

BT-Adapter: Video Conversation is Feasible Without Video Instruction Tuning

Ruyang Liu ^{*1,3}◇ Chen Li ^{*2} Yixiao Ge ² Thomas H. Li ¹ Ying Shan ² Ge Li ^{✉1}

¹School of Electronic and Computer Engineering, Shenzhen Graduate School, Peking University

²Applied Research Center (ARC), Tencent PCG ³Peng Cheng Laboratory

{ruyang@stu, geli@ece, thomas@}.pku.edu.cn {palchenli, yixiaoge, yingsshan}@tencent.com

Abstract

The recent progress in Large Language Models (LLM) has spurred various advancements in image-language conversation agents, while how to build a proficient video-based dialogue system is still under exploration. Considering the extensive scale of LLM and visual backbone, minimal GPU memory is left for facilitating effective temporal modeling, which is crucial for comprehending and providing feedback on videos. To this end, we propose Branching Temporal Adapter (BT-Adapter), a novel method for extending image-language pretrained models into the video domain. Specifically, BT-Adapter serves as a plug-and-use temporal modeling branch alongside the pretrained visual encoder, which is tuned while keeping the backbone frozen. Just pretrained once, BT-Adapter can be seamlessly integrated into all image conversation models using this version of CLIP, enabling video conversations without the need for video instructions. Besides, we develop a unique asymmetric token masking strategy inside the branch with tailor-made training tasks for BT-Adapter, facilitating faster convergence and better results. Thanks to BT-Adapter, we are able to empower existing multimodal dialogue models with strong video understanding capabilities without incurring excessive GPU costs. Without bells and whistles, BT-Adapter achieves (1) state-of-the-art zero-shot results on various video tasks using thousands of fewer GPU hours. (2) better performance than current video chatbots without any video instruction tuning. (3) state-of-the-art results of video chatting using video instruction tuning, outperforming previous SOTAs by a large margin. The code has been available at <https://github.com/farewellthree/BT-Adapter>.

1. Introduction

Since the past year, Large Language Model (LLM) chatbots [25, 34, 37, 45] have emerged as one of the most remarkable advancements within the AI community. Excitingly,

the LLM-driven AI assistants have showcased impressive abilities in comprehension, reasoning, and engaging in conversations. The success of text-only dialogue systems has also sparked the development of image-language conversation agents [7, 9, 19, 50]. These agents combine pretrained image models with LLMs, followed by instruction tuning, to create a fusion of visual and textual understanding for enhanced conversation capabilities.

In comparison to images, videos offer a more comprehensive representation of how humans perceive and interpret the world. Nevertheless, constructing a video-centric dialogue system is a more complex endeavor compared to the image-based one. Firstly, a typically multimodal agent, comprising LLM, a pretrained image encoder (usually CLIP [29]), and additional learnable modules, is already demanding in terms of GPU memory. The incorporation of video dialogue introduces even higher costs, as multiple frames need to be fed as inputs. Secondly, a pretrained encoder as potent and knowledge-enriched as CLIP is currently lacking in the realm of videos, leading to inferior visual understanding. Lastly, gathering video-text instruction data of comparable scale, quality, and diversity to images poses a significant challenge. Rather than building the video agent from scratch, an alternative and promising strategy involves adapting existing pretrained image-centric dialogue models to the video domain [16, 22, 23, 47], which can leverage the abundant knowledge embedded within image models.

Expanding pretrained image models to the video domain is an extensively explored area. The crux of the matter lies in enhancing the capability of 2D models to model temporal dynamics [3, 5, 35]. However, implementing effective temporal modeling within image dialogue models presents notable challenges, primarily due to the trade-off between efficiency and effectiveness in CLIP-based temporal modeling techniques. Methods with strong temporal modeling capabilities, *e.g.*, joint-ST modeling [3, 17, 38, 40] and interpolate-style modeling [24, 26], often necessitate fine-tuning of the entire visual encoder, thereby exacerbating the already significant GPU memory consumption of video

✉Corresponding Author * equal contribution

◇ Work done during internship at ARC Lab, Tencent PCG




Figure 1. The performance overview of our BT-Adapter. On the left, we report zero-shot Recall@1 on MSRVT [39] vs. pretraining GPU hours. On the right, we provide a quantitative comparison of video conversations among existing video dialogue agents.

conversation models. Moreover, full-finetuning would lead to the loss of knowledge encapsulated in pretrained models, resulting in reduced performance for image-based conversations. In contrast, methods that focus on parameter-efficient temporal modeling aim to keep the image encoder frozen and introduce only a few trainable parameters, which are embraced by most video dialogue models. However, these approaches are typically limited temporal modeling and struggle to capture crucial spatiotemporal features.

To tackle the problems in the aforementioned methods, we proposed Branching Temporal Adapter (BT-Adapter), a novel framework to migrate image-text pretrained models to the video domain. As the name implies, BT-Adapter incorporates a branching spatial-temporal module for temporal modeling. Alongside the pretrained image model, BT-Adapter inherits the parameter efficiency advantage from conventional adapters [11] while concurrently achieving effective temporal modeling. Notably, unlike the plug-style adapter, BT-Adapter does not disrupt the forward progression of the pretrained models, thereby safeguarding the integrity of the pretrained multimodal knowledge. After being pretrained with any version of CLIP, BT-Adapter can be seamlessly integrated with all image conversation models using this version of CLIP to activate video conversation, without necessitating video instructions, *e.g.*, openai-CLIP for LLaVa, Eva-CLIP for MiniGPT4 and Instruction-BLIP. Additionally, we have devised a unique asymmetric masking mechanism that exclusively implements tube token masking within the BT-Adapter. Building upon this, we formulate two custom training objectives for the BT-Adapter: Masked Branching Token Alignment (MBTA) and Masked Branching Cross-modal Alignment (MBCA). This approach not only reduces computational demands and accelerates convergence but also yields improved outcomes.

To validate the effectiveness and efficiency of our BT-Adapter, we provide a detailed analysis of various tempo-

ral modeling strategies for video conversation in Sec. 3.2. Furthermore, in Sec. 4, we carry out extensive qualitative and quantitative experiments on BT-Adapter, encompassing both traditional video tasks and video conversations. As depicted in Fig. 1(left), a series of designs ensures that BT-Adapter is highly resource-efficient: our post-pretraining demands just **8 V100(32G)** GPUs in a mere **3 hours**, leading to a reduction in carbon emissions by more than **10240×** and **2687×** compared to CoCa [43] and InternVideo [38] respectively. Building upon this, we still achieve state-of-the-art results in zero-shot video-text retrieval. Regarding video conversation, as shown in Fig. 1(right), we clearly demonstrate that fine-tuned BT-Adapter surpasses the previous state-of-the-art by a significant margin across all benchmarks, and BT-Adapter without instruction tuning has better average performance than fine-tuned SOTAs.

Finally, Our contributions can be summarized as (1) Comprehensive studies on potential temporal modeling methods for video dialogue systems. (2) Branching Temporal Adapter (BT-Adapter) with tail-made masking and training strategies. (3) State-of-the-art results in conventional video tasks and video conversations.

2. Related Work

Temporal Modeling on Image-Language Pretrained Models.

With the wide application of pretrained image-language models [29, 43], how to extend these pretrained multimodal models into the video domain has emerged as a novel yet critical research topic. Existing temporal modeling methods can be broadly categorized into several types. The most straightforward and commonly used one is joint-spatial-temporal modeling [3, 17, 38, 40]. By inputting all video tokens into the image encoder, this approach can effectively model temporal dependencies without other techniques. However, joint-ST modeling necessitates fine-tuning the entire encoder, resulting in significant compu-

tational costs and the degradation of pretrained knowledge. Another representative type is the interpolate-style temporal modeling, including separated-ST modeling [3, 46], message token [24], and ST-Adapter [26]. Inserting temporal modules between the pretrained spatial layers, interpolate-style modeling shares similar cons and pros with joint-ST modeling. Specifically, ST-Adapter is also known as the parameter-efficient temporal module, where the distinctions with our methods are detailed in Sec. 9 in the Appendix. Different from the former two types, concatenate-style temporal modeling [8, 21] attaches temporal modules following the pretrained encoder, which allows for the freezing of the backbone and the preservation of knowledge. Nevertheless, this style offers limited temporal modeling, as it hardly captures crucial low-level spatial-temporal features. In Sec. 3.2, we will experimentally compare the different temporal modeling methods for video dialogue.

In contrast to the preceding types, BT-Adapter adopts a branching temporal modeling structure. Thanks to the design, we circumvent the issues above, enabling proficient temporal modeling, efficient parameter fine-tuning, and multimodal affinity simultaneously. Most similar to our methods, STAN [20] also introduced branch-style temporal modeling. Nonetheless, there are three key distinctions between our work and STAN: (1) Our emphasis is on zero-shot image-to-video transfer and video conversations, areas that have never been explored by STAN. (2) Our performance is notably superior when the backbone is frozen, which is attributed to our unique design like temporal CLS token, zero-initialized temporal projection, and temporal selection. (3) We have introduced the novel training strategy and objectives tailored for the branching temporal structure.

Video Conversation. Recent advancements in multimodal learning have been predominantly propelled by the fusion of visual models with LLM. Yuan et al. [44] initially demonstrated the potential of combining visual models with LLM. Blip-2 [15] further proposed Q-former that maps visual tokens into the text embedding space. Subsequently, several methods presented visual instruction tuning [7, 19, 50] for image-LLM to enable visual conversation. Most closely related to our topic, there have been several developments in the field of video-centric dialogue models over the past few months [16, 22, 23, 47]. Generally, these models consist of the visual encoder, LLM, and temporal module, tuned with video instruction data to realize video conversation. However, in pursuit of efficient training, the temporal modules utilized by these models, such as temporal position embedding [47] and temporal pooling [22, 23], often exhibit limited temporal modeling capabilities. In contrast, our approach maintains parameter-efficient training while simultaneously delivering effective temporal modeling. Thanks to the benchmark raised by Maaz et al. [23], we can quantitatively illustrate the superiority of our methods in Sec. 4.

Moreover, we introduce a novel setting of zero-shot conversation by integrating pretrained temporal modules with image dialogue models, demonstrating the feasibility of video conversation without any video instruction tuning.

3. Methodology

In this section, we elaborate on how to efficiently enhance pretrained image-language models (*e.g.*, CLIP [29] and LLaVA [19]) with temporal modeling capabilities while preserving the pretrained knowledge. The main framework is depicted in Fig. 2. We will introduce the model architectures of BT-Adapter in Sec. 3.1, the exploration of temporal modeling in video dialogue models in Sec. 3.2, and the training strategy and objectives in Sec. 3.3.

3.1. Branching Temporal Adapter

As the pioneer of contrastive image-text pretraining, CLIP [29] has founded widespread application in various domains. Meanwhile, fundamental image dialogue models like BLIP-2 [15] and LLaVA [19] all employ CLIP as visual encoder. Hence, without losing the generalizability, we focus on the adaption of CLIP.

Model Architecture. To enable video input, within CLIP, we treat each frame as an individual image. Given a video with T frames, we divide each frame into N non-overlapping patches, represented as $V = \{V_i\}_{i=1}^T$ and $V_i = \{v_{i,j}\}_{j=0}^N$, where $v_{i,0}$ denotes the [CLS] token. Then, tokens in each frame are added with spatial position and fed independently into the l^{th} CLIP layers:

$$\begin{aligned} V_i'^{(l-1)} &= \text{S_attn}(\text{LN}(V_i^{(l-1)})) + V_i^{(l-1)}, \\ V_i^{(l)} &= \text{FFN}(\text{LN}(V_i'^{(l-1)})) + V_i'^{(l-1)}, \end{aligned} \quad (1)$$

where S_attn, LN and FFN denote spatial attention, layer normalization and feed-forward network.

In contrast to the plug-style adapter that consists of multiple independent modules, BT-Adapter is a continuous network operating as a branch alongside the main backbone. Inside BT-Adapter, we adopt divided space-time attention [3]. Given all video tokens, we first gather patch tokens in the same position across different frames, obtaining $\hat{V} = \{\hat{V}_j\}_{j=1}^N$ and $\hat{V}_j = \{\hat{v}_{i,j}\}_{i=1}^T$. Then, tokens in each position are fed individually into the l^{th} temporal layers:

$$\hat{V}_j'^{(l-1)} = W_t \cdot \text{T_attn}(\text{LN}(\hat{V}_j^{(l-1)})) + \hat{V}_j^{(l-1)}, \quad (2)$$

where T_attn is the self-attention operated on the T dimension. W_t is an additional zero-initialized linear projection, which keeps the training stable during the adaption. Then, the position-wise tokens are reshaped into frame-wise tokens $\{\hat{V}_i'\}_{i=1}^T$, and fed into the spatial layer, following the same procedure as CLIP layer in Eq. 1.

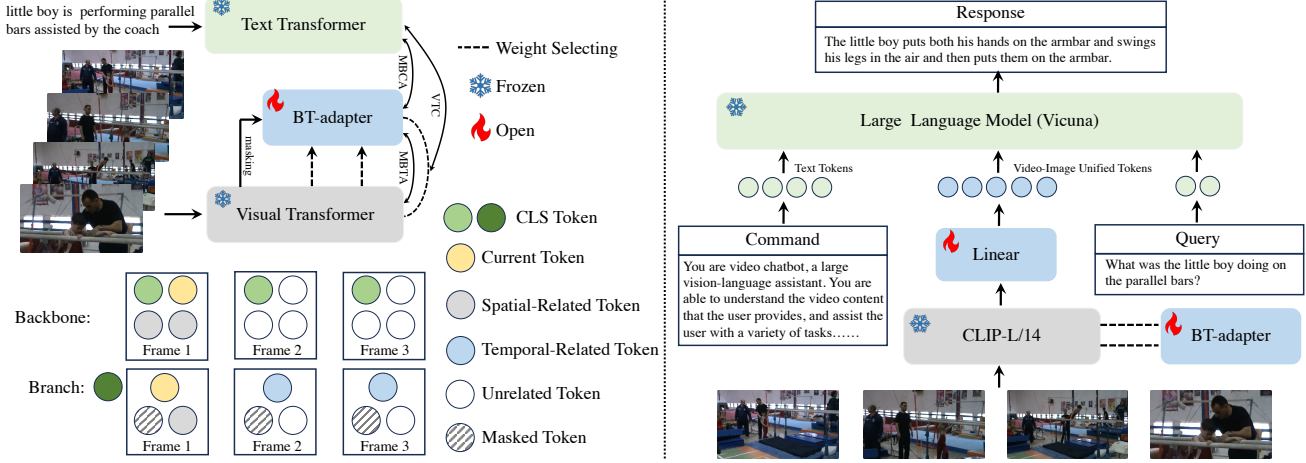


Figure 2. The overview of our model. The left side shows the model architecture and the data flow during pretraining. The right side depicts the pipeline of video conversation.

Backbone-Branch Interaction. Unlike traditional adapters, which are added to every pretrained layer, our branching adapter has significantly fewer layers compared to the main backbone. Assuming we have K layers in the branch, BT-Adapter takes the output of the last $K+1$ CLIP layers as input, each layer on both sides corresponding pairwise. To construct the input of the first BT-Adapter layer from the k^{th} CLIP layer, we first develop a new learnable video [CLS] token $\hat{v}_{0,0}^{(0)}$ to represent the entire video. Then, we concatenate $\hat{v}_{0,0}^{(0)}$ with $V^{(k)}$ and update the patch embeddings with frozen spatial positional embeddings and learnable temporal positional embeddings:

$$\hat{v}_{i,j}^{(0)} = v_{i,j}^{(k)} + P_i^t + P_j^s, \quad (3)$$

where P^s is the spatial position embedding shared with CLIP while P^t is the temporal position embedding. For any other layer of BT-Adapter, we construct its input $\hat{V}^{(l)}$ from the previous branching layer and the CLIP layer at the same level with weight selecting as follows:

$$\hat{v}_{i,j}^{(l)} = \text{Sigmoid}(W_b) \cdot \hat{v}_{i,j}^{(l-1)} + (1 - \text{Sigmoid}(W_b)) \cdot v_{i,j}^{(k+l-1)}, \quad (4)$$

where W_b is the learnable selective weight with zero initialization. At the final stage, the output of the branch and CLIP are combined for the out-of-the-box representation:

$$v = W_{v_proj} (\text{LN}(\text{Sigmoid}(W_b) \cdot \hat{v}_{0,0}^{(-1)} + (1 - \text{Sigmoid}(W_b)) \cdot \frac{1}{T} \sum_{i=1}^T v_{i,0}^{-1})), \quad (5)$$

where W_{v_proj} is the frozen weight in CLIP projecting the visual embedding into joint visual-text feature space. With CLIP kept frozen, the backbone encodes low-level spatial patterns and high-level aligned features, while the branch

focuses on modeling temporal dependencies. In this way, we leverage strong temporal modeling capacities while preserving the pretrained knowledge intact.

3.2. Temporal Modeling for Video Conversation

In this section, we conduct an empirical study to explore potential temporal modeling strategies for video dialogue models, demonstrating the advantages of our approach. We argue that an ideal temporal modeling approach for video-LLM models should meet several criteria: it should be parameter-efficient (P-E), as image dialogue models are already quite large; it should be multimodal-friendly (M-F), preserving as much multimodal alignment knowledge as possible; and it should be temporal-sensitive (T-S), delivering strong performance in time-sensitive scenarios. To measure the ‘‘multimodal-friendly’’ and ‘‘temporal-sensitive’’, we employ the ‘‘Correctness of Information’’ and ‘‘Temporal Understanding’’ metrics from the VideoChatGPT benchmark [23] respectively. ‘‘parameter-efficient’’ is simply decided by whether the CLIP can be frozen. Next, we will explain how we implement various temporal modeling methods for video conversation.

Spatial-Temporal Pooling. Following Maaz et al. [23], we do not integrate any module in CLIP and encode each frame independently. Then, all patch tokens $V \in \mathbf{R}^{T \times N}$ are pooled along the time (T) and spatial (N) dimensions, resulting in a total of $T+N$ tokens. These tokens are subsequently input into the LLM, which has much fewer tokens compared to joint or separate temporal modeling.

Joint-ST modeling. Following Xue et al. [40], we incorporate spatiotemporal positional embeddings to 2D patches and feed all $T * N$ video tokens into CLIP and LLM simultaneously. In this way, tokens in any position or frame can attend to each other, providing a straightforward yet effective method for modeling temporal dependencies. To make training feasible, we utilize the FSDP [48] to facilitate pa-

Table 1. The finetuning results of different temporal modeling on video conversation. * means CLIP is frozen.

Model	P-E	M-F	T-S	Correctness of Information	Temporal Understanding
Baseline*				2.38	1.93
ST Pooling*	✓	✓		2.40	1.98
Joint-ST			✓	1.92	2.11
Separate-ST			✓	2.10	2.13
Separate-ST*	✓			2.29	2.01
BT-Adapter*	✓	✓	✓	2.55	2.26

parameter and gradient sharing between GPUs.

Separate-ST modeling. Unlike joint-ST modeling, separate-ST modeling retains and leverages CLIP’s spatial layer. All $T * N$ patch tokens are fed in LLM. Following Zeng et al. [46], we insert temporal attention before each CLIP layer and add temporal position embeddings to the input. Regarding the backbone, separate-ST modeling offers flexibility, allowing us to experiment with both frozen and unfrozen CLIP layers.

Branch Temporal Adapter. We directly equip CLIP-L/14 with BT-Adapter in Sec. 3.1 to achieve video encoder. In LLaVA, LLM takes the output from the second-to-last layer of the visual encoder as its input. Hence, we also take the second-to-last output from both CLIP and BT-Adapter and combine them with learnable balance weight. Finally, combined patch tokens $V \in \mathbf{R}^{T \times N}$ are pooled along the time (T) dimension, resulting in only N tokens for inputting into LLM, which is the least among all methods.

Results. We replace CLIP-L/14 in LLaVA with the implemented video encoders and then proceed with video instruction tuning. In all methods, we keep LLM frozen while opening the linear projection between the visual encoder and LLM. No pretraining on video-text datasets is included for fair comparison. Results are presented in Table 1. It can be observed that methods with a frozen CLIP tend to be parameter-efficient but perform poorly in terms of temporal modeling. On the other hand, methods with good temporal understanding suffer from high computation costs and knowledge loss. In contrast, our BT-Adapter achieves a balance by being parameter-efficient, multimodal-friendly, and temporal-sensitive simultaneously. The results reveal that BT-Adapter has a clear advantage over other temporal modeling methods for video conversation.

3.3. Pretraining with Asymmetric Masking

As proved by previous studies [38, 40], CLIP-based video encoder can harvest stronger performance on downstream video tasks after being post-pretrained on large-scale video-text data. Hence, we also involved our BT-Adapter with video-text pretraining. However, the expensive computation cost has always been a bottleneck when scaling up the training. Inspired by the recent success of masked modeling in visual-language pretraining [18, 33], we develop a unique asymmetric masking strategy for BT-Adapter. Specifically, we maintain tokens in the frozen CLIP unchanged while

applying a tubular mask to the tokens in BT-Adapter. This mask randomly masks a certain percentage ($\rho\%$) of patch tokens in the same position across different frames. This approach allows us to maximize the retention of pretraining knowledge from CLIP while reducing spatial-temporal redundancy. Thanks to the asymmetric mask, we can maintain a high mask ratio ($\rho \geq 70$) without compromising performance, leading to a reduction of at least half of the computational budget. As a result, with the assistance of frozen backbone and token masking, we can accomplish the costly video-text pretraining in just a few hours. Furthermore, based on asymmetric masking, we have devised two tailor-made training objectives for the branching temporal structure, in addition to the Video-Text Contrastive.

Video-Text Contrastive (VTC). VTC is the most widely-used basic objective for cross-modal alignment. Given the global video feature v in Eq. 5 and global text feature t , we formulate \mathcal{L}_{VTC} as:

$$\mathcal{L}_{nce}(x, y) = -\frac{1}{B} \sum_{m=1}^B \log \frac{\exp(\tau x_m \cdot y_n)}{\sum_{n=1}^B \exp(\tau x_m \cdot y_n)}, \quad (6)$$

$$\mathcal{L}_{VTC} = \mathcal{L}_{nce}(v, t) + \mathcal{L}_{nce}(t, v),$$

where B is the batch size and m, n is the index in a batch and τ is the temperature scale.

Masked Branching Token Alignment (MBTA). Masked video modeling has been proven beneficial for spatial-temporal representation learning [33], but pixel reconstruction in Tong et al. [33] is computationally expensive. In our approach, we have a frozen CLIP and the masked BT-Adapter, where the unmasked CLIP is naturally a teacher for the masked branch. Hence, we can conduct the token alignment in an end-to-end pretraining without extra forward propagation. Specifically, we compute the mean squared error (MSE) between the unmasked tokens in BT-Adapter and the corresponding tokens in CLIP in the last layer, formulating \mathcal{L}_{MBTA} as:

$$\mathcal{L}_{MBTA} = \frac{1}{(1-\rho)NT} \sum_{i=1}^T \sum_{j=1}^N (W_{v,proj} \cdot v_{i,j}^{-1} - W'_{v,proj} \cdot \hat{v}_{i,j}^{-1})^2, \quad (i, j) \notin M, \quad (7)$$

where M is the set of masked tokens, and $W'_{v,proj}$ is a freshly initialized linear projection distinct from $W_{v,proj}$. Utilizing MBTA, we can constrain the masked branch to reconstruct and align the semantic information present in CLIP, leading to better multimodal results.

Masked Branching Cross-Modal Alignment (MBCA). In our model, the backbone and the branch handle different aspects: CLIP encodes static spatial features, while BT-Adapter captures dynamic information. Therefore, we additionally align the branching patch tokens with the text embedding to align the temporal-sensitive information, which

Table 2. The results of video conversation on video-based generative performance benchmarking. FT and ZS mean with and without video instruction tuning respectively.

Method	Temporal Understanding	Correctness of Information	Detail Orientation	Contextual Understanding	Consistency	Mean Score
VideoLLaMA [47]	1.82	1.96	2.18	2.16	1.79	1.98
LLaMA-Adapter [9]	1.98	2.03	2.32	2.30	2.15	2.16
VideoChat [16]	1.94	2.23	2.50	2.53	2.24	2.29
VideoChatGPT [23]	1.98	2.40	2.52	2.62	2.37	2.38
BT-Adapter-LLaVA (ZS)	2.13	2.16	2.46	2.89	2.20	2.46
BT-Adapter-LLaVA (FT)	2.34	2.68	2.69	3.27	2.46	2.69

Table 3. The results of video conversation zero-shot question-answering. FT and ZS mean with and without instruction tuning respectively.

Method	MSVD-QA		MSRVTT-QA		ActivityNet-QA	
	Acc	Score	Acc	Score	Acc	Score
VideoLLaMA	51.6	2.5	29.6	1.8	12.4	1.1
LLaMA-Adapter	54.9	3.1	43.8	2.7	34.2	2.7
VideoChat	56.3	2.8	45.0	2.5	26.5	2.2
VideoChatGPT	64.9	3.3	49.3	2.8	35.2	2.7
Ours (ZS)	67.0	3.6	51.2	2.9	46.1	3.2
Ours (FT)	67.5	3.7	57.0	3.2	45.7	3.2

Table 4. The results of zero-shot video conversation on different image-centric dialogue models.

Method	Temporal	Correctness
LLaVA-Vicuna	1.78	2.06
+BT-Adapter	2.13	2.16
MiniGPT4-Vicuna	1.88	2.48
+BT-Adapter	2.56	2.71
InstructBLIP-Vicuna	2.03	2.90
+BT-Adapter	2.38	2.93

is formulated as:

$$\hat{v} = \frac{1}{(1-\rho)NT} \sum_{i=1}^T \sum_{j=1}^N W_{v\text{-proj}} \cdot \hat{v}_{i,j}^{-1}, \quad (8)$$

$$\mathcal{L}_{MBCA} = \mathcal{L}_{nce}(\hat{v}, t) + \mathcal{L}_{nce}(t, \hat{v}), \quad (i, j) \notin M.$$

Different from \mathcal{L}_{VTC} , we utilize the patch tokens for the alignment. This choice is made because the information is highly centralized in the CLS token for CLIP, whereas downstream applications like multimodal conversation or generation typically rely on the patch tokens as input.

4. Experiments

4.1. Experiment Settings

Tasks and Datasets. We have evaluated our BT-Adapter on two main aspects: traditional video tasks and video-centric dialogue. To begin with, we pretrained the BT-Adapter on WebVid-2M [2]. For traditional video tasks, we consider zero-shot text-to-video retrieval and video recognition, covering seven benchmarks: (a)MSR-VTT [39]. (b)DiDeMo [1]. (c)LSMDC [30]. (d)ActivityNet [4]. (e)Kinetic-400 [5]. (f)HMDB-51 [13]. (g)UCF-101 [31]. We used Recall@K (R@K) for text-to-video retrieval and Top-K Accuracy (A@K) for video recognition as the evaluation metrics.

For video dialogue, we conducted evaluations for the zero-shot video conversation (without instruction tuning) and instruction-tuned video conversation. The VideoChatGPT-100K [23] is employed for supervised instruction tuning. To assess the quality of the responses quantitatively, we employ the VideoChatGPT benchmark,

which consists of five metrics for video-based generative performance benchmarking and six metrics for zero-shot question-answer evaluation. Additional details and settings of each dataset can be found in Sec. 6 in the Appendix.

Implementation Details. We employ openai-CLIP-L/14 and EVA-CLIP-G [32] as the video backbone. For all downstream tasks, we use openai-CLIP-L/14. We adopt 4 layers of BT-Adapter in default. During pretraining, the masking ratio ρ is set as 70%. The temperature scale τ is fixed as 0.01 for contrastive loss. The weight of the three losses is 1 : 1 : 1. It takes 3 hours to train one epoch on WebVid on 8 V100-32G GPUs. For video conversation, we implement the instruction tuning based on BT-Adapter-LLaVA-7B. Unless specifically noted, we use BT-Adapter-LLaVa in all video dialogue tasks. InstructBLIP and MiniGPT4 are also included for zero-shot evaluation. All conversation models are based on Vicuna1.0-7B [49]. During instruction tuning, we update the linear projection and BT-Adapter, while keeping the rest architecture frozen. It takes 3 hours to train three epochs on 8 A100 40GB GPUs. More training details are listed in Appendix.

4.2. Quantitative evaluation

Video Dialogue. Thanks to the benchmarks introduced by Maaz et al. [23], we can quantitatively compare the performance of various video conversation models. The results are presented in Table 2 and Table 3, where we compare BT-Adapter with all existing video-centric dialogue models, including VideoLLaMA [47], LLaMA-Adapter [9], VideoChat [16], and VideoChatGPT [23]. It is observed that our zero-shot model, even without any instruction tun-

Table 5. The zero-shot results of text-to-video retrieval on MSR-VTT, DiDemo, LSMDC, and ActivityNet.

Method	Pretraining Scale	MSR-VTT			DiDeMo			LSMDC			ActivityNet		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
<i>Non-CLIP models</i>													
Frozen [2]	5M	24.7	46.9	57.2	21.1	46.0	56.2	-	-	-	-	-	-
Clover [12]	5M	26.4	49.5	60.0	29.5	55.2	66.3	17.4	29.2	38.2	-	-	-
OmniVL [36]	14M	34.6	58.4	66.6	33.3	58.7	68.5	-	-	-	-	-	-
HiTeA [42]	5M	29.9	54.2	62.9	36.1	60.1	70.3	15.5	31.1	39.8	-	-	-
Singularity [14]	17M	34.0	56.7	66.7	37.1	61.7	69.9	-	-	-	30.6	55.6	66.9
VideoCoCa [41]	100M	34.3	57.8	67.0	-	-	-	-	-	-	34.5	63.2	76.6
<i>CLIP-L/14</i>													
CLIP [29]	-	35.4	58.8	68.1	30.3	54.9	65.4	17.0	31.8	40.3	28.8	57.6	71.8
ImageBind [10]	-	36.8	61.8	70.0	-	-	-	-	-	-	-	-	-
InternVideo [38]	12.8M	40.7	-	-	31.5	-	-	17.6	-	-	30.7	-	-
TVTSv2 [46]	8.5M	38.2	62.4	73.2	34.6	61.9	71.5	17.3	32.5	41.4	-	-	-
UMT-L [17]	5M	33.3	58.1	66.7	34.0	60.4	68.7	20.0	37.2	43.7	31.9	60.2	72.0
BT-Adapter	2M	40.9	64.7	73.5	35.6	61.9	72.6	19.5	35.9	45.0	37.0	66.7	78.9

Table 6. Ablation study on the structures of BT-Adapter. We report the results on zero-shot R@1 of MSRVT and DiDemo and zero-shot video conversation.

Model	MSRVTT	DiDemo	Temporal
CLIP (baseline)	35.4	30.3	1.78
+4 layer separate-ST	35.7	31.0	1.81
+branch modeling	37.4	32.9	1.97
+backbone-branch interaction	38.5	33.9	2.06

Table 7. Ablation study on the training objectives. We report the zero-shot R@1 of retrieval.

MBTA	MBCA	MSRVTT	DiDemo
		38.5	33.9
✓		39.3	34.3
	✓	40.1	34.9
✓	✓	40.9	35.6

ing, outperforms methods that require instruction tuning on average. When fine-tuned using video instruction data, the superiority of our approach becomes even more pronounced. These results underscore the effectiveness of the BT-Adapter as a superior method for temporal modeling in video conversation models compared to existing approaches. Notably, our BT-Adapter exhibits a significantly larger performance margin over other methods on ActivityNet in both Table 3 and 5, highlighting its particular strength in handling long video sequences. Moreover, we have integrated the BT-Adapter with various pretrained image conversation models. As illustrated in Table 4, pretrained BT-Adapter yields consistent advancement on all image-language chatbots without extra instruction tuning, underscoring the broad applicability of our method.

Traditional Video Tasks. The results of zero-shot text-to-video retrieval are presented in Table 5. Compared to the previous SOTAs, BT-Adapter achieves competitive performance across all datasets, obtaining the best results on most metrics. For instance, in comparison to UMT and Singularity, although there is a slight lag in terms of R@1 on DiDeMo and LSMDC respectively, we have surpassed them by more than 5% on MSRVT and ActivityNet. Furthermore, we achieve superior results with significantly fewer pretraining scales and GPU hours than all the mentioned methods. For example, we use 130× fewer GPU hours than

UMT, 560× fewer than TVTSv2, and 2687× fewer than InterVideo. The results of zero-shot action recognition are posted in Table 9 in Appendix.

4.3. Ablation Study

Model Structure. To investigate the components within the BT-Adapter, we initiated our study by exclusively employing VTC for pretraining and subsequently evaluating the zero-shot performance across retrieval tasks and video conversation, as presented in Table 6. Initially, we explored the implementation of separate-ST modeling within the last 4 layers. However, the resulting improvements were marginal. Subsequently, we transitioned the separate-ST network into the branch, where we witnessed notable progress. This shift underscores the efficacy of branching modeling. In the final step, we studied the interaction module, *i.e.*, multi-level selective combination between backbone and branch. This addition led to further enhancements across all benchmark datasets. More fine-grained ablations on model components can be found in Sec. 10 in appendix.

Training Objectives. We investigated the influence of two innovative training objectives in Table 7, where all experiments were conducted using BT-Adapter with a mask rate of 70%. Our observations reveal that both MBTA and MBCA yield improvements in downstream results, with MBCA demonstrating more substantial progress. This out-




Figure 3. Ablation study on the number of BT-Adapter layers.




Figure 4. Ablation study on the masking rate during pretraining.

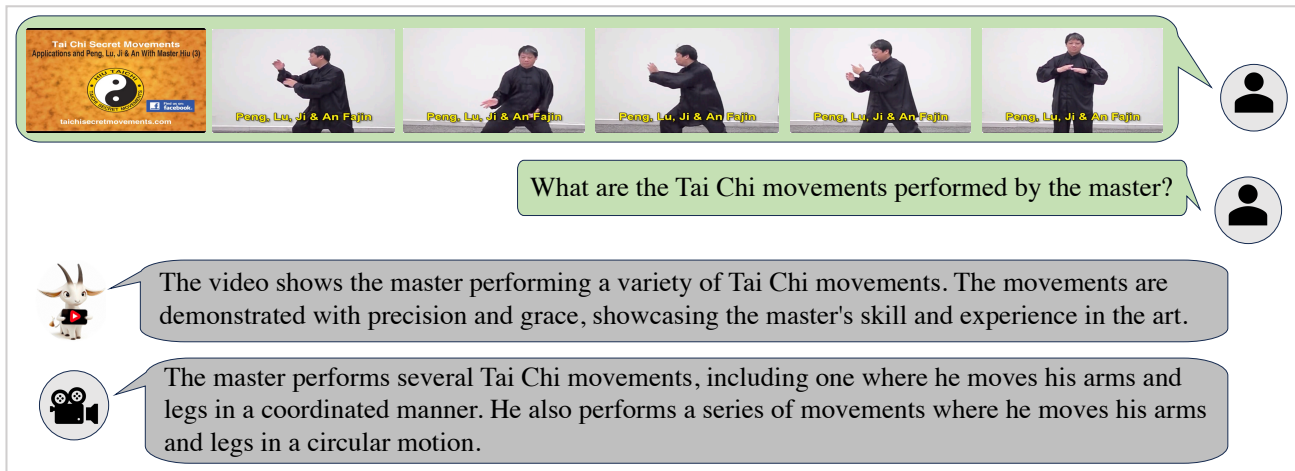


Figure 5. A qualitative result of video conversation. We present the answers from VideoChatGPT [23] (upper) and our BT-Adapter-LLaVA (down).

come suggests that individually aligning temporal modules and textual outputs effectively mitigates the limitations inherent in pretrained image-text models. Moreover, the combination of both objectives results in further advancements.

Number of Branching Layers and Masking Rate. The number of branching layers and the masking rate in our model inherently involve a trade-off between computational resources and performance. To determine the optimal settings, we conducted ablation experiments. Firstly, in Fig. 3, we present the results of zero-shot retrieval and conversation as we vary the number of layers. Notably, we observed that the performance improvement plateaus at around 4-6 layers. Secondly, in Fig. 4, we provide results for zero-shot retrieval and the corresponding convergence GPU time across different masking rates. Our findings indicate that the results remain stable when the masking rate is set at or below 70%, while higher masking rates significantly lower the training time. Therefore, we have chosen to maintain a masking rate of 70%.

4.4. Qualitative Results

In Fig. 5, we present a qualitative example of video conversation. Unlike the general description from VideoChatGPT, our model provides an informative and accurate response

to a video-related question, highlighting the effectiveness of the BT-Adapter in video comprehension. Additional examples of video dialogues covering various aspects can be found in the Appendix.

5. Conclusion

This paper presents a novel approach to achieve parameter-efficient yet effective image-to-video adaptation and video conversation. Our proposed solution, named Branching Temporal Adapter (BT-Adapter), is a branching separate-ST network for temporal modeling. Building upon the BT-Adapter, we introduce an asymmetric masking technique along with two novel training objectives. Extensive experiments demonstrate the superiority of our model over other temporal modeling methods for video conversation. By seamlessly integrating the BT-Adapter with any pretrained image conversation model, we achieve video dialogues without necessitating manual video instruction tuning. With significantly lower computation costs, we attain state-of-the-art results across two zero-shot video tasks and video conversations.

Acknowledgements. This work was supported by National Natural Science Foundation of China (No. 62172021).

References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, pages 5803–5812, 2017. **6, 1**
- [2] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1728–1738, 2021. **6, 7, 1**
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, page 4, 2021. **1, 2, 3**
- [4] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. **6, 1**
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. **1, 6**
- [6] MMAAction2 Contributors. Openmmlab’s next generation video understanding toolbox and benchmark. <https://github.com/open-mmlab/mmaaction2>, 2020. **1**
- [7] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. **1, 3**
- [8] Han Fang, Pengfei Xiong, Luhui Xu, and Yu Chen. Clip2video: Mastering video-text retrieval via image clip. *arXiv preprint arXiv:2106.11097*, 2021. **3**
- [9] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xianguyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023. **1, 6**
- [10] Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190, 2023. **7**
- [11] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019. **2**
- [12] Jingjia Huang, Yanan Li, Jiashi Feng, Xinglong Wu, Xiaoshuai Sun, and Rongrong Ji. Clover: Towards a unified video-language alignment and fusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14856–14866, 2023. **7**
- [13] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 International conference on computer vision*, pages 2556–2563. IEEE, 2011. **6, 1**
- [14] Jie Lei, Tamara L Berg, and Mohit Bansal. Revealing single frame bias for video-and-language learning. *arXiv preprint arXiv:2206.03428*, 2022. **7**
- [15] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. **3**
- [16] KunChang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. **1, 3, 6**
- [17] Kunchang Li, Yali Wang, Yizhuo Li, Yi Wang, Yanan He, Limin Wang, and Yu Qiao. Unmasked teacher: Towards training-efficient video foundation models. *arXiv preprint arXiv:2303.16058*, 2023. **1, 2, 7**
- [18] Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. Scaling language-image pre-training via masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23390–23400, 2023. **5**
- [19] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. **1, 3**
- [20] Ruyang Liu, Jingjia Huang, Ge Li, Jiashi Feng, Xinglong Wu, and Thomas H Li. Revisiting temporal modeling for clip-based image-to-video knowledge transferring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6555–6564, 2023. **3, 2**
- [21] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022. **3**
- [22] Ruiyu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Minghui Qiu, Pengcheng Lu, Tao Wang, and Zhongyu Wei. Valley: Video assistant with large language model enhanced ability. *arXiv preprint arXiv:2306.07207*, 2023. **1, 3**
- [23] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023. **1, 3, 4, 6, 8**
- [24] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding language-image pretrained models for general video recognition. In *European Conference on Computer Vision*, pages 1–18. Springer, 2022. **1, 3, 2**
- [25] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022. **1**
- [26] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning. *Advances in Neural Information Processing Systems*, 35:26462–26477, 2022. **1, 3, 2**

- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [1](#)
- [28] Yijun Qian, Lijun Yu, Wenhe Liu, and Alexander G Hauptmann. Rethinking zero-shot action recognition: Learning from latent atomic actions. In *European Conference on Computer Vision*, pages 104–120. Springer, 2022. [2](#)
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [1](#), [2](#), [3](#), [7](#)
- [30] Anna Rohrbach, Atousa Torabi, Marcus Rohrbach, Niket Tandon, Christopher Pal, Hugo Larochelle, Aaron Courville, and Bernt Schiele. Movie description. *International Journal of Computer Vision*, 123(1):94–120, 2017. [6](#), [1](#)
- [31] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [6](#), [1](#)
- [32] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023. [6](#)
- [33] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022. [5](#)
- [34] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. [1](#)
- [35] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. [1](#)
- [36] Junke Wang, Dongdong Chen, Zuxuan Wu, Chong Luo, Lu-wei Zhou, Yucheng Zhao, Yujia Xie, Ce Liu, Yu-Gang Jiang, and Lu Yuan. Omnivl: One foundation model for image-language and video-language tasks. *Advances in neural information processing systems*, 35:5696–5710, 2022. [7](#)
- [37] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*, 2022. [1](#)
- [38] Yi Wang, Kunchang Li, Yizhuo Li, Yinan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, et al. Internvideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*, 2022. [1](#), [2](#), [5](#), [7](#)
- [39] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296, 2016. [2](#), [6](#), [1](#)
- [40] Hongwei Xue, Yuchong Sun, Bei Liu, Jianlong Fu, Ruihua Song, Houqiang Li, and Jiebo Luo. Clip-vip: Adapting pre-trained image-text model to video-language alignment. In *The Eleventh International Conference on Learning Representations*, 2022. [1](#), [2](#), [4](#), [5](#)
- [41] Shen Yan, Tao Zhu, Zirui Wang, Yuan Cao, Mi Zhang, Soham Ghosh, Yonghui Wu, and Jiahui Yu. Video-text modeling with zero-shot transfer from contrastive captioners. *arXiv preprint arXiv:2212.04979*, 2022. [7](#)
- [42] Qinghao Ye, Guohai Xu, Ming Yan, Haiyang Xu, Qi Qian, Ji Zhang, and Fei Huang. Hitea: Hierarchical temporal-aware video-language pre-training. *arXiv preprint arXiv:2212.14546*, 2022. [7](#)
- [43] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. [2](#)
- [44] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. [3](#)
- [45] Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414*, 2022. [1](#)
- [46] Ziyun Zeng, Yixiao Ge, Zhan Tong, Xihui Liu, Shu-Tao Xia, and Ying Shan. Tvtv2: Learning out-of-the-box spatiotemporal visual representations at scale. *arXiv preprint arXiv:2305.14173*, 2023. [3](#), [5](#), [7](#), [2](#)
- [47] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023. [1](#), [3](#), [6](#)
- [48] Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, et al. Pytorch fsdp: experiences on scaling fully sharded data parallel. *arXiv preprint arXiv:2304.11277*, 2023. [4](#)
- [49] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena, 2023. [6](#)
- [50] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. [1](#), [3](#)