

JRDB-Reasoning: A Difficulty-Graded Benchmark for Visual Reasoning in Robotics

Simindokht Jahangard¹, Mehrzad Mohammadi², Yi Shen¹, Zhixi Cai¹, Hamid Rezatofghi¹

¹Monash University

²Sharif University of Technology
simindokht.jahangard@monash.edu

Abstract

Recent advances in Vision-Language Models (VLMs) and large language models (LLMs) have greatly enhanced visual reasoning, a key capability for embodied AI agents like robots. However, existing visual reasoning benchmarks often suffer from several limitations: they lack a clear definition of reasoning complexity, offer no control to generate questions over varying difficulty and task customization, and fail to provide structured, step-by-step reasoning annotations (workflows). To bridge these gaps, we formalize reasoning complexity, introduce an adaptive query engine that generates customizable questions of varying complexity with detailed intermediate annotations, and extend the JRDB dataset with human-object interaction and geometric relationship annotations to create JRDB-Reasoning, a benchmark tailored for visual reasoning in human-crowded environments. Our engine and benchmark enable fine-grained evaluation of visual reasoning frameworks and dynamic assessment of visual-language models across reasoning levels.

Code — <https://github.com/sjahangard/JRDB-Reasoning>

Datasets — <https://jrdb.erc.monash.edu/>

Introduction

Recent advances in Vision-Language Models (VLMs) and large language models (LLMs) have significantly transformed computer vision, evolving it from basic visual perception to higher-level visual reasoning (Amizadeh et al. 2020). By integrating cognitive capabilities that closely resemble human perception and understanding, visual reasoning is especially crucial for embodied AI agents such as robots. These advances enable robots not only to perceive and interpret visual information but also to engage in cognition and reasoning—allowing them to make decisions, solve problems, and draw meaningful conclusions. Developing models with advanced visual reasoning capabilities requires high-quality benchmarks that not only provide rigorous evaluation but also challenge models with diverse, context-rich, and complex reasoning tasks (Wu et al. 2024; Patraucean et al. 2023; Mangalam, Akshulakov, and Malik 2023; Chandrasegaran et al. 2024). The evolution from

static 2D images (Yu et al. 2016; Chen et al. 2024a; Johnson et al. 2017) to dynamic videos, 3D environments, and multimodal datasets (Li et al. 2024; Shen et al. 2024; Fu et al. 2025) has enabled models to capture temporal dependencies, integrate multimodal cues, and engage in high-level reasoning for real-world decision-making. However, existing visual reasoning benchmarks often group under the broad label of “reasoning,” without distinguishing between different levels of complexity. In reality, visual reasoning spans a continuum—from simple perceptual tasks, such as object recognition, to multi-step logical deductions that require sequential and causal reasoning. This lack of granularity causes benchmarks to evaluate basic recognition tasks in the same way as those requiring complex, multi-stage inference. Consequently, current benchmarks are not able to accurately measure a model’s capacity for advanced reasoning, limiting their effectiveness in evaluating true reasoning capabilities.

Moreover, current datasets lack the ability to adjust question complexity based on user-specified parameters such as the type of task (e.g., VG, VQA), the focus of the question (e.g., humans, objects), and the level of spatial or temporal reasoning (e.g., in images or videos). This limitation hinders the customization of reasoning tasks for diverse evaluation needs, thereby restricting comprehensive model assessment and reducing adaptability to user-specific requirements.

Additionally, these datasets generally provide only input data and corresponding labels, while omitting the intermediate reasoning steps necessary to derive the final answer. Such steps are vital for evaluating models that rely on chain-of-thought (step-by-step) reasoning, such as workflow-based systems, neuro-symbolic frameworks, and program-guided models (Sur’is, Menon, and Vondrick 2023; Ke et al. 2024), where understanding the reasoning trajectory is key to assessing the model’s inference process.

In this paper, we address all these gaps by: (i) defining and formalizing the complexity of perception-to-reasoning tasks, (ii) developing an adaptive query engine that dynamically generates non-predefined questions with adjustable complexity levels, (iii) providing step-by-step solutions (intermediate annotation) for each visual reasoning question within the workflow, and (iv) annotating human-object interactions manually and computing geometric relationships to enhance the robotic JRDB dataset for integration with an

JRDB-Reasoning: A Difficulty-Graded Benchmark for Visual Reasoning in Robotics



Figure 1: Illustration of JRDB-Reasoning on VG and VQA for images and videos. (a) and (e) depict VG tasks on images & HOI. (b) presents a VQA-Counting question with reasoning steps: (1) identify all humans, (2) filter for females, (3) find female looking at the robot. (c) illustrates a VG multi-object tracking task, while (d) shows a VQA-Wh question applied to a video.

adaptive query engine.

First, we formalize the reasoning complexity of a visual reasoning task in terms of the number of reasoning steps required to reach a solution. By representing a scene and its entities as a spatio-temporal graph, we argue that the complexity of a visual reasoning question is closely tied to the number of nodes (entities, S) and edges (relationships, R), as well as their temporal dynamics (time slots, T). As the number of these factors increases, the overall task complexity rises. This structured formulation provides a systematic framework for assessing reasoning difficulty and evaluating model capabilities with greater precision.

Second, we introduce an adaptive query engine that dynamically generates non-predefined questions with adjustable complexity. Users can customize queries across spatial and temporal dimensions in images and videos, tailoring them to specific tasks, including task type (VG or VQA), modality (image or video), subject focus (human, object, or both), and spatial/temporal scope (single, pair, or clique).

Third, the engine generates step-by-step intermediate solutions for each query, providing workflow annotations that capture the entire reasoning trajectory. For example, as illustrated in Figure 1(b), given the query “Count the number of female persons looking at the robot (me),” the engine pro-

duces a structured reasoning sequence: (1) identifying all humans, (2) filtering female individuals, and (3) detecting females looking at the robot. This decomposition enables fine-grained evaluation of reasoning accuracy and facilitates deeper analysis of VLMs or workflow-based framework, including compositional models (Sur’is, Menon, and Vondrick 2023; You et al. 2023; Ke et al. 2024; Cai et al. 2025).

Fourth, we manually annotated human-object interactions with associated confidence levels, Figure 1(e). In addition, we extracted geometric relationships, including spatial relations and distances between each entity (human or object). Through this process, we enriched the JRDB dataset (Martin-Martin et al. 2021), complementing and extending existing annotations (Ehsanpour et al. 2022; Jahangard et al. 2024; Vendrow et al. 2023; Le et al. 2024). All annotations are integrated into our adaptive query engine to generate questions.

Finally, leveraging advances in VLMs (Zhang, Li, and Bing 2023; Wu et al. 2023b), we assess their performance on JRDB-Reasoning, our novel visual reasoning benchmark, to evaluate their capacity for tasks of varying complexity. Despite the engine’s ability to generate highly complex questions, we focus on a simpler subset aligned with current model capabilities. These evaluations provide key insights

into model performance, limitations, and the benchmark’s potential to drive future research.

Related Works

Visual Reasoning Benchmarks. High-quality datasets and benchmarks drive visual reasoning progress, evolving from 2D to 3D and from images to videos, enabling models to tackle complex multimodal tasks like VG and VQA with richer cues and contexts for improved decision-making. Benchmarks evolved from 2D images to videos and 3D environments. Early benchmarks like Flickr30K (Plummer et al. 2015) and DAQUAR Visual7W (Zhu et al. 2016) focused on entity localization, but lacked the scene context necessary for testing deeper reasoning abilities. Datasets like RefCOCO Series (Kazemzadeh et al. 2014), Ref-Adv (Akula et al. 2020), COCO-QA (Ren, Kiros, and Zemel 2015), Visual Genome (Krishna et al. 2017), VQAv2 (Goyal et al. 2017), CLEVR (Johnson et al. 2017) and GQA (Hudson and Manning 2019) introduced more complex tasks, enhancing visual grounding and improving the ability to infer relationships and interactions in more context-rich environments. With video’s rise, benchmarks such as TGIF-QA (Jang et al. 2017), TVQA (Lei et al. 2018), STAR (Wu et al. 2024), VideoQA (Yang et al. 2003), MVBench (Li et al. 2024), Perception Test (Patraucean et al. 2023), VITATECS (Li et al. 2023e), LLAVIDAL (Chakraborty et al. 2024), MLVU (Zhou et al. 2024), Towards Event-oriented (Du et al. 2024), and TempCompass (Liu et al. 2024c) advanced temporal reasoning in dynamic, event-based, and egocentric contexts. Datasets like OSCaR (Nguyen et al. 2024), EgoSchema (Mangalam, Akshulakov, and Malik 2023), and HourVideo (Chandrasegaran et al. 2024) further pushed first-person perspective research in interactive reasoning, action tasks, and large-scale egocentric video understanding. Furthermore, 3D benchmarks like ScanReason (Zhu et al. 2024), LAMM (Yin et al. 2023), SpatialRGPT (Cheng et al. 2025), Mono3DVG (Zhan, Yuan, and Xiong 2024), and M3DBench (Li et al. 2023d), and in autonomous driving like ScanQA-3D (Azuma et al. 2022), NuScenes-QA (Qian et al. 2024) focus on spatial reasoning in realistic 3D environments, addressing the need for models that can navigate and understand immersive spaces. While visual reasoning benchmarks expand the scope to include dynamic and interactive environments that test models on adaptability and complex reasoning in real-world applications. This expansion into more complex, multimodal, and spatially aware datasets has enabled models to address increasingly dynamic and interactive challenges. Visual reasoning benchmarks lack clear definitions of difficulty levels, customization of question complexity, and reasoning steps, hindering accurate model evaluation. In this paper, we define perception-to-reasoning complexity, build an adaptive query engine for dynamic questions with adjustable difficulty, provide step-by-step solutions, and introduce JRDB-Reasoning, a dataset for robots in human-crowded environments.

Vision Language Models. Multimodal models, particularly VLMs, have advanced from perception to complex reasoning, with increasing emphasis on reasoning capabilities, as reviewed here. These advances build on LLM progress

in language and reasoning, as seen in GPT-4 (Achiam et al. 2023), Gemini (Team et al. 2023), Claude, and GPT-4o (Hurst et al. 2024). Meanwhile, models such as VideoChat (Li et al. 2023c), Visual ChatGPT (Wu et al. 2023a), VALLEY (Luo et al. 2023), Otter (Li et al. 2023a), and MiniGPT-4 (Zhu et al. 2023), along with newer methods like LLaVA-NeXT-Video (Zhang et al. 2024b) and LongVU (Shen et al. 2024), have focused on advancements in video understanding and visual processing. Additional efforts from models like mPLUG-Owl2 (Ye et al. 2024), SPHINX (Lin et al. 2023), Intern-VL (Chen et al. 2024b), Yi-VL (AI et al. 2024), VideoChat2 (Li et al. 2023c), Cambrian-1 (Tong et al. 2024), PLLaVA (Xu et al. 2024), Blip2 (Li et al. 2023b), Florence2 (Xiao et al. 2024), and MiniGPT4-Video (Ataallah et al. 2024), as well as recent models such as Emu3 (Wang et al. 2024b) and Pixtral (Agrawal et al. 2024), have significantly contributed to improving video comprehension, contextual learning, and instruction tuning. Notably, models like GLIP (Li et al. 2022), ReCLIP (Subramanian et al. 2022), GroundingDINO (Chen, Anjum, and Gurari 2022), and YOLO-World (Cheng et al. 2024a), Gemini (Team et al. 2023), GPT-4o (Hurst et al. 2024), Qwen2-VL (Wang et al. 2024a), LLaVA-Video (Zhang et al. 2024b), InternVL2 (Chen et al. 2024b), LongVU (Shen et al. 2024), and Flash-VStream (Zhang et al. 2024a) advance more in terms of multimodal understanding, long-context reasoning, and efficient video and language processing. Together, these advances highlight the growing potential of multimodal methods to process and integrate diverse forms of information. However, the ability of these methods to handle reasoning at different levels of difficulty remains unexplored. In this paper, we evaluate some of these models’ reasoning performance across various difficulty levels using the JRDB-Reasoning to analyze their capacity for reasoning in complex visual scenarios.

Visual Reasoning Engine and Dataset

Formalizing Reasoning Complexity

Reasoning tasks vary in complexity, ranging from simple object recognition to more intricate problems requiring multiple reasoning steps. As the number of steps increases, the difficulty correspondingly rises. We use this hypothesis as a foundation to formalize visual reasoning complexity here.

Figure 2 provides an overview of our approach to formalizing reasoning complexity in visual reasoning. We propose that perception and reasoning exist on a continuum (difficulty is vague) rather than as distinct levels. As illustrated in Figure 2(a), reasoning complexity spans from basic perception to advanced multi-step reasoning, typically depending on the number of intermediate steps required to reach a solution. But, this complexity cannot always be quantified before solving a visual reasoning question.

To address this, we propose an alternative complexity measure that strongly correlates with the number of reasoning steps. Specifically, we represent a scene as a dynamic spatio-temporal graph, where nodes correspond to entities such as humans, objects, and surfaces, while edges capture

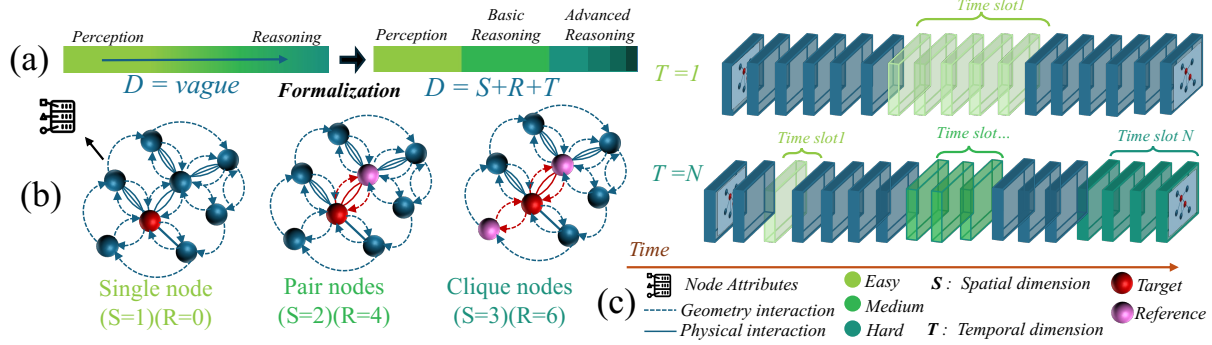


Figure 2: Formalization of complexity involves nodes (S), interactions (R), and time slots (T); greater involvement indicates higher complexity (darker green), while lighter green shows lower complexity.

various types of interactions (e.g. geometric, physical), as shown in Figure 2(b). It can be a target (question asked about it) or a reference (related to the target). Each entity (S) is a unique instance (e.g. a person, an object, or a surface) tracked over time. These entities and their interactions (R) evolve across different time intervals, referred to as time slots (T), as illustrated in Figure 2(c).

As more nodes, edges, and time slots are introduced, reasoning complexity in a visual reasoning question increases, indicated by the color transition from light green to dark green in Figure 2. This structured approach provides a systematic framework for understanding reasoning complexity, enabling a more precise evaluation of model capabilities.

To estimate the complexity of a visual reasoning question, we introduce formulas modeling element involvement, spatial interactions, and temporal spans in the dynamic scene graphs. To validate this formulation, we conducted a user study with 20 participants who rated the difficulty of 500 questions. The results showed strong agreement between human judgments and our computed scores, confirming that questions involving more entities, spatial reasoning, or temporal tracking are perceived as more complex.

(1) Entity (Node) Involvement (S). Question difficulty depends first on how many entities (or nodes) are involved, target or reference in Figure 2. This is quantified by counting the nodes $N = \{n_1, \dots, n_k\}$, where k is the total number of entities. Each node n_i has a unique ID and set of attributes, which can be static (e.g., age, gender) or dynamic (e.g., actions, states). Entity involvement S is expressed as $S = k$.

(2) Spatial Relationships (R). This factor captures spatial interactions between entities within a single frame, represented by edges E_s connecting nodes $n_i, n_j \in N$, i.e.,

$$E_s = \{(n_i, n_j) \mid n_i, n_j \in N\}.$$

These edges encode properties like proximity or orientation. The spatial interaction complexity is defined as

$$R = |E_s|,$$

where $|E_s|$ denotes the number of spatial relationships (edges) at a given time frame.

(3) Temporal Interactions (T). This factor captures cross-temporal interactions via edges E_t connecting entities

$n_i, n_j \in N$ over time intervals t_1, t_2 , i.e.,

$$E_t = \{(n_i, n_j, t_1, t_2) \mid n_i, n_j \in N, t_1, t_2 \in \text{time frames}\}.$$

These edges capture continuity or tracking across time. The temporal interaction complexity is defined as

$$T = |E_t|,$$

where $|E_t|$ is the number of time slots involving each entity.

Overall Difficulty (D). The difficulty of a visual question depends on entity involvement (S), spatial interactions (R), and temporal interactions (T), $D = S + R + T$.

By adopting this approach, we develop a more systematic framework for understanding task difficulty, which allows for a nuanced evaluation of model performance. This structured methodology not only improves benchmarking techniques but also offers deeper insights into a model’s capability to progress from basic perception to advanced reasoning.

Generative Adaptive Query Engine

We developed an adaptive query engine that enables generating open-ended, compositional questions with customizable complexity, based on preferences across spatial and temporal scales in images or videos. It supports diverse tasks such as VG and VQA (Wh and Counting type question, Figure 1(b)(d)). In addition, the engine produces intermediate annotations that provide detailed descriptions at every stage of reasoning. These annotations offer deeper interpretability and are particularly valuable for evaluating VLMs that rely on multi-step reasoning and compositional understanding.

The engine operates based on user-defined preferences, including task type (VG or VQA), modality (image or video), subject focus (human, object, or both), and spatial/temporal scope (single, pair, or clique). Depending on the selected modality, a spatial-temporal graph is constructed. Based on the spatial configuration, all possible combinations of node attributes and edge relationships are computed, and one combination is randomly selected. The search process then proceeds over the spatial-temporal graph using this chosen combination. At each step, corresponding reasoning steps—such as intermediate annotations—are

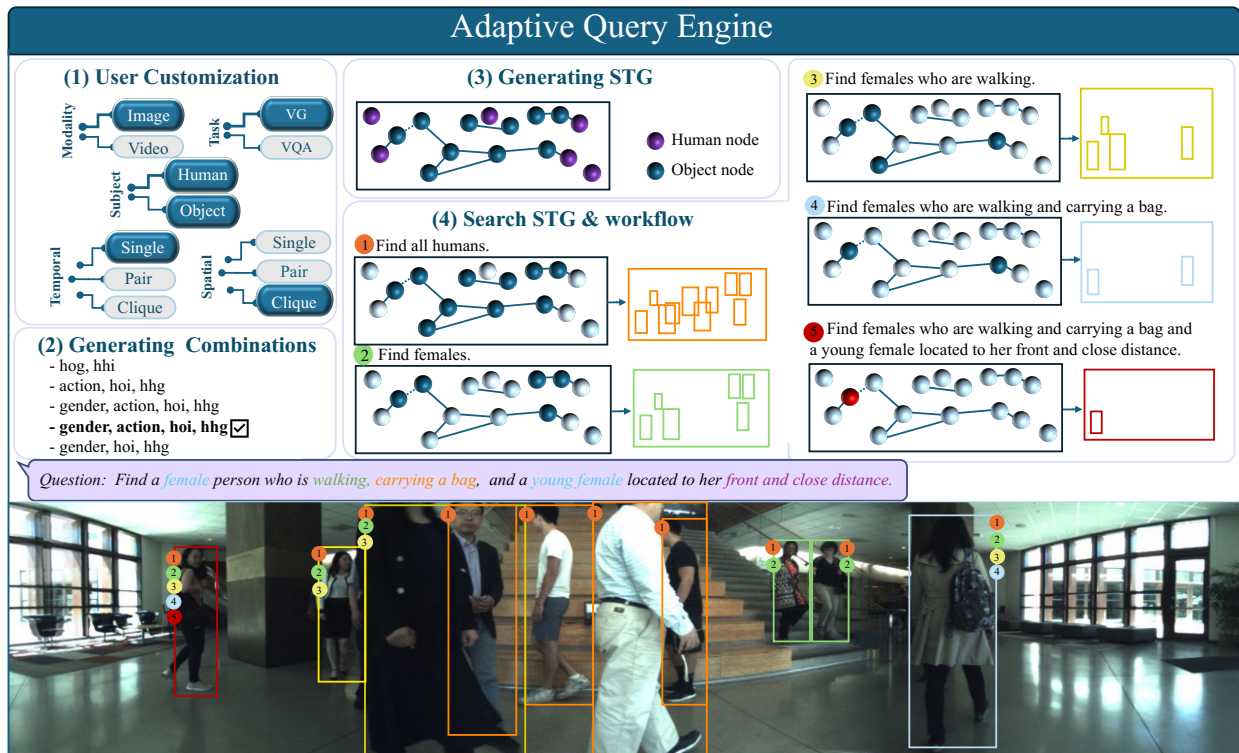


Figure 3: Example of the Generative Query-Based Engine: (1) user selects preferences: modality (image), subject (human&object), task (VG), spatial (clique), temporal (single). (2) All attribute combinations are computed. (3) a spatial-temporal graph (STG) is generated. (4) The STG is searched to generate intermediate annotations and queries.

generated, and the selected attributes or relationships are incorporated into the query. Further details of the generative query engine, along with an illustrative example, are provided below. As shown in Figure 3, the proposed engine operates through the following steps: (1) Users begin by customizing their query preferences via parameters that define the difficulty level. In this example, the user selects the VG task, specifies both human and object as subjects, chooses the image modality ($T = 1$), and sets the spatial preference to a clique configuration with three nodes ($S = 3$). (2) Given the spatial configuration, all possible combinations of node attributes and edge relationships are computed, and one combination is randomly selected. For instance, the attribute set [“gender”, “action”, “hoi”, “hhg”] with ($R = 2$) may be chosen, where “hoi” stands for human-object interaction, and “hhg” refers to human-human geometry. (3) A spatial-temporal graph (STG) is then generated based on the selected modality. (4) Finally, the engine searches through the spatial-temporal graph and incrementally generates both intermediate annotations and the corresponding questions.

Dataset and Annotations

Human Object Interaction. Human-object interaction (HOI) enhances visual reasoning by connecting human actions to objects, improving identification, activity recognition, and scene interpretation. It also extends the range of questions generated in JRDB annotations. In JRDB-

Reasoning, we offer multi-label, fine-grained HOI annotations, shown in Figure 1(e), at the frame level, categorized into four groups, Figure 4. The first category focuses on pose-based activities such as walking, standing, and sitting, shown by dash Figure 1(e), which are the most common in the dataset. The second category involves observational interactions, where individuals visually engage with objects or their surroundings, but these are less frequent. The third category emphasizes physical interactions with objects, like touching and carrying, and is strongly represented. The fourth category includes manipulative interactions, involving precise handling of objects, such as operating and working, which are also notable. These categories provide a detailed framework for analyzing human-object interactions in dynamic environments, with pose-based and physical interactions being the most common, while observational interactions are the least frequent.

Annotation Process. To annotate human-object interactions, we use a specialized toolbox that assigns unique IDs to 2D and 3D bounding box annotations. Our annotation process follows a quality control procedure aligned with previous JRDB benchmarks to ensure consistency. Interaction annotations are carefully aligned with individual actions, and strict guidelines are followed to maintain accuracy. We also address challenges such as robot distance, varying lighting conditions, occlusions, and crowded environments by assigning difficulty levels—Easy (1), Medium (2), and Hard

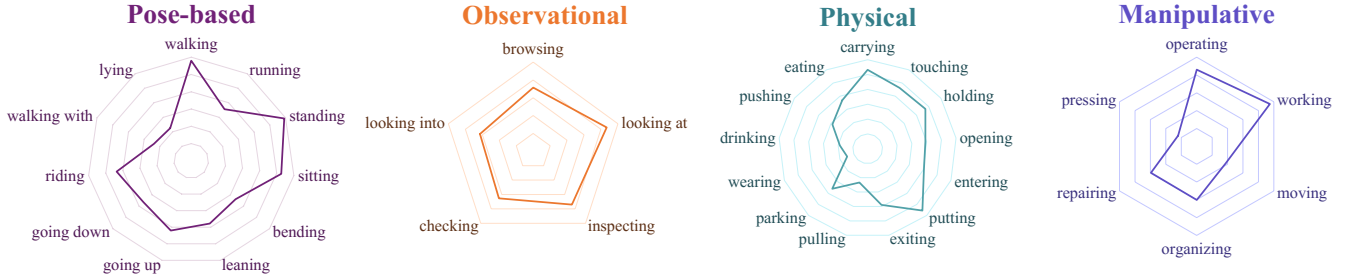


Figure 4: Human-object interaction falls into four types: Pose-based, Observational, Physical, and Manipulative.

Relation	Cond. θ	Distance	Cond. d
(up/down) front	$[-30, 30]$	Very close	$(0, 0.5]$
(up/down) front-left	$(30, 45]$	Close	$[0.5, 1.5]$
(up/down) left	$(45, 135]$	Moderate	$[1.5, 5]$
(up/down) back-left	$(135, 145]$	Far	$[5, 10]$
(up/down) front-right	$[-45, -30]$	Very far	$[10, \infty)$
(up/down) right	$[-135, -45]$		
(up/down) back-right	$[-145, -135]$		
(up/down) back	else		

Table 1: Spatial geometric relationship based on θ & d .

(3)—based on annotator confidence. Each label is reviewed by two additional annotators, with random quality checks to enhance fairness and reliability.

Geometry Relationships. Geometry relationships, including distance and spatial relationship, improve visual reasoning by providing context for object recognition, navigation, and decision-making. They help systems understand object positioning, track movement, predict interactions, and optimize path finding. In robotics and computer vision, this enhances scene understanding, human-robot interaction, and autonomous decision-making.

Spatial relationships. To extract spatial geometry relationships, inspired by (Qian et al. 2024), we compute the spatial relationship relative to a target point, which could be a robot or a person with orientation. The target point serves as the origin for measuring distances and angles in 3D space. As explained in JRDB (Martin-Martin et al. 2021), a robot has a fixed orientation, while for humans, their orientation is also taken into account, shown with red arrow in Figure 1(e). This captures spatial relationships by accounting for both position and facing direction, yielding a more accurate environmental representation. To achieve this, we compute the angle between the vector that connects the centers of two bounding boxes and the forward direction of target. The formula for this angle θ is expressed as:

$$\theta = \cos^{-1} \left(\frac{(B_1[:2] - B_2[:2]) \cdot (B_1[:2] - B_2[:2])}{\|B_1[:2] - B_2[:2]\|^2} \right)$$

where $B_i = [x, y]$ represents the 3D bounding box of entity i . Here, the forward direction of the target point (e.g., robot or human) is defined as 0° , with counterclockwise angles being considered positive. For spatial relationships, we calculate the relative position between two points, B_1

and B_2 , along both the X and Y coordinates. Additionally, we take the Z-coordinate into account and calculate the distance, incorporating this third dimension (up or down) into the analysis which resulted in 16 possible directions. Based on the calculated angle, the relationship between the two objects is categorized in Table 1.

Distance. To enhance interpretability and streamline analysis, we divide distances into five categories based on predefined thresholds (Table 1). This classification offers a clear and consistent framework for analyzing distance measures. For human-human interactions, distances are computed between their 3D bounding boxes (Martin-Martin et al. 2021). In human-object scenarios, we use point cloud data and object annotations (Le et al. 2024), applying efficient nearest neighbor searches via *KD-trees*¹. A KD-tree is built for each point cloud pair to perform k-nearest neighbor queries, and the mean neighbor distance quantifies proximity.

Experiment

With the rapid advancement of VLMs and multi-modal architectures capable of integrating heterogeneous modalities with increasingly sophisticated reasoning abilities, our work aims to provide a systematic and rigorous evaluation of their performance across multiple levels of task complexity. In particular, we focus on visual reasoning tasks, encompassing both VG and VQA, to examine the models' reasoning and comprehension capabilities in multi-modal settings. To complement automated evaluation metrics, we conduct a human evaluation with three expert annotators, reporting the mean of their scores to ensure robustness and reliability. Furthermore, we introduce a standardized evaluation toolkit, carefully designed to maintain consistent difficulty levels and enable a fair comparative analysis across different VLMs.

Evaluation Metrics. For VG, we adopt the evaluation metric from prior research (Sur'is, Menon, and Vondrick 2023; Ke et al. 2024; Cai et al. 2025), where a predicted bounding box is considered correct if its mIoU with the ground-truth box exceeds 0.5. For VQA, we use the accuracy metric as described in (Johnson et al. 2017; Jahangard et al. 2024; Hudson and Manning 2019).

Setup and Details. To evaluate the models, we utilize the JRDB-Reasoning dataset, which contains questions of varying difficulty. We categorize the dataset into three diffi-

¹https://en.wikipedia.org/wiki/K-d_tree

culty levels, denoted by \mathcal{D}_1 , \mathcal{D}_2 , and \mathcal{D}_3 , defined as $\mathcal{D}_1 = \{(S, R, T) \mid S = 1, R = 0, T = 1\}$, $\mathcal{D}_2 = \{(S, R, T) \mid S = 2, R \geq 1, T = 1\}$, $\mathcal{D}_3 = \{(S, R, T) \mid S = 3, R \geq 2, T = 1\}$, where S denotes the number of nodes (humans or objects), R represents the number of interactions (edges), and T indicates the number of time slots. \mathcal{D}_1 includes questions about a single node ($S = 1$) and its attributes, with no interactions ($R = 0$) and a single time slot ($T = 1$). \mathcal{D}_2 increases complexity by involving two nodes ($S = 2$) and at least one interaction ($R \geq 1$), still with $T = 1$. \mathcal{D}_3 requires reasoning over three nodes ($S = 3$) and at least two interactions ($R \geq 2$), with $T = 1$. We exclude levels beyond \mathcal{D}_3 as current models struggle with higher complexities in image and video reasoning. To ensure a focused evaluation, we conduct experiments using a single time slot ($T = 1$). This setup aligns with the design of many existing models, such as video grounding methods, which are typically developed for single-object tracking rather than multi-object temporal interactions.

Visual Grounding. We selected several advanced models and evaluated them in zero-shot model, including GroundingDINO (Liu et al. 2024b), Florence-v2 (Xiao et al. 2023), YOLO-World (Cheng et al. 2024b), and Qwen2.5-VL (Bai et al. 2025). As shown in Table 2, we observe that as the value of D increases—indicating more nodes (entities) and connections—the complexity of the questions also rises, leading to a decline in model performance. Among the models, Florence-V2, Qwen2.5-VL and Grounding DINO consistently demonstrated the best and most stable performance at all difficulty levels. In contrast, YOLO-World dropped in performance with increased complexity, which could be due to their inability to handle more intricate or ambiguous image-text relationships, leading to a decline in performance as the tasks became more challenging.

The human evaluation row shows scores given by human evaluators to assess the performance of multi-modal LLMs. The scores (95.5, 87.6, and 86.6) reflect how well the models are perceived, with the highest score (95.5) being for \mathcal{D}_1 where it is single node ($S = 1$) and $T = 1$ and there is no interaction, which likely involves the easiest questions. The decrease in scores with \mathcal{D}_2 and \mathcal{D}_3 may be due to added complexity in R , $R \geq 1$ and $R \geq 2$, more interactions such as geometric questions and challenges with stitched images, where the asked person being on either side of the image could make some questions harder for the evaluator.

Step-by-Step Evaluation. We evaluate state-of-the-art

Multi-Modal LLM	\mathcal{D}_1	\mathcal{D}_2	\mathcal{D}_3
	(S=1, R=0, T=1)	(S=2, R≥1, T=1)	(S=3, R≥2, T=1)
Florence-V2	37.6	35.7	25.2
GroundingDINO	25.8	22.5	14.7
Qwen2.5-VL	19.8	17.4	14.0
YOLO-World	13.5	5.3	3.1
Human Eval	95.5	87.6	86.6

Table 2: VG Experiment: VLM evaluation on varying difficulty levels using mIoU on JRDB-Reasoning.

Multi-Modal LLM	\mathcal{D}_1	\mathcal{D}_2	\mathcal{D}_3
	(S=1, R=0, T=1)	(S=2, R≥1, T=1)	(S=3, R≥2, T=1)
InternVL 2.5	49.9	49.2	44.6
Paligemma	48.6	47.5	44.6
Qwen2.5-VL	36.3	33.0	27.7
LLaVA-NeXT	35.4	33.2	28.7
Human Eval	96.8	88.5	87.1

Table 3: VQA Experiment: Accuracy of VLMs on JRDB-Reasoning across difficulty levels.

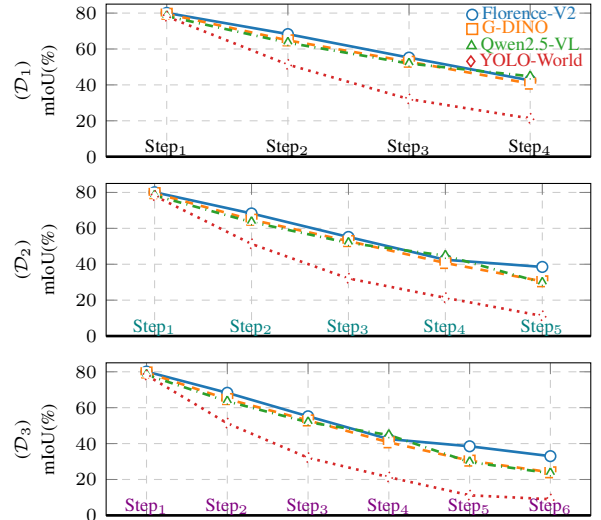


Figure 5: Step-by-Step Evaluation of VLMs on VG Task

models on the VG task using the JRDB-Reasoning benchmark, designed with progressively increasing compositional complexity. In \mathcal{D}_1 , queries involve individual-level attributes—[“human”, “gender”, “age”, “action”]—evaluated in four steps: detecting a person, then filtering by gender, age, and action (e.g., “*Find a young female person who is standing*”). \mathcal{D}_2 introduces relational reasoning to previous steps (e.g., “*...chatting with a man*”), while \mathcal{D}_3 incorporates spatial relations (e.g., “*...with a white man located to her right*”). As complexity increases, model performance degrades consistently, see Figure 5. While all models perform well on basic detection and attribute filtering (steps 1–3), performance steadily declines as later steps introduce increasingly complex reasoning challenges. YOLO-World performs competitively at early stages but falters when relational and spatial comprehension is required. In contrast, models like Qwen 2.5-VL and Florence-2 exhibit a more gradual decline, demonstrating stronger integration of linguistic and visual understanding. Their ability to parse nuanced syntax and semantics allows more faithful alignment with complex human query.

Visual Question Answering. We conduct VQA experiments on images using state-of-the-art models, including Paligemma (Beyer et al. 2024), Qwen2.5-VL (Liu et al. 2024a), LLaVA-NeXT (Liu et al. 2024a) and InternVL

Multi-Modal LLM	\mathcal{D}_1	\mathcal{D}_2	\mathcal{D}_3
	(S=1, R=0, T=1)	(S=2, R \geq 1, T=1)	(S=3, R \geq 2, T=1)
InternVL 2.5	49.1	48.6	46.0
Qwen2.5-VL	47.3	41.8	36.5
LLaVA-NeXT-Video	41.2	36.8	29.7
Human Eval	98.3	92.1	89.1

Table 4: VQA Experiment: VLM evaluation on videos of varying difficulty levels using accuracy on JRDB-Reasoning.

2.5 (Zhang, Li, and Bing 2023), as shown in Table 3. As seen in the results, the performance follows the same trend observed with VG, where as the complexity D increases, model performance decreases. This is due to the increasing difficulty of the questions. Additionally, we observe that InternVL 2.5 performs better than the other models, likely due to its enhanced ability to handle complex image relationships, improving its understanding of object interactions, spatial reasoning, and context within the images. Qwen2.5-VL (Bai et al. 2025), InternVL 2.5 (Zhang, Li, and Bing 2023), and LLaVA-NeXT-Video (Zhang et al. 2024b) are evaluated for video. Consistent with the trends seen in models evaluated on VG, performance drops as question complexity increases. However, model performance on video is better, as they can analyze multiple frames to answer the questions. In summary, model performance decreases with more complex questions, and they still underperform compared to humans.

Conclusion

In this work, we define and formalize perception-to-reasoning task complexity to address the absence of a reasoning complexity definition in existing datasets. We develop an adaptive query engine that dynamically generates non-predefined questions with adjustable complexity, enabling dataset creation with controlled question difficulty and task customization. We also provide structured, step-by-step reasoning annotations. In addition, we labeled human-object interactions manually and computing geometric relationships, thereby enhancing the robotic JRDB dataset for integration with the query engine. Our engine and benchmark support fine-grained evaluation of visual reasoning frameworks and dynamic assessment of VLMs across multiple reasoning levels.

Acknowledgments

This work was supported by the DARPA Assured Neuro-Symbolic Learning and Reasoning (ANSR) program (FA8750-23-2-1016), the ONR Global X-Challenge Grant (N62909-25-1-2067), and the Australian Research Council Discovery Project ARC DP2020102427.

References

Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.;

Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Agrawal, P.; Antoniak, S.; Hanna, E. B.; Bout, B.; Chaplot, D.; Chudnovsky, J.; Costa, D.; De Monicault, B.; Garg, S.; Gervet, T.; et al. 2024. Pixtral 12B. *arXiv preprint arXiv:2410.07073*.

AI, A. Y.; Chen, B.; Li, C.; Huang, C.; Zhang, G.; Zhang, G.; Li, H.; Zhu, J.; Chen, J.; Chang, J.; et al. 2024. Yi: Open Foundation Models by 01. AI. *arXiv preprint arXiv:2403.04652*.

Akula, A.; Gella, S.; Al-Onaizan, Y.; Zhu, S.-C.; and Reddy, S. 2020. Words Aren’t Enough, Their Order Matters: On the Robustness of Grounding Visual Referring Expressions. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 6555–6565. Online: Association for Computational Linguistics.

Amizadeh, S.; Palangi, H.; Polozov, A.; Huang, Y.; and Koishida, K. 2020. Neuro-symbolic visual reasoning: Disentangling. In *International Conference on Machine Learning*, 279–290. PMLR.

Ataallah, K.; Shen, X.; Abdelrahman, E.; Sleiman, E.; Zhu, D.; Ding, J.; and Elhoseiny, M. 2024. Minigt4-video: Advancing multimodal llms for video understanding with interleaved visual-textual tokens. *arXiv preprint arXiv:2404.03413*.

Azuma, D.; Miyanishi, T.; Kurita, S.; and Kawanabe, M. 2022. Scanqa: 3d question answering for spatial scene understanding. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19129–19139.

Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923*.

Beyer, L.; Steiner, A.; Pinto, A. S.; Kolesnikov, A.; Wang, X.; Salz, D.; Neumann, M.; Alabdulmohsin, I.; Tschannen, M.; Bugliarello, E.; et al. 2024. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*.

Cai, Z.; Ke, F.; Jahangard, S.; Garcia de la Banda, M.; Haf-fari, R.; Stuckey, P. J.; and Rezaatofighi, H. 2025. NAVER: A Neuro-Symbolic Compositional Automaton for Visual Grounding with Explicit Logic Reasoning. *arXiv preprint arXiv:2502.00372*.

Chakraborty, R.; Sinha, A.; Reilly, D.; Govind, M. K.; Wang, P.; Bremond, F.; and Das, S. 2024. Llavidal: Benchmarking large language vision models for daily activities of living. *arXiv preprint arXiv:2406.09390*.

Chandrasegaran, K.; Gupta, A.; Hadzic, L. M.; Kota, T.; He, J.; Eyzaguirre, C.; Durante, Z.; Li, M.; Wu, J.; and Fei-Fei, L. 2024. Hourvideo: 1-hour video-language understanding. *arXiv preprint arXiv:2411.04998*.

Chen, C.; Anjum, S.; and Gurari, D. 2022. Grounding answers for visual questions asked by visually impaired people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19098–19107.

- Chen, J.; Wei, F.; Zhao, J.; Song, S.; Wu, B.; Peng, Z.; Chan, S.-H. G.; and Zhang, H. 2024a. Revisiting Referring Expression Comprehension Evaluation in the Era of Large Multimodal Models. *arXiv preprint arXiv:2406.16866*.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24185–24198.
- Cheng, A.-C.; Yin, H.; Fu, Y.; Guo, Q.; Yang, R.; Kautz, J.; Wang, X.; and Liu, S. 2025. Spatialrgpt: Grounded spatial reasoning in vision-language models. *Advances in Neural Information Processing Systems*, 37: 135062–135093.
- Cheng, T.; Song, L.; Ge, Y.; Liu, W.; Wang, X.; and Shan, Y. 2024a. YOLO-World: Real-Time Open-Vocabulary Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16901–16911.
- Cheng, T.; Song, L.; Ge, Y.; Liu, W.; Wang, X.; and Shan, Y. 2024b. Yolo-world: Real-time open-vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16901–16911.
- Du, Y.; Zhou, K.; Huo, Y.; Li, Y.; Zhao, W. X.; Lu, H.; Zhao, Z.; Wang, B.; Chen, W.; and Wen, J.-R. 2024. Towards event-oriented long video understanding. *arXiv preprint arXiv:2406.14129*.
- Ehsanpour, M.; Saleh, F.; Savarese, S.; Reid, I.; and Rezatofighi, H. 2022. Jrd-act: A large-scale dataset for spatio-temporal action, social group and activity detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20983–20992.
- Fu, C.; Dai, Y.; Luo, Y.; Li, L.; Ren, S.; Zhang, R.; Wang, Z.; Zhou, C.; Shen, Y.; Zhang, M.; et al. 2025. Videomme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 24108–24118.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6904–6913.
- Hudson, D. A.; and Manning, C. D. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6700–6709.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Jahangard, S.; Cai, Z.; Wen, S.; and Rezatofighi, H. 2024. JRDB-Social: A Multifaceted Robotic Dataset for Understanding of Context and Dynamics of Human Interactions Within Social Groups. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22087–22097.
- Jang, Y.; Song, Y.; Yu, Y.; Kim, Y.; and Kim, G. 2017. Tgifqa: Toward spatio-temporal reasoning in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2758–2766.
- Johnson, J.; Hariharan, B.; Van Der Maaten, L.; Fei-Fei, L.; Lawrence Zitnick, C.; and Girshick, R. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2901–2910.
- Kazemzadeh, S.; Ordonez, V.; Matten, M.; and Berg, T. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 787–798.
- Ke, F.; Cai, Z.; Jahangard, S.; Wang, W.; Haghighi, P. D.; and Rezatofighi, H. 2024. HYDRA: A Hyper Agent for Dynamic Compositional Visual Reasoning. In *European Conference on Computer Vision*, 132–149. Springer.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123: 32–73.
- Le, D. T.; Gou, C.; Datta, S.; Shi, H.; Reid, I.; Cai, J.; and Rezatofighi, H. 2024. JRDB-PanoTrack: An Open-world Panoptic Segmentation and Tracking Robotic Dataset in Crowded Human Environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22325–22334.
- Lei, J.; Yu, L.; Bansal, M.; and Berg, T. L. 2018. Tvqa: Localized, compositional video question answering. *arXiv preprint arXiv:1809.01696*.
- Li, B.; Zhang, Y.; Chen, L.; Wang, J.; Yang, J.; and Liu, Z. 2023a. Otter: A multi-modal model with in-context instruction tuning. *arXiv preprint arXiv:2305.03726*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023b. 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. 2023c. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*.
- Li, K.; Wang, Y.; He, Y.; Li, Y.; Wang, Y.; Liu, Y.; Wang, Z.; Xu, J.; Chen, G.; Luo, P.; et al. 2024. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22195–22206.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; et al. 2022. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10965–10975.
- Li, M.; Chen, X.; Zhang, C.; Chen, S.; Zhu, H.; Yin, F.; Yu, G.; and Chen, T. 2023d. M3DBench: Let’s Instruct Large Models with Multi-modal 3D Prompts. *arXiv preprint arXiv:2312.10763*.

- Li, S.; Li, L.; Ren, S.; Liu, Y.; Liu, Y.; Gao, R.; Sun, X.; and Hou, L. 2023e. Vitatecs: A diagnostic dataset for temporal concept understanding of video-language models. *arXiv preprint arXiv:2311.17404*.
- Lin, Z.; Liu, C.; Zhang, R.; Gao, P.; Qiu, L.; Xiao, H.; Qiu, H.; Lin, C.; Shao, W.; Chen, K.; et al. 2023. Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models. *arXiv preprint arXiv:2311.07575*.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26296–26306.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024b. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, 38–55. Springer.
- Liu, Y.; Li, S.; Liu, Y.; Wang, Y.; Ren, S.; Li, L.; Chen, S.; Sun, X.; and Hou, L. 2024c. Tempcompass: Do video llms really understand videos? *arXiv preprint arXiv:2403.00476*.
- 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*.
- Mangalam, K.; Akshulakov, R.; and Malik, J. 2023. Egoschema: A diagnostic benchmark for very long-form video language understanding. *Advances in Neural Information Processing Systems*, 36: 46212–46244.
- Martin-Martin, R.; Patel, M.; Rezatofighi, H.; Sheno, A.; Gwak, J.; Frankel, E.; Sadeghian, A.; and Savarese, S. 2021. Jrdp: A dataset and benchmark of egocentric robot visual perception of humans in built environments. *IEEE transactions on pattern analysis and machine intelligence*.
- Nguyen, N.; Bi, J.; Vosoughi, A.; Tian, Y.; Fazli, P.; and Xu, C. 2024. OSCaR: Object State Captioning and State Change Representation. *arXiv preprint arXiv:2402.17128*.
- Patraucean, V.; Smaira, L.; Gupta, A.; Recasens, A.; Markeeva, L.; Banarse, D.; Koppula, S.; Malinowski, M.; Yang, Y.; Doersch, C.; et al. 2023. Perception test: A diagnostic benchmark for multimodal video models. *Advances in Neural Information Processing Systems*, 36: 42748–42761.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, 2641–2649.
- Qian, T.; Chen, J.; Zhuo, L.; Jiao, Y.; and Jiang, Y.-G. 2024. NuScenes-QA: A Multi-Modal Visual Question Answering Benchmark for Autonomous Driving Scenario. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4542–4550.
- Ren, M.; Kiros, R.; and Zemel, R. 2015. Exploring models and data for image question answering. *Advances in neural information processing systems*, 28.
- Shen, X.; Xiong, Y.; Zhao, C.; Wu, L.; Chen, J.; Zhu, C.; Liu, Z.; Xiao, F.; Varadarajan, B.; Bordes, F.; et al. 2024. Longvu: Spatiotemporal adaptive compression for long video-language understanding. *arXiv preprint arXiv:2410.17434*.
- Subramanian, S.; Merrill, W.; Darrell, T.; Gardner, M.; Singh, S.; and Rohrbach, A. 2022. ReCLIP: A Strong Zero-Shot Baseline for Referring Expression Comprehension. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5198–5215.
- Sur’s, D.; Menon, S.; and Vondrick, C. 2023. ViperGPT: Visual Inference via Python Execution for Reasoning. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 11854–11864.
- Team, G.; Anil, R.; Borgeaud, S.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; Millican, K.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Tong, S.; Brown, E.; Wu, P.; Woo, S.; Middepogu, M.; Akula, S. C.; Yang, J.; Yang, S.; Iyer, A.; Pan, X.; et al. 2024. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*.
- Vendrow, E.; Le, D. T.; Cai, J.; and Rezatofighi, H. 2023. Jrdp-pose: A large-scale dataset for multi-person pose estimation and tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4811–4820.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; et al. 2024a. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Wang, X.; Zhang, X.; Luo, Z.; Sun, Q.; Cui, Y.; Wang, J.; Zhang, F.; Wang, Y.; Li, Z.; Yu, Q.; et al. 2024b. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*.
- Wu, B.; Yu, S.; Chen, Z.; Tenenbaum, J. B.; and Gan, C. 2024. Star: A benchmark for situated reasoning in real-world videos. *arXiv preprint arXiv:2405.09711*.
- Wu, C.; Yin, S.; Qi, W.; Wang, X.; Tang, Z.; and Duan, N. 2023a. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*.
- Wu, S.; Fei, H.; Qu, L.; Ji, W.; and Chua, T.-S. 2023b. NExT-GPT: Any-to-Any Multimodal LLM. *arXiv preprint arXiv:2309.05519*.
- Xiao, B.; Wu, H.; Xu, W.; Dai, X.; Hu, H.; Lu, Y.; Zeng, M.; Liu, C.; and Yuan, L. 2023. Florence-2: Advancing a unified representation for a variety of vision tasks (2023). URL <https://arxiv.org/abs/2311.06242>.
- Xiao, B.; Wu, H.; Xu, W.; Dai, X.; Hu, H.; Lu, Y.; Zeng, M.; Liu, C.; and Yuan, L. 2024. Florence-2: Advancing a Unified Representation for a Variety of Vision Tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4818–4829.

- Xu, L.; Zhao, Y.; Zhou, D.; Lin, Z.; Ng, S. K.; and Feng, J. 2024. Pllava: Parameter-free llava extension from images to videos for video dense captioning. *arXiv preprint arXiv:2404.16994*.
- Yang, H.; Chaisorn, L.; Zhao, Y.; Neo, S.-Y.; and Chua, T.-S. 2003. VideoQA: question answering on news video. In *Proceedings of the eleventh ACM international conference on Multimedia*, 632–641.
- Ye, Q.; Xu, H.; Ye, J.; Yan, M.; Hu, A.; Liu, H.; Qian, Q.; Zhang, J.; and Huang, F. 2024. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13040–13051.
- Yin, Z.; Wang, J.; Cao, J.; Shi, Z.; Liu, D.; Li, M.; Huang, X.; Wang, Z.; Sheng, L.; Bai, L.; et al. 2023. Lamm: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark. *Advances in Neural Information Processing Systems*, 36: 26650–26685.
- You, H.; Sun, R.; Wang, Z.; Chen, L.; Wang, G.; Ayyubi, H. A.; Chang, K.-W.; and Chang, S.-F. 2023. IdealGPT: Iteratively Decomposing Vision and Language Reasoning via Large Language Models. *arXiv preprint arXiv:2305.14985*.
- Yu, L.; Poirson, P.; Yang, S.; Berg, A. C.; and Berg, T. L. 2016. Modeling context in referring expressions. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, 69–85. Springer.
- Zhan, Y.; Yuan, Y.; and Xiong, Z. 2024. Mono3DVG: 3D Visual Grounding in Monocular Images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6988–6996.
- Zhang, H.; Li, X.; and Bing, L. 2023. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*.
- Zhang, H.; Wang, Y.; Tang, Y.; Liu, Y.; Feng, J.; Dai, J.; and Jin, X. 2024a. Flash-vstream: Memory-based real-time understanding for long video streams. *arXiv preprint arXiv:2406.08085*.
- Zhang, Y.; Li, B.; Liu, H.; Lee, Y.; Gui, L.; Fu, D.; Feng, J.; Liu, Z.; and Li, C. 2024b. LLaVA-Next: A Strong Zero-Shot Video Understanding Model. *arXiv preprint arXiv:2406.06579*.
- Zhou, J.; Shu, Y.; Zhao, B.; Wu, B.; Xiao, S.; Yang, X.; Xiong, Y.; Zhang, B.; Huang, T.; and Liu, Z. 2024. Mlvu: A comprehensive benchmark for multi-task long video understanding. *arXiv preprint arXiv:2406.04264*.
- Zhu, C.; Wang, T.; Zhang, W.; Chen, K.; and Liu, X. 2024. ScanReason: Empowering 3D Visual Grounding with Reasoning Capabilities. In *European Conference on Computer Vision (ECCV)*.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.
- Zhu, Y.; Groth, O.; Bernstein, M.; and Fei-Fei, L. 2016. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4995–5004.