set, and the bottom part is on 500 randomly sampled cases. We get the SOTA results among MLLM based models, we also have advantage for justification compared to ADAPT-8 even it takes extra signal.

Dataset				DRAMA				BDD-X			
Module	Paradigm	GFlops	Params	Caption		Detection		Description		Justification	
				B4↑	C↑	mIoU↑	Acc↑	B4↑	C↑	B4↑	C↑
Image	None	–	–	0.513	3.002	64.0	70.8	0.244	1.349	0.072	0.709
Average Pool	None	–	–	0.526	3.107	64.7	73.3	0.274	1.776	0.078	0.803
Transformer	Time-only	2.7	13.7M	0.526	3.126	65.6	73.7	0.278	1.809	0.080	0.850
T Mamba	Time-only	2.2	7.7M	0.533	3.207	66.8	75.4	0.284	1.907	0.085	0.857
Timesformer	S-T divided	116.9	71.4M	0.533	3.218	66.3	74.3	0.285	1.898	0.082	0.852
DST Mamba	S-T divided	12.0	14.4M	0.527	3.187	66.6	75.0	0.283	1.883	0.087	0.947
JST Mamba	S-T Joint	8.7	7.7M	0.536	3.221	66.3	74.8	0.298	1.957	0.086	0.928

Table 3: **Ablation for different temporal information processing module.** Acc refers to the proportion that IoU>0.5.

LCP (Malla et al. 2023) even with optical flow. Among MLLM based methods, HiLM-D (Ding et al. 2023), the previous SOTA, which uses ST-Adapter to process temporal information, but gets lower mIoU in video format than its image format, while our H-MBA is both helpful for the caption and detection performance, it shows the effectiveness.

The understanding of BDD-X dataset requires models to describe the ego car’s action and explain the intention. Many previous methods utilize extra control signals, such as driving speed and angle, to help make the prediction. To keep consistent with the former task, we just use 8 frames inputs as a uniform setting. Our H-MBA gets the SOTA results among MLLM based models in Table 2, including BLIP-2 (Li et al. 2023a), Video-Chat (Li et al. 2023b), Video-LLaMA (Zhang et al. 2023) and DriveGPT-4 (Xu et al. 2023) with extra signal. Directly using general video foundation models for AD tasks may not be superior due to domain bias, which also proves the effectiveness of our proposed method. And for the previous non-MLLM SOTA ADAPT (Jin et al. 2023), though the extra signal helps to get higher description scores, our H-MBA still gains better performance for the justification of the action, which means that MLLMs learn good understanding of the driving scenario.

## Ablations

**Choice of Temporal processing module.** In this part, we first ablate which temporal module is suitable for driving tasks with both consideration of performance and computation cost. We compare our designed mamba-series modules

with corresponding mainstream attention temporal modules in Table 3. “Image” means we only use the current frame feature of the video, which is also the Shikra baseline. And “Average Pool” means we mean pool the video features in the frame dimension, the performance has slight improvement but still not satisfying without temporal sequence modeling. Then we compare our T Mamba module with Transformer (Vaswani et al. 2017) block, both of them directly learn the relationships between frames. We can see that, Mamba block has advantages in both performance and Flops cost, in such small scale and lightweight modeling. Next, we consider spatio-temporal divided modeling, corresponding to the Timesformer (Bertasius, Wang, and Torresani 2021) and DST Mamba in the table. Though the spatial interaction further improves the results on these tasks, the extra computation cost of Timesformer is the highest. Compared to Timesformer, our Mamba-based methods seem to have more advantages in risk object detection, which needs to integrate temporal changes while retaining the current positions of the objects in current frame. Finally, for our JST Mamba, we fully leverage the capability of Mamba block to process long sequences, with added temporal embedding to explicitly denote temporal distinctions. It has better performance in caption tasks. Note the corresponding attention-based operation requires excessive computational resources, which would cause Out of Memory (OOM) in our machine. On the whole, the Mamba-based modules achieve higher performance with lower computation cost, while different Mamba granularities have advantages in different tasks, and our H-

Adapter	DRAMA				BDD-X			
	Caption		Detection		Description		Justification	
	B4 ↑	C ↑	mIoU ↑	Acc ↑	B4 ↑	C ↑	B4 ↑	C ↑
Direct Addition (DA)	0.520	3.157	66.6	75.1	0.301	1.908	0.081	0.892
Inception Concat (IC)	0.515	3.027	64.3	72.5	0.246	1.461	0.072	0.712
Q-Mamba	0.536	3.227	66.9	75.6	0.301	1.996	0.088	0.949

Table 4: Ablation for the choice of adapter. Q-Mamba keeps the optimal performance across multiple tasks.

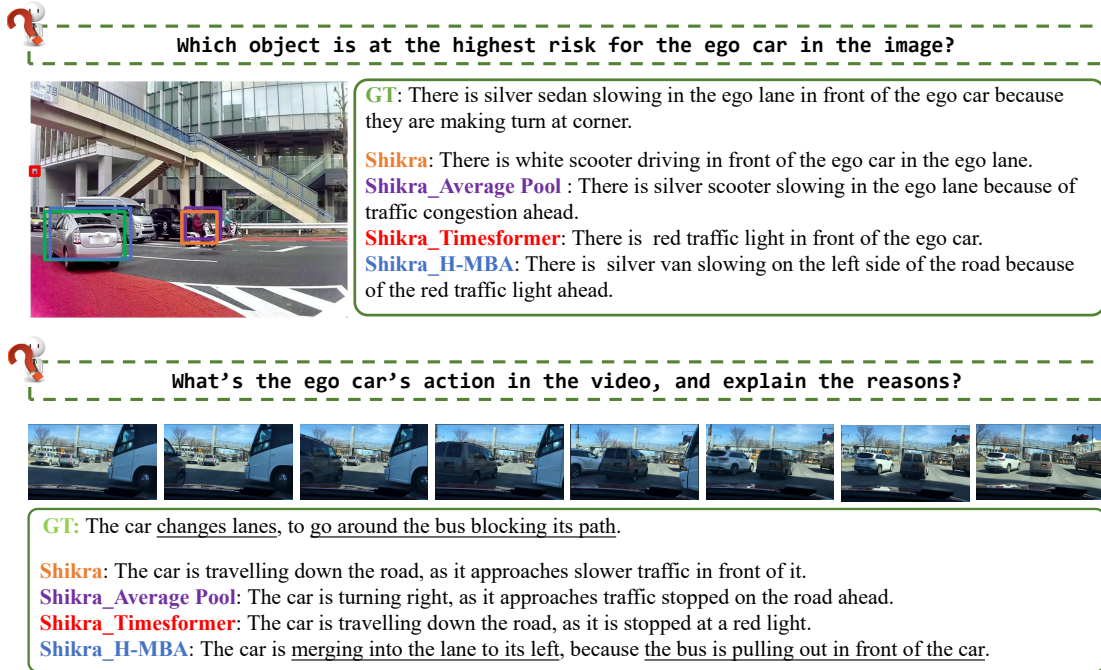


Figure 4: Visualization comparison for the output of different temporal processing modules. For some rare scenarios, such as lane changing caused by a bus occupying the road ahead, our H-MBA can recognize the action and provide reasonable explanations, while others focus on the wrong points.

MBA adaptively assembles them to benefit all the tasks.

**Choice of Adapter.** Next we figure out the optimal setting of three feature adapters. As shown in Table 4, DA aggregates all the features by learnable weights, yet the results indicate that it merely yields a somewhat averaged outcome. The results of IC are relatively low, it may be because the weights are difficult to converge for too many concatenated features. Then for our Q-Mamba, the results have achieved optimal results on each metric, some even higher than single one, which demonstrates the model could effectively learn the most relevant knowledge from the multi-scale features. More ablations could be seen in the Appendix.

**Visualization.** We provide numerous visual examples and compare the results obtained using different time processing modules, as shown in Fig 4. In addition to normal scenes like “the car is driving” and “the car is stopped”, our model can also accurately identify “turning” and “merging lane” behaviors. For some rare scenarios, such as lane changing caused by a bus occupying the road ahead in the example, the model can also recognize the action and provide reason-

able explanations, while others focus on the wrong points, e.g., traffic lights. These examples vividly demonstrate the advantages of hierarchical feature processing. Moreover, we show some real-world accident cases in the appendix, demonstrating the generalization of H-MBA for practical application to enhance driving safety.

## 5 Conclusion

Recent MLLM based approaches have primarily focused on data-labeling and fine-tuning for driving tasks, while we focus on the complicated scene changes in driving videos, and put forward H-MBA for multiple driving tasks. H-MBA utilizes high and low temporal resolution branches, integrated with three different Mamba-based processors to learn features of different granularity. The hierarchical features are fused by the Q-Mamba adapter and sent to LLMs to boost the performance when handling driving videos, with little extra computation. We achieve SOTA results in several tasks such as risk object localization, demonstrating the effectiveness of our approach, and showing the potential for practical application and improving driving safety.

## Acknowledgements

This work was supported by the National Key R&D Program of China(NO.2022ZD0160505), the National Natural Science Foundation of China under Grant(62272450), and the Joint Lab of CAS-HK.

## References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736.
- Atakishiyev, S.; Salameh, M.; Yao, H.; and Goebel, R. 2021. Explainable artificial intelligence for autonomous driving: A comprehensive overview and field guide for future research directions. *arXiv preprint arXiv:2112.11561*.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding? In *ICML*, volume 2, 4.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chen, K.; Zhang, Z.; Zeng, W.; Zhang, R.; Zhu, F.; and Zhao, R. 2023. Shikra: Unleashing Multimodal LLM’s Referential Dialogue Magic. *arXiv preprint arXiv:2306.15195*.
- Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H. W.; Sutton, C.; Gehrmann, S.; et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113.
- Dai, W.; Li, J.; Li, D.; Tiong, A. M. H.; Zhao, J.; Wang, W.; Li, B.; Fung, P. N.; and Hoi, S. 2024. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in Neural Information Processing Systems*, 36.
- Dao, T.; Fu, D. Y.; Saab, K. K.; Thomas, A. W.; Rudra, A.; and Ré, C. 2023. Hungry Hungry Hippos: Towards Language Modeling with State Space Models. In *Proceedings of the 11th International Conference on Learning Representations (ICLR)*.
- Deruyttere, T.; Vandenhende, S.; Grujicic, D.; Van Gool, L.; and Moens, M.-F. 2019. Talk2car: Taking control of your self-driving car. *arXiv preprint arXiv:1909.10838*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ding, X.; Han, J.; Xu, H.; Zhang, W.; and Li, X. 2023. HiLM-D: Towards High-Resolution Understanding in Multimodal Large Language Models for Autonomous Driving. *arXiv preprint arXiv:2309.05186*.
- Driess, D.; Xia, F.; Sajjadi, M. S.; Lynch, C.; Chowdhery, A.; Ichter, B.; Wahid, A.; Tompson, J.; Vuong, Q.; Yu, T.; et al. 2023. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*.
- Fu, D.; Li, X.; Wen, L.; Dou, M.; Cai, P.; Shi, B.; and Qiao, Y. 2024. Drive like a human: Rethinking autonomous driving with large language models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 910–919.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Goel, K.; and Ré, C. 2021. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Jaegle, A.; Gimeno, F.; Brock, A.; Vinyals, O.; Zisserman, A.; and Carreira, J. 2021. Perceiver: General perception with iterative attention. In *International conference on machine learning*, 4651–4664. PMLR.
- Jin, B.; Liu, X.; Zheng, Y.; Li, P.; Zhao, H.; Zhang, T.; Zheng, Y.; Zhou, G.; and Liu, J. 2023. Adapt: Action-aware driving caption transformer. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 7554–7561. IEEE.
- Karabacak, M.; and Margetis, K. 2023. Embracing large language models for medical applications: opportunities and challenges. *Cureus*, 15(5).
- Kim, J.; Misu, T.; Chen, Y.-T.; Tawari, A.; and Canny, J. 2019. Grounding human-to-vehicle advice for self-driving vehicles. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10591–10599.
- Kim, J.; Rohrbach, A.; Darrell, T.; Canny, J.; and Akata, Z. 2018. Textual explanations for self-driving vehicles. In *Proceedings of the European conference on computer vision (ECCV)*, 563–578.
- Li, C.; Wong, C.; Zhang, S.; Usuyama, N.; Liu, H.; Yang, J.; Naumann, T.; Poon, H.; and Gao, J. 2024a. Llavamed: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, K.; He, Y.; Wang, Y.; Li, Y.; Wang, W.; Luo, P.; Wang, Y.; Wang, L.; and Qiao, Y. 2023b. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*.
- Li, K.; Li, X.; Wang, Y.; He, Y.; Wang, Y.; Wang, L.; and Qiao, Y. 2024b. Videomamba: State space model for efficient video understanding. *arXiv preprint arXiv:2403.06977*.
- Liang, J.; Huang, W.; Xia, F.; Xu, P.; Hausman, K.; Ichter, B.; Florence, P.; and Zeng, A. 2023. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 9493–9500. IEEE.

- Lin, H.; Bai, H.; Liu, Z.; Hou, L.; Sun, M.; Song, L.; Wei, Y.; and Sun, Z. 2024. Mope-clip: Structured pruning for efficient vision-language models with module-wise pruning error metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27370–27380.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024a. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024b. VMamba: Visual State Space Model. *arXiv preprint arXiv:2401.10166*.
- Loshchilov, I.; and Hutter, F. 2016. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.
- Luo, R.; Zhao, Z.; Yang, M.; Dong, J.; Qiu, M.; Lu, P.; Wang, T.; and Wei, Z. 2023. Valley: Video assistant with large language model enhanced ability. *arXiv preprint arXiv:2306.07207*.
- Maaz, M.; Rasheed, H.; Khan, S.; and Khan, F. S. 2023. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*.
- Malla, S.; Choi, C.; Dwivedi, I.; Choi, J. H.; and Li, J. 2023. Drama: Joint risk localization and captioning in driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1043–1052.
- Pan, J.; Lin, Z.; Zhu, X.; Shao, J.; and Li, H. 2022. St-adapter: Parameter-efficient image-to-video transfer learning. *Advances in Neural Information Processing Systems*, 35: 26462–26477.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Smith, J. T.; Warrington, A.; and Linderman, S. W. 2022. Simplified state space layers for sequence modeling. *arXiv preprint arXiv:2208.04933*.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Venugopalan, S.; Rohrbach, M.; Donahue, J.; Mooney, R.; Darrell, T.; and Saenko, K. 2015. Sequence to sequence-video to text. In *Proceedings of the IEEE international conference on computer vision*, 4534–4542.
- Wang, H.; Cai, P.; Sun, Y.; Wang, L.; and Liu, M. 2021. Learning interpretable end-to-end vision-based motion planning for autonomous driving with optical flow distillation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 13731–13737. IEEE.
- Wei, J.; Bosma, M.; Zhao, V. Y.; Guu, K.; Yu, A. W.; Lester, B.; Du, N.; Dai, A. M.; and Le, Q. V. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, 2048–2057. PMLR.
- Xu, Z.; Zhang, Y.; Xie, E.; Zhao, Z.; Guo, Y.; Wong, K. K.; Li, Z.; and Zhao, H. 2023. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *arXiv preprint arXiv:2310.01412*.
- Yang, Y.; Xing, Z.; and Zhu, L. 2024. Vivim: a video vision mamba for medical video object segmentation. *arXiv preprint arXiv:2401.14168*.
- Zang, Y.; Li, W.; Han, J.; Zhou, K.; and Loy, C. C. 2023. Contextual object detection with multimodal large language models. *arXiv preprint arXiv:2305.18279*.
- Zhang, H.; Xin, Y.; Li, X.; and Bing, L. 2023. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*.
- Zhou, X.; Liu, M.; Zagar, B. L.; Yurtsever, E.; and Knoll, A. C. 2023. Vision language models in autonomous driving and intelligent transportation systems. *arXiv preprint arXiv:2310.14414*.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*.