Video-CoT: A Comprehensive Dataset for Spatiotemporal **Understanding of Videos Based on Chain-of-Thought**

Shuyi Zhang* Institute of Automation, CAS School of Artifcial Intelligence, UCAS Beijing, China zhangshuyi2024@ia.ac.cn

Lingfeng Zhang Shenzhen International GraduateSchool,Tsinghua University Shenzhen, Guangdong, China lfzhang715@gmail.com

Xiaoshuai Hao* Beijing Academy of Artificial Intelligence (BAAI) Beijing, China xshao@baai.ac.cn

Pengwei Wang Beijing Academy of Artificial Intelligence (BAAI) Beijing, China pwei@baai.ac.cn

Hongxuan Ma[†] Institute of Automation, CAS School of Artifcial Intelligence, UCAS Beijing, China hongxuan.ma@ia.ac.cn

Abstract

1 2

3

4 5

8

9

10

11

13

14

15

16 17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

49

Video content comprehension is essential for various applications, ranging from video analysis to interactive systems. Despite advancements in large-scale vision-language models (VLMs), these models often struggle to capture the nuanced, spatiotemporal details essential for thorough video analysis. To address this gap, we introduce *Video-CoT*, a groundbreaking dataset designed to enhance spatiotemporal understanding using Chain-of-Thought (CoT) methodologies. Video-CoT contains 192,000 fine-grained spatiotemporal question-answer pairs and 23,000 high-quality CoTannotated samples, providing a solid foundation for evaluating spatiotemporal understanding in video comprehension. Additionally, we provide a comprehensive benchmark for assessing these tasks, with each task featuring 750 images and tailored evaluation metrics. Our extensive experiments reveal that current VLMs face significant challenges in achieving satisfactory performance, highlighting the difficulties of effective spatiotemporal understanding. Overall, the Video-CoT dataset and benchmark open new avenues for research in multimedia understanding and support future innovations in intelligent systems requiring advanced video analysis capabilities. By making these resources publicly available, we aim to

57 https://doi.org/10.1145/nnnnnnnnnnnn

58

Beijing, China

Yingbo Tang Institute of Automation, CAS School of Artifcial Intelligence, UCAS Beijing, China

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97

98

99

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

116

tangyingbo2020@ia.ac.cn

Zhongyuan Wang Beijing Academy of Artificial Intelligence (BAAI) Beijing, China zyuan@baai.ac.cn

Shanghang Zhang[†] State Key Laboratory of Multimedia Information Processing, School of Computer Science, Peking University, shanghang@pku.edu.cn

encourage further exploration in this critical area. Project website: https://video-cot.github.io/.

CCS Concepts

• Computing methodologies → Visual content-based indexing and retrieval; Video summarization.

Keywords

Video Understanding, Spatiotemporal Alignment, Multimedia

1 Introduction

Multimodal content understanding and reasoning are essential for advancing universal artificial intelligence, particularly in the video modality [9, 11, 30, 36, 44], where effective comprehension is crucial for applications ranging from automated analysis to interactive systems [10, 12, 37, 44]. However, existing video datasets typically focus on content summarization or specific temporal and spatial dimensions, which limits their effectiveness for fine-grained spatiotemporal reasoning. This highlights the need for a task-rich spatiotemporal dataset that facilitates comprehensive video understanding.

Recent advancements in rule-based reinforcement learning have significantly improved the reasoning capabilities of large language models (LLMs), such as GPT-40 [17], Deepseek-R1 [8], and Kimi-1.5 [35]. This progress has sparked interest in Chain-of-Thought (CoT) data [40], which incorporates intermediate reasoning steps for problem-solving [24, 50], aiming to enhance complex reasoning abilities by capturing the logical chains of human thought. Although some recent works have generated CoT reasoning data in the video modality [5, 14, 23, 47], they often overlook fine-grained spatiotemporal information-such as event start and end times and

^{*}Both authors contributed equally to this research. 48

[†]Corresponding Authors.

Permission to make digital or hard copies of all or part of this work for personal or 50 classroom use is granted without fee provided that copies are not made or distributed 51 for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the 52 author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or 53 republish, to post on servers or to redistribute to lists, requires prior specific permission 54 and/or a fee. Request permissions from permissions@acm.org.

ACM MM, Dublin, Ireland 55

^{© 2025} Copyright held by the owner/author(s). Publication rights licensed to ACM. 56 ACM ISBN 978-x-xxxx-x/YYYY/MM

ACM MM, October 27-31, 2025, Dublin, Ireland



Figure 1: Overview of the Video-CoT dataset. The Video-CoT dataset encompasses three key dimensions: spatio-temporal localization and captioning, spatio-temporal grounding, and spatio-temporal reference.

the positions of subjects within videos. Therefore, there is a pressing need to develop a video dataset based on CoT methodologies, addressing these gaps and supporting advanced reasoning in video comprehension.

To address this gap, we introduce *Video-CoT*, a groundbreaking dataset designed to enhance spatiotemporal understanding using Chain-of-Thought (CoT) methodologies, one example of which is shown in Fig. 1. Video-CoT contains 192,000 fine-grained spatiotemporal question-answer pairs and 23,000 high-quality CoT-annotated samples, providing a robust resource for training and evaluating VLMs in video comprehension. The dataset is organized into three components: Spatio-Temporal Localization and Captioning, which includes Temporal Video Localization (TVL) and Video Captioning (VC) tasks for temporal localization and content summarization; Spatio-Temporal Grounding, consisting of Spatial Video Grounding (SVG) and Spatio-Temporal Video Grounding (STVG), delivering accurate temporal and spatial coordinates based on event descriptions; and Spatio-Temporal Reference, comprising Spatial Relationship Reference (SRR) and Temporal Video Reference (TVR), addressing relational descriptions of event subjects in time and space. Through these diverse tasks, we aim to enhance the model's reasoning capibilities to spatiotemporal information. Additionally, we provide a comprehensive benchmark for assessing these tasks, each featuring 750 images and tailored evaluation metrics. Extensive experiments reveal that current VLMs face significant challenges in achieving satisfactory performance, underscoring the complex-ities of effective spatiotemporal understanding. In summary, our contributions are as follows:

• Proposal of Video-CoT Dataset. Introduction of a novel dataset designed to enhance spatiotemporal understanding

through Chain-of-Thought (CoT) methodologies, offering a solid foundation for evaluating video comprehension.

- **Comprehensive Benchmark and Evaluation Metrics.** Development of a robust benchmark that includes tailored evaluation metrics for assessing the performance of various tasks within the dataset. This benchmark supports three key dimensions: spatiotemporal localization and captioning, spatiotemporal grounding, and spatiotemporal reference, encompassing six distinct subtasks.
- Extensive Experiments. Conducting thorough experiments that highlight the challenges faced by current visual language models (VLMs) in achieving effective spatiotemporal comprehension, offering critical insights that can drive innovation and advancement in future research.

2 Related Work

Video Datasets Video understanding is a crucial capability for Multimodal Large Language Models (MLLMs), with accurate video datasets serving as their foundation. Early datasets like UCF-101 [31] and Sports-1M [18] initiated action recognition by emphasizing visual feature learning. As the field matured, more complex datasets emerged, including Kinetics-400/700 [19], Something-Something [7], and ActivityNet-Captions [21]. However, existing video datasets often fall short in capturing comprehensive spatiotemporal information. Many focus solely on either temporal or spatial features, lacking effective integration. For example, the Jester dataset [26] emphasizes gesture recognition over time but offers minimal spatial annotations for hand and body positions. Similarly, the Fine-Gym dataset [29] provides temporal annotations for actions in gym videos but lacks detailed spatial context regarding the relationships between athletes and equipment. These gaps underscore the need

S.Zhang, X.Hao and Y.Tang et al.

Video-CoT: A Comprehensive Dataset for Spatiotemporal Understanding of Videos Based on Chain-of-Thought



Video Length Distribution Across 6 tasks

Figure 3: The video length statistic.

Figure 2: The distribution of tasks across three primary categories.

for new datasets that effectively combine both spatial and temporal dimensions to enhance video understanding research.

Chain-of-Thought (CoT) Chain-of-Thought (CoT) prompting [40] has emerged as a pivotal technique for enhancing the complex reasoning capabilities of large language models (LLMs). By explicitly decomposing multi-step problems into intermediate reasoning pathways, CoT significantly improves performance on tasks such as mathematical reasoning [4] and commonsense QA [32]. Subsequent advances focus on three directions: (1) Minimizing demonstration dependency, as seen in Zero-Shot CoT activated by trigger phrases like "Let's think step by step" [20]; (2) Optimizing reasoning robustness through methods like self-consistency (aggregating multiple reasoning paths) [38] and automated demonstration generation [49]; (3) Integrating external tools such as program interpreters [6] and symbolic solvers to address computational limitations. Despite these innovations, critical challenges persist in ensuring reasoning fidelity-generated chains often contain logical inconsistencies [15] and factual hallucinations [25], motivating ongoing research into verifiable reasoning frameworks and errorcorrection mechanisms.

3 Dataset Construction and Analysis

In this section, we provide an in-depth exploration of each spatiotemporal task within our dataset and outline the Chain-of-Thought (CoT) data generation process in detail.

3.1 Task Definition

Our Video-CoT dataset consists of six tasks categorized into three key aspects:

Spatio-Temporal Localization and Captioning connects visual perception with language understanding by detecting specific temporal events in videos and generating meaningful textual descriptions. This task establishes the foundation for identifying and articulating key events within video content. It comprises two subtasks: *Temporal Video Localization (TVL)*, which focuses on pinpointing the precise temporal boundaries of actions or events in untrimmed videos, such as identifying the exact duration of a "basketball dunk" in a sports clip; and *Video Captioning (VC)*, which extends beyond event detection to generate natural language descriptions of video content, enabling more intuitive interaction with video data across various contexts.

Spatio-Temporal Grounding aims to align textual queries with corresponding spatio-temporal regions in videos, integrating spatial and temporal reasoning to enhance multimodal understanding. This task includes two components: *Spatial Video Grounding (SVG)*, which focuses on locating specific spatial regions in video frames that correspond to a textual query, such as identifying a "black SUV car leaving an adult in pink clothes"; and *Spatio-Temporal Video Grounding (STVG)*, which extends SVG by predicting a spatio-temporal sequence of bounding boxes across multiple frames that match a text query, like tracking "a dog chasing a ball". STVG integrates spatial localization with temporal tracking, addressing challenges like object occlusion and appearance variations to ensure robust alignment between textual and visual representations.

Spatio-Temporal Reference focuses on understanding object relationships and temporal dynamics within video content. By modeling spatial and temporal dependencies, this task enhances the interpretation of complex interactions in dynamic scenes. It comprises two key components: *Spatial Relationship Reference (SRR)*, which identifies spatial relationships between objects, such as "above", "behind", or "next to" within video scenes— for example, detecting "a book on a table" during a room scan, crucial for precise scene understanding; and *Temporal Video Reference (TVR)*, which describes video content within specific time intervals. Given a defined segment, TVR generates concise and contextually accurate textual descriptions of events and actions occurring during that interval, requiring a deep understanding of temporal dependencies and event dynamics to produce temporally coherent and meaningful descriptions.

3.2 CoT Dataset Building

Data Collection and Curation for our Video-CoT dataset primarily relies on publicly available datasets. For the tasks of Temporal Video Localization (TVL), Video Captioning (VC), and Temporal Video Reference (TVR), we directly utilize existing datasets (VTimeLLM-stage2 [16], MSR-VTT [43], MSVD [3] and WebVid [2]). In contrast, the tasks of Spatial Relationship Reference (SRR), Spatial Video Grounding (SVG), and Spatio-Temporal Video Grounding (STVG) involve additional processing based on the HCSTVG-V1 [33] and VidSTG [51] datasets, which provide dense annotations for each frame.



Figure 4: Pipeline for Constructing the Video-CoT Dataset.

Our filtering pipeline selects subjects or objects as query targets and ensures their consecutive presence across video frames, discarding objects that appear for less than 2 seconds. We derive the start and end timestamps of objects with continuous presence, along with their pixel locations at initial, final, and intermediate integer-second timestamps based on the video's frame counts and per-frame annotations. By maintaining the spatial relationships between primary event subjects and other objects in the environment, we construct spatiotemporal relational question-answering data, resulting in a total of 192K fine-grained spatial-temporal QA pairs.

The distribution of our dataset across six tasks is illustrated in Fig. 2. Additionally, the dataset demonstrates a balanced distribution across various video length intervals, as shown in Fig. 3. Short videos (0–40 seconds) comprise the majority, providing ample data for tasks that require rapid event detection and description. Simultaneously, the dataset includes a significant number of longer videos exceeding 160 seconds, enabling models to train on complex long-term temporal sequences while maintaining performance. This distribution reflects real-world video content characteristics, encompassing both short-term events and extended sequences, thereby supporting the development of robust models capable of generalizing across diverse video lengths and inference tasks.

CoT Data Generation We construct a pipeline for Chain-of-Thought (CoT) data generation in the Video-CoT dataset, as illustrated in Fig. 4. Leveraging existing video datasets and annotations, we employ a two-stage methodology utilizing a vision-language model (VLM) to analyze spatiotemporal relationships. Specifically, we utilize the Qwen2.5-VL-72B-Instruct model, known for its ex-ceptional ability to generate coherent, long-form reasoning chains. The VLM processes input prompts alongside corresponding video data while referencing relevant examples, producing structured rea-soning outputs encapsulated within <think> and </think> tags to represent the complete thought process.

To ensure the quality of the generated CoT data, we apply taskspecific evaluation metrics for assessing the answers. We establish thresholds based on the average accuracy of individual tasks and the best-performing results from prior studies, ensuring that the model's outputs surpass the average performance of state-of-theart models. Only samples with answers exceeding these thresholds are retained in the Video-CoT dataset, effectively filtering out lowquality chains of thought that may contain logical inconsistencies or factual errors. To further enhance dataset quality, we implement a manual expert review process to eliminate any remaining lowquality data. This two-stage quality control mechanism—combining automated screening based on metrics and manual expert verification—ensures that our Video-CoT dataset comprises only highquality chain-of-thought samples, thereby effectively improving the model's spatiotemporal reasoning capabilities.

4 Method

To enhance the spatiotemporal reasoning capabilities of the visionlanguage model (VLM) and validate the effectiveness of the *Video-CoT dataset*, we propose a comprehensive fine-tuning framework comprising two methods: Answer Supervised Fine-tuning (Ans-SFT) and Chain of Thoughts Supervised Fine-tuning (CoT-SFT). These methods utilize distinct data types from the original dataset and the *Video-CoT dataset* to fine-tune the same model, optimizing performance across various levels of spatiotemporal understanding. In the following sections, we provide detailed descriptions of these two fine-tuning methods, including their model architectures, inputoutput configurations, and loss function formulations.

4.1 Answer-Supervised Fine-Tuning (Ans-SFT)

Ans-SFT focuses on direct answer alignment through end-to-end optimization of answer generation probability. The input-output formulation is expressed as:

Input:
$$V, Q \to \text{Output: } A$$
, (1)

where $V \in \mathbb{R}^{T \times H \times W \times C}$ represents the input video with *T* frames, and *Q* denotes the natural language question. Given this input pair, the model \mathcal{M} generates the target answer *A* by maximizing the conditional likelihood $P(A \mid V, Q)$:

$$A^* = \arg\max_{A} P(A \mid V, Q; \theta).$$
(2)

The method employs standard cross-entropy loss for sequence modeling:

$$\mathcal{L}_{\text{Ans}} = -\sum_{t=1}^{|A|} \log P(a_t \mid a_{< t}, V, Q).$$
(3)

Video-CoT: A Comprehensive Dataset for Spatiotemporal Understanding of Videos Based on Chain-of-Thought

Models	Params.	TVL	VC	SRR TVR		SVG STVG		VG
		tIoU↑	MENTOR↑	EM↑	MENTOR↑	sIoU↑	tIoU↑	sIoU↑
		C	losed-Source MLL	Ms				
GPT-40 [17]	-	21.7	26.9	14.6	25.0	15.9	14.3	7.4
Gemini-1.5-pro [34]	-	24.2	29.0	12.8	27.3	20.9	15.8	9.9
		C	pen-Source MLL	Ms				
Qwen2.5-VL [1]	3B	4.4	4.6	10.0	14.8	10.3	8.5	5.7
Qwen2.5-VL [1]	7B	20.0	19.1	12.4	20.1	12.5	12.9	5.9
InternVideo2_5_Chat [39]	8B	11.8	20.1	7.1	3.4	9.4	10.4	5.4
Video-UTR [45]	7B	10.3	3.7	-	-	2.9	-	-
LLaVA-NeXT-Video [48]	7B	8.5	4.3	6.6	-	-	-	-
LLaVA-onevision-qwen2 [22]	7B	10.4	14.2	9.5	12.8	4.5	13.0	4.9

Table 1: Performance Comparison of VLMs on Video-CoT Benchmark Tasks.

Table 2: Performance of Our Proposed Models on Video-CoT Benchmark Tasks.

Models	Params.	TVL	VC	SRR	TVR	SVG	ST	STVG	
		tIoU↑	MENTOR↑	EM↑	MENTOR↑	sIoU↑	tIoU↑	sIoU↑	
Qwen2.5-VL [1] (Baseline)	3B	4.4	4.6	10.0	14.8	10.3	8.5	5.7	
Video-Ans-SFT (Ours)	3B	8.3 (+3.9↑)	12.3 (+7.7↑)	14.1 (+4.1↑)	22 . 8 (+8 . 0↑)	14.2 (+3.9↑)	12.2 (+3.7↑)	8.1 (+2.4↑)	
Video-CoT-SFT (Ours)	3B	19.7 (+15.3 ↑)	15.4 (+10.8 ↑)	16.9 (+6.9 ↑)	21.1 (+6.3↑)	17.0 (+6.7↑)	14.1 (5.6↑)	9.2 (3.5↑)	

This approach is effective for spatiotemporal localization tasks where answers correspond to specific spatiotemporal coordinates (*e.g.*, timestamps or bounding boxes). Direct SFT enables rapid convergence and efficient training for tasks requiring precise localization. However, it may face limitations in handling tasks that demand fine-grained analysis or multi-step reasoning, such as Spatio-Temporal Video Grounding (STVG) or Spatial Relationship Reference (SRR), where intermediate reasoning steps are crucial for achieving accurate results.

4.2 Chain-of-Thought-Supervised Fine-Tuning (CoT-SFT)

CoT-SFT introduces hierarchical reasoning supervision by leveraging structured reasoning chains from our *Video-CoT dataset*. The input-output formulation expands the basic framework to include intermediate reasoning steps:

Input:
$$V, Q \rightarrow \text{Output: } R, A,$$
 (4)

where $R = \{r_1, r_2, ..., r_n\}$ represents the annotated reasoning chain with *n* intermediate steps, and *A* denotes the final answer. The model learns to generate both the reasoning process and the ultimate solution through a two-phase optimization objective:

$$\mathcal{L}_{\text{CoT}} = -\sum_{t=1}^{|R|} \log P(r_t \mid r_{< t}, V, Q; \theta) - \lambda \sum_{t=1}^{|A|} \log P(a_t \mid a_{< t}, R, V, Q; \theta),$$
(5)

where λ is a hyperparameter that balances the importance of reasoning step generation and final answer prediction. This dualobjective formulation enables the model to simultaneously learn explicit reasoning patterns and accurate answer generation.

This approach is particularly effective for complex multi-step tasks where intermediate reasoning processes—such as object trajectory analysis, event sequencing, or spatial relationship inference—are crucial for achieving correct solutions. To enhance training efficiency and model robustness, we implement a curriculum learning strategy that progressively increases reasoning chain complexity. The curriculum begins with simple 2-step reasoning scenarios (*e.g.*, temporal localization followed by spatial identification) and advances to sophisticated 5-step chains involving complex spatial relationship deduction and causal inference.

Experiments

5.1 Environment Setup

Benchmark To evaluate the effectiveness of our proposed methods and dataset, we constructed the Video-CoT-Benchmark using the same methodology as the Video-CoT dataset. This comprehensive benchmark encompasses all six spatiotemporal understanding subtasks: Video Captioning (VC), Temporal Video Localization (TVL), Spatial Video Grounding (SVG), Spatio-Temporal Video Grounding (STVG), Spatial Relationship Reference (SRR), and Temporal Video Reference (TVR). The benchmark includes 4,500 video questionanswer pairs that are entirely distinct from the training dataset, ensuring a thorough and unbiased evaluation of model performance across various spatiotemporal reasoning tasks.

Implementation Details We conduct our experiments using 581 the Qwen2.5-VL-3B-Instruct model as the baseline, fine-tuning 582 it through two distinct methodologies: Answer-Supervised Fine-583 Tuning (Ans-SFT) and Chain-of-Thought-Supervised Fine-Tuning 584 (CoT-SFT). The fine-tuning process utilizes the AdamW optimizer 585 with a learning rate of 10^{-6} , employing mixed-precision training 586 587 to enhance efficiency on a single GPU. Training is performed for 1 epoch to ensure effective learning from the dataset while minimiz-588 589 ing the risk of overfitting.

Evaluation Metrics We established distinct evaluation metrics 590 for each task. For Temporal Video Localization (TVL), we used 591 temporal Intersection over Union (tIoU) to assess localization ac-592 curacy. In Video Captioning (VC) and Temporal Video Reference 593 (TVR), we employed the MENTOR metric to evaluate the quality 594 and relevance of generated captions and textual descriptions. For 595 Spatial Video Grounding (SVG) and Spatio-Temporal Video Ground-596 ing (STVG), we utilized spatial Intersection over Union (sIoU) to 597 measure precision in spatial localization within single frames and 598 599 across frames, respectively, with STVG also incorporating tIoU 600 to ensure temporal accuracy. For Spatial Relationship Reference (SRR), Exact Match (EM) was used to verify if predicted spatial 601 relationships align with ground truth labels. These criteria were 602 selected based on their relevance to each task's requirements and 603 their effectiveness in measuring model performance in capturing 604 spatiotemporal information. 605

5.2 Results

606

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

The evaluation on Video-CoT benchmark reveals significant limitations in current state-of-the-art VLMs when performing finegrained spatiotemporal reasoning for videos, as shown in Table 1. Closed-source models generally outperform open-source counterparts across all benchmark tasks, with Gemini-1.5-pro achieving the highest scores in most metrics. However, even the advanced models demonstrate limited performance, particularly in STVG task where the best score of 9.9 sIoU remains relatively low.

Table 2 demonstrates the effectiveness of our proposed training approaches compared to the baseline Qwen2.5-VL (3B) model. Both the Video-Ans-SFT and Video-CoT-SFT methods yield substantial improvements across all tasks, with the CoT-SFT variant showing particularly strong gains in temporally-oriented metrics (tIoU) for TVL and TVR tasks. The Video-CoT-SFT achieves a remarkable 19.7 tIoU in TVL, outperforming not only its 3B baseline but also all other open-source 7B models listed in Table 1. This performance advantage confirms that Chain-of-Thought fine-tuning specifically enhances a model's capacity for long-term spatiotemporal reasoning compared to standard answer-based supervision, providing a promising direction for developing more capable models.

6 Potential Applications

The Video-CoT dataset, with its dense spatiotemporal annotations, balanced distribution of short and long videos, and structured Chainof-Thought (CoT) annotations, enables advanced research across multiple domains. We highlight four key application areas that leverage its unique characteristics:

Training and Benchmarking Spatiotemporal Reasoning The Video-CoT dataset offers dense spatiotemporal annotations and 639

640

641

642

643

644

645

646

647

648

649

650

651

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

object trajectory information, providing high-quality benchmark data for video spatiotemporal relation reasoning tasks [39, 42, 45, 48]. It supports the training of vision-language models (VLMs) with spatiotemporal reasoning capabilities, including understanding relationships and event progression. The structured Chain-of-Thought annotations (<think>...</think>) enhance the development of interpretable video reasoning models, applicable in areas such as autonomous driving (analyzing vehicle-pedestrian interactions) and surveillance systems (detecting anomalous spatiotemporal patterns). Additionally, its per-frame object localization and duration labels enable precise evaluation of relational reasoning accuracy.

Advancing Long-term Video Understanding The dataset features videos exceeding 160 seconds in duration, allowing researchers to develop innovative long-term temporal modeling approaches [41, 46] with memory capabilities. These extended sequences support research on long-range dependency modeling and event detection, particularly in applications requiring state tracking, such as sports commentary. Moreover, the balanced mix of short and long videos enhances studies on temporal generalization.

Cross-modal Retrieval Enhancement The comprehensive annotations in Video-CoT significantly advance video-text retrieval systems [9–13, 28] by enabling cross-modal search capabilities. The dataset supports fine-grained queries that incorporate spatial, temporal, and relational constraints, such as identifying specific object interactions within defined time windows. This level of detail facilitates accurate temporal grounding of textual descriptions to specific video segments and enhances compositional cross-modal retrieval through advanced relational reasoning. These capabilities hold transformative potential in forensic investigations, allowing analysts to efficiently search through extensive video evidence.

Intelligent Human-Computer Interaction Systems The precise spatiotemporal annotations in Video-CoT create new opportunities for intelligent interactive systems [27] requiring nuanced video understanding. In augmented and virtual reality applications, the dataset's detailed object trajectory data facilitates more natural interactions between virtual and real-world objects, ensuring proper occlusion effects. For robotic systems, timestamped object positions and interaction patterns enhance dynamic environment mapping, enabling robots to anticipate human actions effectively. The structured <think> annotations are especially valuable for developing explainable AI assistants, allowing them to articulate their spatiotemporal reasoning processes during human-computer interactions and bridging the gap between complex visual understanding and intuitive communication.

7 Conclusion

This paper introduces *Video-CoT*, a novel dataset aimed at enhancing spatiotemporal understanding in video comprehension using Chain-of-Thought (CoT) methodologies. It contains 192,000 finegrained spatiotemporal question-answer pairs and 23,000 highquality CoT-annotated samples, organized into three key components. A comprehensive benchmark for evaluation is also provided. Our experiments reveal significant challenges faced by current vision-language models (VLMs). In summary, Video-CoT represents a valuable contribution to the community, advancing research in multimodal content understanding. Video-CoT: A Comprehensive Dataset for Spatiotemporal Understanding of Videos Based on Chain-of-Thought

755

756

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810

811 812

References

697

698

699

700

701

702

703

704

705

706

707

708

709

710

711

712

713

714

715

716

718

719

720

721

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. 2025. Qwen2. 5-vl technical report. arXiv preprint arXiv:2502.13923 (2025).
- [2] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. 2021. 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. 1728–1738.
- [3] David Chen and William B Dolan. 2011. Collecting highly parallel data for paraphrase evaluation. In Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies. 190–200.
- [4] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. arXiv preprint arXiv:2110.14168 (2021).
- [5] Kaituo Feng, Kaixiong Gong, Bohao Li, Zonghao Guo, Yibing Wang, Tianshuo Peng, Benyou Wang, and Xiangyu Yue. 2025. Video-r1: Reinforcing video reasoning in mllms. arXiv preprint arXiv:2503.21776 (2025).
- [6] Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Pal: Program-aided language models. In International Conference on Machine Learning. 10764–10799.
- [7] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. 2017. The" something something" video database for learning and evaluating visual common sense. In Proceedings of the IEEE international conference on computer vision. 5842–5850.
- [8] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948 (2025).
- Xiaoshuai Hao and Wanqian Zhang. 2023. Uncertainty-aware alignment network for cross-domain video-text retrieval. Advances in Neural Information Processing Systems 36 (2023), 38284–38296.
- [10] Xiaoshuai Hao, Wanqian Zhang, Dayan Wu, Fei Zhu, and Bo Li. 2022. Listen and look: Multi-modal aggregation and co-attention network for video-audio retrieval. In IEEE International Conference on Multimedia and Expo. 1–6.
- [11] Xiaoshuai Hao, Wanqian Zhang, Dayan Wu, Fei Zhu, and Bo Li. 2023. Dual alignment unsupervised domain adaptation for video-text retrieval. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 18962–18972.
- [12] Xiaoshuai Hao, Yucan Zhou, Dayan Wu, Wanqian Zhang, Bo Li, and Weiping Wang. 2021. Multi-Feature Graph Attention Network for Cross-Modal Video-Text Retrieval. In Proceedings of the 2021 International Conference on Multimedia Retrieval. 135–143.
- [13] Xiaoshuai Hao, Yucan Zhou, Dayan Wu, Wanqian Zhang, Bo Li, Weiping Wang, and Dan Meng. 2021. What matters: Attentive and relational feature aggregation network for video-text retrieval. In *IEEE International Conference on Multimedia* and Expo. 1–6.
- [14] Xiaoshuai Hao, Yi Zhu, Srikar Appalaraju, Aston Zhang, Wanqian Zhang, Bo Li, and Mu Li. 2023. Mixgen: A new multi-modal data augmentation. In Proceedings of the IEEE/CVF winter conference on applications of computer vision. 379–389.
- [15] Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. arXiv preprint arXiv:2305.02301 (2023).
- [16] Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. 2024. Vtimellm: Empower llm to grasp video moments. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 14271–14280.
- [17] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-40 system card. arXiv preprint arXiv:2410.21276 (2024).
- [18] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. 2014. Large-scale video classification with convolutional neural networks. In Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. 1725–1732.
- [19] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. 2017. The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017).
- [20] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. Advances in neural information processing systems 35 (2022), 22199–22213.
- [21] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In Proceedings of the IEEE International Conference on Computer Vision. 706–715.
- [22] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. 2024. Llava-onevision: Easy visual task transfer. arXiv preprint arXiv:2408.03326 (2024).

- [23] Xinhao Li, Ziang Yan, Desen Meng, Lu Dong, Xiangyu Zeng, Yinan He, Yali Wang, Yu Qiao, Yi Wang, and Limin Wang. 2025. Videochat-r1: Enhancing spatio-temporal perception via reinforcement fine-tuning. arXiv preprint arXiv:2504.06958 (2025).
- [24] Ruilin Luo, Zhuofan Zheng, Yifan Wang, Yiyao Yu, Xinzhe Ni, Zicheng Lin, Jin Zeng, and Yujiu Yang. 2025. URSA: Understanding and Verifying Chain-ofthought Reasoning in Multimodal Mathematics. arXiv preprint arXiv:2501.04686 (2025).
- [25] Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. In The 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics.
- [26] Joanna Materzynska, Guillaume Berger, Ingo Bax, and Roland Memisevic. 2019. The jester dataset: A large-scale video dataset of human gestures. In Proceedings of the IEEE/CVF International Conference on Computer vision Workshops.
- [27] Jinhui Pang, Jinglin He, Noureldin Mohamed Abdelaal Ahmed Mohamed, Changqing Lin, Zhihui Zhang, and Xiaoshuai Hao. 2025. A Hierarchical Reinforcement Learning Framework for Multi-UAV Combat Using Leader-Follower Strategy. arXiv preprint arXiv:2501.13132 (2025).
- [28] Wanqian Zhang Jinchao Zhang Zexian Yang Xiaoshuai Hao Bo Li Sa Zhu, Huashan Chen. 2025. Uneven Event Modeling for Partially Relevant Video Retrieval. In *IEEE International Conference on Multimedia and Expo.* 1–6.
- [29] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. 2020. Finegym: A hierarchical video dataset for fine-grained action understanding. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2616–2625.
- [30] Fangxun Shu, Lei Zhang, Hao Jiang, and Cihang Xie. 2023. Audio-visual llm for video understanding. arXiv preprint arXiv:2312.06720 (2023).
- [31] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012).
- [32] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2018. Commonsenseqa: A question answering challenge targeting commonsense knowledge. arXiv preprint arXiv:1811.00937 (2018).
- [33] Zongheng Tang, Yue Liao, Si Liu, Guanbin Li, Xiaojie Jin, Hongxu Jiang, Qian Yu, and Dong Xu. 2021. Human-centric spatio-temporal video grounding with visual transformers. *IEEE Transactions on Circuits and Systems for Video Technology* 32, 12 (2021), 8238–8249.
- [34] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv preprint arXiv:2403.05530 (2024).
- [35] Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. 2025. Kimi k1. 5: Scaling reinforcement learning with llms. arXiv preprint arXiv:2501.12599 (2025).
- [36] Junke Wang, Dongdong Chen, Chong Luo, Xiyang Dai, Lu Yuan, Zuxuan Wu, and Yu-Gang Jiang. 2023. Chatvideo: A tracklet-centric multimodal and versatile video understanding system. arXiv preprint arXiv:2304.14407 (2023).
- [37] Shuai Cheng Lin WangXiaoshuai Hao Wanqian Zhang Xiaohua Chen Wei Wang. 2025. Multi-Granularity Based Collaborative Learning for Semi-Supervised Hashing. In IEEE International Conference on Multimedia and Expo (ICME). 1–6.
- [38] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. arXiv preprint arXiv:2203.11171 (2022).
- [39] Yi Wang, Xinhao Li, Ziang Yan, Yinan He, Jiashuo Yu, Xiangyu Zeng, Chenting Wang, Changlian Ma, Haian Huang, Jianfei Gao, et al. 2025. InternVideo2. 5: Empowering Video MLLMs with Long and Rich Context Modeling. arXiv preprint arXiv:2501.12386 (2025).
- [40] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems 35 (2022), 24824–24837.
- [41] Yuetian Weng, Mingfei Han, Haoyu He, Xiaojun Chang, and Bohan Zhuang. 2024. Longvlm: Efficient long video understanding via large language models. In European Conference on Computer Vision. 453–470.
- [42] Yujie Wu, Huaihai Lyu, Yingbo Tang, Lingfeng Zhang, Zhihui Zhang, Wei Zhou, and Siqi Hao. 2025. Evaluating GPT-4o's Embodied Intelligence: A Comprehensive Empirical Study. *TechRxiv preprint techrxiv*.174495686.69962588/v1 (2025).
- [43] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In Proceedings of the IEEE conference on computer vision and pattern recognition. 5288–5296.
- [44] Cilin Yan, Haochen Wang, Shilin Yan, Xiaolong Jiang, Yao Hu, Guoliang Kang, Weidi Xie, and Efstratios Gavves. 2024. Visa: Reasoning video object segmentation via large language models. In *European Conference on Computer Vision*. Springer, 98–115.
- [45] En Yu, Kangheng Lin, Liang Zhao, Yana Wei, Zining Zhu, Haoran Wei, Jianjian Sun, Zheng Ge, Xiangyu Zhang, Jingyu Wang, et al. 2025. Unhackable temporal rewarding for scalable video mllms. arXiv preprint arXiv:2502.12081 (2025).

ACM MM, October 27-31, 2025, Dublin, Ireland

S.Zhang, X.Hao and Y.Tang et al.

- [46] Lingfeng Zhang, Xiaoshuai Hao, Qinwen Xu, Qiang Zhang, Xinyao Zhang, Pengwei Wang, Jing Zhang, Zhongyuan Wang, Shanghang Zhang, and Renjing Xu. 2025. Mapnav: A novel memory representation via annotated semantic maps for vlm-based vision-and-language navigation. arXiv preprint arXiv:2502.13451
 (2025).
- [47] Xingjian Zhang, Siwei Wen, Wenjun Wu, and Lei Huang. 2025. Tinyllava-videor1: Towards smaller lmms for video reasoning. arXiv preprint arXiv:2504.09641 (2025).
 - [48] Y Zhang, B Li, H Liu, Y Lee, L Gui, D Fu, J Feng, Z Liu, and C Li. 2024. Llava-next: A strong zero-shot video understanding model. (2024).
- [49] Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. arXiv preprint arXiv:2210.03493 (2022).
- [50] Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2023. Multimodal chain-of-thought reasoning in language models. arXiv preprint arXiv:2302.00923 (2023).
- [51] Zhu Zhang, Zhou Zhao, Yang Zhao, Qi Wang, Huasheng Liu, and Lianli Gao. 2020. Where does it exist: Spatio-temporal video grounding for multi-form sentences. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 10668–10677.