

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ZEBRA-CoT: A DATASET FOR INTERLEAVED VISION-LANGUAGE REASONING

Anonymous authors

Paper under double-blind review

ABSTRACT

Humans often rely on visual aids, such as diagrams or sketches, when tackling complex problems. Teaching multimodal models to adopt similar strategies, a process known as Visual Chain of Thought (visual CoT), is much more difficult. The main challenges are: (1) weak performance of off-the-shelf visual CoT, which hinders reinforcement learning, and (2) the lack of high-quality visual CoT training data. We introduce **ZEBRA-CoT**, a diverse large-scale interleaved text–image reasoning dataset with 182,384 reasoning traces across 18 domains with over 50 distinct tasks. This dataset is specifically designed to train models to natively perform visual CoT. We emphasize four categories of tasks where sketching or visual reasoning is especially natural, spanning (a) *scientific questions* such as geometry, physics, and algorithms; (b) *2D visual reasoning tasks* like visual search and jigsaw puzzles; (c) *3D reasoning tasks* including 3D multi-hop inference, embodied and robot planning; and (d) *visual logic problems and strategic games* like chess. Fine-tuning Anole-7B model on ZEBRA-CoT yields a +12% improvement in our test-set accuracy and up to +13% performance gains on standard VLM benchmarks. Similarly, fine-tuning Bagel-7B produces models capable of generating high-quality interleaved visual reasoning chains, underscoring ZEBRA-CoT’s effectiveness in advancing multimodal reasoning.

1 INTRODUCTION

Human cognition naturally integrates multimodal thought processes when solving complex problems. For example, a high school student sketches diagrams to solve geometry or physics problems, an engineer creates diagrams to design and debug workflows, and a data scientist generates plots to better understand data. These visual aids are central to effective problem solving. While recent vision-language models (VLMs) have shown strong performance on multimodal tasks like visual question answering, their reasoning traces remain predominantly textual. Enabling models to explicitly reason in the visual space, Visual Chain of Thought (visual CoT), remains a fundamental open challenge. Unlocking visual CoT may improve reasoning performance in domains where visual intuition is relevant and may make the reasoning patterns expressed by models more interpretable to humans.

Recent advances in frontier multimodal models (Team et al., 2023; Hurst et al., 2024; Bai et al., 2025; OpenAI, 2025a; Team, 2024; Chern et al., 2024; Sun et al., 2024; Deng et al., 2025) have made visual CoT feasible primarily through agentic pipelines that leverage external tools (e.g., Python functions, or expert vision models) for visual programming (Surfs et al., 2023), such as generating sketches for geometry, algorithms, and spatial reasoning tasks (Hu et al., 2024; OpenAI, 2025b), or bounding boxes for fine-grained visual tasks (Shao et al., 2024a; Wu and Xie, 2024; Zheng et al., 2025). An emerging possibility is innate visual reasoning, where models directly generate explicit visual tokens during their thinking process (Li et al., 2025; Chern et al., 2025; Xu et al., 2025b). However, current VLMs with interleaved text and image generation capabilities (Team, 2024; Chern et al., 2024) either fail to generate useful visual aids for reasoning or are not inherently trained for such multimodal generation during the reasoning process (Deng et al., 2025), making reinforcement learning approaches to reasoning infeasible. Li et al. (2025) demonstrate visual CoT in synthetic mazes by training specialist models, but we remain far from foundation models capable of general high-quality visual CoT, largely due to the lack of large-scale diverse interleaved text and image reasoning training datasets.

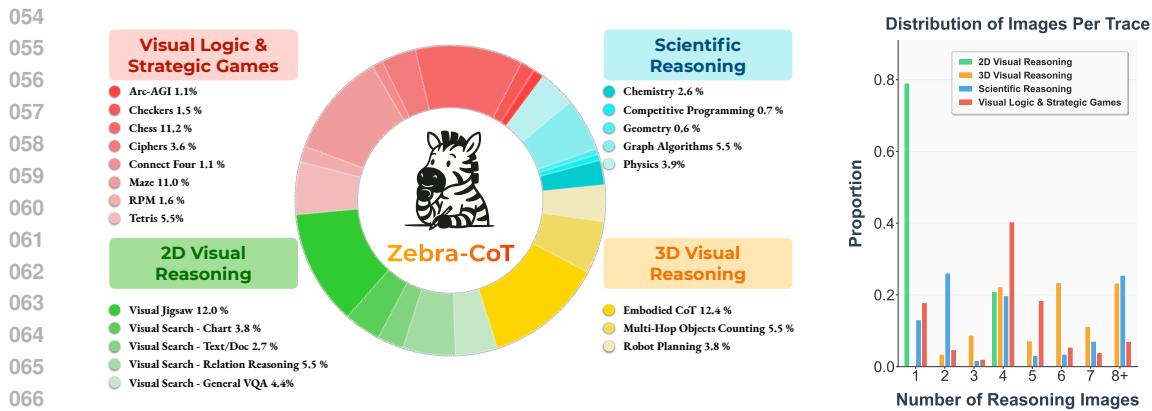


Figure 1: We curate a large-scale multimodal dataset by sourcing and cleaning raw traces from real-world domains, and generating synthetic examples using templated reasoning filled in by VLMs. ZEBRA-CoT comprises 4 major categories and 18 subcategories, encompassing over **182K** instances in total. A detailed breakdown of the data statistics appears in Table 3.

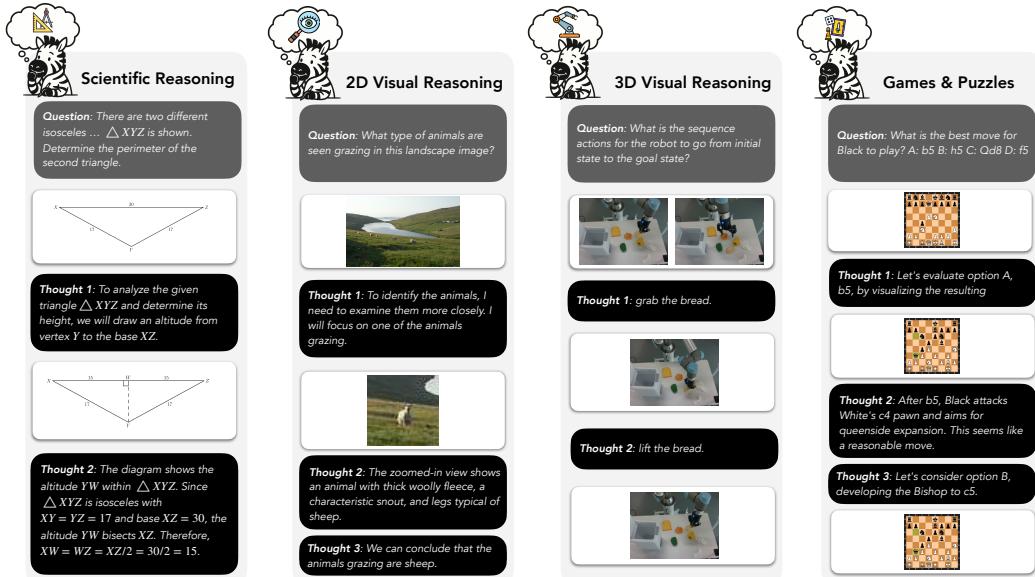


Figure 2: Visual CoT helps answer complex visual reasoning questions, as illustrated by examples from ZEBRA-CoT.

To support the development of next generation vision language models that can explicitly reason with both text and visual modalities, we present **ZEBRA-CoT**, a high quality dataset of interleaved text and image reasoning traces. Our dataset covers four main categories: scientific questions, 2D visual reasoning, 3D visual reasoning, and visual logic and strategic games, each containing multiple subdomains and task types, as exemplified in Figure 2. To the best of our knowledge, ZEBRA-CoT is the first dataset to provide diverse and logically coherent multimodal reasoning traces across such a wide range of domains. Unlike prior large-scale interleaved datasets that are primarily composed of web-scraped image-text pairs with weak semantic alignment and no explicit reasoning structure (Li et al., 2024b; Awadalla et al., 2024; Zhu et al., 2023), ZEBRA-CoT is carefully curated as a training resource in the spirit of high-quality text-based reasoning datasets. At the same time, compared to the only existing open-source interleaved text visual reasoning dataset we are aware of, VISUAL-CoT (Shao et al., 2024a), which focuses on a single task of visual search, ZEBRA-CoT introduces a much broader and more diverse set of tasks with richer reasoning trajectories. We provide a detailed comparison with other datasets below in Table 1.

Dataset	Primary Task	CoT Modality	Suitability for visual CoT Training
GQA	Compositional visual QA	Text	No visual CoT
ScienceQA	Multimodal science QA	Text	No visual CoT
CLEVR	Synthetic compositional visual QA	Text	No visual CoT
VCR	Visual commonsense QA with rationale	Text	No visual CoT
VideoCoT	Video QA	Text	No visual CoT
EgoCOT	Embodied planning	Text	No visual CoT
LLaVA-CoT	Multimodal reasoning QA	Text	No visual CoT
MAmmoTH-VL	Large scale multimodal instruction tuning	Text	No visual CoT
MM-Verify	Multimodal reasoning with verification	Text	No visual CoT
R1-Onevision	A SFT and RL multimodal reasoning dataset	Text	No visual CoT
Visual CoT	Visual-search QA with bbox CoT	Image, Text	Limited to visual search tasks
MM-PhyQA	Physics visual CoT	Image, Text	Physics data only, not open sourced
CoT VLA	Robotics visual CoT	Image, Action	No text reasoning
OmniCorpus	10 B-level interleaved corpus	None	Noisy pretraining data without CoT
MINT-1T	1 T-token web-scale interleaved data	None	Noisy pretraining data without CoT
ZEBRA-COT	Diverse and high quality visual CoT	Image, Text	Diverse interleaved vision-language CoT

Table 1: ZEBRA-COT introduces a broader set of high quality visual CoT traces compared with prior datasets and pipelines.

Our contributions are summarized as follows:

1. We release ZEBRA-COT, a high quality and diverse dataset with interleaved text and visual CoT that contains 182,384 samples for training models to natively perform visual CoT for problem solving. Details regarding the dataset are shown in Section 3
2. We evaluate three frontier LLMs, including GPT-5, Claude Sonnet 4, and Gemini 2.5 Pro, on the tasks in ZEBRA-COT in Section 4. Despite their advanced multimodal reasoning capabilities, these models perform poorly on those challenging tasks, with an average of 31.51%. Moreover, to demonstrate the effectiveness and value of visual CoT, we construct a scaffolding experiment that provides the first one or two multimodal CoT steps in context. Accuracy rises to 47.99% after one step (**+16.48** pts) and 56.70% after two steps (**+25.19** pts) overall, with gains of up to **+43.77** pts in specific domains. These findings highlight the challenging nature of our dataset, the quality of our reasoning traces, and the value of visual CoT.
3. After fine-tuning ANOLE-7B (Chern et al., 2024) on our training set, we improved the accuracy on our in-distribution test set from 4.2% to 16.9%. When evaluating the resulting model on benchmarks requiring visual reasoning, our ANOLE-ZEBRA-COT-7B model achieves an average improvement of **4.9%** across seven challenging datasets, with a maximum gain of **13.1%** on a visual logic benchmark, as shown in Table 2.
4. We fine-tune BAGEL-7B (Deng et al., 2025), a high-quality multimodal model that cannot natively generate interleaved text and images on our dataset. After fine-tuning, the model is able to inherently generate high-quality visual CoT during its own reasoning process, making it well-suited for future RL training, as shown qualitatively in the examples in Figure 4 and Appendix B.

2 RELATED WORK

Visual chain of thought. The community has predominantly been tackling visual CoT by using visual programming to generate images (Surís et al., 2023; Zhang et al., 2023; Mitra et al., 2024; Yang* et al., 2023; Wu and Xie, 2024; Hu et al., 2024; Menon et al., 2024; OpenAI, 2025b; Zheng et al., 2025). In particular, VISUAL SKETCHPAD (Hu et al., 2024) presents the most versatile open-source visual reasoning agents among existing works, handling a wide range of tasks. Another line of work explores model-generated images: for example, Rose et al. (2023) uses a diffusion model to bridge gaps in storytelling, and Chern et al. (2025) generates intermediate images to improve image generation tasks; Zhao et al. (2025) generates intermediate images as subgoal predictions and derives actions based on them for robotic planning; Li et al. (2025) and Xu et al. (2025b) explore spatial reasoning tasks like mazes by visualizing each temporal step. However, these model-generated

162 image approaches are mostly specialists, and developments are still primitive compared to visual
 163 programming methods that leverage external tools.

164 **Visual reasoning datasets.** Many multimodal visual reasoning datasets have been proposed, such
 165 as GQA (Hudson and Manning, 2019), SCIENCEQA (Lu et al., 2022), VIDEOCoT (Wang et al.,
 166 2024c), EGOCOT (Mu et al., 2023), LLAVA-COT (Xu et al., 2024), MAMMOTH-VL (Guo et al.,
 167 2024), MM-VERIFY (Sun et al., 2025), R1-ONEVISION (Yang et al., 2025), CLEVR (Johnson
 168 et al., 2017), VCR (Zellers et al., 2019), although most focus on multi-modality only in the in-
 169 put question, leaving the reasoning traces purely textual. Among them, VISUAL-CoT (Shao et al.,
 170 2024a) stands out as the only open-source dataset featuring interleaved text and image reasoning.
 171 MM-PHYQA (Anand et al., 2024) on the other hand, introduces a paradigm for incorporating im-
 172 ages into the reasoning process for physics problems, though the dataset is not publicly available.
 173 Several vision-centric benchmarks (Fu et al., 2024b; Hao et al., 2025a) present diverse and challeng-
 174 ing tasks, but they lack annotated reasoning traces.

175 **Interleaved text and image datasets.** Large-scale corpora with interleaved text and images have
 176 become essential for pretraining VLMs with reasoning capabilities (Alayrac et al., 2022; Chen and
 177 Wang, 2022; Sun et al., 2024; Wang et al., 2024b; Hurst et al., 2024; Li et al., 2024a; Bai et al., 2025;
 178 Team et al., 2025). However, in most existing interleaved text and image datasets MULTIMODAL
 179 C4 (Zhu et al., 2023), OBELICS (Laurençon et al., 2023), OMNICORPUS (Li et al., 2024b), images
 180 are primarily used for recognition, captioning, or as supplementary context in text-based reasoning,
 181 rather than serving as explicit visual aids that contribute meaningfully to the reasoning process.
 182 While MINT-1T (Awadalla et al., 2024) includes some scientific content from arXiv where images
 183 may aid reasoning, both the text traces and visual content are often noisy and not well-suited for
 184 post-training or fine-grained reasoning tasks. Instead, our ZEBRA-CoT introduces a broader and
 185 higher-quality set of visual CoT examples, enabling effective training for visual reasoning.

186 3 DATA CURATION DETAILS AND COMPOSITIONS

189 3.1 CURATING A DIVERSE AND HIGH QUALITY INTERLEAVED VISION AND LANGUAGE 190 REASONING DATASET

191 A key challenge in training multimodal generation models to output visual CoT natively is the lack
 192 of datasets with strong logical coherence between text and visual modalities, and diverse categories
 193 of such visual CoT. Existing interleaved datasets often fail to provide clear, meaningful connections
 194 that demonstrate when and why visual reasoning is necessary for problem-solving, while current
 195 visual CoT datasets are confined to a few domains, limiting the model’s ability to learn generalizable
 196 visual CoT capabilities when faced with out-of-distribution problems.

197 To address these gaps, we developed a comprehensive data curation pipeline that bridges logical
 198 connections across modalities, as shown in Figure 5. For logical coherence across modalities, we
 199 leverage frontier vision-language models (Gemini-2.5 Pro) to enrich reasoning traces and ensure a
 200 clear logical flow between textual reasoning and visual aids. For diversity, we combine real-world
 201 problems from multiple domains (mathematics, physics, chemistry, coding, chess, visual question
 202 answering, robotics) with synthetic examples generated through computer programming, simula-
 203 tion, and graphic rendering. This pipeline enabled us to curate over 182 K high-quality interleaved
 204 text and visual reasoning traces spanning four major categories: scientific reasoning, 2D visual rea-
 205 soning, 3D visual reasoning, and visual logic and strategic games. Unlike existing limited datasets
 206 that focus primarily on visual search or spatial reasoning, our curated dataset provides the breadth
 207 and diversity necessary for training models that can generalize across domains. For details regard-
 208 ing our data curation pipeline, please refer to Appendix A.2. In the following sections, we provide
 209 a brief introduction to the tasks of each broad category. And for the details regarding the subcat-
 210 egory and domains, please refer to Appendices A.3 to A.6. For prompt templates, please refer to
 211 Appendix F.

213 3.2 SCIENTIFIC QUESTIONS

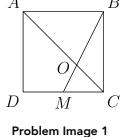
215 Visual reasoning is particularly valuable in STEM domains, as it enables the visualization of abstract
 216 concepts such as auxiliary lines, free-body diagrams, and sketches, which clarify ideas that are hard

216
217
218
219
220
221
222
223
224
225
226
227


Geometry

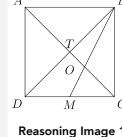
Question: In square $ABCD$, AD is 4 centimeters, and M is the midpoint of \overline{CD} . Let O be the intersection of \overline{AC} and \overline{BM} . What is the ratio of OC to OA ?

Problem Image 1



<think> Thought 1: To better visualize the geometric relationships, especially within triangle BCD , let's add the diagonal BD to our diagram. Let T be the intersection of AC and BD . This will help us identify if any segment related to AC can be considered a median of triangle BCD . </think>

Reasoning Image 1



<think> Thought 2: The updated diagram shows the square $ABCD$ with both diagonals AC and BD intersecting at point T . The segment BM intersects AC at O . Now, let's focus on triangle BCD ... </think>

to describe in language and support step-by-step problem solving in ways that mirror human cognition. In ZEBRA-CoT, this category spans subdomains including geometry, physics, chemistry, algorithmic problem solving, and graph problems. For geometry, physics, and chemistry, we leverage openly licensed datasets and textbooks, using Gemini-2.5 (Comanici et al., 2025) to denoise and parse them into clean, logically structured visual CoT. For graph problems, we employ computer programs to generate images and text templates, which are then diversified using Gemini-2.5. For algorithmic problems, we use a GPT-4.1 agent built upon Hu et al. (2024) to produce detailed traces for solving competitive programming tasks. For details regarding all tasks in this domain, see Appendix A.3.

3.3 2D VISUAL REASONING



Visual Jigsaw

Question: A jigsaw puzzle shows a picture of pandas but has several missing pieces, indicated by gray areas. Identify which set of pieces from the options provided will correctly complete the puzzle.

Problem Image 1



<think> Thought 0: Analyze the original puzzle image to understand the missing sections and the surrounding details. **Thought 1:** To determine the correct missing pieces, visualize the complete image or find the original source image. </think>

Reasoning Image 1



<think> Thought 2: Compare the full image with the options provided to see which set of pieces matches the content and position of the gray areas in the original puzzle. **Thought 3:** Option D contains the correct parts of the image needed to fill the missing sections based on the full image. </think>

In 2D visual reasoning, visual aids support the manipulation and interpretation of 2D visual information, enabling tasks that involve spatial arrangement, pattern recognition, and fine-grained inspection. For this category, we include tasks such as visual search and visual jigsaw. For visual search, we adapt datasets from Shao et al. (2024a) and incorporate two types of visual aids: drawing bounding boxes and zooming into focal regions. We apply those visual CoT broadly across data categories, such as charts, documents, relations, and general VQA. For visual jigsaw tasks, we crop images from ImageNet (Deng et al., 2009) to create puzzles with a random number of missing pieces in diverse shapes. The visual CoT is either iteratively filling in the pieces or reconstructing the original image directly. Further details are provided in Appendix A.4.

3.4 3D VISUAL REASONING

For 3D visual reasoning tasks, we focus on two domains: (1) embodied reasoning and robotic planning in the physical world, and (2) understanding 3D transformations from different viewpoints. For the first domain, prior work has shown that generating visual predictions of the physical world and extracting inverse dynamics can improve performance in long-horizon decision-making in robotics (Zhao et al., 2025; Yang et al., 2024). To capture this, we reformulate the ALFRED benchmark (Shridhar et al., 2020) into an image goal-conditioned planning task in which models generate detailed step-by-step plans to transition from an initial state to a goal state. We also adapt RoboMIND (Wu et al., 2024) for real-world robot planning, where models receive initial and goal images, along with descriptions of robot embodiment, and must produce precise high-level action plans. For reasoning about 3D transformations, we design multi-hop object counting tasks inspired

270
271
272
273

Robot Planning

274
275
276
277
278
279
280
281
282

Question:
Please describe the sequence of robot actions from the initial state (the first image) to the goal state (the second image).



Problem Image 1 & 2

<think> Thought 1:
move the robot's arm towards the bread.
</think>



Reasoning Image 1

<think> Thought 2:
use the robot's arm to grab the bread.
</think>



Reasoning Image 2

<think> Thought 3:
move the bread towards the table.
</think>



Reasoning Image 3

283
284

by CLEVR (Johnson et al., 2017), where scenes undergo sequential modifications, such as adding or removing objects, requiring models to visually reason through each transformation step. For details, see Appendix A.5

288



3D Multi Hop Object Counting

289
290
291
292
293
294
295
296
297

Question: Remove zero CAPSULE objects. Eliminate all PYRAMID objects. Remove zero PURPLE objects. Discard all CUBE objects. Count how many blue-colored objects remaining.



Problem Image 1

<think> Thought 2: Proceeding to Step 2, the instruction is to "Remove all PYRAMID objects." Upon reviewing the current scene, there is one orange pyramid present. This object must be removed from the arrangement. The updated scene after this removal is depicted in the following image.
</think>



Reasoning Image 1

<think> Thought 3: For Step 3, the directive states to "Remove 0 PURPLE object." Similar to Step 1, this operation specifies that no purple objects are to be removed. The objects remaining in the scene after Step 2 are preserved...
</think>

298
299

3.5 VISUAL LOGIC AND STRATEGIC GAMES

300

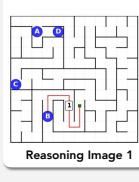


Maze

301
302
303
304
305
306
307
308
309
310
311

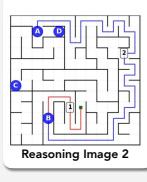
Question: Determine the smallest number of movements required to cover all letters without repetition, starting at the green marker (1-step transitions in four directions around obstacles) Options: (A) 88, (B) 99, (C) 90, (D) 86

<think> Thought 3: The first segment of the optimal path takes us from the initial green marker to the letter B. This journey covers a distance of 10 moves. The image below illustrates this initial part of the route.
</think>



Reasoning Image 1

<think> Thought 4: Following the visit to B, the next leg of the journey leads to letter A. This particular segment is quite lengthy, requiring 41 moves due to the winding nature of the maze. The next visual highlights this path.
</think>



Reasoning Image 2

312
313
314
315
316
317
318
319
320
321
322
323

For visual logic puzzles (IQ matrices, Tetris, ciphers, ARC-AGI (Chollet et al., 2024)), previous VLMs tended to solve problems primarily using text reasoning. They first verbalize visual inputs into text, which causes information loss and makes visually salient patterns, such as spatial relationships, difficult to capture. In contrast, humans solve these directly and efficiently via visual imagination and manipulation, even for babies who have not yet acquired language capabilities (Zhu et al., 2020). To bridge the gap, we construct visual CoT traces that include explicit intermediate visual transformations to encourage models to solve these problems visually. Similarly, for strategic games (chess, checkers, Connect Four), decision making typically involves searching and generating counterfactual rollouts. While LLMs can simulate this by symbolizing board states into text, much of the spatial structure is lost, and rollouts in text space are difficult for problems with large visual information. Thus, we render those search and simulation steps into images so that models trained on this data can perform long-horizon planning in the visual space inherently. Finally, we generate a diverse suite of maze tasks and visual CoT traces that require a combination of capabil-



ties, including high-level symbolic search and low-level perception. For details of those tasks, see Appendix A.6.

4 ANALYSIS OF ZEBRA-COT AND THE VALUE OF VISUAL CoT

Proprietary frontier models (GPT-5 (OpenAI, 2025c), Gemini-2.5 Pro (Comanici et al., 2025), Claude-4 Sonnet (Anthropic, 2025)) have achieved state-of-the-art performance on multimodal reasoning benchmarks. Despite their advanced multimodal capabilities, we show that they struggle significantly with the tasks in ZEBRA-COT. To explore these limitations and demonstrate the challenging nature of our dataset alongside the effectiveness of visual reasoning traces, we design a scaffolding experiment. Specifically, our dataset consists of structured reasoning chains: <question> → <text-reasoning-1> → <visual-reasoning-1> → <text-reasoning-2> → <visual-reasoning-2> → ... → <answer>.

In the zero-shot setting, we provide models only with the <question> (containing both image and text). For scaffolding experiments, we incrementally provide the first k multimodal reasoning steps as context:

- **1MT** ($k = 1$): <question> + <text-reasoning-1> + <visual-reasoning-1>
- **2MT** ($k = 2$): <question> + <text-reasoning-1> + <visual-reasoning-1> + <text-reasoning-2> + <visual-reasoning-2>

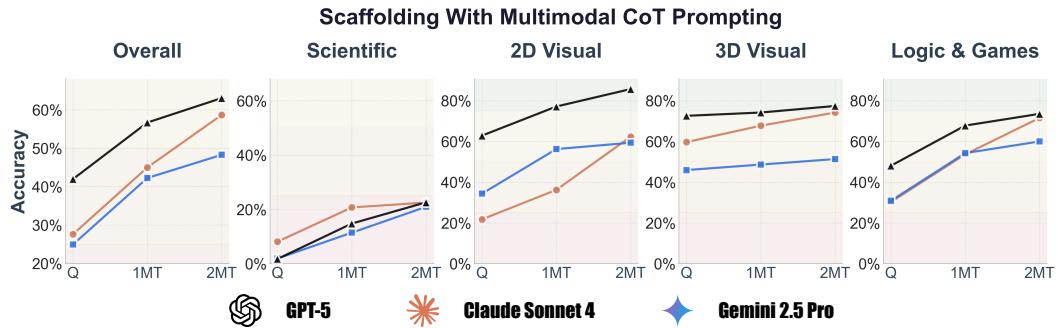


Figure 3: Scaffolding experiment with frontier models. **Q** represents zero-shot question-only evaluation, **1MT** denotes a question with the first multimodal reasoning step provided, and **2MT** indicates a question with the first two multimodal reasoning steps. We show that **even frontier models with the best multimodal reasoning capabilities perform poorly overall on tasks in ZEBRA-COT**. However, as we provide the first one or two multimodal steps to those models, the accuracy improves significantly.

Importantly, most tasks in ZEBRA-COT require various multimodal reasoning steps (which can involve as many as 20 images) to reach the final answer. By providing only the first two steps as scaffolding, we ensure that models must still perform substantial reasoning to solve the task. We can safely assume that the provided steps serve as guidance rather than revealing the solution. Since our dataset comprises diverse tasks, some of which extend beyond traditional QA formats (e.g., robotic planning and embodied CoT) that are not suitable for evaluation, we select the most challenging and

representative examples for evaluation: graph questions for scientific reasoning, visual jigsaw for 2D spatial reasoning, multihop object counting for 3D reasoning, and maze/chess/tetris for visual logic and strategic games.

We plot the results for three evaluation settings across each task domain in Figure 3. We observe that frontier models achieve poor zero-shot performance: GPT-5 reaches 41.98% accuracy, while Claude-4 Sonnet and Gemini-2.5 Pro achieve only 27.61% and 24.93% respectively. However, with multimodal CoT scaffolding, we observe substantial improvements: average accuracy across the three models increases to 47.99% (**+16.48%**) with one reasoning step and 56.70% (**+25.19%**) with two steps.

Performance gains vary across task types, but we generally see an improvement trend. Maze tasks show the most dramatic improvements, which jump from 52.59% to 76.60% (+24.01%) and to 96.36% (+43.77%) on average, while challenging tasks such as graph reasoning improve from 3.92% to 22.03% (+18.11%) with two multimodal reasoning steps on average. Even tasks with higher baseline performance, such as multihop object counting (with an initial accuracy of 59.40%), benefit from visual CoT, eventually reaching 67.65% accuracy on average. Detailed statistics are shown in Table 8.

To isolate the contribution of visual reasoning aids from text CoT in our traces, we conduct an ablation task where we remove all visual aids from the reasoning traces and retain only the textual steps. We observe that text-only CoT yields substantially smaller performance gains compared to full visual CoT, and in some cases even degrades performance. This is expected: in our dataset, the visual and textual components are highly complementary. Many reasoning steps reference visual elements that, once removed, leave the text chain logically incomplete or incoherent. Model even requests for the missing visual aids that are referred in the text cot. These results indicate that the majority of the performance improvements stem from the visual reasoning steps, or the combined visual + text reasoning, rather than from textual CoT alone. The statistics for text only results are shown here:

5 TRAINING MODELS ON ZEBRA-COT

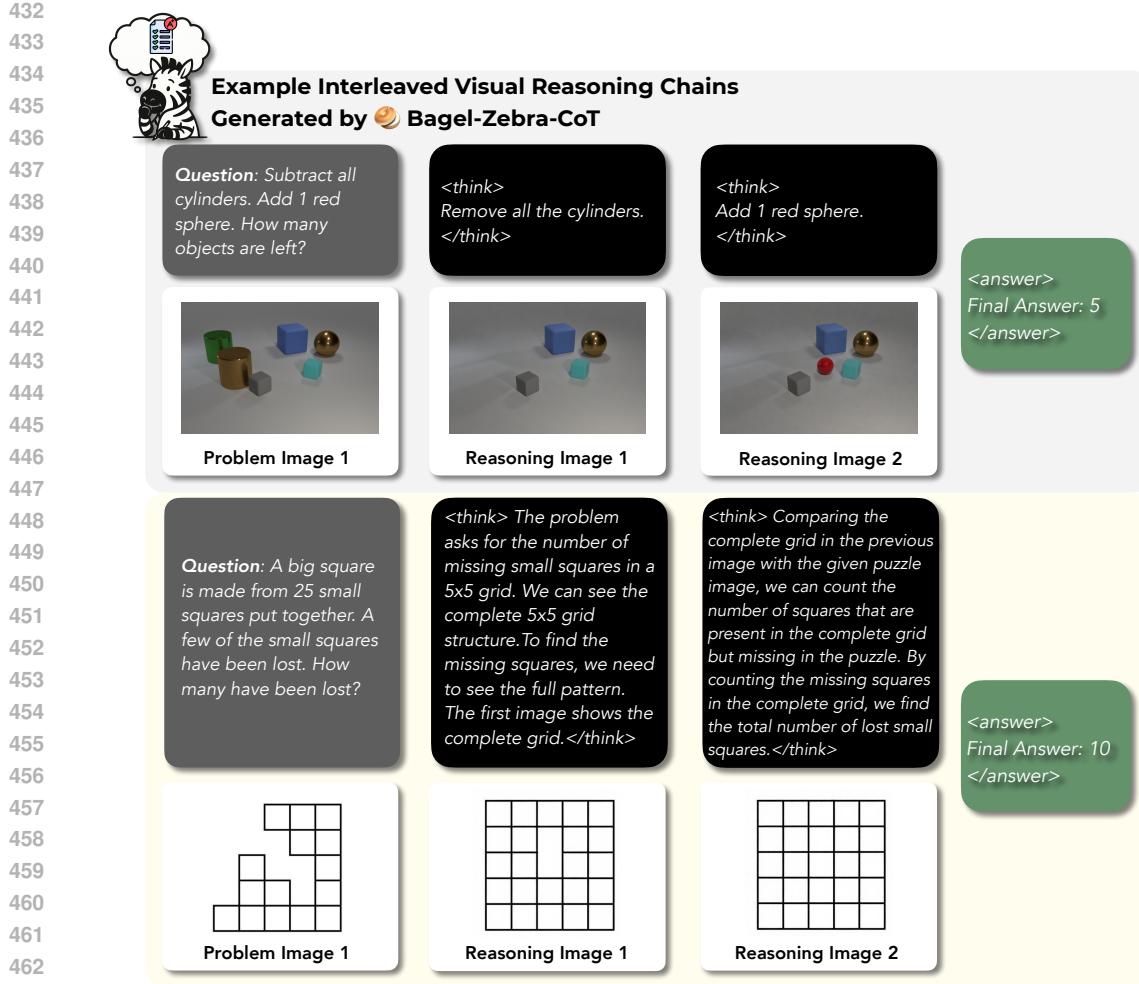
Model	MathVision*	MathVista*	VisuLogic	EMMA	MMVP	Blink	Vstar
Anole with CoT prompting	13.80	22.80	8.50	12.80	10.00	26.46	23.60
Anole-Zebra-CoT (Ours)	16.45	25.30	21.80	15.02	15.33	31.25	27.20

Table 2: Overall performance (%) across eight datasets for the base Anole model with chain-of-thought prompting vs. the same Anole model further trained on ZEBRA-CoT. *We evaluate on the mini versions of MathVision and MathVista because interleaved generation is time consuming. A full breakdown of each evaluation set is presented in Appendix C.

Anole-Zebra-CoT. We fine-tune Anole (Chern et al., 2024) on our dataset, which builds on Chameleon (Team, 2024), using the codebase from Chern et al. (2025). We finetune the model fully end-to-end on a node with $8 \times$ H200 GPUs for 12 hours, with a learning rate of 1×10^{-5} , cosine decay, a batch size of 8, and a max token length of 12288. We train the model for 10k steps. To evaluate our trained model, we set the maximum generation length to 16384. After fine-tuning Anole on our ZEBRA-CoT corpus, the accuracy increased from 4.2% to 16.9%, delivering a 4 times relative performance improvement and a 12% gain in accuracy.

Furthermore, we evaluate seven challenging benchmarks that require visual reasoning, including MathVision (Wang et al., 2024a), MathVista (Lu et al., 2024), VisuLogic (Xu et al., 2025a), EMMA (Hao et al., 2025b), MMVP (Tong et al., 2024), BLINK (Fu et al., 2024b), and Vstar (Wang et al., 2023). All the evaluations are done using VLMEvalKit (Duan et al., 2024). To ensure a fair comparison, we use chain-of-thought prompting (Wei et al., 2022) when evaluating the base Anole model. As shown in Table 2, training with ZEBRA-CoT significantly improves the Anole model across all benchmarks. Most notably, it could improve the Anole model’s visual logical reasoning capabilities by 13.3 points.

Bagel-Zebra-CoT. To further test whether ZEBRA-CoT can enhance a stronger backbone, we fine-tune the BAGEL-7B model (Deng et al., 2025) end-to-end on a node with $8 \times$ H200 GPUs for 1,000



464
 465
 466
 467
 468
 469
 470
 471
 472
 473
 474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485

Figure 4: Example interleaved reasoning chains generated by Bagel-Zebra-CoT, a Bagel-7B model finetuned on ZEBRA-CoT. These traces demonstrate ZEBRA-CoT’s for instilling intrinsic visual reasoning capability in complex multimodal models.

steps using packed sequences with 60,000 tokens, a learning rate of 2×10^{-5} , and cosine decay. We cap all images at a resolution of 512 on the minimum side, resulting in approximately 1,024+ visual tokens per image. Because the original Bagel implementation cannot natively generate interleaved text-image outputs, we revise the training loop to include a loss term at the `<|vision_start|>` token, enabling seamless visual token generation. We additionally wrap text reasoning tokens between `<think>` and `</think>`, and the final answer within `<answer>` and `</answer>`. At inference time, when encountering `<im_end>`, we sample one additional token to check whether the next token is `<|vision_start|>`; if so, the model itself seamlessly switches to image generation mode to generate visual aids. The entire interleaved generation process only stops if the model generates the `<answer>` token.

We observe that our trained model can inherently generate visual CoT when solving problems, even on tasks outside its training distribution. This suggests its potential as a strong initialization for future reinforcement learning fine-tuning. In Figure 4, we include representative reasoning traces produced by the model. We further include more reasoning traces in Appendix B, as well as a model performance analysis in Appendix D

486 6 CONCLUSION & FUTURE DIRECTIONS 487

488 In this paper, we introduced ZEBRA-COT, a large-scale dataset of 182K interleaved text-image
489 reasoning traces spanning 4 major categories across 18 domains with over 50 distinct tasks. Fine-
490 tuning experiments demonstrate substantial improvements: Anole-7B achieves an average 4.9 %
491 gain across seven challenging benchmarks, with up to 13.1% on visual logic tasks, while Bagel-
492 7B learns to inherently generate visual aids during problem solving, a capability absent in the base
493 model.

494 This work opens several exciting avenues for future research. Most immediately, models trained on
495 ZEBRA-COT, particularly our Bagel variant that natively generates visual thoughts, provide strong
496 initializations for reinforcement learning. Just as text-based reasoning models have benefited from
497 RL fine-tuning to improve logical consistency and correctness, we envision similar gains for visual
498 reasoning through RL with verifiable rewards (Shao et al., 2024b; Guo et al., 2025) or fine-grained
499 rewards (Zeng et al., 2024; Fu et al., 2025).

500 We believe ZEBRA-COT represents a crucial step toward AI systems that think visually as naturally
501 as humans sketch diagrams, generate graphs, and use spatial reasoning to solve complex problems.
502 With our dataset and fine-tuned model, we hope to accelerate progress toward this goal.
503

504 7 LLM USAGE DISCLOSURE 505

506 We used LLM for two purposes. The first one is for improving grammar and wording when writing
507 the paper. The second usage is synthetic data generation, where details can be found in Section 3
508 and Appendix A.2
509

510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 REFERENCES

541

542 AgileX Robotics. Cobot magic. <https://global.agilex.ai/products/cobot-magic>, 2023.

543

544 Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel
545 Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan
546 Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian
547 Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo
548 Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language
549 model for few-shot learning. 2022. URL <https://arxiv.org/abs/2204.14198>.

550

551 Avinash Anand, Janak Kapuriya, Apoorv Singh, Jay Saraf, Naman Lal, Astha Verma, Rushali
552 Gupta, and Rajiv Shah. Mm-phyqa: Multimodal physics question-answering with multi-image
553 cot prompting. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages
554 53–64. Springer, 2024.

555

556 Anthropic. System card: Claude opus 4 & claude sonnet 4. System card, Anthropic, May 2025. URL
557 <https://www-cdn.anthropic.com/6d8a8055020700718b0c49369f60816ba2a7c285.pdf>. Original May 2025;
558 changelog entries July 16, 2025 and Sept 2, 2025.

559

560 Anas Awadalla, Le Xue, Oscar Lo, Manli Shu, Hannah Lee, Etash Guha, Sheng Shen, Mohamed
561 Awadalla, Silvio Savarese, Caiming Xiong, et al. Mint-1t: Scaling open-source multimodal data
562 by 10x: A multimodal dataset with one trillion tokens. *Advances in Neural Information Processing
Systems*, 37:36805–36828, 2024.

563

564 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
565 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,
566 2025.

567

568 Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and
569 Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

570

571 Xi Chen and Xiao Wang. Pali: Scaling language-image learning in 100+ languages. In *Conference
on Neural Information Processing Systems (NeurIPS)*, 2022.

572

573 Ethan Chern, Jiadi Su, Yan Ma, and Pengfei Liu. Anole: An open, autoregressive, native large mul-
574 timodal models for interleaved image-text generation. *arXiv preprint arXiv:2407.06135*, 2024.

575

576 Ethan Chern, Zhulin Hu, Steffi Chern, Siqi Kou, Jiadi Su, Yan Ma, Zhijie Deng, and Pengfei Liu.
577 Thinking with generated images. *arXiv preprint arXiv:2505.22525*, 2025.

578

579 François Chollet. On the measure of intelligence. *arXiv preprint arXiv:1911.01547*, 2019.

580

581 Francois Chollet, Mike Knoop, Gregory Kamradt, and Bryan Landers. Arc prize 2024: Technical
582 report. *arXiv preprint arXiv:2412.04604*, 2024.

583

584 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit
585 Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the
586 frontier with advanced reasoning, multimodality, long context, and next generation agentic capa-
587 bilities. *arXiv preprint arXiv:2507.06261*, 2025.

588

589 Joost CF de Winter, Dimitra Dodou, and Yke Bauke Eisma. Responses to raven matrices: Governed
590 by visual complexity and centrality. *Perception*, 52(9):645–661, 2023.

591

592 Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao
593 Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv
preprint arXiv:2505.14683*, 2025.

594

595 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-
596 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,
597 pages 248–255. Ieee, 2009.

594 Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong,
 595 Yuhang Zang, Pan Zhang, Jiaqi Wang, et al. Vlmevalkit: An open-source toolkit for evaluat-
 596 ing large multi-modality models. In *Proceedings of the 32nd ACM International Conference on*
 597 *Multimedia*, pages 11198–11201, 2024.

598 Franka Emika GmbH. *Franka Emika Panda Robot Arm*, 2018. <https://www.franka.de>.

600 Deqing Fu, Ruohao Guo, Ghazal Khalighinejad, Ollie Liu, Bhuwan Dhingra, Dani Yogatama, Robin
 601 Jia, and Willie Neiswanger. IsoBench: Benchmarking multimodal foundation models on isomor-
 602 phic representations. In *First Conference on Language Modeling (COLM)*, 2024a.

603 Deqing Fu, Tong Xiao, Rui Wang, Wang Zhu, Pengchuan Zhang, Guan Pang, Robin Jia, and
 604 Lawrence Chen. TLDR: Token-level detective reward model for large vision language mod-
 605 els. In *The Thirteenth International Conference on Learning Representations*, 2025. URL
 606 <https://openreview.net/forum?id=Zy2XgaGpDw>.

607 Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A
 608 Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but
 609 not perceive. In *European Conference on Computer Vision*, pages 148–166. Springer, 2024b.

610 Daya Guo, Dejian Yang, Huawei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
 611 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
 612 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

613 Jarvis Guo, Tuney Zheng, Yuelin Bai, Bo Li, Yubo Wang, King Zhu, Yizhi Li, Graham Neubig,
 614 Wenhui Chen, and Xiang Yue. Mammoth-vl: Eliciting multimodal reasoning with instruction
 615 tuning at scale. *arXiv preprint arXiv:2412.05237*, 2024.

616 Yunzhuo Hao, Jiawei Gu, Huichen Will Wang, Linjie Li, Zhengyuan Yang, Lijuan Wang, and
 617 Yu Cheng. Can mllms reason in multimodality? emma: An enhanced multimodal reasoning
 618 benchmark. *arXiv preprint arXiv:2501.05444*, 2025a.

619 Yunzhuo Hao, Jiawei Gu, Huichen Will Wang, Linjie Li, Zhengyuan Yang, Lijuan Wang, and
 620 Yu Cheng. Can mllms reason in multimodality? emma: An enhanced multimodal reasoning
 621 benchmark. *arXiv preprint arXiv:2501.05444*, 2025b.

622 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn
 623 Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset.
 624 In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks*
 625 *Track (Round 2)*, 2021. URL <https://openreview.net/forum?id=7Bywt2mQsCe>.

626 Yushi Hu, Weijia Shi, Xingyu Fu, Dan Roth, Mari Ostendorf, Luke Zettlemoyer, Noah A Smith,
 627 and Ranjay Krishna. Visual sketchpad: Sketching as a visual chain of thought for multimodal
 628 language models. *arXiv preprint arXiv:2406.09403*, 2024.

629 Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning
 630 and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer*
 631 *vision and pattern recognition*, pages 6700–6709, 2019.

632 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
 633 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*
 634 *arXiv:2410.21276*, 2024.

635 Michael Igorevich Ivanitskiy, Rusheb Shah, Alex F. Spies, Tilman Räuker, Dan Valentine, Can
 636 Rager, Lucia Quirke, Chris Mathwin, Guillaume Corlouer, Cecilia Diniz Behn, and Samy Wu
 637 Fung. A configurable library for generating and manipulating maze datasets, 2023. URL <https://arxiv.org/abs/2309.10498>.

638 Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and
 639 Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual
 640 reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 641 pages 2901–2910, 2017.

648 Alex Lau-Zhu, Emily A Holmes, Sally Butterfield, and Joni Holmes. Selective association between
 649 tetris game play and visuospatial working memory: A preliminary investigation. *Applied cognitive*
 650 *psychology*, 31(4):438–445, 2017.

651

652 Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov,
 653 Thomas Wang, Siddharth Karamcheti, Alexander M. Rush, Douwe Kiela, Matthieu Cord, and
 654 Victor Sanh. Obelics: An open web-scale filtered dataset of interleaved image-text documents,
 655 2023.

656 Chengzu Li, Wenshan Wu, Huanyu Zhang, Yan Xia, Shaoguang Mao, Li Dong, Ivan Vulić, and
 657 Furu Wei. Imagine while reasoning in space: Multimodal visualization-of-thought. *arXiv preprint*
 658 *arXiv:2501.07542*, 2025.

659

660 Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zeyun Ma, and Chunyuan Li.
 661 Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv*
 662 *preprint arXiv:2407.07895*, 2024a.

663

664 Qingyun Li, Zhe Chen, Weiyun Wang, Wenhui Wang, Shenglong Ye, Zhenjiang Jin, Guanzhou
 665 Chen, Yinan He, Zhangwei Gao, Erfei Cui, et al. Omnicorpus: A unified multimodal corpus of
 10 billion-level images interleaved with text. *arXiv preprint arXiv:2406.08418*, 2024b.

666

667 Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord,
 668 Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for
 669 science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521,
 2022.

670

671 Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-
 672 Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of
 673 foundation models in visual contexts. In *International Conference on Learning Representations*
 (ICLR), 2024.

674

675 Sachit Menon, Richard Zemel, and Carl Vondrick. Whiteboard-of-thought: Thinking step-by-step
 676 across modalities. *arXiv preprint arXiv:2406.14562*, 2024.

677

678 MIT OpenCourseWare. [Course Title]. <https://ocw.mit.edu/>, 2022. MIT OpenCourse-
 679 Ware: Massachusetts Institute of Technology. License: Creative Commons BY-NC-SA.

680

681 Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. Compositional chain-of-
 682 thought prompting for large multimodal models. In *Proceedings of the IEEE/CVF Conference*
 683 *on Computer Vision and Pattern Recognition*, pages 14420–14431, 2024.

684

685 William Moebs, Samuel J. Ling, and Jeff Sanny. *University Physics Volume 1*. OpenStax,
 686 Houston, Texas, 2016. URL <https://openstax.org/books/university-physics-volume-1/pages/1-introduction>. Licensed under
 687 CC BY 4.0.

688

689 Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhui Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng
 690 Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of
 691 thought. *Advances in Neural Information Processing Systems*, 36:25081–25094, 2023.

692

693 OpenAI. Openai o3 and o4-mini system card. Technical report, OpenAI, April 2025a. URL
 694 <https://cdn.openai.com/pdf/2221c875-02dc-4789-800b-e7758f3722c1/o3-and-o4-mini-system-card.pdf>.

695

696 OpenAI. Thinking with images. <https://openai.com/index/thinking-with-images/>, April 2025b. Accessed: 2025-07-21.

697

698 OpenAI. Gpt-5 system card, August 2025c. URL <https://openai.com/index/gpt-5-system-card/>. System card for GPT-5.

699

700 Bharath Ramsundar, Peter Eastman, Patrick Walters, Vijay Pande, Karl Leswing, and Zhenqin Wu.
 701 *Deep Learning for the Life Sciences*. O'Reilly Media, 2019. <https://www.amazon.com/Deep-Learning-Life-Sciences-Microscopy/dp/1492039837>.

702 Daniel Rose, Vaishnavi Himakunthal, Andy Ouyang, Ryan He, Alex Mei, Yujie Lu, Michael
 703 Saxon, Chinmay Sonar, Diba Mirza, and William Yang Wang. Visual chain of thought: bridging
 704 logical gaps with multimodal infillings. *arXiv preprint arXiv:2305.02317*, 2023.

705

706 Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hong-
 707 sheng Li. Visual cot: Advancing multi-modal language models with a comprehensive dataset and
 708 benchmark for chain-of-thought reasoning. *Advances in Neural Information Processing Systems*,
 709 37:8612–8642, 2024a.

710 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 711 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of
 712 mathematical reasoning in open language models, 2024b. URL <https://arxiv.org/abs/2402.03300>.

713

714 Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi,
 715 Luke Zettlemoyer, and Dieter Fox. ALFRED: A Benchmark for Interpreting Grounded Instruc-
 716 tions for Everyday Tasks. In *The IEEE Conference on Computer Vision and Pattern Recognition*
 717 (*CVPR*), 2020. URL <https://arxiv.org/abs/1912.01734>.

718

719 Linzhuang Sun, Hao Liang, Jingxuan Wei, Bihui Yu, Tianpeng Li, Fan Yang, Zenan Zhou, and
 720 Wentao Zhang. Mm-verify: Enhancing multimodal reasoning with chain-of-thought verification.
 721 *arXiv preprint arXiv:2502.13383*, 2025.

722

723 Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiying Yu, Yueze Wang, Yongming Rao,
 724 Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative multimodal models are in-context
 725 learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-
 726 nition*, pages 14398–14409, 2024.

727

728 Dídac Surís, Sachit Menon, and Carl Vondrick. Vipergpt: Visual inference via python execution for
 729 reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages
 11888–11898, 2023.

730

731 Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint
 732 arXiv:2405.09818*, 2024.

733

734 Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut,
 735 Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly
 736 capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

737

738 Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej,
 739 Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical
 report. *arXiv preprint arXiv:2503.19786*, 2025.

740

741 Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide
 742 shut? exploring the visual shortcomings of multimodal llms, 2024.

743

744 Universal Robots A/S. *UR5e Collaborative Robot Arm*, 2018. <https://www.universal-robots.com/products/ur5e/>.

745

746 Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and
 747 Hongsheng Li. Measuring multimodal mathematical reasoning with math-vision dataset. In *The
 748 Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks
 749 Track*, 2024a. URL <https://openreview.net/forum?id=QWTCCxMpPA>.

750

751 Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan
 752 Zhang, Yueze Wang, Zhen Li, Qiying Yu, et al. Emu3: Next-token prediction is all you need.
 753 *arXiv preprint arXiv:2409.18869*, 2024b.

754

755 Yan Wang, Yawen Zeng, Jingsheng Zheng, Xiaofen Xing, Jin Xu, and Xiangmin Xu. Videocot:
 756 A video chain-of-thought dataset with active annotation tool. *arXiv preprint arXiv:2407.05355*,
 757 2024c.

756 Yuxuan Wang, Zilong Zheng, Xueliang Zhao, Jinpeng Li, Yueqian Wang, and Dongyan Zhao. VS-
 757 TAR: A video-grounded dialogue dataset for situated semantic understanding with scene and topic
 758 transitions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Lin-
 759 guistics (Volume 1: Long Papers)*, pages 5036–5048, Toronto, Canada, July 2023. Association
 760 for Computational Linguistics. URL [https://aclanthology.org/2023.acl-long.
 761 276](https://aclanthology.org/2023.acl-long.276).

762 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 763 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
 764 neural information processing systems*, 35:24824–24837, 2022.

765 Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhiqin Yang, Meng Li, Yinuo
 766 Zhao, Zhiyuan Xu, Guang Yang, Zhen Zhao, Guangyu Li, Zhao Jin, Lecheng Wang, Jilei Mao,
 767 Xinhua Wang, Shichao Fan, Ning Liu, Pei Ren, Qiang Zhang, Yaoxu Lyu, Mengzhen Liu,
 768 Jingyang He, Yulin Luo, Zeyu Gao, Chenxuan Li, Chenyang Gu, Yankai Fu, Di Wu, Xingyu
 769 Wang, Sixiang Chen, Zhenyu Wang, Pengju An, Siyuan Qian, Shanghang Zhang, and Jian Tang.
 770 Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation.
 771 *arXiv preprint arXiv:2412.13877*, 2024.

772 Penghao Wu and Saining Xie. V*: Guided visual search as a core mechanism in multimodal llms.
 773 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages
 774 13084–13094, 2024.

775 Guowei Xu, Peng Jin, Li Hao, Yibing Song, Lichao Sun, and Li Yuan. Llava-o1: Let vision language
 776 models reason step-by-step. *arXiv preprint arXiv:2411.10440*, 2024.

777 Weiyue Xu, Jiahao Wang, Weiyun Wang, Zhe Chen, Wengang Zhou, Aijun Yang, Lewei Lu,
 778 Houqiang Li, Xiaohua Wang, Xizhou Zhu, Wenhui Wang, Jifeng Dai, and Jinguo Zhu. Visulogic:
 779 A benchmark for evaluating visual reasoning in multi-modal large language models, 2025a. URL
 780 <https://arxiv.org/abs/2504.15279>.

781 Yi Xu, Chengzu Li, Han Zhou, Xingchen Wan, Caiqi Zhang, Anna Korhonen, and Ivan Vulić. Visual
 782 planning: Let's think only with images. *arXiv preprint arXiv:2505.11409*, 2025b.

783 Sherry Yang, Jacob Walker, Jack Parker-Holder, Yilun Du, Jake Bruce, Andre Barreto, Pieter
 784 Abbeel, and Dale Schuurmans. Video as the new language for real-world decision making. *arXiv
 785 preprint arXiv:2402.17139*, 2024.

786 Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng
 787 Yin, Fengyun Rao, Minfeng Zhu, et al. R1-onevision: Advancing generalized multimodal rea-
 788 soning through cross-modal formalization. *arXiv preprint arXiv:2503.10615*, 2025.

789 Zhengyuan Yang*, Linjie Li*, Jianfeng Wang*, Kevin Lin*, Ehsan Azarnasab*, Faisal Ahmed*,
 790 Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. Mm-react: Prompting chatgpt for multi-
 791 modal reasoning and action. 2023.

792 Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual
 793 commonsense reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and
 794 pattern recognition*, pages 6720–6731, 2019.

795 Yongcheng Zeng, Guoqing Liu, Weiyu Ma, Ning Yang, Haifeng Zhang, and Jun Wang. Token-level
 796 direct preference optimization. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian
 797 Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st
 798 International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learn-
 799 ing Research*, pages 58348–58365. PMLR, 21–27 Jul 2024. URL [https://proceedings.
 800 mlr.press/v235/zeng24c.html](https://proceedings.mlr.press/v235/zeng24c.html).

801 Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, and Song-Chun Zhu. Raven: A dataset for relational
 802 and analogical visual reasoning. In *Proceedings of the IEEE/CVF conference on computer vision
 803 and pattern recognition*, pages 5317–5327, 2019.

804 Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. Multimodal
 805 chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*, 2023.

810 Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li,
811 Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-
812 language-action models. *arXiv preprint arXiv:2503.22020*, 2025.

813

814 Ziwei Zheng, Michael Yang, Jack Hong, Chenxiao Zhao, Guohai Xu, Le Yang, Chao Shen, and
815 Xing Yu. Deepeyes: Incentivizing” thinking with images” via reinforcement learning. *arXiv*
816 *preprint arXiv:2505.14362*, 2025.

817 Wang Zhu, Hexiang Hu, Jiacheng Chen, Zhiwei Deng, Vihan Jain, Eugene Ie, and Fei Sha. Baby-
818 Walk: Going farther in vision-and-language navigation by taking baby steps. In *Proceedings of*
819 *the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2539–2556.
820 *Association for Computational Linguistics*, 2020.

821

822 Wanrong Zhu, Jack Hessel, Anas Awadalla, Samir Yitzhak Gadre, Jesse Dodge, Alex Fang, Young-
823 jae Yu, Ludwig Schmidt, William Yang Wang, and Yejin Choi. Multimodal c4: An open, billion-
824 scale corpus of images interleaved with text. *Advances in Neural Information Processing Systems*,
36:8958–8974, 2023.

825

826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864 **A DATASET DETAILS**
865866 **A.1 DATA STATISTICS.**
867868 Here we show detailed statistics about ZEBRA-CoT’s categories.
869870 Table 3: Statistics of ZEBRA-CoT.
871

872 General Category	873 Sub Category	874 Count	875 Percentage (%)
876 2D Visual Reasoning	877 Visual Jigsaw	878 21,899	879 12.0
	880 Visual Search	881 30,000	882 16.4
	883 Subtotal	884 51,899	885 28.5
886 3D Visual Reasoning	887 Embodied Cot	888 22,666	889 12.4
	890 Multi-Hop Objects Counting	891 10,000	892 5.5
	893 Robot Planning	894 6,944	895 3.8
	896 Subtotal	897 39,610	898 21.7
899 Scientific Reasoning	900 Chemistry	901 4,666	902 2.6
	903 Competitive Programming	904 1,207	905 0.7
	906 Geometry	907 1,058	908 0.6
	909 Graph Algorithms	910 10,000	911 5.5
	912 Physics	913 7,090	914 3.9
	915 Subtotal	916 24,021	917 13.2
918 Visual Logic Strategic Games	919 Arc-Agi	920 2,000	921 1.1
	922 Checkers	923 2,753	924 1.5
	925 Chess	926 20,483	927 11.2
	928 Ciphers	929 6,589	930 3.6
	931 Connect Four	932 2,029	933 1.1
	934 Maze	935 20,000	936 11.0
	937 RPM	938 3,000	939 1.6
	940 Tetris	941 10,000	942 5.5
943 Total	944 Subtotal	945 66,854	946 36.7
		947 182,384	948 100.0

918
919

A.2 CURATING DIVERSE AND HIGH QUALITY VISUAL COT

920
921
922
923
924
925
926

Bridging logical connections across modalities. A key challenge in training multimodal generation models to output visual CoT natively is the lack of datasets with strong logical coherence between text and visual modalities, and diverse categories of such visual CoT. Existing interleaved datasets often fail to provide clear, meaningful connections that demonstrate when and why visual reasoning is necessary for problem-solving, while current visual CoT datasets are confined to a few domains, limiting the model’s ability to learn generalizable visual CoT capabilities when faced with out-of-distribution problems.

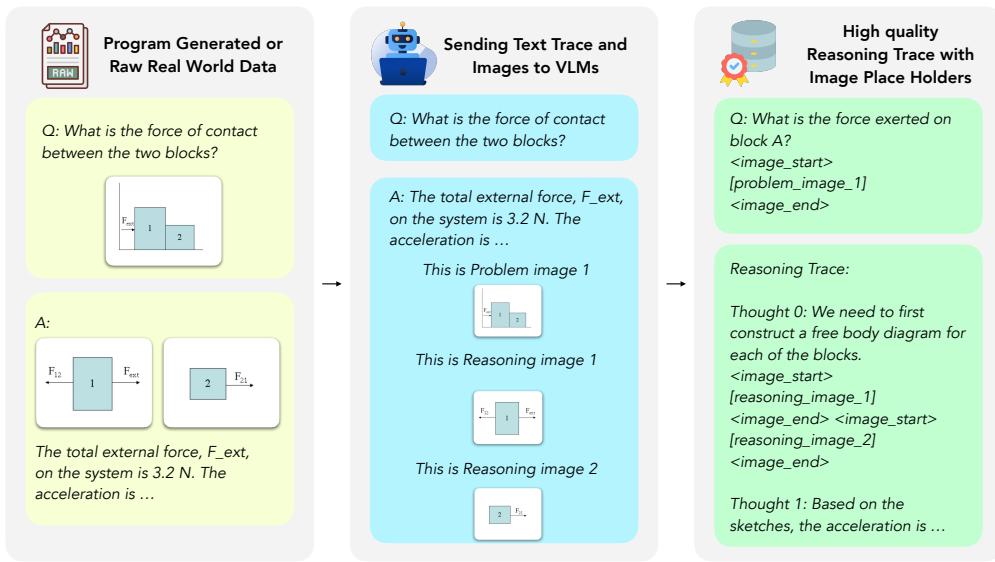
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947

Figure 5: An overview of our data curation pipeline.

To address these requirements, we first source a diverse range of question types and domains. For real world data, we source high-quality problems from online resources such as math, physics, coding, and chess competition datasets. We then extract and clean the available raw reasoning traces that contain text and images. However, even from high quality sources, traces can still lack clear logical connections between modalities, as well as clear references to the images for automatic parsing into interleaved text and image data ready for training. For example, most geometry data uses reference labels such as “Figure x ”, which makes it hard to find the mapping between the actual image and the text reference. For synthetic data, we create our own examples by generating images or utilizing real images from online sources, then crafting corresponding reasoning templates. This procedure raises a clear issue, namely that we lack diversity and expressiveness in textual reasoning regarding templated data. For instance, in visual search tasks, it is crucial to elucidate the rationale behind drawing specific bounding boxes, and in chess, generating reflections and descriptions of move visualizations is key.

We address both of these issues using frontier VLMs (Gemini-2.5 and GPT-4.1) to fill in the template placeholders, enhance the reasoning traces, and complete the textual reasoning narrative. We feed both images and raw text reasoning traces into the language model and ask the language model to output pure text traces with image placeholders. We further filter out invalid cases, such as multiple image placeholders referring to the same image and unreferenced image placeholders, to ensure that the data can be automatically parsed into a training dataset.

Broadening breadth and diversity of interleaved visual language reasoning dataset. Furthermore, existing multimodal rationale datasets are also limited in their breadth. The only available datasets focus on either visual search (Wu and Xie, 2024; Shao et al., 2024a) or spatial reasoning, such as maze navigation (Li et al., 2025). Such limited datasets are unlikely to enable training visual reasoning models that can generalize across domains more broadly. Visual Sketchpad (Hu et al., 2024) offers a diverse range of VLM agents to tackle a wider variety of questions. Though Sketchpad offers a powerful and significant contribution to generating visual aids, the pipeline is

972 not designed for collecting post-training datasets. First, the reasoning traces generated by agentic
973 pipelines often involve tool call errors and debug information, which degrade their quality. Second,
974 the scalability and diversity of the dataset are fundamentally constrained by the limited number of
975 agent tool designs and the high cost, as each reasoning trace may require many API calls. To tackle
976 these issues, we curate a total of over **182K** high-quality interleaved text and visual reasoning traces,
977 spanning four major categories: scientific reasoning, 2D visual reasoning, 3D visual reasoning, and
978 visual logic and strategic games. We provide the details in the section below and example traces
979 from our dataset.

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026
1027

A.3 SCIENTIFIC QUESTIONS

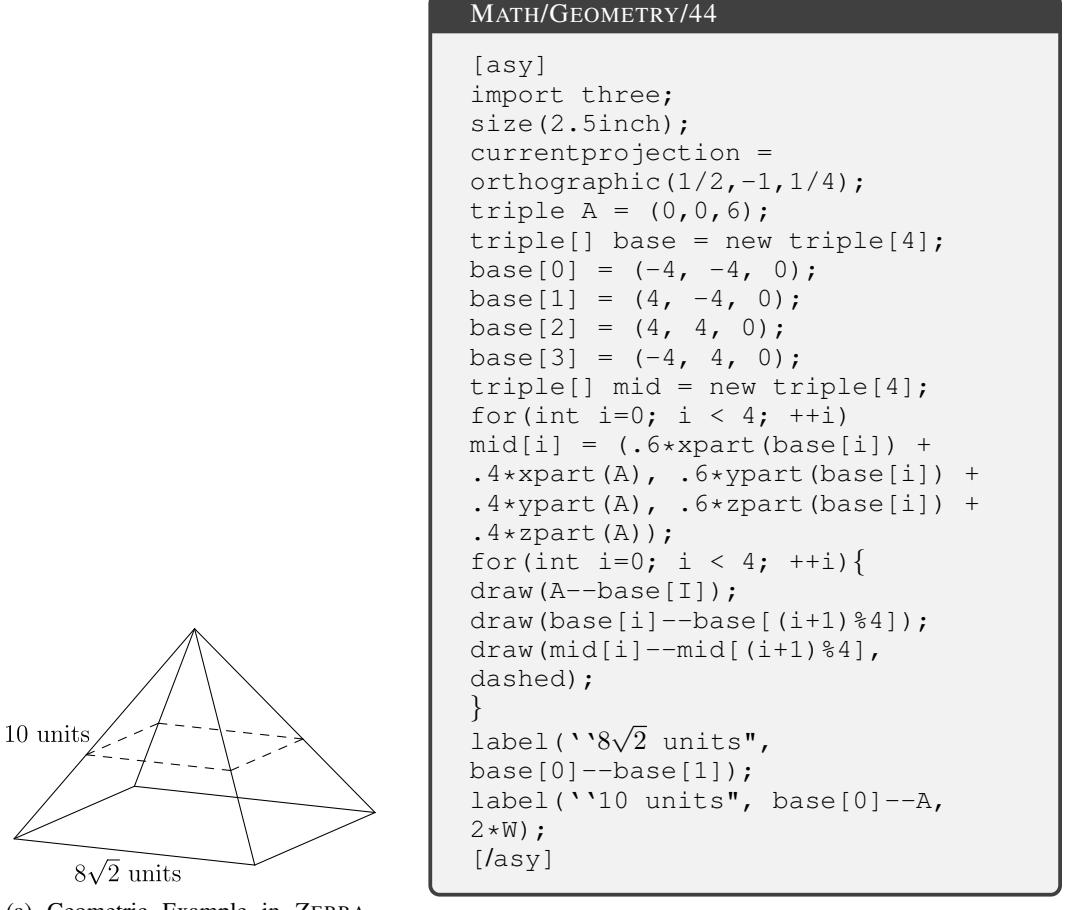
1028
1029
1030
1031
1032
1033

Geometry. Geometric understanding is a core ability for multimodal models to ground reasoning over complicated mathematical tasks. Many datasets have been proposed to evaluate mathematical capabilities, including geometry. The MATH dataset (Hendrycks et al., 2021) is widely used for evaluating the mathematical performance of LLMs. Although the MATH dataset includes numerous geometry competition problems, their geometric elements are provided as plotting code rather than rendered images (see Figure 6).

1034
1035

Here, we provide example code for geometry sketch generation.

1036

1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064

(a) Geometric Example in ZEBRA-CoT

(b) Geometric Example in MATH Dataset (Hendrycks et al., 2021)

1065
1066
1067

Figure 6: Comparison of the same geometric figure in our ZEBRA-CoT dataset and the MATH dataset. Ours focus on multimodal reasoning and explicitly plot the geometry problem than using the text-only plotting codes.

1071
1072

In ZEBRA-CoT, we convert every piece of plotting code into figure renderings, producing both the problem diagram and its solution illustration to serve as an explicit visual reasoning chain for model training.

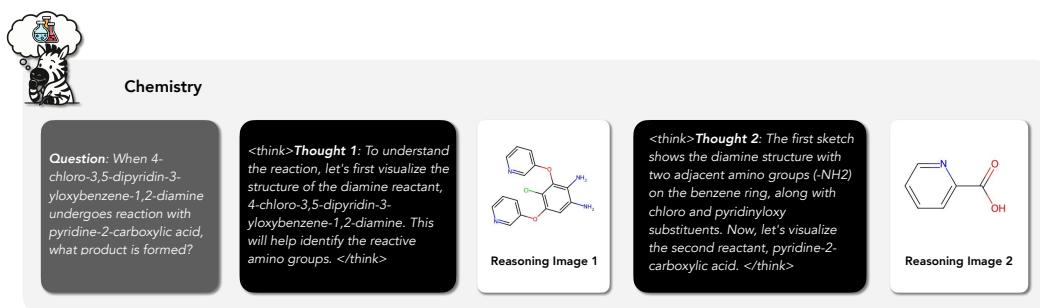
1076
1077
1078
1079

In total, we collect 1,061 samples from the MATH dataset’s train split. Our data provides only rendered images for both the problem and solution reasoning chains, with no plotting code included. Solving these problems requires generating images to assist. The problems are not restricted to the geometry subcategory but also include some problems from counting and probability, pre-algebra, pre-calculus, etc.

1080 **Physics.** A variety of physics problems benefit from sketches, such as free body diagrams for
 1081 force analysis, motion diagrams for kinematics, circuit diagrams for electricity, and ray diagrams in
 1082 optics. We construct samples of classical mechanics problems programmatically. Problem instances
 1083 are generated from parametric Python templates (e.g., Atwood machines, inclined planes, elastic
 1084 collisions, pendulums), with physically plausible parameters sampled from predefined ranges. For
 1085 each sample, we render free-body diagrams, kinematic visuals, and structured CoT traces capturing
 1086 the full solution process.

1087 We also leverage openly licensed resources such as OpenStax (MIT OpenCourseWare, 2022) and
 1088 MIT OCW (Moebs et al., 2016) to generate more diverse and complex physics problems, ultimately
 1089 achieving scalable and legally clear dataset generation while ensuring diverse, high-quality exam-
 1090 ples.

1091 **Chemistry.** Organic reaction prediction is a classic multimodal reasoning task, typically framed as
 1092 symbolic input and structural output. We include a chemistry subset of 4,700 two-to-one reactions
 1093 from the **USPTO-50K** dataset (Ramsundar et al., 2019), filtered for distinct reactants and single
 1094 products. Each reaction trace includes three visual artifacts: individual molecular depictions of each
 1095 reactant, a combined schematic of both reactants side-by-side, and the resulting product structure.
 1096 Molecules are rendered with `RDKit`, and names are retrieved from PubChem when available. Text
 1097 prompts use randomized templates (e.g., “What is formed by combining acetic acid and ethanol?”),
 1098 and PubChem names are included when available. This visual progression helps models learn com-
 1099 positional chemical structure without SMILES or reaction templates.



1100
 1101
 1102
 1103
 1104 **Algorithmic problem solving.** Humans naturally create visual diagrams when solving complex
 1105 problems, transforming abstract concepts into spatial representations for deeper reasoning. We
 1106 formalize this by interpreting coding problems through compact visual scaffolds: one or two diagrams
 1107 depicting graph structure, edge weights, etc. To build traces, we run an iterative “visual sketchpad”
 1108 loop: GPT-4.1 receives a prompt and returns THOUGHT statements plus VIS_SPEC blocks when
 1109 sketches are needed; we render specs with `networkx/matplotlib`, feed images back to the
 1110 model, and repeat until complete, then clean transcripts with post-processing.

1111
 1112
 1113
 1114 Problem samples come from competitive programming, prioritizing real-world abstractions like lo-
 1115 gistics, network routing, and flow optimization. The orchestrator produces simple visual structures
 1116 emphasizing clarity over style. Each trace contains the problem prompt, 1–3 reference diagrams, and
 1117 polished explanations, supporting grounded reasoning in discrete structures while mirroring how al-
 1118 gorithms are taught. The final corpus comprises 1,200 diverse algorithm-based problems spanning
 1119 competitive programming.

1120
 1121 **Graph problems.** Graph algorithms are useful for large language model applications because
 1122 they efficiently organize and traverse structured relationships, for example in search and retrieval
 1123 applications. Methods like shortest-path and subgraph matching enable multi-step reasoning by
 1124 connecting relevant concepts across knowledge graphs. Recent work by Fu et al. (2024a) shows
 1125 that although LLMs can solve graph problems such as connectivity and maximum flow to some
 1126 extent when a textual description of the graph is given, *multimodal* LLMs suffer when solving graph
 1127 problems. This finding suggests potential for improving multimodal models’ graph-understanding
 1128 abilities by guiding their reasoning over images.

1134 We create 10,000 graph problems with full reasoning traces spanning over four tasks: graph
 1135 connectivity, shortest path, minimum spanning tree, and topological sort. Each task has about 2,500
 1136 samples, with one problem image and at most 19 reasoning images per sample. Each reasoning
 1137 image is coupled with an explanation for the underlying algorithms, for example, Dijkstra for the
 1138 shortest path, BFS for connectivity, etc.

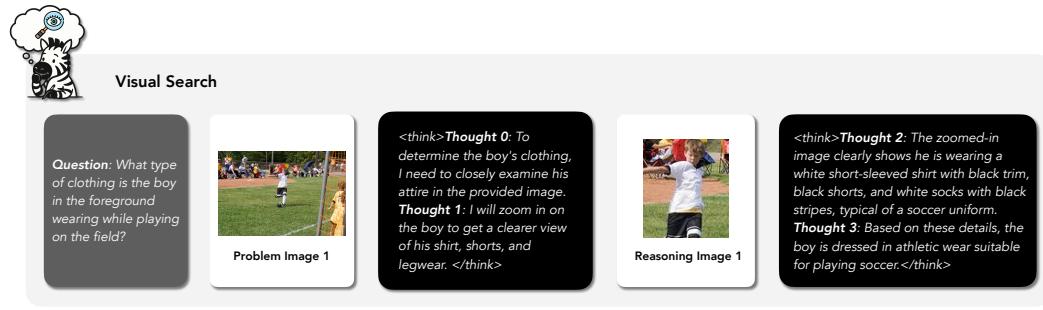
1139

1140 A.4 2D VISUAL REASONING

1141

1142 **Visual search.** Previous research has shown that drawing bounding boxes and zooming can im-
 1143 prove accuracy on visual search tasks (Wu and Xie, 2024; Shao et al., 2024a). We follow such tasks
 1144 by creating two types of traces, one for drawing bounding boxes and one for zooming. We use data
 1145 from Shao et al. (2024a) to generate our traces covering four categories of visual search tasks: chart,
 1146 text/doc, relation study, and general VQA.

1147



1148

1149 **Visual jigsaw.** Visual jigsaw refers to filling in missing pieces of an image, as in a jigsaw puzzle.
 1150 Each puzzle is constructed from an ImageNet (Deng et al., 2009) image, with 1 to 4 missing pieces
 1151 of varying shapes, including rectangles and irregular regions. Each puzzle includes four multiple-
 1152 choice options, where each option presents a set of candidate missing pieces. Only one set correctly
 1153 matches the pieces removed from the original ImageNet image. We generate two types of visual
 1154 CoT traces for solving each puzzle. In the first type, we iteratively fill in the missing patches using
 1155 the pieces from each multiple-choice option and identify the one that produces a coherent image.
 1156 In the second type, we imagine what the original image would look like and then select the option
 1157 whose pieces best match the imagined reconstruction.

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

A.5 3D VISUAL REASONING

1170

1171

1172

1173

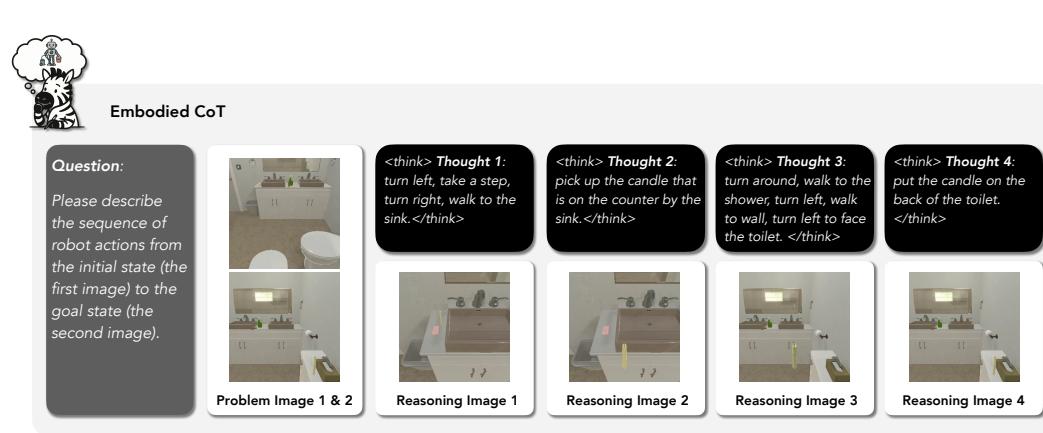
1174

1175

1176

1177

1178



1188 In this new task, the model receives two images: the initial and goal states. Then the model is
 1189 tasked with generating a textual description of the high-level planning steps required to transition
 1190 from the initial to the goal state. To emphasize the role of visual reasoning, we require the generated
 1191 descriptions to be detailed and step-by-step (e.g., *“turn and go to the TV; pick up the bowl that is on*
 1192 *the TV stand in front of the TV; with the bowl in hand...”*) rather than brief summaries (e.g., *“move*
 1193 *bowl to coffee table”*), which can often be produced through shortcut reasoning without capturing
 1194 intermediate visual steps.

1195 We compile the entire training set, as well as the seen and unseen validation sets from ALFRED,
 1196 resulting in a total of 7,080 examples spanning diverse visual reasoning trajectories. When multiple
 1197 textual reasoning annotations exist for a single visual trajectory, we include all of them, resulting in
 1198 22,666 textual reasoning traces.

1199 **Robot planning.** While low-level manipulation may rely on reactive control, continuous planning
 1200 for complex tasks often requires *high-level visual guidance*, making visual CoT essential for
 1201 bridging perception and long-horizon decision-making in robot planning. Similarly, we reformulate
 1202 **RoboMIND** (Wu et al., 2024), a multi-embodiment dataset of real-world robot manipulation, into
 1203 an image goal-conditioned planning task. In this setting, a model is provided with the initial and goal
 1204 states images, along with a textual description of the robot setup (e.g., AgileX (AgileX Robotics,
 1205 2023), Franka (Franka Emika GmbH, 2018), or UR5e (Universal Robots A/S, 2018)), and is tasked
 1206 with generating a detailed textual plan outlining the high-level steps required to transition from the
 1207 initial to the goal state.

1208 Unlike embodied planning tasks that often involve partial observability and require agents to infer
 1209 unobserved states, this robot planning task is fully observable. Therefore, the challenge lies not
 1210 in imagining the visual trajectory but in articulating precise movements for each arm or gripper to
 1211 accomplish the task (e.g., *“[left] move towards the oven door and [right] grab the corn.”*).

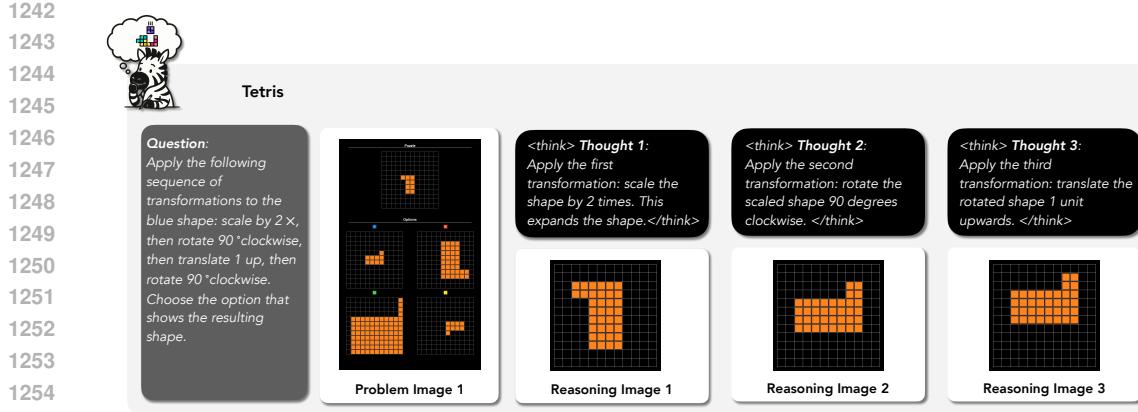
1212 To control degrees of freedom, we exclude the humanoid robot examples from the original Robo-
 1213 MIND dataset, focusing solely on tasks involving robotic arms. This results in a curated subset of
 1214 6,945 robot planning tasks, each annotated with human-generated high-level actions that serve as
 1215 visual reasoning trajectories.

1216 **3D multi-hop objects counting.** A core aspect of human visual-spatial reasoning is understanding
 1217 transformations and imagining scenes from different viewpoints. For this task, our setup follows
 1218 a structure similar to that of Johnson et al. (2017), using 10 predefined shape types (e.g., sphere,
 1219 cylinder, donut) in various colors. At each step, we randomly apply one of three operations: remove
 1220 all instances of an attribute (e.g., all red objects), remove a subset (e.g., 5 red objects), or add new
 1221 objects (e.g., 2 blue prisms, 1 red sphere). We then create questions that ask about the quantity
 1222 of specific attributes or what objects are left in the field. To increase difficulty, the initial scenes
 1223 are rendered from varying viewpoints (front, back, left, right), where some objects may be partially
 1224 occluded by those in front. The first visual reasoning step involves generating a top-down 45° view
 1225 to reconstruct the full scene, allowing the model to see potentially blocked objects. The subsequent
 1226 visual sketches correspond to each transformation step in the instruction. We also improve upon the
 1227 data from Johnson et al. (2017) by adding in different materials, backgrounds, and floor designs.

1228 A.6 VISUAL LOGIC AND STRATEGIC GAMES

1229 **Visual logic puzzles.** Humans approach logic puzzles such as Tetris, Raven’s Progressive Matrices
 1230 (RPM, Zhang et al., 2019), and the Abstraction & Reasoning Corpus (ARC-AGI, Chollet, 2019;
 1231 Chollet et al., 2024) primarily through visuospatial reasoning: we see how pieces combine, trans-
 1232 form, or complete a pattern before committing to an answer. These logic games rely heavily on
 1233 visuospatial working memory, which is correlated with general intelligence level (Lau-Zhu et al.,
 1234 2017; de Winter et al., 2023).

1235 To enhance models with such cognitive ability, we include the following tasks. For *Tetris*, we collect
 1236 three types of tasks: a) shape assembly: given a silhouette and candidate tetromino sets, select the
 1237 one that perfectly tiles the shape; b) grid completion: fill a partially occupied grid using a specified
 1238 set of tetrominoes; c) spatial transformation: apply a sequence of geometric operations (translate,
 1239 rotate, mirror, scale) to an irregular shape in the grids. The visual CoT involves visualizing each
 1240 transformation step. For *RPM* (IQ matrix), we include three types from Zhang et al. (2019) that



involve compositional reasoning. The reasoning trace identifies visual patterns for each compositional component across rows or columns. For *ARC-AGI*, while prior models often rely on textual reasoning, humans typically solve these tasks through visual pattern recognition. To better align with human strategies, we construct two types of visual CoT. The first begins with matrix representations of the training examples and test input; the reasoning trace first visualizes the training examples, the test input, and finally the predicted output. The second type directly uses visual representations in the task instruction, thus the model only has to generate a visual sketch of the predicted output as part of its reasoning process. For all data, we use VLM to generate accompanying textual descriptions to enrich interleaved text-image rationales.

Mazes. Mazes serve as a canonical testbed for visual CoT reasoning, bridging low-level perception with high-level symbolic search. Unlike purely pixel-based 2D visual tasks such as visual search and visual jigsaw, mazes possess explicit graph structure yet remain visually intuitive, letting us disentangle vision errors from planning errors.

We adopt the `maze-dataset` library to procedurally generate thousands of grid mazes with diverse topologies (lattice type, branch factor, loop density).¹ Each instance is exported in two complementary formats: a) `m.as_pixels()`, an RGB raster that encodes walls, free cells, start ■, and goal ■, suitable for visual perception; b) `MazePlot`, a vector overlay that can superimpose solution paths, candidate trajectories, heat-maps, or landmark nodes for human-readable walk-throughs. To increase maze diversity, we also use OpenAI Gym’s `FrozenLake-v1` environment (Brockman et al., 2016).

We evaluate a broad spectrum of spatial reasoning skills across multiple question types: *I. topological analysis* (e.g., counting isolated regions, identifying connected components under 4- or 8-connectivity, finding the largest connected area), *II. pathfinding* (e.g., determining reachable endpoints, computing shortest paths, enumerating all optimal routes), *III. navigation planning* (e.g., selecting correct paths from alternatives, calculating minimal moves to reach targets), and *IV. coverage problems* (e.g., visiting all marked locations, identifying the farthest reachable position). This diverse task suite goes beyond simple start-to-goal navigation, encompassing the full range of spatial reasoning strategies that humans use to interpret complex environments. We also introduce varying complexity of the matrix, including different maze side lengths ranging from (5, 15), different branching factors b , loop probability ℓ , and number of distractor endpoints k . Larger n exponentially increases the search space, while higher b and ℓ degrade heuristic admissibility. Both of those require genuine planning rather than rote memorization.

Chess.

Strategic planning in chess involves simulating multiple futures and selecting moves that maximize long-term advantage. To support counterfactual reasoning, we construct a dataset of mid-game positions from rated Lichess games², each with structured visual traces. Given a position, Stockfish

¹`maze-dataset` supports recursive-backtracker, randomized Prim, Wilson, and Kruskal generators; see (Ivanitskiy et al., 2023).

²<https://lichess.org/>

1296
1297
1298
1299
1300
1301



1302 What's White's best move? Options: A: Ba2, B: Na4, C: Qf5, D: Bb3.

1303
1304
1305
1306
1307
1308
1309



(A) **Ba2**: Safe position (B) **Na4**: Poorly placed (C) **Qf5**: Exposed queen (D) **Bb3**: Vulnerable position
supports central pawns weak attack on b6 vulnerable to g6 weaker than Ba2

1310
1311
1312 Figure 7: Traces showing reasoning for each move option. Option A (Ba2) is evaluated as strongest,
1313 providing safe bishop placement while supporting potential central pawn advances.

1314
1315
1316
1317
1318
1319
1320
1321
1322

identifies the optimal move, and three alternates are sampled randomly from legal moves. Each candidate is visualized independently for comparative evaluation. By rendering possibilities in isolation, move consequences, tempo gain, structural weakening, and tactical motifs become legible, enabling better strategic reasoning. Traces are formatted as multiple-choice tasks with visual sketches, encouraging tactical awareness and pattern recognition. Postprocessing with Gemini 2.5 Flash refines traces into coherent multimodal reasoning sequences for short and long-horizon planning, yielding 11,500 chess traces across diverse positions and hypothetical move states.

1323
1324
1325
1326
1327

Connect 4. We also include 2,100 Connect-Four traces that teach short-horizon grid planning. Each starts from a mid-game position obtained after 4–20 random moves; a depth-4 α – β minimax (center-control, open-three heuristic) then plays the forced win. Every drop is rendered on a 7×6 board with a glowing overlay, and an LLM condenses the raw log into a clear, step-by-step “How does Red win?” narrative pairing text and visuals.

1328
1329
1330
1331
1332
1333

Checkers. Our checkers traces, initiated from randomized mid-game boards and continued through depth-4 minimax, are designed to capture key tactical motifs such as forced captures, multi-jump chains, and king-row advancement, providing dense supervision for spatially grounded decision making. To ensure data quality, we apply rule-based validators that discard invalid traces. Final traces are passed through a language model to synthesize coherent visual reasoning aligned with *turn-based planning*, ending with 2,800 full traces.

1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Ciphers. Classical encryption schemes convert symbolic rules into spatial transformations, enabling visual reasoning over grids and coordinates. We generate examples of *Scytale*, *Polybius*, *Rail Fence* and *Playfair* ciphers, each with multimodal traces showing encoding steps. Visual sketches highlight relevant regions such as column positions, grid lookups, or digraph geometry, while text describes applied rules (e.g., “*Letter ‘P’ goes on the second rail*,” “*A’ is at coordinates (1, 1)*.”) Early steps use visuals, later steps use language, mirroring real cryptanalysis. The 6,600 traces align symbolic manipulation with visual structure.

1350 B EXTRA BAGEL ZEBRA CoT GENERATED VISUAL CoT
1351

1352

1353 Visual CoT Example: EMMA
1354

1355

1356 **Prompt:**

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

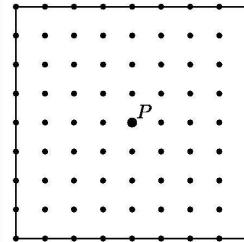
1402

1403

Visual CoT Example: EMMA

Prompt:

There are 81 grid points (uniformly spaced) in the square shown in the diagram below, including the points on the edges. Point P is the center of the square. Given that point Q is randomly chosen from among the other 80 points, what is the probability that line PQ is a line of symmetry for the square?

**Visual CoT:**

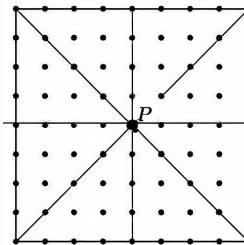
<think>

The problem asks for the probability that a randomly chosen point Q from the 80 points other than the center P results in the line PQ being a line of symmetry for the square. We need to determine the number of such points Q and divide by the total number of possible points (80).

First, let's visualize the grid and the center point P . The grid has 9 points along each side, so there are 80 points in total (excluding P). A line of symmetry for a square passes through the center P and divides the square into two congruent halves. For PQ to be a line of symmetry, Q must be the reflection of P across this line. This means Q must be on the grid line that forms a line of symmetry with P .

Let's consider the lines of symmetry for the square. A square has 4 lines of symmetry: 2 diagonals, the vertical line through the center, and the horizontal line through the center. Let's visualize these lines.

</think>



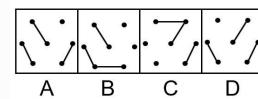
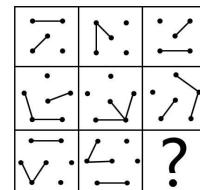
[Visual Thought]

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413

1414 Visual CoT Example: VisuLogic

1415 **Prompt:**

1416 Solve the complex visual logical reasoning problem
1417 through step-by-step reasoning.

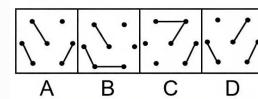
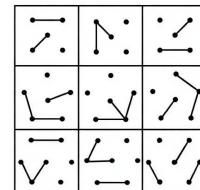


1429 **Visual CoT:**

1430 <think>

1431 Observe the grid to identify patterns in the shapes and their
1432 components. The grid is divided into 3x3 cells. Each cell
1433 contains two distinct shapes: a line segment and a dot. I
1434 will analyze the patterns for the line segment and the dot
1435 separately. For the line segment, the pattern is: (Top Left,
1436 Top Right) = (Bottom Left, Bottom Right). For the dot, the
1437 pattern is: (Top Left, Top Right) = (Bottom Left, Bottom
1438 Right). Now, I will apply these patterns to the third row and
1439 third column to find the missing shape. The missing shape
1440 in the third row and third column should have a line segment
1441 like the one in the third row, third column (which is the same
1442 as the third row, second column line segment). The missing
1443 shape in the third row and third column should also have a
1444 dot like the one in the third row, third column (which is the same as the third row,
1445 second column dot). Combining these, the missing shape should be a line segment
1446 from the third row, second column and a dot from the third row, second column.

1447 </think>



1449 **[Visual Thought]**

1450
1451
1452
1453
1454
1455
1456
1457

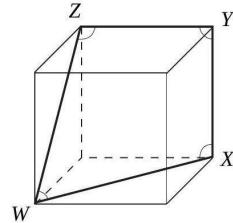
1458
1459
1460
1461
1462
1463
1464
1465

Visual CoT Example: EMMA

Prompt:

1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

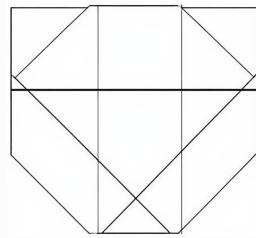
The picture shows a cube with four marked angles: $\angle WXY$, $\angle XYZ$, $\angle YZW$, and $\angle ZWX$. What is the sum of these angles?



Visual CoT:

<think>

The problem asks for the sum of four specific angles in a cube. The angles are $\angle WXY$, $\angle XYZ$, $\angle YZW$, and $\angle ZWX$. These angles are formed by the edges of the cube meeting at vertices X, Y, Z, W .



[Visual Thought]

In a cube, each vertex is formed by three edges that are mutually perpendicular. This means that at each vertex, the three edges meet to form right angles.

Let's consider vertex X . The edges WX , XY , and XZ meet at vertex X . Since the edges are perpendicular, the angle $\angle WXY = 90^\circ$.

Similarly, we can analyze the other vertices:

- At vertex Y : edges XY , YZ are perpendicular, so $\angle XYZ = 90^\circ$
- At vertex Z : edges YZ , ZW are perpendicular, so $\angle YZW = 90^\circ$
- At vertex W : edges ZW , WX are perpendicular, so $\angle ZWX = 90^\circ$

We have identified the measures of all four angles: $\angle WXY = 90^\circ$, $\angle XYZ = 90^\circ$, $\angle YZW = 90^\circ$, $\angle ZWX = 90^\circ$.

To find the sum of these angles, we add their measures:

$$90^\circ + 90^\circ + 90^\circ + 90^\circ = 360^\circ$$

The sum of the angles $\angle WXY$, $\angle XYZ$, $\angle YZW$, and $\angle ZWX$ is **360°**.

</think>

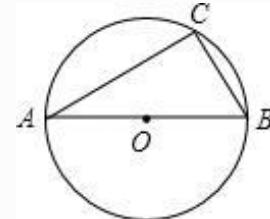
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523

1524 **Visual CoT Example: MathVista**

1525
1526
1527

Prompt:

1528 Question: As shown in the figure, AB is the diameter of
1529 $\odot O$, and point C is on $\odot O$. If $\angle A = 40.0$, then the degree
1530 of $\angle B$ is?



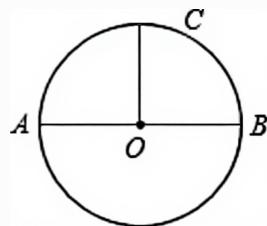
1531
1532
1533
1534
1535
1536

Visual CoT:

1537
1538
1539
1540
1541
1542
1543

<think>

First, let's analyze the geometry of the problem. Since AB is the diameter of the circle, any angle inscribed in the semicircle formed by AB will be a right angle. Thus, angle ACB is a right angle, measuring 90 degrees. Now, let's draw an auxiliary line connecting point O (the center of the circle) to point C . This line segment OC is a crucial part of our reasoning. The line segment OC can be labeled as shown in the image.



[Visual Thought]

Since AB is the diameter of the circle, O is the center. Therefore, AO and BO are radii of the circle, making them equal in length. By the properties of an isosceles triangle ($AO = BO$), and having an angle of 90 degrees at C (i.e., angle ACB), we can deduce that triangle ABC is an isosceles right-angled triangle. This means that angle A is equal to angle CBA , and since angle ACB is 90 degrees, the two other angles must add up to 90 degrees. Given that angle A is 40 degrees, angle B must be 50 degrees. Therefore, angle B is 50 degrees.

</think>

1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580

1581 Visual CoT Example: BLINK

1582 **Prompt:**

1583 How many kids are sitting on the bench in the image?



1593 **Visual CoT:**

1594 **<think>**



1595 Let's count the number of kids sitting on the bench in the
1596 image. To get a clearer view of the kids on the bench, I
1597 will zoom in on the relevant area. The zoomed-in image
1598 clearly shows four children sitting on the bench. Based on
1599 this visual evidence, the number of kids on the bench is four.
1600 **</think>**

1601 **[Visual Thought]**

1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

1620 C ANOLE-ZEBRA PERFORMANCE BREAKDOWN
1621

Split	Anole	Anole-Zebra-CoT (Ours)
Overall	12.80	15.03
Chemistry	12.84	15.48
Coding	9.75	16.31
Math	13.12	14.35
Physics	21.79	10.90

1629
1630 Table 4: EMMA: breakdown by subject (%).
1631

Subtask	Anole	Anole-Zebra-CoT (Ours)
Overall	22.80	24.90
Scientific reasoning	30.33	32.79
Textbook question answering	36.08	29.75
Numeric commonsense	16.67	17.36
Arithmetic reasoning	15.58	18.98
Visual question answering	24.58	29.61
Geometry reasoning	20.50	23.01
Algebraic reasoning	25.27	24.56
Geometry problem solving	21.15	24.04
Math word problem	9.14	12.37
Logical reasoning	29.73	10.81
Figure question answering	24.54	28.25
Statistical reasoning	20.27	26.58

1646
1647 Table 5: MathVista: breakdown by subtask for base vs. our model (%).
1648

Subtask	Anole	Anole-Zebra-CoT (Ours)
Overall	8.50	21.80
Quantitative reasoning	8.78	21.81
Spatial reasoning	8.23	22.08
Positional reasoning	8.82	19.85
Attribute reasoning	9.76	25.61
Stylistic reasoning	10.00	24.44
Other	5.56	18.52

1658
1659 Table 6: Visual Logic: breakdown by subtask (%).
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

1674
 1675
 1676
 1677
 1678
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692

Category	Anole	Anole-Zebra-CoT (Ours)
Overall	26.46	31.25
Art Style	19.66	35.04
Counting	19.17	15.00
Forensic detection	0.00	20.45
Functional correspondence	17.69	22.31
IQ test	26.67	23.33
Jigsaw	11.33	39.33
Multi-view reasoning	48.12	21.05
Object localization	50.82	45.90
Relative depth	38.71	41.94
Relative reflectance	29.10	27.61
Semantic correspondence	19.42	17.99
Spatial relation	41.26	57.34
Visual correspondence	21.51	26.16
Visual similarity	30.37	44.44

Table 7: Blink: breakdown by category (%).

1708
 1709
 1710
 1711
 1712
 1713
 1714
 1715
 1716
 1717
 1718
 1719
 1720
 1721
 1722
 1723
 1724
 1725
 1726
 1727

D BAGEL PERFORMANCE ANALYSIS

We evaluate our Bagel model trained on ZEBRA-CoT across several benchmarks but did not observe substantial improvements over the original model, where the original generates pure text responses. In fact, we even saw slight performance drops on some tasks such as MathVista. A detailed analysis revealed a likely cause of this decline. The Bagel model employs two visual encoders: a ViT-based understanding encoder and a VAE-based generation encoder. For generated images, the model often produces hallucinations. For example, when instructed to remove all red balls from a scene, the generated image may also remove yellow balls. When this corrupted image is passed back into the ViT encoder, the representation correctly reflects that both red and yellow balls are missing, leading the model to reason over inaccurate visual information, ultimately reducing accuracy. Instead generating pure text responses avoids such image generation hallucinations.

1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

E SCAFFOLDING RESULTS BREAKDOWN

Chess					
Model	Q (%)	1MT (%)	2MT (%)	Δ 1MT (%)	Δ 2MT (%)
Claude-4 Sonnet	32.95	57.95	67.05	25.00	34.09
Gemini-2.5 Pro	15.07	39.73	39.73	24.66	24.66
GPT-5	45.78	62.65	61.45	16.87	15.66
Graph					
Model	Q (%)	1MT (%)	2MT (%)	Δ 1MT (%)	Δ 2MT (%)
Claude-4 Sonnet	8.11	20.72	22.52	12.61	14.41
Gemini-2.5 Pro	1.90	11.43	20.95	9.52	19.05
GPT-5	1.74	14.78	22.61	13.04	20.87
2D Visual Jigsaw					
Model	Q (%)	1MT (%)	2MT (%)	Δ 1MT (%)	Δ 2MT (%)
Claude-4 Sonnet	21.74	36.23	62.32	14.49	40.58
Gemini-2.5 Pro	34.38	56.25	59.38	21.88	25.00
GPT-5	62.86	77.14	85.71	14.29	22.86
Maze					
Model	Q (%)	1MT (%)	2MT (%)	Δ 1MT (%)	Δ 2MT (%)
Claude-4 Sonnet	35.06	58.44	94.81	23.38	59.74
Gemini-2.5 Pro	59.70	85.07	97.01	25.37	37.31
GPT-5	63.01	86.30	97.26	23.29	34.25
3D Multi-Hop Counting					
Model	Q (%)	1MT (%)	2MT (%)	Δ 1MT (%)	Δ 2MT (%)
Claude-4 Sonnet	59.68	67.74	74.19	8.06	14.52
Gemini-2.5 Pro	45.95	48.65	51.35	2.70	5.41
GPT-5	72.58	74.19	77.42	1.61	4.84
Tetris					
Model	Q (%)	1MT (%)	2MT (%)	Δ 1MT (%)	Δ 2MT (%)
Claude-4 Sonnet	18.87	39.62	45.28	20.75	26.42
Gemini-2.5 Pro	8.57	25.71	31.43	17.14	22.86
GPT-5	30.77	50.00	59.62	19.23	28.85

Table 8: Scaffolding evaluation results across task domains. Q: zero-shot question-only; 1MT: question with first multimodal reasoning step; 2MT: question with first two multimodal reasoning steps. Δ columns show absolute improvement over baseline (Q).

Chess					
Model	Q (%)	1TT (%)	2TT (%)	Δ 1TT (%)	Δ 2TT (%)
Claude-4 Sonnet	32.95	42.05	40.91	9.10	7.96
Gemini-2.5 Pro	15.07	13.70	19.86	-1.37	4.79
GPT-5	45.78	51.81	62.65	6.03	16.87
Graph					
Model	Q (%)	1TT (%)	2TT (%)	Δ 1TT (%)	Δ 2TT (%)
Claude-4 Sonnet	8.11	20.72	15.32	12.61	7.21
Gemini-2.5 Pro	1.90	5.71	5.71	3.81	3.81
GPT-5	1.74	11.30	19.13	9.56	17.39
2D Visual Jigsaw					
Model	Q (%)	1TT (%)	2TT (%)	Δ 1TT (%)	Δ 2TT (%)
Claude-4 Sonnet	21.74	39.13	30.43	17.39	8.69
Gemini-2.5 Pro	34.38	24.22	25.00	-10.16	-9.38
GPT-5	62.86	68.57	71.43	5.71	8.57
Maze					
Model	Q (%)	1TT (%)	2TT (%)	Δ 1TT (%)	Δ 2TT (%)
Claude-4 Sonnet	35.06	49.35	55.84	14.29	20.78
Gemini-2.5 Pro	59.70	23.29	39.73	-36.41	-19.97
GPT-5	63.01	63.01	75.34	0.00	12.33
3D Multi-Hop Counting					
Model	Q (%)	1TT (%)	2TT (%)	Δ 1TT (%)	Δ 2TT (%)
Claude-4 Sonnet	59.68	69.35	69.35	9.67	9.67
Gemini-2.5 Pro	45.95	54.84	56.45	8.89	10.50
GPT-5	72.58	72.58	74.19	0.00	1.61
Tetris					
Model	Q (%)	1TT (%)	2TT (%)	Δ 1TT (%)	Δ 2TT (%)
Claude-4 Sonnet	18.87	26.42	32.08	7.55	13.21
Gemini-2.5 Pro	8.57	11.54	7.69	2.97	-0.88
GPT-5	30.77	28.85	42.31	-1.92	11.54

Table 9: Text-only CoT evaluation results across task domains. Q: zero-shot question; 1TT: first text reasoning step; 2TT: first two text reasoning steps. Δ columns show absolute improvement over Q.

1890 **F PROMPT TEMPLATES**
1891
1892
1893 **F.1 PROMPT FOR ENHANCING RAW REASONING TRACES FOR ONLINE AND AGENTIC DATA**
1894
1895 **Prompt Template 1**
1896
1897 You are an expert in creating clean and logically coherent
1898 → multimodal chain of thought traces. Your task is to
1899 → analyze
1900 and comprehend a raw reasoning trace with interleaved text
1901 → and images, then transform it into a clean, step-by-step
1902 → multimodal
1903 reasoning trace that correctly solves the original problem.
1904 ===== INPUT =====
1905 1. Problem & Noisy Trace: A raw interleaved text and image
1906 → reasoning trace. Images in this trace are represented by
1907 → placeholders:
1908 → - `'[problem_image_X]'` for original problem images (e.g.,
1909 → → `'[problem_image_1]'`, `'[problem_image_2]'`)
1910 → - `'[reasoning_image_X]'` for images generated during
1911 → → reasoning (e.g., `'[reasoning_image_1]'`,
1912 → → `'[reasoning_image_2]'`)
1913 2. Image Data: The actual image data corresponding to the
1914 → placeholders, provided separately.
1915 ===== Your Task =====
1916 Generate a clean, logical multimodal reasoning trace as
1917 → **plain text** that represents the *ideal* reasoning
1918 → process to solve the problem.
1919 ===== OUTPUT FORMAT =====
1920 You MUST generate the formatted reasoning trace with the
1921 → following structure:
1922
1923 **QUESTION:**
1924 <The original problem statement with text and image
1925 → placeholders: <image_start>[problem_image_1]<image_end>,
1926 → <image_start>[problem_image_2]<image_end>, etc. Stay as
1927 → close to the original problem statement as possible but
1928 → remove noise to ensure clarity>
1929
1930 **REASONING TRACE:**
1931 THOUGHT 0: <Clear description of initial reasoning step that
1932 → identifies key elements of the problem>
1933 THOUGHT 1: <Next reasoning step, often explaining why an
1934 → image will be created>
1935 <image_start>[reasoning_image_1]<image_end>
1936 THOUGHT 2: <Further reasoning step based on the image,
1937 → explaining insights gained>
1938 <image_start>[reasoning_image_2]<image_end>
1939 // Additional thoughts and images as needed
1940 <image_start>[reasoning_image_X]<image_end>
1941 THOUGHT N: <Final reasoning step before the answer,
1942 → summarizing key insights>
1943
1944 **FINAL ANSWER:**

```

1944
1945 <The final calculated answer based on the reasoning>
1946
1947 ===== Guidelines =====
1948
1949 1. Enhancing Original Trace Rather than Generating New Trace:
1950   - Instead of generating a new trace, your task is to
1951     ↳ enhance the original trace (which is a correct trace
1952     ↳ but rather concise and lacks coherent multimodal
1953     ↳ reasoning) by adding more details and explanations, see
1954     ↳ the following sections of guidelines for more details.
1955   - You MUST use all the images provided in the original
1956     ↳ trace.
1957   - You should use the original trace as a reference rather
1958     ↳ than copying it verbatim.
1959
1960 2. Multimodal Reasoning Flow:
1961   - Develop a coherent, step-by-step chain of thought that
1962     ↳ seamlessly integrates textual and visual reasoning.
1963   - Clearly explain the necessity of generating a sketch /
1964     ↳ visual thought / image before introducing its
1965     ↳ placeholder.
1966   - After each image placeholder, describe the insights
1967     ↳ gained from the sketch / visual thought / image, and
1968     ↳ how it contributes to advancing the solution.
1969   - Ensure each step logically builds on the previous ones,
1970     ↳ especially between text reasoning and visual reasoning
1971     steps.
1972
1973 3. Image Placeholders and References:
1974   - Use placeholder tags ONLY when you want to actually
1975     ↳ insert/show/generate an image in your trace. When
1976     ↳ doing so, write the corresponding placeholder tag
1977     ↳ exactly as shown, including the <image_start> and
1978     ↳ <image_end> tags.
1979   - Each unique image in the original problem and the
1980     ↳ reasoning trace should be represented by a unique
1981     ↳ placeholder tag, and each unique placeholder tag
1982     ↳ should only show up once in the trace.
1983   - When referring to images in your explanations, use
1984     ↳ natural language descriptions (e.g., "the diagram in
1985     ↳ the question", "the first sketch", "the visual thought
1986     ↳ X I created") instead of using placeholder tags. This
1987     ↳ is important because it helps us to parse into
1988     ↳ interleaved text and image sequences.
1989   - For images from the original problem, use:
1990     ↳ <image_start>[problem_image_X]<image_end>
1991   - For sketches or visuals generated during reasoning, use:
1992     ↳ <image_start>[reasoning_image_X]<image_end>
1993
1994 4. Narrative Style:
1995   - Remove irrelevant technical details such as debugging
1996     ↳ info, code snippets, and LaTeX package imports.
1997   - Eliminate verbose language that do not contribute to
1998     ↳ solving the problem.
1999   - Focus on the essential reasoning path that leads to the
2000     ↳ correct solution, using concise and clear language to
2001     ↳ describe the overall reasoning process.
2002

```

1998

F.2 PROMPT FOR ENHANCING PROGRAM GENERATED TEMPLATE DATA

1999

2000

2001

Prompt Template 1

2002

You are an expert in enhancing multimodal reasoning traces.
 → Your task is to transform a template reasoning trace into
 → a diverse multimodal reasoning trace that correctly
 → solves the problem, while staying close to the original
 → template and final answer.

2008

===== INPUT =====

1. Problem & Template Trace: A template with interleaved text
 → and image placeholders:
 → - `'[problem_image_X]'` for original problem images (e.g.,
 → → `'[problem_image_1]'`)
 → - `'[reasoning_image_X]'` for images generated during
 → → reasoning (e.g., `'[reasoning_image_1]'`)
 2. Image Data: The actual image data corresponding to the
 → placeholders, provided separately.

2017

===== Your Task =====

Generate a concise multimodal reasoning trace as ****plain**
 → **text****.

2021

===== OUTPUT FORMAT =====

You MUST generate the formatted reasoning trace with the
 → following structure:

2024

QUESTION:

<Rewrite the problem statement in your own words while
 → maintaining all key information. Do not change key
 → information. Include image placeholders:
 → - <image_start>[problem_image_1]<image_end>,
 → - <image_start>[problem_image_2]<image_end>, etc.>

2031

REASONING TRACE:

THOUGHT 0: <Identify key elements of the problem>
 THOUGHT 1: <Explain reasoning step, often why an image /
 → sketch / visual thought is needed>
 <image_start>[reasoning_image_1]<image_end>
 THOUGHT 2: <Explain insights from the image>
 <image_start>[reasoning_image_2]<image_end>
 // Additional thoughts and images as needed
 <image_start>[reasoning_image_X]<image_end>
 THOUGHT N: <Summarize key insights before answer>

2041

FINAL ANSWER:

<The original final answer in the template, do not change it>

2044

===== Guidelines =====

2045

1. Diversifying the Template:

- Rewrite the problem statement and reasoning steps in
 → your own words while preserving all key information.
- Avoid deviating from the original template reasoning
 → structure. Your job is to diversify the text of the
 → original trace, not the logic.

2052
 2053 - Vary the language and phrasing to avoid repetitive
 2054 ↳ patterns.
 2055 - You MUST use all the images provided in the original
 2056 ↳ trace.
 2057 - You MUST keep the original final answer.
 2058 - Maintain the original template's core reasoning
 2059 ↳ structure and rationale while introducing textual
 2060 ↳ reasoning refinements rather than substantial changes
 2061 ↳ to the logical flow.

2062 2. Multimodal Reasoning Flow:
 2063 - Develop a coherent, step-by-step chain of thought that
 2064 ↳ seamlessly integrates textual and visual reasoning.
 2065 - Clearly explain the necessity of generating a sketch /
 2066 ↳ visual thought / image before introducing its
 2067 ↳ placeholder.
 2068 - After each image placeholder, describe the insights
 2069 ↳ gained from the sketch / visual thought / image, and
 2070 ↳ how it contributes to advancing the solution.
 2071 - Ensure each step logically builds on the previous ones,
 2072 ↳ especially between text reasoning and visual reasoning
 2073 ↳ steps.

2074 3. Image Placeholders and References:
 2075 - Use placeholder tags ONLY when you want to actually
 2076 ↳ insert/show/generate an image in your trace. When
 2077 ↳ doing so, write the corresponding placeholder tag
 2078 ↳ exactly as shown, including the <image_start> and
 2079 ↳ <image_end> tags.
 2080 - Each unique image in the original problem and the
 2081 ↳ reasoning trace should be represented by a unique
 2082 ↳ placeholder tag, and each unique placeholder tag
 2083 ↳ should only show up once in the trace.
 2084 - When referring to images in your explanations, use
 2085 ↳ natural language descriptions (e.g., "the diagram in
 2086 ↳ the question", "the first sketch", "the visual thought
 2087 ↳ X I created") instead of using placeholder tags. This
 2088 ↳ is important because it helps us to parse into
 2089 ↳ interleaved text and image sequences.
 2090 - For images from the original problem, use:
 2091 ↳ <image_start>[problem_image_X]<image_end>
 2092 - For sketches or visuals generated during reasoning, use:
 2093 ↳ <image_start>[reasoning_image_X]<image_end>

F.3 PROMPT FOR ALGORITHMIC PROBLEMS

Prompt Template 2

2100 You are an expert in mathematical problem solving,
 2101 ↳ algorithmic reasoning, visual explanation, and creating
 2102 ↳ multimodal reasoning traces.

2103 ---

2104 1. STRICT VISUALIZATION POLICY (IMPORTANT):

2106
 2107 You are only allowed to produce at most 3 [VIS_SPEC]
 2108 → visualizations, and they must all appear at the very
 2109 → beginning of your reasoning (within the first 3--4
 2110 → thoughts). You may only use the following types for these
 2111 → visualizations:
 2112 - graph
 2113 - flow_network
 2114 - tree_from_dict
 2115 - tree_from_root
 2116 - grid
 2117
 2118 After these initial visualizations, you must do all further
 2119 → reasoning purely mentally/textually or with
 2120 → pseudocode--NO MORE [VIS_SPEC] blocks are allowed after
 2121 → the first 3. Any attempt to include more than 3
 2122 → visualizations or use a disallowed type will be ignored.
 2123 The visual reasoning should only be used to understand the
 2124 → setup of the question - humans visualize at the beginning
 2125 → to ``set the board.'' The actual problem solving is done
 2126 → purely textually.
 2127
 2128 ****General Rules:****
 2129 - Interleave THOUGHT steps and [VIS_SPEC] image requests.
 2130 - Your final reasoned solution must match the logic of the
 2131 → given solution code.
 2132 - Prefix THOUGHT 0 with REASONING TRACE in the previous line.
 2133 - Prefix each reasoning step with ``THOUGHT n:'' (n starts at
 2134 → 0, less than 50 words each).
 2135 - Max 3 [VIS_SPEC] blocks, all within the first 3--4
 2136 → thoughts.
 2137 - Diagram #1: raw structural sketch (graph topology, blank
 2138 → grid, etc.).
 2139 - Diagram #2--3: showcase pivotal elements if helpful.
 2140 - ****Internal self-check (no output):**** ``Would a human
 2141 → scribble this as a quick setup sketch?'' If the answer is
 2142 → no, ****do not**** emit a VIS_SPEC.
 2143 - Strictly do not regenerate the same image - simply refer to
 2144 → the previous images in text if needed.
 2145 - Max of 10 thoughts.
 2146 - Every visualization request ****must**** use a minimal
 2147 → [VIS_SPEC] block with the correct type specified. Do not
 2148 → use any other format.
 2149 - Do ****not**** include file names, imports, or drawing code.
 2150 → The orchestrator will handle image generation.
 2151 - If you cannot meaningfully visualize or correctly visualize
 2152 → a thought using the provided tools and inputs, then do
 2153 → not generate an image.
 2154 - Images are meant to be simple and visually cohesive - do
 2155 → not make grandiose images with titles and axis - it's
 2156 → simply for a baseline understanding of the question.
 2157 - The first line of the trace should be QUESTION: followed by
 2158 → a detailed in depth recap of the problem, specifying all
 2159 → the important aspects, without mentioning the solution.

2. Validity Rules:

```

2160
2161     - All [VIS_SPEC] parameters must be valid, fully-formed
2162     → Python literals.
2163     - For [VIS_SPEC] type "grid", the values must be a valid
2164     → Python list of lists with exactly rows rows and cols
2165     → columns (or a flat list of length rows * cols), and each
2166     → value should be a number or string.
2167     - For type graph, tree_from_dict, tree_from_root, and
2168     → similar, node and edge labels may be strings or integers,
2169     → but all structures must be valid Python literals.
2170     - Never output incomplete or empty lists/arrays/dicts in
2171     → [VIS_SPEC] blocks. All lists must be fully closed and
2172     → contain at least one value, unless an empty structure is
2173     → explicitly required by the problem.
2174     - Do not use variable names, symbolic labels, ellipses, or
2175     → placeholders (e.g., a1, x, \ldots, an) anywhere in the
2176     → [VIS_SPEC].
2177
2178     ---
2179
2180     **[VIS_SPEC] Reference Examples: Your blocks must follow the
2181     → same format as these.**
2182
2183     [VIS_SPEC]
2184     type: graph
2185     nodes: [A,B,C]
2186     edges: [(A,B),(B,C)]
2187     [/VIS_SPEC]
2188
2189     [VIS_SPEC]
2190     type: flow_network
2191     nodes: [A,B,C]
2192     edges: [(A,B),(B,C)]
2193     flows (optional): {(A,B): 2, (B,C): 1}
2194     capacities (optional): {(A,B): 3, (B,C): 2}
2195     [/VIS_SPEC]
2196
2197     \ldots
2198     \ldots
2199     \ldots
2200
2201     3. Reflection step immediately after each VIS_SPEC
2202     - Write a new THOUGHT that:
2203         a. Describes what you see in the previous generated
2204             → `reasoning_image_N.png`.
2205         b. Explains how it informs your next reasoning move.
2206
2207     4. FINAL ANSWER
2208     - After all reasoning, output ``FINAL ANSWER:`` and your
2209         → concise solution (pseudocode is sufficient)
2210
2211     5. Formatting and Output Requirements
2212     - Everything must be plain text with only the full
2213         → QUESTION (just the problem itself, not the name of the
2214             → problem), FINAL ANSWER, REASONING TRACE marker,
2215             → THOUGHT lines and VIS_SPEC markers.

```

2214 G IMPACT STATEMENT
2215

2216 All data sourced in this work were either publicly available under open licenses or generated syn-
2217 thetically. We ensured that all original content and assets used in the dataset creation process respect
2218 copyright and licensing terms. No human subjects were involved, and we do not foresee any di-
2219 rect harm to individuals or communities as a result of this work. The dataset is intended solely for
2220 academic research to improve multimodal reasoning capabilities in AI systems.

2221 H LICENSES
2222

2223 We list the licenses involved in this work as follows,

2224

- 2225 • Anole-7B model is under *Chameleon Research License*.
- 2226 • BAGEL-7B-MoT model is licensed under the *Apache 2.0 license*. It is finetuned from
2227 *Qwen2.5-7B-Instruct* and *siglip-so400m-14-384-flash-attn2* model, and uses the *FLUX.1-2228 schnell VAE model*, all under *Apache 2.0*.
- 2229 • ImageNet dataset in under *BSD 3* license.
- 2230 • Visual CoT dataset is licensed under *CC BY 4.0*
- 2231 • MATH dataset (Hendrycks et al., 2021) is under *MIT License*.
- 2232 • OpenStax Physics books are license under *CC BY 4.0*.
- 2233 • MIT OCW Physics lecture notes under *CC BY 4.0*.
- 2234 • Maze datasets is licensed under *CC BY 4.0*.

2235