

CoSPlan: Corrective Sequence Planning via Scene Graph Incremental Updates

Shresth Grover[†], Priyank Pathak, Akash Kumar, and Yogesh S Rawat

UCF Institute of Artificial Intelligence, University of Central Florida (UCF)

shgrover@ucsd.edu {priyank, akash.kumar, yogesh}@ucf.edu

Project Page: https://shroglck.github.io/cos_plan/

Dataset: <https://huggingface.co/datasets/shrg7/COSPLAN>

Abstract. Vision Language Models (VLMs) have shown promising planning capabilities, yet their success remains confined to the text domain, leaving visual decision-making relatively underexplored. Addressing this gap, we introduce **Corrective Sequence Planning (CoSPlan)** benchmark, where VLMs must plan a sequence of visual actions from an initial scene to a target scene. **CoSPlan** evaluates models on their ability to imagine and execute a coherent set of visual steps required to reach the goal (**Step Completion**). To prevent any shortcuts that simply describe the final scene, we introduce an erroneous action in decision making, which must be *detected* (**Error Detection**) and *corrected* to reach the goal, enabling a deeper understanding of the task. **CoSPlan** spans across 4 tasks: maze navigation, block re-arrangement, image reconstruction, and object re-organization. Despite using advanced reasoning strategies such as *Chain-of-Thought* and *Scene Graphs*, VLMs struggle on **CoSPlan**, while still showing promising performance in the text domain. Addressing this, we propose **Scene Graph Incremental updates (SGI)**, a novel training-free method to transform images into ‘*textual*’ scene graphs, enabling step-by-step reasoning through iterative scene graph refinement. SGI yields an average of $\simeq 4.4\%$ $\uparrow$ on **CoSPlan** w/ generalization on PlanBench and VQA. *Link for solving puzzles on the project page.*

Keywords: Benchmark · Reasoning · LLMs · VLMs · Planning

1 Introduction

Humans are uniquely equipped to handle visual-language information to solve complex problems by predicting future steps under physical constraints. In AI, Large-scale Vision Language Models (VLMs) [1, 8] replicate a similar mechanism in the textual domain [21, 55]. Using techniques such as Chain-of-Thought (CoT) and deep-thinking paradigms, VLMs can generate a sequence of textual steps enabling reasoning over long sequences of instructions. This success in text-based decision making raises a critical but **vastly unexplored question**: “*What happens when this sequential thought process must occur in the visual domain?*”

[†]work done as intern




Fig. 1: CoSPlan : (Left) Given the initial *multi-step instructions* (yellow), VLM is asked to complete the remaining steps to reach blue dot. Errors include starting at the wrong position (green dot shifted by 1 cell), and choosing the path because it’s the first option among all possible paths. (Right) The model sees the initial & final states, along with an initial context of already performed actions (w/ an **error**). VLM identifies the error and picks the set of remaining actions (**correcting error**) to reach goal.

VLM’s strong zero-shot generalization is ideal for complex visual reasoning, a prominent feature in real-world decision-making (*e.g.* robotics, autonomous navigation *etc.*). These models can accurately predict (imagine) some future scenes, yet they struggle to sustain coherent multi-step visual reasoning (Fig. 1 (left)). Moreover, explicitly visualizing (generating) future scenes is computationally expensive, making naive simulation impractical for real-world deployment.

Recent benchmarks focus on simple visual puzzles for evaluating logic reasoning [26, 29, 64], but lack complex decision making. In contrast, decision-making problems like *sequential planning* require a step-by-step plan to reach a goal [32, 37, 56]. However, research on such planning tasks is overwhelmingly confined to the textual domain [39, 55], with limited-to-no exploration in vision [47]. Bridging this gap, we focus on *vision*-based logical reasoning guided by *textual* instructions for decision making, involving *temporal* sequence of actions. Our **Corrective Sequence Planning** (CoSPlan) is designed to evaluate VLMs’ reasoning abilities to complete tasks through multi-step visual actions. Starting from an initial scene and a sequence of actions, the model must understand the context and determine a plan that correctly completes the remaining steps to reach the goal (**Step Completion**). This requires not only understanding the scene and its constraints, but also updating the scene after each action.

Unlike ideal settings with perfect instructions [43], CoSPlan reflects realistic conditions by introducing an *intentional erroneous step* in the action sequence. Models are then judged on their ability to perform **Error Detection**, requiring models to identify the incorrect step, correct it, and reach the goal. This not only *ensures a genuine understanding of the task*, where models must *scrutinize each and every action*, but also **prevents cheating** [15, 45], where models can’t simply select an option that matches the final state (bypassing true reasoning). Errors also reflect practical scenarios where long sequences of actions may introduce errors. An example of CoSPlan is shown in Fig. 1 (right).

Table 1: Reasoning Benchmarks: Temporal updates the scene with action. The multi-step tasks don’t complete after one action. Error is erroneous decision-making.

Benchmark	Modality			Testing Environment		Task	
	Vision	Text	Temporal	Synthetic	Real World	Multi-step	Error
ALFWorld [48]		✓	✓	✓		✓	
PlanBench [55]		✓	✓	✓		✓	
WebArena [72]		✓	✓	✓		✓	
SpatialEval [58]	✓	✓		✓	✓		
SURDS [19]	✓	✓	✓		✓		
CoSPlan (<i>Ours</i>)	✓	✓	✓	✓	✓	✓	✓

CoSPlan introduces 4 error-prone planning tasks: 1) **Maze-E** for navigation, 2) **Blocks-World-E** for block re-arrangement, 3) **Shuffle-E** for image reconstruction, and 4) **Robo-VQA-E** for real-world re-organization. We evaluate several VLMs, including *GPT-4o* and GPT-5.1 [1], *CoG-VLM* [61], *Intern-VLM* variants [9, 62], *Janus-Pro-7B* [8], *Qwen* variants [5, 60], and Llama3-8B [17]. Their vanilla performance is near random guessing, indicating the difficulty of the benchmark. Hence, two deeper reasoning strategies are applied: Chain-of-Thought (CoT) [63] to encourage step-wise reasoning, and Scene Graph (SG) [7] to structure object positions and attributes. This setup highlights the benchmark’s complexity and provides a rigorous test of visual reasoning ability.

While CoT & SG are known to perform well on error-free planning tasks, we show that they struggle on CoSPlan. They compress reasoning from initial to final state into a single step. This forces models to implicitly simulate (*internalize*) long action sequences without explicit intermediate representations. To address this, we introduce **SGI** (**Scene Graph Incremental updates**), a training-free method that represents *images as text-based Scene Graphs* and *refines them step-by-step*, explicitly generating intermediate states. Thus, SGI i) reduces long horizon reasoning burden, ii) tracks evolving scenes, and iii) detects and corrects errors, making corrective sequence planning significantly more robust.

In summary, we make the following contributions: i) **CoSPlan** (**Corrective Sequence Planning**) is the first multimodal benchmark with temporal sequences of actions in vision + language domain to evaluate VLMs visual reasoning in error-prone scenarios. CoSPlan includes four planning tasks to test the abilities of *Error Detection* and *Step Completion* to reach a desired goal. ii) We evaluate VLMs and reveal weaknesses in error handling, vision-based sequence planning, and understanding. iii) We propose **SGI**, a **Scene Graph Incremental** technique, that refines representations of scene for every action, enhancing robustness in CoSPlan and datasets such as VQA [58], and PlanBench [55].

2 Related work

Reasoning in VLMs Enhancing Reasoning without fine-tuning [10] is of great interest in VLMs. Wei *et al.* [63] introduced the Chain-of-Thought (CoT) via a

Table 2: CoSPlan Dataset Details: ‘Size’ are the number of image-text pairs, ‘Initial Context’ is the *average* number of actions already performed, and ‘Remaining Steps’ is the *average* number of additional steps on top of the initial context, required to reach the goal. The ‘Source’ is where images (or text) are taken from.

Dataset	Task	Type	Size	Initial Context	Remaining Steps	Source
Maze-E	Navigation	Path Planning	5000	2.0	4.6	Synthetic
Blocks-World-E	Re-arrange	Blocks	5000	2.0	3.8	Synthetic
Shuffle-E	Re-construct	Puzzle	1000	3.7	7.1	ImageNet [14]
Robo-VQA-E	Re-organize	Real-world	350	5.5	4.1	ROM [44]

series of reasoning steps to guide LLMs. CoT shows significant gain on mathematical reasoning tasks [51]; however, it *lacks spatial relationships*, making it unfit for visual reasoning. Chen *et al.* [7] proposed using Scene Graphs (SG) as structured representations, to improve VLM’s reasoning abilities on visual grounding tasks like VQA [13], image generation [24], spatial reasoning [28], *etc.* However, SG tries to hallucinate an entire sequence in one go, facing challenges in erroneous multi-step visual reasoning. Our **SGI** approach converts images into textual graphs, which VLMs can then iterate over to generate intermediate scenes like CoT, without relying on external annotations (unlike [40]).

Sequential Planning Valmeekam *et al.* [55] proposed variations of Sequence Planning, including completing steps based on a partial context. Most sequential planning datasets [3, 31, 43, 44, 67] assume ideal instructions, which may not hold outside the lab environment. Many rely on human / video-based supervision [6, 12, 69], limiting scalability. Most benchmarks focus on textual domain [4, 39, 66, 71], with limited exploration in the vision domain [11, 47, 54] (Tab. 1). SpatialEval [58] only evaluates static images with no temporal change of states. Recent works have also examined VLMs for planning [20, 22, 25, 41, 68]. CoSPlan is the first benchmark to evaluate VLMs on sequential planning under *vision-language and temporal domain* with error-prone instructions.

Explainable AI & Analysis Explaining behavior VLMs in the real-world has been quite a well-researched topic [27, 35, 36, 49, 50]. However, agentic models like GPT [1] and Qwen [60] are fairly recent, and their explanations of decision-making are still in their infancy [18, 30, 34]. In this work, we aim to study how these AI models make reasoning decisions when faced with erroneous step completion tasks. We leverage *single-choice MCQ* for this, as it aids controlled analysis of VLM decision-making. This further aids easy analysis of failure cases and serves as motivation for our SGI algorithm,

3 CoSPlan Benchmark

Corrective Sequence Planning (CoSPlan) mimics general decision-making by evaluating the model’s ability to navigate a complex challenge of detecting and correcting non-optimal (error) steps in a sequential planning task. In this setup,




Fig. 2: CoSPlan : We provide the initial ($\mathcal{I}_0$) and the final state ($\mathcal{I}_g$), with an initial set of already performed actions (initial context in orange, visualized as black arrows). We introduce *one error* that either violates the rules of the environment or is sub-optimal; however, one error in Shuffle-E would have a cascading effect, requiring multiple backtracks for correction, hence ignored (Sec. 3.1). VLMs perform: *Step Completion*, choose the correct set of future steps (green) to reach $\mathcal{I}_g$, and *Error Detection*, detecting the erroneous step (red) in the initial context. Coordinates (row, column) are 0-indexed. Safeguard against cheating (Sec. 3.3) highlighted in pink for Blocks-World-E.

- Model $\mathcal{M}$ progresses from an initial state $\mathcal{I}_0$ to a goal state $\mathcal{I}_g$ through a sequence of N actions: $(\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_N)$.
- We introduce an intentional non-optimal (error) action $\mathcal{A}_{\mathcal{E}}$ within the *initial context (already performed k actions)* $(\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_{\mathcal{E}}, \dots, \mathcal{A}_N)$.
- The model must *detect* this erroneous action $\mathcal{A}_{\mathcal{E}}$ and course-correct to *complete* the remaining actions steps $(\mathcal{A}_{k+1}, \mathcal{A}_{k+2}, \dots, \mathcal{A}_N)$ towards the final goal. Mathematically, it can be shown as

$$\begin{array}{l}
 \text{CoSPlan} \\
 \mathcal{M}(\mathcal{A}_{1,\dots,\mathcal{E},\dots,k}; \mathcal{I}_0; \mathcal{I}_g) \\
 \text{Initial State: } \mathcal{I}_0 \\
 \text{Goal: } \mathcal{I}_g \\
 \text{Performed actions: } \mathcal{A}_{1,\dots,\mathcal{E},\dots,k}
 \end{array}
 \rightarrow
 \begin{cases}
 \text{Error Detection} \\
 \text{Identify } \mathcal{A}_{\mathcal{E}} \\
 \text{Step Complete} \\
 \mathcal{A}_{k+1,k+2,\dots,N}
 \end{cases}
 \quad (1)$$

This setup is used to solve diverse scenarios, such as *re-constructing* a correct image from shuffled image tiles, *re-arranging* & *re-organizing* the objects / blocks into a coherent order (obeying physics), and *navigating* through a maze. Success relies on addressing and resolving non-optimal (errors) steps encountered along the way (Fig. 2). Sec. 3.1 describes each dataset with proposed planning tasks.

What’s an Error? We loosely define ‘error’ as a plausible but suboptimal action that deviates from the optimal path to the goal, potentially resulting in longer sequences. An error can also be a purely wrong action that makes it impossible to reach the goal without correction, *e.g.* referencing non-existent objects, violating task/physics constraints *etc.* Presence of erroneous action forces models to reason over *all actions* / *intermediate states* to identify and correct the error.

3.1 Benchmark Datasets

We introduce four sequential planning datasets, each featuring diverse tasks with intentional sub-optimal errors (except ‘Shuffle-E’) posing unique challenges in corrective sequence planning. Similar to previous planning benchmarks [16, 18, 53, 57], we opt for multiple-choice questions [59] to test: i) **Error Detection**: Identifying the non-optimal erroneous action from initial context (already performed actions), or selecting “none of the above”. ii) **Step Completion**: Selecting the correct answer among 5 options that would correct the mistake and lead to the final goal. The use of synthetic datasets [38, 58] has been shown to test reasoning vulnerabilities in VLMs. An overview is provided in Tab. 2 & Fig. 2, respectively. Intuition behind this formulation explained in *Supplementary*.

Maze-E (Fig. 2a): The *goal* is to solve a maze while navigating from the start cell (green, $\mathcal{I}_0$) to the goal cell (blue $\mathcal{I}_g$). *Input* is a maze layout with $\mathcal{I}_0$, $\mathcal{I}_g$, initial sequence of moves under the constraints of moving *only* up, down, left, or right, within maze boundaries, while *avoiding* red cells. Dataset is constructed via randomly sampled grid of size $\in [3 \times 3, 8 \times 8]$, and up to 5 obstacles¹.

Blocks-World-E (Fig. 2b): It’s a block rearrangement task where the objective is to transform an initial configuration ($\mathcal{I}_0$) into a target configuration ($\mathcal{I}_g$). *Input* includes $\mathcal{I}_0$, $\mathcal{I}_g$ arrangements along with a partially executed stacking sequence under the constraint of moving *only* top blocks (nothing on top) and place it either on another top block or empty column. The dataset is generated using OpenCV, with 3 to 8 blocks randomly distributed across columns.

Shuffle-E (Fig. 2c): Here, the objective is to recover (reconstruct) the original image by swapping shuffled image patches. *Input* is a starting shuffled image ($\mathcal{I}_0$), the final restored image ($\mathcal{I}_g$), and the initial sequence of image patch swaps. Dataset contains 1000 uniformly sampled ImageNet [14] images across all classes.

Robo-VQA-E (Fig. 2d): Goal is to perform a sequence of actions to re-organize real-world objects in an initial scene ($\mathcal{I}_0$) to reach desired arrangement ($\mathcal{I}_g$). The dataset (modified ROM [44]) contains 350 curated image-text pairs. *Inputs*: starting scene, target scene, and a sequence of initial object placement actions.

3.2 Erroneous Step $\mathcal{A}_\epsilon$

We introduce one error step in the initial context. For **Maze-E**, we randomly select one of the following incorrect moves: i) move out of bounds, ii) move

¹ Black & white pattern helps distinguish cells/navigate, using OpenCV




Fig. 3: Error Correction: Initial context has erroneous move from (1,0) to red cell (2,1). True option corrects this error (yellow), while the one without (pink) isn't correct.

into an obstacle (red cell), and iii) move diagonally. For **Blocks-World-E** error involves randomly picking *one block* and randomly picking one of the wrong columns, either i) stacking inefficiently ii) placing blocks in physically impossible positions, such as floating in the air or between blocks, or iii) moving non-top blocks. In **Robo-VQA-E**, we use Qwen 2 VL-8B (curated by us) to annotate the scene and pick two random objects/containers, *e.g.* pen, dustbin, to generate errors. These include: 1) Suboptimal rearrangement *e.g.* “place obj1 (pen) in cont2 (dustbin)” 2) Arbitrary interaction with objects outside the scene *e.g.* “break obj1 using obj2” 3) Arbitrary or incorrect placements, “cover cont1”. 3) Invalid actions *e.g.* “opening an already open door”. For **Shuffle-E** an erroneous patch swap will cascade into subsequent wrong swaps, making it *infeasible* to maintain uniformity across other one-error settings.

3.3 Safeguard against Shortcuts and Cheating

Revealing the final scene I_g enables shortcutting: the model can simply choose the option that matches I_g without reasoning about the initial state or constraints. To prevent this, we add a distractor that also reaches I_g but **omits** the required error-correction step. Conversely, correct selection requires understanding the scene to rectify actions, rather than relying on final-state visual similarity. Fig. 2b ‘Step Completion’ shows option 2 (green) and option 3 (pink) both match I_g , but only option 2 corrects the error of moving block 4 from the 2nd column to 1st column. Similarly, in Fig. 3, the option 2 (yellow) reverts the illegal diagonal move, while option 1 simply reaches the target.

3.4 Evaluation

CoSPlan is evaluated using a multiple-choice question (MCQ) framework (shown in Fig. 2), with Top-1(%) accuracy as the evaluation metric². We *independently* evaluate **Step Completion**: the model chooses *one correct option from 5 options* (random accuracy $\frac{1}{5}$), and **Error Detection**: MCQ setup presents *initial context actions as choices*, with an additional option of ‘none of the above’ denoting no error present (random accuracy = $\mathbb{E}[\frac{1}{\text{Initial context length}+1}]$).

² Higher number implies better performance $\uparrow$.




Fig. 4: (Left) *Chain of Thought*: All VLMs are asked to iterate actions step-by-step (context) and verify if they reach the goal. (Right) *Scene Graph*: GPT-4o describes objects as *nodes*, location as *edges*, state as *attributes*. Detailed description in Sec. 3.5.

3.5 Models & Techniques

Following the OpenVLM leaderboard and the recent *Capture* [38] benchmark, we evaluate CoSP1an on GPT-5.1, GPT-4o, CoG-VLM, Intern-VLM 2, Intern-VLM 3, Janus-Pro-7B, Llama3-8B, Qwen2-VL-8B, and Qwen3-VL-8B. As GPT-4o and GPT-5.1 are paid models, we use them for limited experiments. We further examine enhanced reasoning strategies Chain-of-Thought (CoT) [63] and Scene Graphs (SG) [7] to improve VLM performance. Detailed comparisons between SG and CoT, along with model specifics, are provided in *Supplementary*.

Chain-of-Thought (CoT) CoT [46, 63] has 3 steps i) *Identify*: VLMs are provided with ‘additional’ details of the problem and constraints; ii) *Context*: Step-by-step description (enforcing) of each action (from the initial context); iii) *Verify*: ‘additional’ verification that the VLMs follow all constraints and steps. This approach is model-agnostic, and the same set of instructions is provided to all models. An example is shown in Fig. 4 (left) and in *Supplementary*.

Scene Graphs (SG) SG adds spatial awareness to CoT, aiding reasoning abilities. We ‘QUERY’ (prompt) VLM to construct SG for the initial ($\mathcal{I}_0$) and goal states ($\mathcal{I}_g$), capturing three key components: a) Objects in the scene (*nodes*), b) Attributes (*properties*) and interactions between objects, and c) Spatial relationships (*edges*) capturing relative positions. An example SG is shown in Fig. 4 (right) and *Supplementary*. Different VLM’s have different structural representations (model-specific SG), however, we ensure uniform prompt for all VLMs.

3.6 Results & Benchmark Analysis

We conduct our analysis on the **Step Completion** task using CoT unless mentioned otherwise. Key insights are highlighted. Additional analyses and extended discussions are provided in the *Supplementary*.

Benchmark: Tab. 3 & 4 compares VLMs on CoSP1an, via Vanilla method (raw image-text input), CoT and SG. SG outperforms CoT, which in turn outperforms vanilla models (few exceptions), highlighting need for structured representations in visual reasoning. Proprietary very large reasoning models like GPT-4o

Table 3: CoSPlan Step Completion: ‘V’ denotes vanilla (no CoT or SG). Both initial & final SG(s) are input to VLM. † means vanilla evaluation skipped because of financial constraints (CoT & SG generally outperform vanilla anyways); **Lighter shade** indicates $\pm 2\%$ close to random choice; **Darker** is intentionally choosing wrong ($<$ random).

VLM	Robo-VQA-E (%)			Shuffle-E (%)			Maze-E (%)			Blocks-World-E (%)		
	V	CoT	SG	V	CoT	SG	V	CoT	SG	V	CoT	SG
Random	20			20			20			20		
Human	42.1			53.7			95.7			81.8		
Llama3-8B [17]	18.1	18.3	19.1	17.3	17.7	18.5	19.5	20.1	21.3	21.3	22.7	23.2
CoG-VLM [61]	13.1	12.5	21.5	23.1	27.1	23.7	25.1	25.9	26.5	25.5	25.2	26.7
Janus-pro-7B [8]	14.1	14.7	21.3	23.2	23.1	23.5	20.4	20.2	21.7	24.2	23.1	25.1
Qwen2 VL-8B [60]	17.1	17.6	18.9	24.1	24.9	25.1	26.5	27.9	28.3	18.1	18.6	18.8
Qwen3 VL-8B [5]	20.3	21.2	23.4	28.3	29.4	30.3	35.1	34.6	36.2	26.5	27.9	28.3
Intern-VLM 2 [9]	22.1	23.5	25.1	20.1	23.2	23.4	21.6	35.8	41.2	18.3	21.2	18.9
Intern-VLM 3 [62]	27.1	29.4	31.4	24.1	25.6	27.1	48.6	48.8	50.1	25.7	27.1	29.4
GPT-4o [1] †	-	48.2	52.2	-	27.6	30.1	-	45.6	46.1	-	49.7	54.3
GPT-5.1 [1] †	-	49.3	51.3	-	31.5	33.4	-	49.1	48.3	-	48.1	52.3
Gemini-3-pro [52] †	-	61.6	67.3	-	57.4	61.6	-	68.2	70.4	-	69.5	71.3

Table 4: CoSPlan Error Detection: Format same as Tab. 3. Models look at the initial context and try to identify the erroneous action.

VLM	Robo-VQA-E			Maze-E			Blocks-World-E		
	V	CoT	SG	V	CoT	SG	V	CoT	SG
Random	25.4			26.1			26.1		
Human	60.0			90.0			80.0		
Llama3-8B [17]	11.2	10.7	13.6	18.7	18.3	19.7	27.3	27.6	28.2
CoG-VLM [61]	32.1	33.4	35.3	6.4	8.4	13.3	41.3	43.1	44.5
Janus-pro-7B [8]	17.5	18.1	26.1	20.5	19.1	21.0	29.3	31.0	27.6
Qwen2 VL-8B [60]	9.2	9.1	9.6	20.5	20.8	20.7	32.3	30.6	35.2
Qwen3 VL-8B [5]	13.1	15.6	18.1	23.8	24.1	25.8	34.1	35.3	36.7
Intern-VLM 2 [9]	24.3	25.2	26.1	32.8	33.1	33.4	36.5	37.9	37.3
Intern-VLM 3 [62]	25.1	26.6	28.1	33.3	34.7	35.1	42.4	43.1	44.3
GPT-4o [1] †	-	45.3	44.2	-	40.3	35.3	-	35.1	42.1
GPT-5.1 [1] †	-	54.9	46.3	-	39.8	37.7	-	38.3	44.6
Gemini-3-pro [52] †	-	57.3	62.5	-	62.6	67.8	-	67.5	71.8

and GPT-5.1 make reasonably informed decisions while relatively light-weight & open-source ones **perform near or below random chance**, indicating the inability to visualize (imagine, predict) intermediate steps. Accuracy less than random (dark red) can partially be explained by overwhelmingly picking certain options [70] (Janus and Qwen 2 select ‘option A’ 90%+ and 75%+ times) or cheats to reach goal without error correction; For SG, InternVLM-2 cheats 43%, InternVLM-3 41%, CoG-VLM 27%, and Qwen3 35%. Based on human scores, the difficulty for tasks follows: Robo-VQA-E > Shuffle-E > Blocks-World-E > Maze-E. The near-random performance is *only* on one error. Generating problems with multiple cascading errors will have exponential complexity, limited by automation (*e.g.* GPT-4o) not equipped to handle even one error.




Fig. 5: (Left) Error-free Eval VLMs excel in error-free settings, but struggle on error-prone ones. (Mid) Out vs In-context: Errors from within context (scene) are harder than random ones (out-context). **(Right) Multi-modal vs Text-only** VLMs perform exceptionally well when simulating text-based steps but struggle on visual reasoning.




Fig. 6: (Left) MCQ options count Intern-VLM (SG) accuracy drops with the # of MCQ options, hinting at randomness. **(Mid) MCQ context** Step Completion with 9 remaining steps shows only k steps towards the goal. Constant accuracy indicates VLMs are not simulating the MCQ option to reach the target. **(Right) Length of Initial Context** VLM accuracy shows a positive correlation with the number of already performed steps.

Impact of Error: Fig. 5 left) shows VLMs excel in error-free settings (GPT-4o near-perfect) but struggle on error-prone ones, hinting at lack of reasoning, possibly picking options that describes the target (cheating Sec. 3.3). These highlight the need for challenging, error-prone benchmarks like CoSPlan to expose the gap between clean training and practical error-prone scenarios, while demanding logical visual reasoning (CoG-VLM and Janus-pro-7B near random predictions).

Effect of Context: CoSPlan includes a random mix of two types of errors: i) *In-Context*: Erroneous step involving objects present in the scene, ii) *Out-Context*: Error uses random objects not in scene (e.g. basketball, hammer, etc). Figure 5 mid) shows lower performance on In-Context errors, suggesting VLMs struggle more when erroneous actions involve *plausible* objects from within the scene while they can handle the out-of-context errors with relative ease.

Visual vs Text: Similar to previous works [23, 42], describing our tasks (vision+text) in text-only format significantly boosts reasoning accuracy (Fig. 5 right)). This further exposes models inability to iteratively visualize intermediate reasoning steps, despite performing well in purely text based ones (Fig. 7). Qwen2’s near-random prediction in Blocks-World-E reveals its limitations in logical reasoning. Tasks like Robo-VQA-E and Shuffle-E cannot be faithfully represented as text-only without visual aid.




Fig. 7: Failure in Simulating Visual Steps: Proprietary VLMs like Grok [65], Claude [2], ChatGPT [33] fail to generate intermediate states starting from the initial scene. On Blocks-World-E, Grok misplaces the red block, Claude moves purple block to the left instead of right, and ChatGPT fails to generate correct steps.

MCQ Options: Fig. 6 left) shows that increasing the number of MCQ options drops the performance exposing the randomness in picking options (first reported by [18]). This reveals the added logical complexity of MCQ in CoSPlan.

Context from MCQ: Fig. 6 mid) evaluates step completion with 9 remaining steps to reach $\mathcal{I}_g$, under a fixed context length of 2. The MCQ reveals only the next k steps toward $\mathcal{I}_g$ (e.g., $k = 2$ exposes only $\frac{2}{9}$ steps and does not reach the goal). All models (including GPT-4o) maintain nearly constant accuracy across k , indicating an inability to iterate over visual steps in MCQ options. Ideally, accuracy should improve as more steps toward the goal are revealed. Instead, performance suggests blind option bias (e.g., Janus selects ‘A’ 94% of the time) and decisions driven solely by the initial context.

Initial Context importance: Fig. 6 right) shows that model accuracy increases with longer initial context. We hypothesize that including more already-performed, *non-negotiable* steps brings the intermediate state closer to the goal, thereby reducing the burden of predicting the remaining steps.

Can models visualize actions: Given initial, and final state, VLMs are expected to internalize steps between the two states. Fig. 7 visualizes how models imagine future steps (states). Models may succeed on certain steps but fail on subsequent ones, revealing brittle incoherent long-horizon reasoning, partially explaining the random prediction observed on CoSPlan.

4 Scene Graph Incremental update (SGI)

VLMs near random performance on CoSPlan (Tab. 3) fails to leverage the intermediate steps (Fig. 6 (mid)). With only initial and final images, models internally interpolate missing states (CoT & SG), a process they struggle in the visual domain, even though they perform with relative ease in the text domain




Fig. 8: SGI: 1) Initial and Goal Scene Graphs (SG) are generated. 2) Incremental Scene Update sequentially modifies SG for each action A_i 3) Similarity Comparison matches the resultant SG with Goal graph for searching for the best-aligned sequence.

(Fig. 5 (right)). Incorporating these, we propose **Scene Graph Incremental update (SGI)** to extend Scene Graph reasoning. We leverage SG to transform visual images into textual representations, and apply an *iterative step-by-step reasoning on text space instead of the visual domain*. Instead of reasoning within a single static scene, SGI explicitly derives next-time-frame scene graphs as actions unfold in evolving scenes (Fig. 8). This incremental formulation bridges the gap between the initial and final states via *explicit intermediate states*, and decomposing reasoning into smaller transitions, thereby improving corrective sequence planning and error detection. The hierarchy can be visualized as $\text{CoT} \subseteq \text{SG} \subseteq \text{SGI}$.

4.1 Algorithm

An overview (Fig. 8) and pseudo code for step completion is shown in Algorithm 1. More details and SGI for error detection in *Supplementary*.

1) Vanilla Scene Graphs (SG): We ‘QUERY’ VLMs to generate the Scene Graphs for the initial state I_0 as S_0 and the final goal I_g as S_g (Sec. 3.5). We evaluated the performance of these vanilla SGs in Tab. 3 as $SG = [S_0, S_g]$

2) Incremental Scene Update: Starting from the initial SG (S_0), we feed a textual description of each action ($\mathcal{A}_1, \dots, \mathcal{A}_k$) to VLM and ask it ‘SIMULATE’ the action on the SG producing intermediate SG (S_c). The ‘SIMULATE’ prompt to VLM: “*Simulate the given action sequence from the initial state, incrementally updating the scene graph.*” modifying nodes, attributes, and edges of SG. We use this initial context SG S_C (same for all MCQ options) to ‘SIMULATE’ each MCQ option independently producing for final SG S_m for the ‘m-th’ option.

3) Similarity Comparison: After simulating each action, VLM is asked to compare ‘SIMILARITY’ between resultant SG S_m and goal SG S_g , via “*Compare the resulting scene graph with the goal scene graph to identify incorrect relationships, misplaced objects, or unmet constraints and score them between 0-100.*”. The option that the VLMs think is most similar (S_m and S_g) is chosen as its prediction. *Note:* SG(s) are model-specific internal representations; hence compared via model-as-a-judge, as there are no universal metrics for all models.

Algorithm 1 SGI for Step Completion (Sec. 4)**Input:** Initial state $\mathcal{I}_0$, Goal state $\mathcal{I}_g$, Initial Context actions $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_{k-1}, \mathcal{A}_k, \dots, \mathcal{A}_N$ **Objective:** Best option from MCQ for $\mathcal{A}_{k+1}, \mathcal{A}_{k+2}, \dots, \mathcal{A}_N$ **Require:** VLM $\mathcal{M}$, Step Completion MCQ m options.

```

1) Initial and goal Vanilla Scene Graph, (ref Section 3.5)
1:  $S_0 \leftarrow \text{QUERY}[\mathcal{M}(\mathcal{I}_0)] \parallel S_g \leftarrow \text{QUERY}[\mathcal{M}(\mathcal{I}_g)]$ 

2) Incremental Scene Update ( $S_0 \rightarrow S_c \rightarrow S_m$ )
2:  $S_c \leftarrow S_0$ 
3: for  $\mathcal{A}_i$  in  $[\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_{k-1}, \mathcal{A}_k, \dots, \mathcal{A}_N]$  do
4:    $S_c \leftarrow \text{SIMULATE}[\mathcal{M}(S_c, \mathcal{A}_i)]$  //  $\mathcal{M}$  simulates i-th action  $\mathcal{A}_i$  to iteratively
                                     update intermediate context SG  $S_c$ 
5: end for
6: for each option  $m \in \text{MCQ}$  do
7:    $\mathcal{A}_{k+1}^m, \mathcal{A}_{k+2}^m, \dots, \mathcal{A}_N^m \leftarrow m \parallel S_m \leftarrow S_c$  // actions & context SG  $S_c$  for option  $m$ 
8:   for  $\mathcal{A}_i^m$  in  $[\mathcal{A}_{k+1}^m, \mathcal{A}_{k+2}^m, \dots, \mathcal{A}_N^m]$  do
9:      $S_m \leftarrow \text{SIMULATE}[\mathcal{M}(S_m, \mathcal{A}_i^m)]$  // iteratively simulate  $m^{th}$  option actions
10:  end for
11: end for

3) Similarity : VLM determines similarity between MCQs derived SG  $S_m$  and  $S_g$ 
12:  $m' \leftarrow \arg \max_{m \in \text{MCQ}} \text{SIMILARITY}[\mathcal{M}(S_m, S_g)]$ 
13: Output:  $m'$ 

```

4.2 Results

CoSPlan Results: Tab. 5 shows that SGI consistently outperforms vanilla Scene Graph across all tasks. For Step Completion, SGI yields an average gain of 4.4%, with improvements up to 10.3%, 4.8%, and 10.0% for Intern-VLM 2, Intern-VLM-3, and GPT-4o. For Error Detection, gains average 4.1%, 1.7%, and 9.2% respectively. On Blocks-World-E, *Gemini-2.5-pro* [52] improves from 67% with CoT to 70% with SG to 71.5% with SGI. Although SGI requires one VLM call per step, iterating over text-based SG is far cheaper than synthesizing intermediate images. Total compute cost is the length of the initial context + number of MCQ options $\times$ Avg. number of steps per option. Without stepwise updates, performance is near random for Intern-VLM 2. The additional compute yields up to 13% improvement in error detection, demonstrating its effectiveness on CoSPlan. SGI has fewer cheating instances: Intern-VLM-2 (43% vs 39%), Intern-VLM-3 (41% vs 35%), Qwen-3 (35% vs 30%), and CoG-VLM (27% vs 23%).

Scene Graph Quality: Tab. 6 analyzes the impact of scene graph quality on SGI. We extract ground-truth (GT, ‘Oracle’) graphs for Block-World-E and Maze-E and compare them against a noisy setting, where the VLM perturbs object positions to introduce structural errors in scene graph. The results show that Oracle SG consistently outperforms VLM-generated and noisy inputs, highlighting the perception challenge: the difficulty in accurately understanding scenes. However, even with Oracle SG, performance remains below text-only reasoning,

Table 5: Scene Graph Incremental update: SGI improves SG, same naming convention as Tab. 3. Blocks-World-E shortened to Blocks-E, and Robo-VQA-E to Robo-E. Other models with SGI are evaluated in the *Supplementary*.

Method	Step Completion (% $\uparrow$)								Error Detection (% $\uparrow$)					
	Robo-E		Shuffle-E		Maze-E		Blocks-E		Robo-E		Maze-E		Blocks-E	
	SG	SGI	SG	SGI	SG	SGI	SG	SGI	SG	SGI	SG	SGI	SG	SGI
Random	20		20		20		20		25.4		26.1		26.1	
Intern-VLM 2	25.1	32.1	23.4	25.2	41.2	43.2	18.9	29.2	26.1	31.5	33.4	34.8	37.3	42.9
Intern-VLM 3	31.4	33.7	27.1	28.6	50.1	54.8	29.4	30.6	28.1	29.5	35.1	36.3	44.3	45.1
GPT-4o	52.2	56.4	30.1	37.0	46.1	56.1	54.3	55.3	44.2	57.4	35.3	41.1	42.1	50.7

Table 6: Oracle and Noisy SG: Robo-VQA-E has no Oracle (GT) SG.

	SG	Shuffle-E		Maze-E		Blocks-E	
		SG	SGI	SG	SGI	SG	SGI
Intern-2	Noisy	20.3	23.4	38.5	41.3	15.3	25.4
	VLM	23.4	25.2	41.2	43.2	18.9	29.2
	GT	23.7	26.5	42.0	59.0	28.1	34.1
Qwen-3	Noisy	27.3	30.6	30.5	43.6	24.3	28.7
	VLM	30.3	31.4	36.2	44.6	28.3	29.5
	GT	30.7	31.8	40.5	55.4	30.3	35.6

Table 7: SGI on VQA [58]. SGI generalization on VQA tasks.

Method	Spatial-Map			Maze-Nav			Spatial-Grid		
	CoT	SG	SGI	CoT	SG	SGI	CoT	SG	SGI
CoG-VLM	25.1	36.7	35.8	32.3	32.4	31.2	30.1	34.3	38.2
Janus-pro-7B	42.4	47.4	47.8	20.8	27.3	29.3	34.4	35.8	36.3
Intern-VLM	36.3	41.3	44.3	28.6	40.5	42.1	33.3	33.8	35.1

indicating that **simulating visual steps** is inherently difficult. Furthermore, the higher robustness of SGI for noisy SG in Qwen 3 demonstrates that the model is not blindly picking option A anymore. This is because each option is independently iterated upon to derive the best option, rather than simply asking VLM to pick one.

External Dataset: Unlike sequence planning, VQA [58] is static and lacks temporal structure. In our setup, the same image serves as both initial and final state, and the VLM simulates each MCQ option for feasibility. Iteratively evaluating options outperforms single-shot SG or CoT (Tab. 7). We also evaluate the text-only PlanBench benchmark [55], specifically Blocksworld Task 8, where the goal changes after the initial context, requiring the model to adapt its plan. Unlike MCQ settings, PlanBench requires generating a full action sequence that is validated by the planning environment. Qwen2-VL-8B with SGI on textual scene graphs achieves the best performance (Tab. 8). Other PlanBench tasks focus on plan generation and are out of scope. Details in *Supplementary*.

MCQ Alternative: Tab. 9 shows the performance of **error detection task** in an open text generation. Given initial context, instead of picking an MCQ option, VLMs are prompted to write the erroneous step as a question-answering task, judged by Qwen-2 8B for similarity between ground truth and generated text. **SGI** outperforms baselines, indicating the generalizability of the algorithm.


Error-Free Scenarios: Fig. 9 shows SGI outperforming SG on both error-prone and error-free sequence planning, highlighting the generalizability of SGI.

Table 8: Planbench [55]
Constrained plan generation

Method	Variant	Task 8 Score ($\uparrow$)
Qwen2	Vanilla	13.8
VL-8B	CoT	14.1
	SG	13.9
<i>(our)</i>	SGI	14.7

Table 9: Text vs
MCQ

Intern-VLM-2	Blocks-World-E		
	Variant	MCQ	Text
Intern-VLM-2	CoT	37.9	40.0
	SG	37.3	38.0
	SGI	42.9	44.0

**Fig. 9: Error Free**

5 Conclusion and Future Work

We introduce **CoSPlan**, a benchmark for evaluating VLMs on error-prone sequential planning from images with text instructions. **CoSPlan** exposes critical weaknesses in VLMs: near-random behavior with even a single error, poor handling of in-context mistakes, and a bias toward text-based reasoning over grounded multimodal understanding. Even strong models such as **GPT-4o** and **GPT-5.1** fail to reliably use visual context to recover correct plans. We propose **SGI** (Scene Graph Incremental Update), which refines scene representations step-by-step after each action. SGI substantially boosts performance across error-prone, error-free, and VQA settings compared to vanilla scene graphs, highlighting its robustness and ability to enhance corrective sequence planning in VLMs.

Our current setup uses static image pairs. Extending **CoSPlan** to video-based planning and to an interactive setting where agents execute actions and observe updated states are natural future directions.

6 Limitations

Our work isolates exactly one erroneous step per sequence to facilitate controlled evaluation. While this design aids in failure analysis, we acknowledge that the single error constraint is somewhat artificial. Real-world alignment requires evaluating models in more complex, multi-error settings.

Acknowledgements

The authors thank Steven Dick (UCF High-Performance Computing) and Vibhav Vineet (Microsoft Research) for their help and contributions to this project. This research has benefited from the Microsoft Accelerating Foundation Models Research (AFMR) grant program.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)

2. Anthropic: Claude sonnet 4.6. <https://www.anthropic.com> (2024), large language model
3. Asai, M.: Photo-realistic blocksworld dataset (2018), <https://arxiv.org/abs/1812.01818>
4. Asai, M., Kajino, H., Fukunaga, A., Muise, C.: Classical planning in deep latent space (2022), <https://arxiv.org/abs/2107.00110>
5. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., Ge, W., Guo, Z., Huang, Q., Huang, J., Huang, F., Hui, B., Jiang, S., Li, Z., Li, M., Li, M., Li, K., Lin, Z., Lin, J., Liu, X., Liu, J., Liu, C., Liu, Y., Liu, D., Liu, S., Lu, D., Luo, R., Lv, C., Men, R., Meng, L., Ren, X., Ren, X., Song, S., Sun, Y., Tang, J., Tu, J., Wan, J., Wang, P., Wang, P., Wang, Q., Wang, Y., Xie, T., Xu, Y., Xu, H., Xu, J., Yang, Z., Yang, M., Yang, J., Yang, A., Yu, B., Zhang, F., Zhang, H., Zhang, X., Zheng, B., Zhong, H., Zhou, J., Zhou, F., Zhou, J., Zhu, Y., Zhu, K.: Qwen3-vl technical report (2025), <https://arxiv.org/abs/2511.21631>
6. Bouhsain, S.A., Alami, R., Simeon, T.: Learning to predict action feasibility for task and motion planning in 3d environments. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 3736–3742. IEEE (2023)
7. Chen, M., Li, Y., et al.: Measuring and improving chain-of-thought reasoning in vision-language models. arXiv preprint arXiv:2309.04461 (2023), <https://arxiv.org/abs/2309.04461>
8. Chen, X., Wu, Z., Liu, X., Pan, Z., Liu, W., Xie, Z., Yu, X., Ruan, C.: Janus-pro: Unified multimodal understanding and generation with data and model scaling. arXiv preprint arXiv:2501.17811 (2025)
9. Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., et al.: 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. pp. 24185–24198 (2024)
10. Cheng, A., et al.: Spatialrgpt: Grounded spatial reasoning in vision language model. OpenReview (2024)
11. Chow, W., Mao, J., Li, B., Seita, D., Guizilini, V., Wang, Y.: Physbench: Benchmarking and enhancing vision-language models for physical world understanding. arXiv preprint arXiv:2501.16411 (2025)
12. Crockett, B., Mueller, C.L., Hayes, B.: Human demonstrations enable efficient solutions to sequential manifold planning problems. In: Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction. pp. 800–809 (2025)
13. Damodaran, V., Chakravarthy, S., Kumar, A., Umapathy, A., Mitamura, T., Nakashima, Y., Garcia, N., Chu, C.: Understanding the role of scene graphs in visual question answering (2021), <https://arxiv.org/abs/2101.05479>
14. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on. pp. 248–255. IEEE (2009), <https://ieeexplore.ieee.org/abstract/document/5206848/>
15. Eshuijs, L., Wang, S., Fokkens, A.: Short-circuiting shortcuts: Mechanistic investigation of shortcuts in text classification (2025), <https://arxiv.org/abs/2505.06032>
16. Fan, C., Liang, Y., Bhardwaj, S., Cobbina, K., Li, M., Zhou, T.: V-rex: Benchmarking exploratory visual reasoning via chain-of-questions (2025), <https://arxiv.org/abs/2512.11995>

17. Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Vaughan, A., Yang, A., Fan, A., Goyal, A., Hartshorn, A., Yang, A., Mitra, A., Sravankumar, A., Korenev, A., Hinsvark, A., Rao, A., Zhang, A., Rodriguez, A., Gregerson, A., Spataru, A., Roziere, B., Biron, B., Tang, B., Chern, B., Caucheteux, C., Nayak, C., Bi, C., Marra, C., McConnell, C., Keller, C., Touret, C., Wu, C., Wong, C., Ferrer, C.C., Nikolaidis, C., Allonius, D., Song, D., Pintz, D., Livshits, D., Wyatt, D., Esiobu, D., Choudhary, D., Mahajan, D., Garcia-Olano, D., Perino, D., Hupkes, D., Lakomkin, E., AlBadawy, E., Lobanova, E., Dinan, E., Smith, E.M., Radenovic, F., Guzmán, F., Zhang, F., Synnaeve, G., Lee, G., Anderson, G.L., Thattai, G., Nail, G., Mialon, G., Pang, G., Cucurell, G., Nguyen, H., Korevaar, H., Xu, H., Touvron, H., Zarov, I., Ibarra, I.A., Kloumann, I., Misra, I., Evtimov, I., Zhang, J., Copet, J., Lee, J., Geffert, J., Vranes, J., Park, J., Mahadeokar, J., Shah, J., van der Linde, J., Billock, J., Hong, J., Lee, J., Fu, J., Chi, J., Huang, J., Liu, J., Wang, J., Yu, J., Bitton, J., Spisak, J., Park, J., Rocca, J., Johnstun, J., Saxe, J., Jia, J., Alwala, K.V., Prasad, K., Upasani, K., Plawiak, K., Li, K., Heafield, K., Stone, K., El-Arini, K., Iyer, K., Malik, K., Chiu, K., Bhalla, K., Lakhota, K., Rantala-Yearly, L., van der Maaten, L., Chen, L., Tan, L., Jenkins, L., Martin, L., Madaan, L., Malo, L., Blecher, L., Landzaat, L., de Oliveira, L., Muzzi, M., Pasupuleti, M., Singh, M., Paluri, M., Kardaş, M., Tsimpoukelli, M., Oldham, M., Rita, M., Pavlova, M., Kambadur, M., Lewis, M., Si, M., Singh, M.K., Hassan, M., Goyal, N., Torabi, N., Bashlykov, N., Bogoychev, N., Chatterji, N., Zhang, N., Duchenne, O., Çelebi, O., Alrassy, P., Zhang, P., Li, P., Vasic, P., Weng, P., Bhargava, P., Dubal, P., Krishnan, P., Koura, P.S., Xu, P., He, Q., Dong, Q., Srinivasan, R., Ganapathy, R., Calderer, R., Cabral, R.S., Stojnic, R., Raileanu, R., Maheswari, R., Girdhar, R., Patel, R., Sauvestre, R., Polidoro, R., Sumbaly, R., Taylor, R., Silva, R., Hou, R., Wang, R., Hosseini, S., Chennabasappa, S., Singh, S., Bell, S., Kim, S.S., Edunov, S., Nie, S., Narang, S., Raparthy, S., Shen, S., Wan, S., Bhosale, S., Zhang, S., Vandenhende, S., Batra, S., Whitman, S., Sootla, S., Collot, S., Gururangan, S., Borodinsky, S., Herman, T., Fowler, T., Sheasha, T., Georgiou, T., Scialom, T., Speckbacher, T., Mihaylov, T., Xiao, T., Karn, U., Goswami, V., Gupta, V., Ramanathan, V., Kerkez, V., Gonguet, V., Do, V., Vogeti, V., Albiero, V., Petrovic, V., Chu, W., Xiong, W., Fu, W., Meers, W., Martinet, X., Wang, X., Wang, X., Tan, X.E., Xia, X., Xie, X., Jia, X., Wang, X., Goldschlag, Y., Gaur, Y., Babaei, Y., Wen, Y., Song, Y., Zhang, Y., Li, Y., Mao, Y., Coudert, Z.D., Yan, Z., Chen, Z., Papakipos, Z., Singh, A., Srivastava, A., Jain, A., Kelsey, A., Shajnfeld, A., Gangidi, A., Victoria, A., Goldstand, A., Menon, A., Sharma, A., Boesenberg, A., Baevski, A., Feinstein, A., Kallet, A., Sangani, A., Teo, A., Yunus, A., Lupu, A., Alvarado, A., Caples, A., Gu, A., Ho, A., Poulton, A., Ryan, A., Ramchandani, A., Dong, A., Franco, A., Goyal, A., Saraf, A., Chowdhury, A., Gabriel, A., Bharambe, A., Eisenman, A., Yazdan, A., James, B., Maurer, B., Leonhardi, B., Huang, B., Loyd, B., Paola, B.D., Paranjape, B., Liu, B., Wu, B., Ni, B., Hancock, B., Wasti, B., Spence, B., Stojkovic, B., Gamido, B., Montalvo, B., Parker, C., Burton, C., Mejia, C., Liu, C., Wang, C., Kim, C., Zhou, C., Hu, C., Chu, C.H., Cai, C., Tindal, C., Feichtenhofer, C., Gao, C., Civin, D., Beaty, D., Kreymer, D., Li, D., Adkins, D., Xu, D., Testuggine, D., David, D., Parikh, D., Liskovich, D., Foss, D., Wang, D., Le, D., Holland, D., Dowling, E., Jamil, E., Montgomery, E., Presani, E., Hahn, E., Wood, E., Le, E.T., Brinkman, E., Arcaute, E., Dunbar, E., Smothers, E., Sun, F., Kreuk, F., Tian, F., Kokkinos, F., Ozgenel, F., Caggioni, F., Kanayet, F., Seide, F., Florez, G.M., Schwarz, G., Badeer, G., Sweet, G., Halpern, G., Herman, G.,

- Sizov, G., Guangyi, Zhang, Lakshminarayanan, G., Inan, H., Shojanazeri, H., Zou, H., Wang, H., Zha, H., Habeeb, H., Rudolph, H., Suk, H., Aspegren, H., Goldman, H., Zhan, H., Damla, I., Molybog, I., Tufanov, I., Leontiadis, I., Veliche, I.E., Gat, I., Weissman, J., Geboski, J., Kohli, J., Lam, J., Asher, J., Gaya, J.B., Marcus, J., Tang, J., Chan, J., Zhen, J., Reizenstein, J., Teboul, J., Zhong, J., Jin, J., Yang, J., Cummings, J., Carvill, J., Shepard, J., McPhie, J., Torres, J., Ginsburg, J., Wang, J., Wu, K., U, K.H., Saxena, K., Khandelwal, K., Zand, K., Matosich, K., Veeraraghavan, K., Michelena, K., Li, K., Jagadeesh, K., Huang, K., Chawla, K., Huang, K., Chen, L., Garg, L., A, L., Silva, L., Bell, L., Zhang, L., Guo, L., Yu, L., Moshkovich, L., Wehrstedt, L., Khabsa, M., Avalani, M., Bhatt, M., Mankus, M., Hasson, M., Lennie, M., Reso, M., Groshev, M., Naumov, M., Lathi, M., Keneally, M., Liu, M., Seltzer, M.L., Valko, M., Restrepo, M., Patel, M., Vyatskov, M., Samvelyan, M., Clark, M., Macey, M., Wang, M., Hermoso, M.J., Metanat, M., Rastegari, M., Bansal, M., Santhanam, N., Parks, N., White, N., Bawa, N., Singhal, N., Egebo, N., Usunier, N., Mehta, N., Laptev, N.P., Dong, N., Cheng, N., Chernoguz, O., Hart, O., Salpekar, O., Kalinli, O., Kent, P., Parekh, P., Saab, P., Balaji, P., Rittner, P., Bontrager, P., Roux, P., Dollar, P., Zvyagina, P., Ratanchandani, P., Yuvraj, P., Liang, Q., Alao, R., Rodriguez, R., Ayub, R., Murthy, R., Nayani, R., Mitra, R., Parthasarathy, R., Li, R., Hogan, R., Battey, R., Wang, R., Howes, R., Rinott, R., Mehta, S., Siby, S., Bondu, S.J., Datta, S., Chugh, S., Hunt, S., Dhillon, S., Sidorov, S., Pan, S., Mahajan, S., Verma, S., Yamamoto, S., Ramaswamy, S., Lindsay, S., Lindsay, S., Feng, S., Lin, S., Zha, S.C., Patil, S., Shankar, S., Zhang, S., Zhang, S., Wang, S., Agarwal, S., Sajuyigbe, S., Chintala, S., Max, S., Chen, S., Kehoe, S., Satterfield, S., Govindaprasad, S., Gupta, S., Deng, S., Cho, S., Virk, S., Subramanian, S., Choudhury, S., Goldman, S., Remez, T., Glaser, T., Best, T., Koehler, T., Robinson, T., Li, T., Zhang, T., Matthews, T., Chou, T., Shaked, T., Vontimitta, V., Ajayi, V., Montanez, V., Mohan, V., Kumar, V.S., Mangla, V., Ionescu, V., Poenaru, V., Mihailescu, V.T., Ivanov, V., Li, W., Wang, W., Jiang, W., Bouaziz, W., Constable, W., Tang, X., Wu, X., Wang, X., Wu, X., Gao, X., Kleinman, Y., Chen, Y., Hu, Y., Jia, Y., Qi, Y., Li, Y., Zhang, Y., Zhang, Y., Adi, Y., Nam, Y., Yu, Wang, Zhao, Y., Hao, Y., Qian, Y., Li, Y., He, Y., Rait, Z., DeVito, Z., Rosnbrick, Z., Wen, Z., Yang, Z., Zhao, Z., Ma, Z.: The llama 3 herd of models (2024), <https://arxiv.org/abs/2407.21783>
18. Grover, S., Vineet, V., Rawat, Y.S.: Navigating hallucinations for reasoning of unintentional activities (2024), <https://arxiv.org/abs/2402.19405>
 19. Guo, X., Zhang, R., Duan, Y., He, Y., Nie, D., Huang, W., Zhang, C., Liu, S., Zhao, H., Chen, L.: Surds: Benchmarking spatial understanding and reasoning in driving scenarios with vision language models. In: NeurIPS (2025)
 20. Hao, S., Gu, Y., Ma, H., Hong, J.J., Wang, Z., Wang, D.Z., Hu, Z.: Reasoning with language model is planning with world model. arXiv preprint arXiv:2305.14992 (2023)
 21. Huang, J., Chen, X., Mishra, S., Zheng, H.S., Yu, A.W., Song, X., Zhou, D.: Large language models cannot self-correct reasoning yet. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=IkMD3fKBPQ>
 22. Huang, X., Liu, W., Chen, X., Wang, X., Wang, H., Lian, D., Wang, Y., Tang, R., Chen, E.: Understanding the planning of llm agents: A survey. arXiv preprint arXiv:2402.02716 (2024)
 23. Jain, A., Vatsa, M., Singh, R.: Words over pixels? rethinking vision in multimodal large language models. In: Kwok, J. (ed.) Proceedings of the Thirty-Fourth Inter-

- national Joint Conference on Artificial Intelligence, IJCAI-25. pp. 10481–10489. International Joint Conferences on Artificial Intelligence Organization (8 2025). <https://doi.org/10.24963/ijcai.2025/1164>, <https://doi.org/10.24963/ijcai.2025/1164>, survey Track
24. Johnson, J., Gupta, A., Fei-Fei, L.: Image generation from scene graphs (2018), <https://arxiv.org/abs/1804.01622>
 25. Kambhampati, S., Valmeekam, K., Guan, L., Verma, M., Stechly, K., Bhambri, S., Saldyt, L.P., Murthy, A.B.: Position: Llms can’t plan, but can help planning in llm-modulo frameworks. In: Forty-first International Conference on Machine Learning (2024)
 26. Kamoi, R., Zhang, Y., Das, S.S.S., Zhang, R.H., Zhang, R.: Visonlyqa: Large vision language models still struggle with visual perception of geometric information. arXiv preprint arXiv:2412.00947 (2024)
 27. Kazmierczak, R., Berthier, E., Frehse, G., Franchi, G.: Explainability and vision foundation models: A survey. *Information Fusion* **122**, 103184 (2025)
 28. Li, X., Guo, D., Liu, H., Sun, F.: Embodied semantic scene graph generation. In: 5th Annual Conference on Robot Learning (2021), <https://openreview.net/forum?id=FCoh40LZ1Gg>
 29. Lu, P., Bansal, H., Xia, T., Liu, J., Li, C., Hajishirzi, H., Cheng, H., Chang, K.W., Galley, M., Gao, J.: Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. arXiv preprint arXiv:2310.02255 (2023)
 30. Luo, H., Specia, L.: From understanding to utilization: A survey on explainability for large language models. arXiv preprint arXiv:2401.12874 (2024)
 31. Nagpal, K., Mehr, N.: Optimal robotic assembly sequence planning: A sequential decision-making approach (2025), <https://arxiv.org/abs/2310.17115>
 32. Nayak, S., Orozco, A.M., Have, M.T., Thirumalai, V., Zhang, J., Chen, D., Kapoor, A., Robinson, E., Gopalakrishnan, K., Ichter, B., Harrison, J., Mahajan, A., Balakrishnan, H.: MAP-THOR: Benchmarking long-horizon multi-agent planning frameworks in partially observable environments. In: Multi-modal Foundation Model meets Embodied AI Workshop @ ICML2024 (2024), <https://openreview.net/forum?id=ZygZN5egzy>
 33. OpenAI, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., Avila, R., Babuschkin, I., Balaji, S., Balcom, V., Baltescu, P., Bao, H., Bavarian, M., Belgum, J., Bello, I., Berdine, J., Bernadett-Shapiro, G., Berner, C., Bogdonoff, L., Boiko, O., Boyd, M., Brakman, A.L., Brockman, G., Brooks, T., Brundage, M., Button, K., Cai, T., Campbell, R., Cann, A., Carey, B., Carlson, C., Carmichael, R., Chan, B., Chang, C., Chantzis, F., Chen, D., Chen, S., Chen, R., Chen, J., Chen, M., Chess, B., Cho, C., Chu, C., Chung, H.W., Cummings, D., Currier, J., Dai, Y., Decareaux, C., Degry, T., Deutsch, N., Deville, D., Dhar, A., Dohan, D., Dowling, S., Dunning, S., Ecoffet, A., Eleti, A., Eloundou, T., Farhi, D., Fedus, L., Felix, N., Fishman, S.P., Forte, J., Fulford, I., Gao, L., Georges, E., Gibson, C., Goel, V., Gogineni, T., Goh, G., Gontijo-Lopes, R., Gordon, J., Grafstein, M., Gray, S., Greene, R., Gross, J., Gu, S.S., Guo, Y., Hallacy, C., Han, J., Harris, J., He, Y., Heaton, M., Heidecke, J., Hesse, C., Hickey, A., Hickey, W., Hoeschele, P., Houghton, B., Hsu, K., Hu, S., Hu, X., Huizinga, J., Jain, S., Jain, S., Jang, J., Jiang, A., Jiang, R., Jin, H., Jin, D., Jomoto, S., Jonn, B., Jun, H., Kaftan, T., Łukasz Kaiser, Kamali, A., Kanitscheider, I., Keskar, N.S., Khan, T., Kilpatrick, L., Kim, J.W., Kim, C., Kim, Y., Kirchner, J.H., Kiros, J., Knight, M., Kokotajlo, D., Łukasz Kondraciuk, Kondrich, A., Konstantinidis, A., Kosic, K., Krueger, G., Kuo, V., Lampe, M.,

- Lan, I., Lee, T., Leike, J., Leung, J., Levy, D., Li, C.M., Lim, R., Lin, M., Lin, S., Litwin, M., Lopez, T., Lowe, R., Lue, P., Makanju, A., Malfacini, K., Manning, S., Markov, T., Markovski, Y., Martin, B., Mayer, K., Mayne, A., McGrew, B., McKinney, S.M., McLeavey, C., McMillan, P., McNeil, J., Medina, D., Mehta, A., Menick, J., Metz, L., Mishchenko, A., Mishkin, P., Monaco, V., Morikawa, E., Mossing, D., Mu, T., Murati, M., Murk, O., Mély, D., Nair, A., Nakano, R., Nayak, R., Neelakantan, A., Ngo, R., Noh, H., Ouyang, L., O’Keefe, C., Pachocki, J., Paino, A., Palermo, J., Pantuliano, A., Parascandolo, G., Parish, J., Parparita, E., Passos, A., Pavlov, M., Peng, A., Perelman, A., de Avila Belbute Peres, F., Petrov, M., de Oliveira Pinto, H.P., Michael, Pokorny, Pokrass, M., Pong, V.H., Powell, T., Power, A., Power, B., Proehl, E., Puri, R., Radford, A., Rae, J., Ramesh, A., Raymond, C., Real, F., Rimbach, K., Ross, C., Rotsted, B., Roussez, H., Ryder, N., Saltarelli, M., Sanders, T., Santurkar, S., Sastry, G., Schmidt, H., Schnurr, D., Schulman, J., Selsam, D., Sheppard, K., Sherbakov, T., Shieh, J., Shoker, S., Shyam, P., Sidor, S., Sigler, E., Simens, M., Sitkin, J., Slama, K., Sohl, I., Sokolowsky, B., Song, Y., Staudacher, N., Such, F.P., Summers, N., Sutskever, I., Tang, J., Tezak, N., Thompson, M.B., Tillet, P., Tootoonchian, A., Tseng, E., Tuggle, P., Turley, N., Tworek, J., Uribe, J.F.C., Vallone, A., Vijayvergiya, A., Voss, C., Wainwright, C., Wang, J.J., Wang, A., Wang, B., Ward, J., Wei, J., Weinmann, C., Welihinda, A., Welinder, P., Weng, J., Weng, L., Wiethoff, M., Willner, D., Winter, C., Wolrich, S., Wong, H., Workman, L., Wu, S., Wu, J., Wu, M., Xiao, K., Xu, T., Yoo, S., Yu, K., Yuan, Q., Zaremba, W., Zellers, R., Zhang, C., Zhang, M., Zhao, S., Zheng, T., Zhuang, J., Zhuk, W., Zoph, B.: Gpt-4 technical report (2024), <https://arxiv.org/abs/2303.08774>
34. Palikhe, A., Yu, Z., Wang, Z., Zhang, W.: Towards transparent ai: A survey on explainable large language models. arXiv preprint arXiv:2506.21812 (2025)
 35. Pathak, P., Karuppasamy, M., Baranwal, A., Vyas, S., Rawat, Y.S.: Robust onion: Peeling open vocab object detectors under noise (2026), <https://arxiv.org/abs/2606.26734>
 36. Pathak, P., Marjit, S., Vyas, S., Rawat, Y.S.: LR0.FM: Low-Res Benchmark and Improving robustness for Zero-Shot Classification in Foundation Models. In: The Thirteenth International Conference on Learning Representations (2025), <https://openreview.net/forum?id=AsFxRSLtqR>
 37. Paul, S.K.: Sequential planning in large partially observable environments guided by llms (2023), <https://arxiv.org/abs/2312.07368>
 38. Pothiraj, A., Stengel-Eskin, E., Cho, J., Bansal, M.: Capture: Evaluating spatial reasoning in vision language models via occluded object counting. arXiv preprint arXiv:2504.15485 (2025)
 39. Ramakrishnan, S.K., Wijmans, E., Kraehenbuehl, P., Koltun, V.: Does spatial cognition emerge in frontier models? arXiv preprint arXiv:2410.06468 (2024)
 40. Rana, K., Haviland, J., Garg, S., Abou-Chakra, J., Reid, I., Suenderhauf, N.: Sayplan: Grounding large language models using 3d scene graphs for scalable task planning. In: 7th Annual Conference on Robot Learning (2023), <https://openreview.net/forum?id=wMpOM0OSs7a>
 41. Rossetti, N., Tummolo, M., Gerevini, A.E., Putelli, L., Serina, I., Chiari, M., Olivato, M.: Learning general policies for planning through gpt models. Proceedings of the International Conference on Automated Planning and Scheduling **34**(1), 500–508 (May 2024). <https://doi.org/10.1609/icaps.v34i1.31510>, <https://ojs.aaai.org/index.php/ICAPS/article/view/31510>

42. Schulze Buschoff, L.M., Akata, E., Bethge, M., Schulz, E.: Visual cognition in multimodal large language models. *Nature Machine Intelligence* **7**(1), 96–106 (2025)
43. Sener, F., Chatterjee, D., Shelepov, D., He, K., Singhanian, D., Wang, R., Yao, A.: Assembly101: A large-scale multi-view video dataset for understanding procedural activities. *CVPR 2022* (2022)
44. Sermanet, P., Ding, T., Zhao, J., Xia, F., Dwibedi, D., Gopalakrishnan, K., Chan, C., Dulac-Arnold, G., Maddineni, S., Joshi, N.J., et al.: Robovqa: Multimodal long-horizon reasoning for robotics. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 645–652. IEEE (2024)
45. Shaib, C., Suriyakumar, V.M., Sagun, L., Wallace, B.C., Ghassemi, M.: Learning the wrong lessons: Syntactic-domain spurious correlations in language models (2026), <https://arxiv.org/abs/2509.21155>
46. Shao, H., Qian, S., Xiao, H., Song, G., Zong, Z., Wang, L., Liu, Y., Li, H.: Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning (2024), <https://arxiv.org/abs/2403.16999>
47. Shojaei, P., Mirzadeh, I., Alizadeh, K., Horton, M., Bengio, S., Farajtabar, M.: The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity (2025), <https://ml-site.cdn-apple.com/papers/the-illusion-of-thinking.pdf>
48. Shridhar, M., Yuan, X., Côté, M.A., Bisk, Y., Trischler, A., Hausknecht, M.: Alf-world: Aligning text and embodied environments for interactive learning. *arXiv preprint arXiv:2010.03768* (2020)
49. Shu, D., Zhao, H., Hu, J., Liu, W., Payani, A., Cheng, L., Du, M.: Large vision-language model alignment and misalignment: A survey through the lens of explainability. In: Christodoulopoulos, C., Chakraborty, T., Rose, C., Peng, V. (eds.) *Findings of the Association for Computational Linguistics: EMNLP 2025*. pp. 1713–1735. Association for Computational Linguistics, Suzhou, China (Nov 2025). <https://doi.org/10.18653/v1/2025.findings-emnlp.90>, <https://aclanthology.org/2025.findings-emnlp.90/>
50. Sim, M.Y., Zhang, W.E., Dai, X., Fang, B.: Can VLMs actually see and read? a survey on modality collapse in vision-language models. In: Che, W., Nabende, J., Shutova, E., Pilehvar, M.T. (eds.) *Findings of the Association for Computational Linguistics: ACL 2025*. pp. 24452–24470. Association for Computational Linguistics, Vienna, Austria (Jul 2025). <https://doi.org/10.18653/v1/2025.findings-acl.1256>, <https://aclanthology.org/2025.findings-acl.1256/>
51. Sprague, Z., Yin, F., Rodriguez, J.D., Jiang, D., Wadhwa, M., Singhal, P., Zhao, X., Ye, X., Mahowald, K., Durrett, G.: To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning. *arXiv preprint arXiv:2409.12183* (2024)
52. Team, G., Anil, R., Borgeaud, S., Alayrac, J.B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A.M., Hauth, A., Millican, K., et al.: Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* (2023)
53. Tran, D.T., Tran, T.K., Hauswirth, M., Le Phuoc, D.: Reasonvqa: A multi-hop reasoning benchmark with structural knowledge for visual question answering. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 18793–18803 (October 2025)
54. Valmeekam, K., Marquez, M., Sreedharan, S., Kambhampati, S.: On the planning abilities of large language models - a critical investigation. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) *Advances in Neural Information Processing Systems*. vol. 36, pp. 75993–76005. Curran Associates,

- Inc. (2023), https://proceedings.neurips.cc/paper_files/paper/2023/file/efb2072a358cefb75886a315a6fcf880-Paper-Conference.pdf
55. Valmeekam, K., Marquez, M., Sreedharan, S., Kambhampati, S.: On the planning abilities of large language models-a critical investigation. *Advances in Neural Information Processing Systems* **36**, 75993–76005 (2023)
 56. Wang, C., Deng, Y., Lyu, Z., Zeng, L., He, J., Yan, S., An, B.: Q*: Improving multi-step reasoning for llms with deliberative planning (2024), <https://arxiv.org/abs/2406.14283>
 57. Wang, J., Ming, Y., Shi, Z., Vineet, V., Wang, X., Li, S., Joshi, N.: Is a picture worth a thousand words? delving into spatial reasoning for vision language models. *Advances in Neural Information Processing Systems* **37**, 75392–75421 (2024)
 58. Wang, J., Ming, Y., Shi, Z., Vineet, V., Wang, X., Li, Y., Joshi, N.: Is a picture worth a thousand words? delving into spatial reasoning for vision language models (2024), <https://arxiv.org/abs/2406.14852>
 59. Wang, J., Ming, Y., Shi, Z., Vineet, V., Wang, X., Li, Y., Joshi, N.: Is a picture worth a thousand words? delving into spatial reasoning for vision language models. In: Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., Zhang, C. (eds.) *Advances in Neural Information Processing Systems*. vol. 37, pp. 75392–75421. Curran Associates, Inc. (2024). <https://doi.org/10.52202/079017-2400>, https://proceedings.neurips.cc/paper_files/paper/2024/file/89cc5e613d34f90de90c21e996e60b30-Paper-Conference.pdf
 60. Wang, P., Bai, S., Tan, S., Wang, S., Fan, Z., Bai, J., Chen, K., Liu, X., Wang, J., Ge, W., et al.: Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191* (2024)
 61. Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., Ji, J., Yang, Z., Zhao, L., Song, X., et al.: Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079* (2023)
 62. Wang, W., Gao, Z., Gu, L., Pu, H., Cui, L., Wei, X., Liu, Z., Jing, L., Ye, S., Shao, J., Wang, Z., Chen, Z., Zhang, H., Yang, G., Wang, H., Wei, Q., Yin, J., Li, W., Cui, E., Chen, G., Ding, Z., Tian, C., Wu, Z., Xie, J., Li, Z., Yang, B., Duan, Y., Wang, X., Hou, Z., Hao, H., Zhang, T., Li, S., Zhao, X., Duan, H., Deng, N., Fu, B., He, Y., Wang, Y., He, C., Shi, B., He, J., Xiong, Y., Lv, H., Wu, L., Shao, W., Zhang, K., Deng, H., Qi, B., Ge, J., Guo, Q., Zhang, W., Zhang, S., Cao, M., Lin, J., Tang, K., Gao, J., Huang, H., Gu, Y., Lyu, C., Tang, H., Wang, R., Lv, H., Ouyang, W., Wang, L., Dou, M., Zhu, X., Lu, T., Lin, D., Dai, J., Su, W., Zhou, B., Chen, K., Qiao, Y., Wang, W., Luo, G.: Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency (2025), <https://arxiv.org/abs/2508.18265>
 63. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* **35**, 24824–24837 (2022)
 64. Wu, Q., Zhao, H., Saxon, M., Bui, T., Wang, W.Y., Zhang, Y., Chang, S.: Vsp: Assessing the dual challenges of perception and reasoning in spatial planning tasks for vlms (2024), <https://arxiv.org/abs/2407.01863>
 65. xAI: Grok 4 model card. Tech. rep., xAI (August 2025), <https://data.x.ai/2025-08-20-grok-4-model-card.pdf>
 66. Xiao, R., Ma, W., Wang, K., Wu, Y., Zhao, J., Wang, H., Huang, F., Li, Y.: Flow-Bench: Revisiting and benchmarking workflow-guided planning for LLM-based agents. In: Al-Onaizan, Y., Bansal, M., Chen, Y.N. (eds.) *Findings of the Association for Computational Linguistics: EMNLP 2024*. pp. 10883–10900. Association for Computational Linguistics, Miami, Florida, USA (Nov 2024). <https://arxiv.org/abs/2408.10100>

- [//doi.org/10.18653/v1/2024.findings-emnlp.638](https://doi.org/10.18653/v1/2024.findings-emnlp.638), <https://aclanthology.org/2024.findings-emnlp.638/>
67. Zhang, H., Guo, H., Guo, S., Cao, M., Huang, W., Liu, J., Zhang, G.: Ing-vp: Mllms cannot play easy vision-based games yet (2024), <https://arxiv.org/abs/2410.06555>
 68. Zhang, J., Tang, L., Song, Y., Meng, Q., Qian, H., Shao, J., Song, W., Zhu, S., Gu, J.: Fltrnn: Faithful long-horizon task planning for robotics with large language models. In: 2024 IEEE International Conference on Robotics and Automation (ICRA). pp. 6680–6686. IEEE (2024)
 69. Zhao, H., Hadji, I., Dvornik, N., Derpanis, K.G., Wildes, R.P., Jepson, A.D.: P3iv: Probabilistic procedure planning from instructional videos with weak supervision (2022), <https://arxiv.org/abs/2205.02300>
 70. Zheng, C., Zhou, H., Meng, F., Zhou, J., Huang, M.: Large language models are not robust multiple choice selectors. arXiv preprint arXiv:2309.03882 (2023)
 71. Zheng, H.S., Mishra, S., Zhang, H., Chen, X., Chen, M., Nova, A., Hou, L., Cheng, H.T., Le, Q.V., Chi, E.H., Zhou, D.: Natural plan: Benchmarking llms on natural language planning (2024), <https://arxiv.org/abs/2406.04520>
 72. Zhou, S., Xu, F.F., Zhu, H., Zhou, X., Lo, R., Sridhar, A., Cheng, X., Ou, T., Bisk, Y., Fried, D., Alon, U., Neubig, G.: Webarena: A realistic web environment for building autonomous agents (2024), <https://arxiv.org/abs/2307.13854>