Think360°

Beyond Depth: Evaluating the Width-centric Reasoning Capability of MLLMs

Supplementary Material

Appendix

A. More Experiment Settings

All experiments are conducted on A800 GPUs. Additionally, we list the maximum output length settings for different models in Table S1.

Table S1. Maximum response length settings for different models. Models with † are evaluated using `vLLM` (Qwen series, MiMo, Kimi, LLama, and GLM-V) or `LMDeploy` (InternVL series).

Model	Max Response Length
♣*Closed-source MLLMs*	
♣GPT-4o	16384
♣GPT-4v-preview	16384
♣GPT-4.1	32768
♣o1	100000
♣o3	100000
♣o4-mini	100000
♣Gemini-2.0-flash	8192
♣Gemini-2.0-flash-thinking-exp-01-21	8192
♣Gemini-2.5-flash-thinking	65536
♣Gemini-2.5-pro	65536
♣Doubao-1-5-vision-pro-250328	16384
♣Doubao-1.5-thinking-vision-pro-250428	16384
♣Doubao-1.5-thinking-vision-pro-250428-nothinking	16384
♣Claude-4-Opus-20250514	32000
♣Claude-4-Opus-20250514-Thinking	32000
♣Claude-3.7-Sonnet-Thinking	64000
♣Claude-4-Sonnet	64000
♣Grok-2-vision-1212	16384
♦*Open-source MLLMs*	
♦LLaVA-Onevision-7B	8192
♦Llama-3.2-Vision-Instruct-11B†	32768
♦GLM-4.1V-Thinking-9B†	32768
♦Kimi-VL-Instruct†	32768
♦Kimi-VL-Thinking†	65536
♦InternVL-2.5-8B†	8192
♦InternVL-3-8B†	16384
♦InternVL-3-14B†	16384
♦InternVL-3.5-8B†	16384
♦Bee-RL-7B†	16384
♦MiMo-VL-RL-7B†	32768
♦MiMo-VL-SFT-7B†	32768
♦Qwen2.5-VL-Instruct-7B†	8192
♦Qwen2.5-VL-Instruct-32B	8192
♦Qwen2.5-VL-Instruct-72B	8192

B. Demonstration of the Scaling of Reasoning Depth & Width

To illustrate how reasoning depth and width scale with problem complexity, we present two representative examples in Fig.S1.

Reasoning Depth. The left demonstrates depth scaling through a clock time-reading task with progressively complex dial markings. In the simplest version, clock positions are directly labeled with numerical values, requiring only visual recognition. As complexity increases, markings transition to arithmetic expressions (addition, subtraction, multiplication, division), then to advanced mathematical operations (integrals, determinants). This progression systematically extends the sequential reasoning chain: models must parse mathematical notation, execute calculations, map results to clock positions, and determine the time—each step dependent on its predecessor. Thus, reasoning depth scales with the computational complexity embedded in visual elements.

Reasoning Width. The right demonstrates width scaling through maze navigation with increasing grid sizes. As the maze expands from small to medium to large configurations, the exploration space grows exponentially. Larger mazes introduce more branching points, dead ends, and alternative routes requiring parallel consideration. Models must explore multiple candidate paths simultaneously, evaluate their viability, and backtrack when necessary. The reasoning width scales directly with spatial complexity and the number of viable paths requiring concurrent evaluation.

C. Error Analysis

We conducted a fine-grained error diagnosis on two advanced thinking models: Doubao-1.5-Thinking-Vision-Pro and Claude-Opus-4-Thinking. As shown in Fig.S2, our analysis reveals that **width-centric errors consistently dominate**, accounting for 56% (Doubao-1.5) and 55% (Claude-Opus-4) of all failures. These significantly outweigh both depth-centric errors (36%-38%) and basic perception or calculation mistakes (<8%).

A closer examination of these failure modes highlights the structural limitations driving these errors. The most prevalent issue across both models is *Incomplete Branch Expansion* (accounting for ~30% of total errors). This indicates that the models arbitrarily collapse multi-path problems into a single path, failing to enumerate alternative branches (often resulting in tunnel vision). Furthermore, *Ineffective Pruning* contributes to ~20% of errors, demonstrating a systematic failure to verify candidate solutions against global constraints.

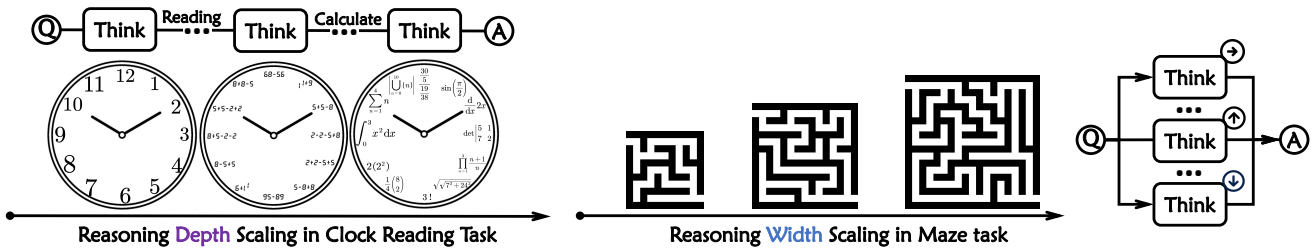

Figure S1. Demonstration of the Scaling of Reasoning Depth & Width

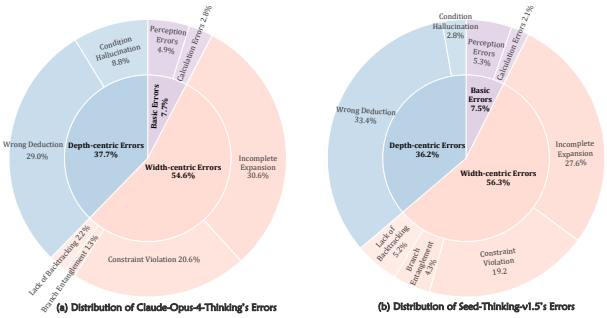

Figure S2. Distribution of error categories for Claude-Opus-4-Thinking (left) and Seed-Thinking-v1.5 (right). The pie charts illustrate that width-centric errors (e.g., incomplete branch expansion and ineffective pruning) dominate the failure modes across different architectures, substantiating that current models are primarily bottlenecked by multi-path reasoning rather than single-step depth.

D. Prompts for Response & Caption Generation

In this section, we present the details of the prompts used for model evaluation and caption generation. Specifically, we employ two different prompting strategies to obtain model responses: one directly requests answers to questions, while the other first requires step-by-step reasoning before providing the final answer. The two prompting strategies are shown below:

Direct Prompt: Please answer this question.
CoT Prompt: Please first think about this question step by step, and then output the final answer.

The caption generation prompt is designed to create highly detailed and accurate descriptions of visual content that serves as a bridge between visual and textual modalities, ensuring that all critical visual information required for mathematical and logical reasoning is preserved in the textual description.

Specially, we emphasizes four key principles: completeness ensures no essential visual elements are omitted; precision requires the use of exact mathematical terminology and notation; comprehensiveness mandates inclusion of every detail necessary for understanding the complete visual context; and clarity ensures logical organization of information to facilitate subsequent reasoning tasks.

Caption Prompt: You are an expert mathematical and logical reasoning analyst.
Create an image caption so detailed and accurate that another model could reconstruct all essential visual information needed for reasoning, using only your description.
Critical Guidelines:

- **Completeness**: Describe objectivel ONLY what is visually present. Do not infer, solve, interpret, or add information not explicitly shown
- **Precision**: Use exact mathematical terminology and standard notation
- **Comprehensiveness**: Include every detail necessary for another model to understand the complete visual context
- **Clarity**: Organize information logically to enable effective reasoning by subsequent models.

E. Question Distribution Analysis

As shown in Fig.S3, the word cloud reveals prominent terms such as "square", "grid", "cell", and "region" indicating strong emphasis on spatial reasoning and "number", "which", "all", "how many" and "times" suggesting the demand for systematic exploration and constraint satisfaction tasks that align with our benchmark's focus on reasoning depth and width.

The distribution of question lengths reveals an average of 73.72 words, with a median of 60. What's more, it exhibits a right-skewed pattern, with the majority of questions (48%) falling within 30-69 words. Notably, 163 questions exceed 100 words, with the longest reaching 298 words, indicating significant variation in problem context length and complexity across the dataset.

F. Detailed Workflow for Game-based and Proof-based Question

This section provides detailed workflows for processing two challenging question types: game-based and proof-based problems. These categories require specialized handling due to their unique characteristics—game-based problems lack explicit question-answer pairs, while proof-based

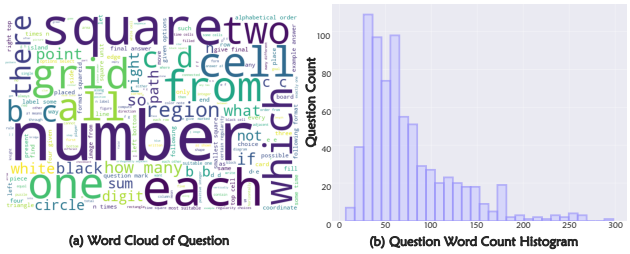

Figure S3. Word Cloud and Length Distribution Histogram of Questions in Think360.

problems contain non-verifiable reasoning processes that complicate objective evaluation.

F.1. Game-based Question Processing

Game-based problems originate from interactive online puzzle games that present visual challenges without predefined questions or explicit answers. As shown in Fig.S7, the processing workflow involves three key stages:

Stage 1: Image Preparation. We capture the initial game state screenshot that contains sufficient conditions for solving the puzzle. This image serves as the primary visual input, preserving all relevant spatial relationships and constraints. To facilitate unambiguous reference to specific regions or positions, we overlay alphabetical labels (A, B, C, etc.) onto the image, creating clearly identifiable reference points.

Stage 2: State Enumeration and Question Design. We systematically enumerate all possible states or values for each labeled position in the puzzle. For example, in a map coloring game, we identify the finite set of colors (e.g., green, red, brown, yellow) that can be assigned to each region. We then design questions that reference specific labeled positions (e.g., "What color is region A?"), transforming the interactive game into a well-defined question-answer format.

Stage 3: Format Standardization. To ensure consistent response formats, we incorporate explicit instructions and examples directly into the problem statement. This includes specifying the answer format (e.g., "Give final answer in following format: [RegionIndex][Color]"), providing notation explanations (e.g., "g(green), r(red), b(brown), y(yellow)"), and including concrete examples (e.g., "Ag means region A is green"). This standardization enables objective verification of model responses.

Throughout this process, we filter out game instructions and UI elements that are irrelevant to the core reasoning task, retaining only the essential puzzle constraints and the designed question.

F.2. Proof-based Question Processing

Proof-based problems from competition and textbook sources typically require demonstrating mathematical statements through logical arguments (see Fig.S8). Since complete proofs are difficult to verify objectively and may not align with our focus on visual reasoning, we adopt a redesign strategy:

Stage 1: Proof Structure Analysis. We carefully analyze the original proof to identify key intermediate results, numerical relationships, or specific conclusions that are objectively verifiable. For instance, in a proof showing that certain polygons cannot be tiled by dominoes, we extract intermediate results like "the relation between boundary lengths and square counts satisfies 4(b-w) = B-W" and the final conclusion "the polygon cannot be tiled."

Stage 2: Context Extraction. We extract the essential context from the original problem statement and proof setup that is necessary for understanding the redesigned questions. This includes definitions (e.g., "a polygon is orthogonal if all angles are 90° or 270°"), notation (e.g., "b and w denote black and white square counts"), and setup procedures (e.g., "give the polygon a chessboard coloring"). This context is incorporated into the new problem statement to ensure self-contained questions.

Stage 3: Question Redesign. Based on the extracted verifiable information, we redesign questions that test understanding of key proof steps or conclusions without requiring the complete proof. For example, instead of "Show that the polygon cannot be tiled," we ask "Determine the relation of b, w, B, W" or "Could such a polygon be tiled? Answer yes or no." These redesigned questions maintain mathematical rigor while enabling objective answer verification.

Stage 4: Quality Control. We filter out portions of the original proof process that cannot be reliably verified or that do not contribute to visual reasoning assessment. Only the redesigned questions with clear, verifiable answers are retained in the final benchmark.

This workflow transforms proof-oriented problems into evaluation-friendly formats while preserving the core mathematical reasoning required, ensuring that our benchmark remains focused on verifiable visual reasoning rather than unconstrained proof generation.

G. More Details for ToT-Eval

In the stage of tree construction, each node preserves the original wording and is classified by step type (e.g., calculation, deduction, conclusion). For solutions consisting only of a final answer, we create a single root node to maintain structural consistency.

ToT Extraction Prompt: You are an expert in decomposing mathematical reasoning into tree structures. Extract key reasoning steps from solutions into a hierarchical tree where depth represents sequential reasoning steps that depend on previous conclusions, and breadth captures parallel exploration of different possibilities at the same depth level.
Critical Guidelines:

- **Verbatim Extraction**: Each node content must be directly extracted from the original solution text, preserving the original wording without paraphrasing
- **Critical Steps Only**: Focus on major logical leaps, calculations, and key deductions rather than simple listings or obvious observations. Keep complete calculation steps as one node (e.g., "ans = 4 + 7 = 11" should not be split)
- **Tree Structure**: Use node ID format {depth}.{sequence}, where child depth = parent depth + 1. Depth 1 nodes must have parent = None (root nodes), and nodes at the same depth are siblings, not parent-child
- **Special Case**: For solutions containing only a simple final answer (e.g., "D", "42"), create a single root node with parent = None to maintain structural consistency

After constructing the reasoning tree and judging each node's correctness, we compute two complementary metrics to quantify reasoning depth and breadth. ToT-Depth measures the quality of the deepest reasoning chains by averaging the accuracy along all paths from root to maximum-depth leaves, thereby rewarding long, correct chains while penalizing early logical breakdowns. ToT-Width measures the quality of parallel reasoning by averaging the accuracy across sibling groups at each depth level, crediting models that successfully explore multiple valid branches. Together, these process-based metrics provide fine-grained assessment of long chain-of-thought responses beyond simple outcome-based accuracy.

ToT Judgement Prompt: You are an expert in evaluating mathematical and logical reasoning steps. Judge the correctness of a single reasoning step within a larger reasoning tree, considering the problem context, reference answer, parent nodes for context, and the image if relevant.
Evaluation Criteria:

- **Correctness**: Is the reasoning in this step logically sound and factually correct?
- **Validity**: Does it follow properly from the parent nodes?
- **Accuracy**: For calculations, are the results mathematically correct?
- **Relevance**: Does it contribute meaningfully to solving the problem?
- **Final Answer Check**: If this is the final answer node, does it match the reference answer?

Output: Respond with only "True" (if correct) or "False" (if incorrect/flawed). Note that intermediate steps can be correct even if the final answer is wrong, and conversely, a step can be flawed even if it leads to the correct final answer.

ToT-Depth measures the quality of the deepest reasoning chains. We first identify all leaf nodes at the maximum depth $d_{\max}$, then trace back from each leaf to the root to form complete reasoning paths. For each path P_i containing nodes $\{n_1, n_2, \ldots, n_{|P_i|}\}$, we calculate its accuracy as the proportion of correct nodes. The final depth score is the average accuracy across all deepest paths:

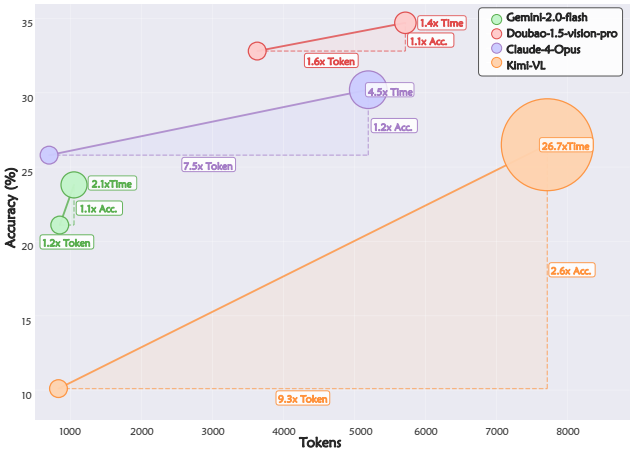

Figure S4. Thinking Mode Ablation. The x-axis shows accuracy improvement, and the y-axis shows token increase. Bubble size visualizes the time cost expansion. For each model, we fix the radius of the non-thinking circle to 1 and scale the radius of the thinking circle by the empirical multiplier of inference time (×Time).

$$\text{ToT-Depth} = \frac{1}{|P|}\sum_{i=1}^{|P|}\frac{\sum_{n\in P_i}\mathbb{1}[\text{correct}(n)]}{|P_i|} \tag{S1}$$

where P is the set of all paths from root to maximum-depth leaves, and $\mathbb{1}[\text{correct}(n)]$ indicates whether node n is judged correct.

ToT-Width measures the quality of parallel reasoning exploration. We group all nodes by their parent, identifying sibling groups that represent alternative reasoning branches at the same depth level. For each parent node p with children $C_p = \{c_1, c_2, \ldots, c_{|C_p|}\}$, we compute the accuracy of this sibling group. The width score is the average accuracy across all such groups:

$$\text{ToT-Width} = \frac{1}{|G|}\sum_{p\in G}\frac{\sum_{c\in C_p}\mathbb{1}[\text{correct}(c)]}{|C_p|} \tag{S2}$$

where G is the set of all parent nodes that have at least one child. These metrics together provide a comprehensive view of both the vertical depth and horizontal breadth of the model's reasoning capabilities.

H. More Experiments Results

In this section, we provide more experiments results.

Performance Patterns Across Subject: As shown in Fig. S6, models demonstrate above-average performance in Algebra and Number Theory subsets, but underperform in Combinatorics, Geometry, and Probability & Statistics.

This gap suggests that current MLLMs are more comfortable with problems that can be reduced to relatively direct symbolic manipulation or formula-based computation, while they struggle with tasks that require exploring large combinatorial spaces, handling spatial relations, or modeling uncertainty. In particular, the deficits in Combinatorics and Probability & Statistics are consistent with the difficulty of width-oriented reasoning, where models must juggle multiple cases, scenarios, or distributions rather than follow a single dominant derivation path.

Thinking vs. No Thinking: As shown in the Fig.S4, enabling thinking mode yields markedly different trade-offs across models. Kimi-VL sits at the "heavy-thinking" extreme: it gains 2.6× accuracy but at the cost of 9.3× tokens and 26.7× inference time, making it the most expensive option in test-time scaling. Claude-4-Opus shows a milder pattern, with 1.2× accuracy, 7.5× tokens, and 4.5× time, indicating more controlled but still substantial overhead.

By contrast, Doubao-1.5-vision-pro and Gemini-2.0-flash behave like "lightweight thinking" models: they achieve around 1.1× accuracy with only 1.4×–2.1× time and 1.2×–1.6× tokens. Overall, these results suggest that some models (e.g., Kimi-VL) aggressively trade latency for gains, while others (e.g., Doubao, Gemini) offer a more balanced accuracy–efficiency compromise.

What's more, Tab.S2 demonstrates more fine-grained influence of Chain-of-Thought prompting on the splits requiring different cognitive capabilities.

Image vs. Text: To further enrich our benchmarks and disentangle visual perception and reasoning, we design four different input settings based on the vision or language modalities. A question can be represented by four different ways: Image+Question (IQ), Image-Only (IO), Caption+Question (CQ) and Image+Question+Caption (ICQ). Specifically, Image-Only version problems are generated by adaptively overlaying the original question text onto the image according to the original image's aspect ratio. The caption with detailed descriptive information about the original image is generated by advanced o4-mini.

The ablation results in Fig.S5 point to a consistent pattern: Image-Only (IO) tends to be the hardest setting, underperforming Image+Question (IQ) in a clear majority of cases (60%) and underscoring persistent challenges in visual perception within current MLLMs. By contrast, Caption+Question (CQ) generally yields small but reliable improvements over IQ, with substantial improvements in certain models such as Llama-3.2-Vision (2.9%) and MiMo-VL-RL (1.7%). Interestingly, Image+Question+Caption (ICQ) typically achieves comparable performance to CQ rather than decisively surpassing it, implying that once a high-quality caption is available, additional raw visual inputs offer limited incremental value, likely due to information redundancy or multimodal fusion overhead. These trends highlight two key takeaways: (i) a persistent perception bottleneck under image-only inputs, and (ii) the value of textual captions as an effective bridge between vision and language, even when the marginal benefit of reintroducing images remains modest.

I. More Examples

In this section, we provide more data points in our benchmarks for intuitive understanding.

Table S2. Fine-grained Influence of Chain-of-Thought prompting on model performances.

Model	#Para.	CoT	ALL			Perceive-and-Comprehend			Trial-and-Error			Divide-and-Conquer			Branch-and-Bound			Hypothesize-and-Test		
			Acc./%	Time/s	Token	Acc./%	Time/s	Token	Acc./%	Time/s	Token	Acc./%	Time/s	Token	Acc./%	Time/s	Token	Acc./%	Time/s	Token
♣Close-source MLLMs																				
♣GPT-4o	-	✗	16.0	13.28	309.03	17.2	12.69	287.16	15.3	10.72	268.64	9.9	12.21	331.46	16.8	13.17	322.85	14.3	12.89	313.87
		✓	16.4	28.41	388.91	17.5	27.53	321.81	15.1	20.88	314.74	12.2	33.59	432.74	16.3	22.56	458.20	13.8	28.07	384.97
		Δ	+0.4	+15.1	+79.9	+0.3	+14.8	+34.6	−0.2	+10.2	+46.1	+2.3	+21.4	+101.3	−0.5	+9.4	+135.4	−0.5	+15.2	+71.1
♣o4-mini	-	✗	42.1	84.61	6736.37	42.8	81.37	6391.21	34.3	106.52	8067.57	38.7	89.62	7460.26	48.0	76.56	6401.12	37.9	91.37	7195.80
		✓	43.4	78.35	7079.01	44.6	75.00	6760.91	37.1	94.76	8651.76	37.2	84.31	7836.95	52.8	72.72	6071.11	37.8	82.15	7511.60
		Δ	+1.3	−6.3	+342.6	+1.8	−6.4	+369.7	+2.8	−11.8	+584.2	−1.5	−5.3	+376.7	+4.8	−3.8	−330.0	−0.1	−9.2	+315.8
♣Claude-4-Opus-20250514	-	✗	25.8	41.45	696.93	26.7	40.50	671.71	22.3	42.40	694.89	18.2	44.50	714.85	28.7	42.66	714.23	21.9	42.57	710.77
		✓	30.4	19.89	722.09	32.6	19.31	685.16	26.6	20.52	729.83	19.6	19.94	727.74	33.3	20.13	758.95	25.5	20.11	731.80
		Δ	+4.6	−21.6	+25.2	+5.9	−21.2	+13.4	+4.3	−21.9	+34.9	+1.4	−24.6	+12.9	+4.6	−22.5	+44.7	+3.6	−22.5	+21.0
♣Claude-4-Sonnet	-	✗	28.2	18.59	785.22	29.7	18.04	747.34	23.2	18.58	752.49	20.9	18.79	830.18	30.9	18.81	824.35	24.1	18.93	796.31
		✓	28.2	15.75	797.74	29.2	15.28	762.82	23.1	15.66	778.02	23.0	15.96	818.09	35.0	16.09	826.26	23.8	15.80	797.85
		Δ	+0.0	−2.8	+12.5	−0.5	−2.8	+15.5	−0.1	−2.9	+25.5	+2.1	−2.8	−12.1	+4.1	−2.7	+1.9	−0.3	−3.1	+1.5
♣Grok-2-vision-1212	-	✗	15.7	15.81	763.63	16.3	15.26	728.68	13.0	16.65	790.89	8.1	15.47	760.14	17.9	14.38	673.29	14.0	15.83	775.08
		✓	17.3	30.35	764.20	19.6	27.52	667.04	14.8	34.94	773.16	12.8	39.25	699.59	9.8	24.35	734.44	15.5	29.86	783.83
		Δ	+1.60	+14.54	+0.57	+3.30	+12.26	−61.64	+1.80	+18.29	−17.73	+4.70	+23.78	−60.55	−8.10	+9.97	+61.15	+1.50	+14.03	+8.75
♦Open-source MLLMs																				
♦Kimi-VL-Instruct†	16A3B	✗	10.1	39.79	829.45	11.1	38.14	750.18	7.7	49.11	1029.10	5.9	43.66	793.32	9.8	39.22	860.11	9.3	41.07	852.28
		✓	10.9	30.20	775.96	12.6	27.94	679.89	8.1	35.44	960.65	5.4	34.09	815.00	10.9	28.80	728.97	10.2	31.52	814.84
		Δ	+0.8	−9.6	−53.5	+1.5	−10.2	−70.3	+0.5	−13.7	−68.4	−0.5	−9.6	+21.7	+1.1	−10.4	−131.1	+0.9	−9.5	−37.4
♦Qwen2.5-VL-Instruct†	7B	✗	11.0	20.31	788.44	13.0	18.62	711.00	8.7	23.43	924.76	5.8	21.03	837.41	10.6	18.05	692.89	10.3	20.95	815.10
		✓	11.1	14.21	649.44	12.8	13.52	605.39	8.2	15.43	723.83	8.6	14.16	662.28	10.6	14.46	634.16	10.0	14.42	662.73
		Δ	+0.1	−6.1	−139.0	−0.2	−5.1	−105.6	−0.6	−8.0	−200.9	+2.8	−6.9	−175.1	+0.0	−3.6	−58.7	−0.3	−6.5	−152.4
♦Qwen2.5-VL-Instruct	72B	✗	16.2	24.20	615.09	18.2	23.63	587.27	14.4	25.87	652.28	9.0	24.46	618.13	13.6	23.65	608.68	13.9	24.39	623.97
		✓	15.9	24.61	627.19	18.3	23.85	592.46	14.0	26.32	667.55	8.8	25.78	641.36	16.3	26.36	629.20	13.6	23.54	622.19
		Δ	−0.3	+0.4	+12.1	+0.1	+0.2	+5.2	−0.4	+0.4	+15.3	−0.2	+1.3	+23.2	+2.7	+2.7	+20.5	−0.3	−0.9	−1.8
♦InternVL-3†	8B	✗	13.1	4.67	379.43	14.2	4.45	360.55	12.3	4.98	405.78	8.6	5.06	413.24	15.4	4.81	391.61	11.4	4.74	385.42
		✓	12.3	5.31	432.41	13.1	5.17	420.59	10.9	5.47	447.35	4.7	5.49	449.65	12.5	5.22	426.41	10.9	5.41	441.31
		Δ	−0.8	+0.6	+53.0	−1.1	+0.7	+60.0	−1.3	+0.5	+41.6	−3.8	+0.4	+36.4	−3.0	+0.4	+34.8	−0.5	+0.7	+55.9
♦MiMo-VL-RL†	7B	✗	28.3	334.21	7380.78	29.7	314.61	6870.68	24.9	348.01	7761.05	19.1	394.47	8446.60	27.9	281.06	7226.50	24.4	352.59	7692.95
		✓	29.9	345.61	7941.18	31.3	325.85	7393.75	25.6	376.58	8255.09	21.4	395.02	8772.46	31.4	306.17	8115.06	25.9	369.66	8338.44
		Δ	+1.6	+11.4	+560.4	+1.6	+11.2	+523.1	+0.7	+28.6	+494.0	+2.3	+0.5	+325.9	+3.5	+25.1	+888.6	+1.4	+17.1	+645.5

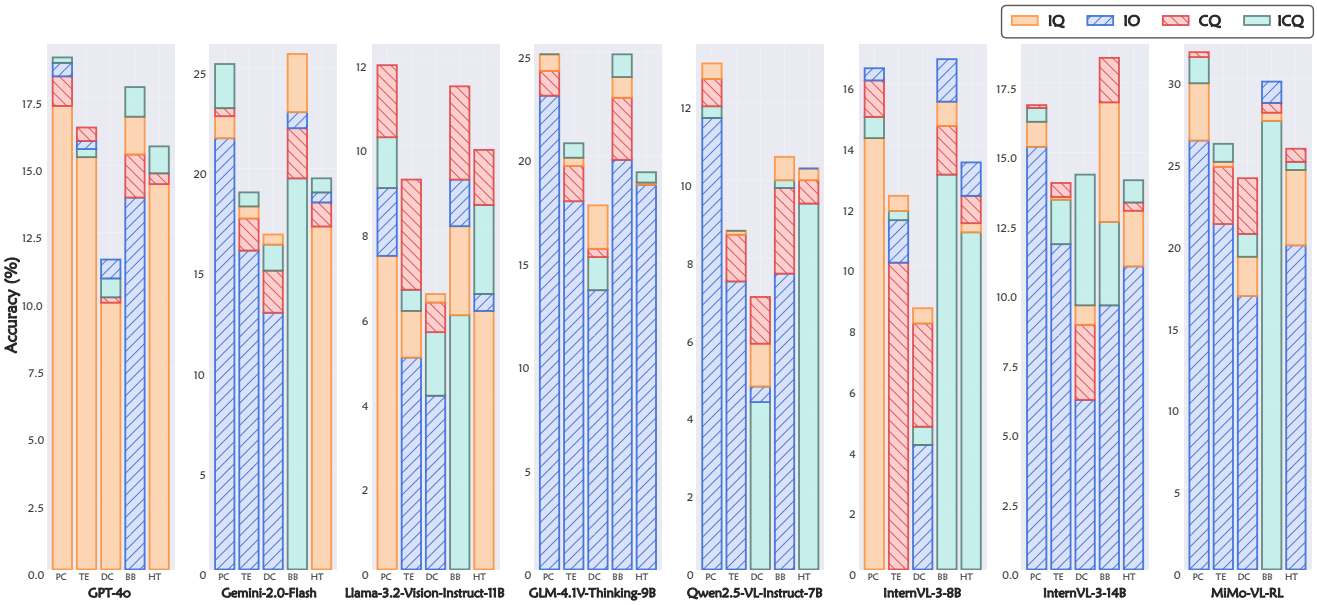

Figure S5. Ablation Study of Input Settings.

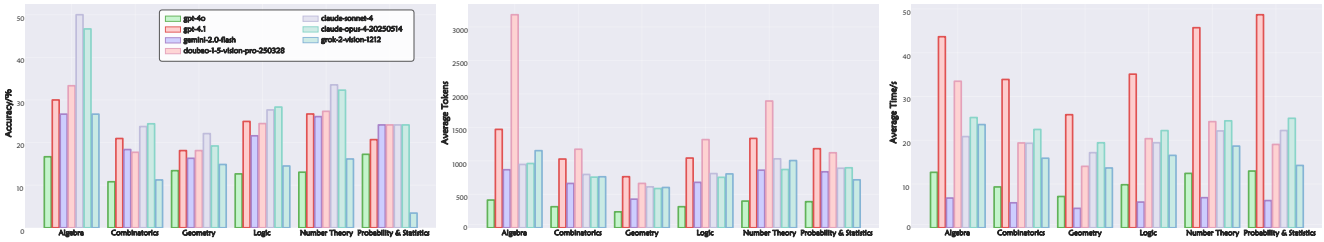

(a) System-1 MLLMs with fast and single-pass reasoning (average time< $50s$, token< 3000)

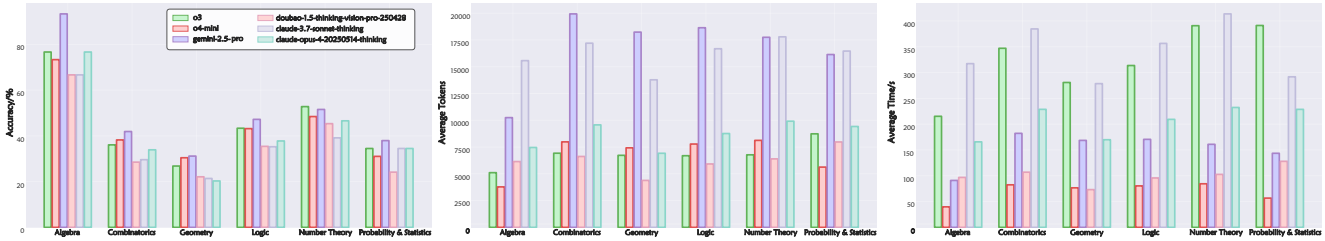

(b) System-2 MLLMs with slow and iterative long CoT reasoning (average time> $50s$, token> 3000).

Figure S6. Reasoning performances comparison across different question types. To account for the substantial disparities in reasoning time and token usage across models, we categorize them using predefined time/token thresholds to better highlight the performance profiles of system-1 and system-2 models. Please zoom in for a better view.

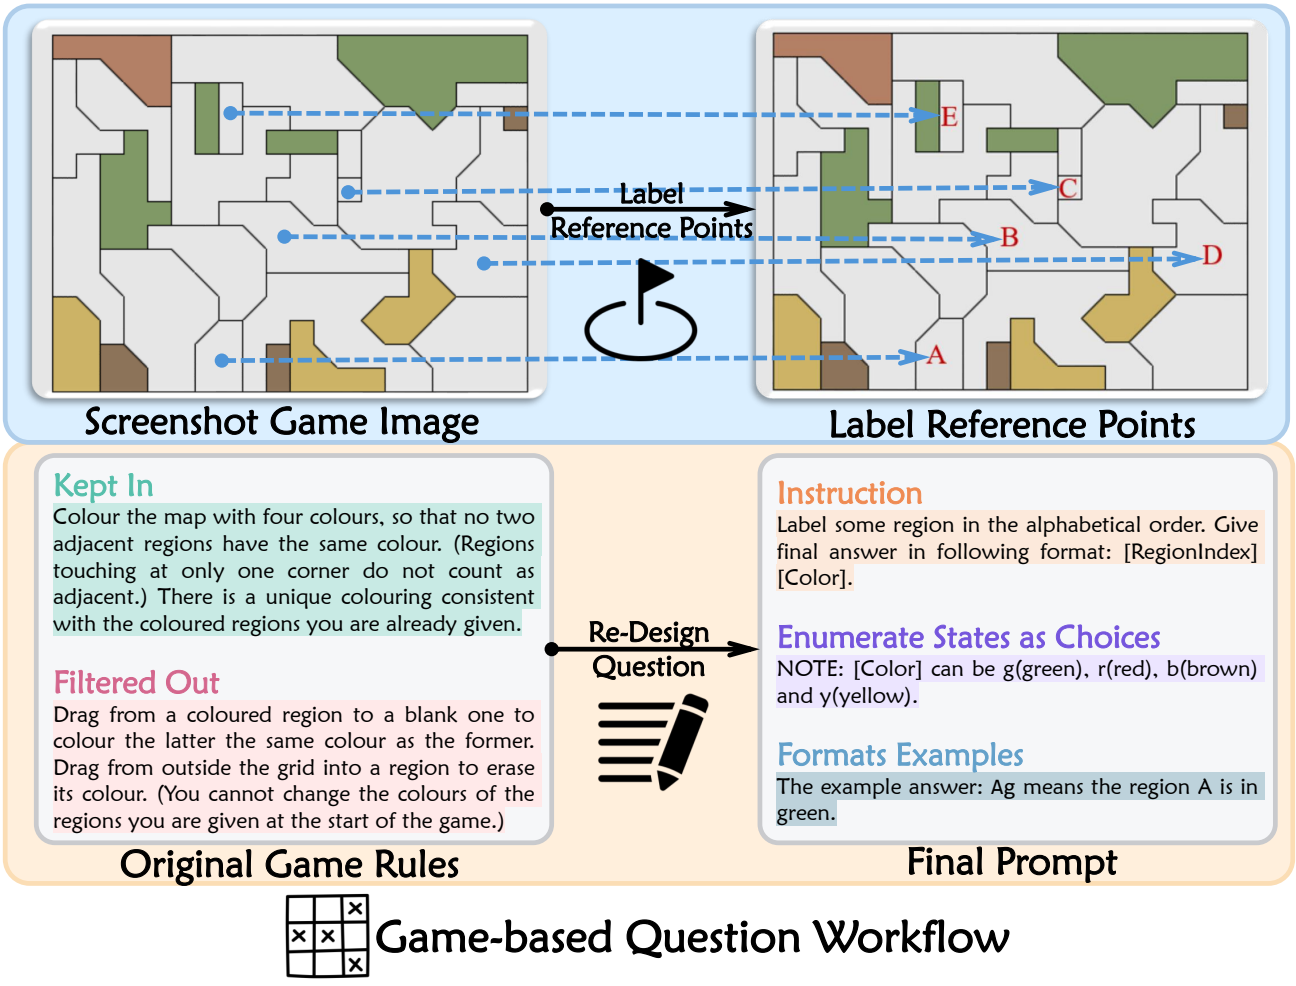

Figure S7. Workflow for Game-based Questions

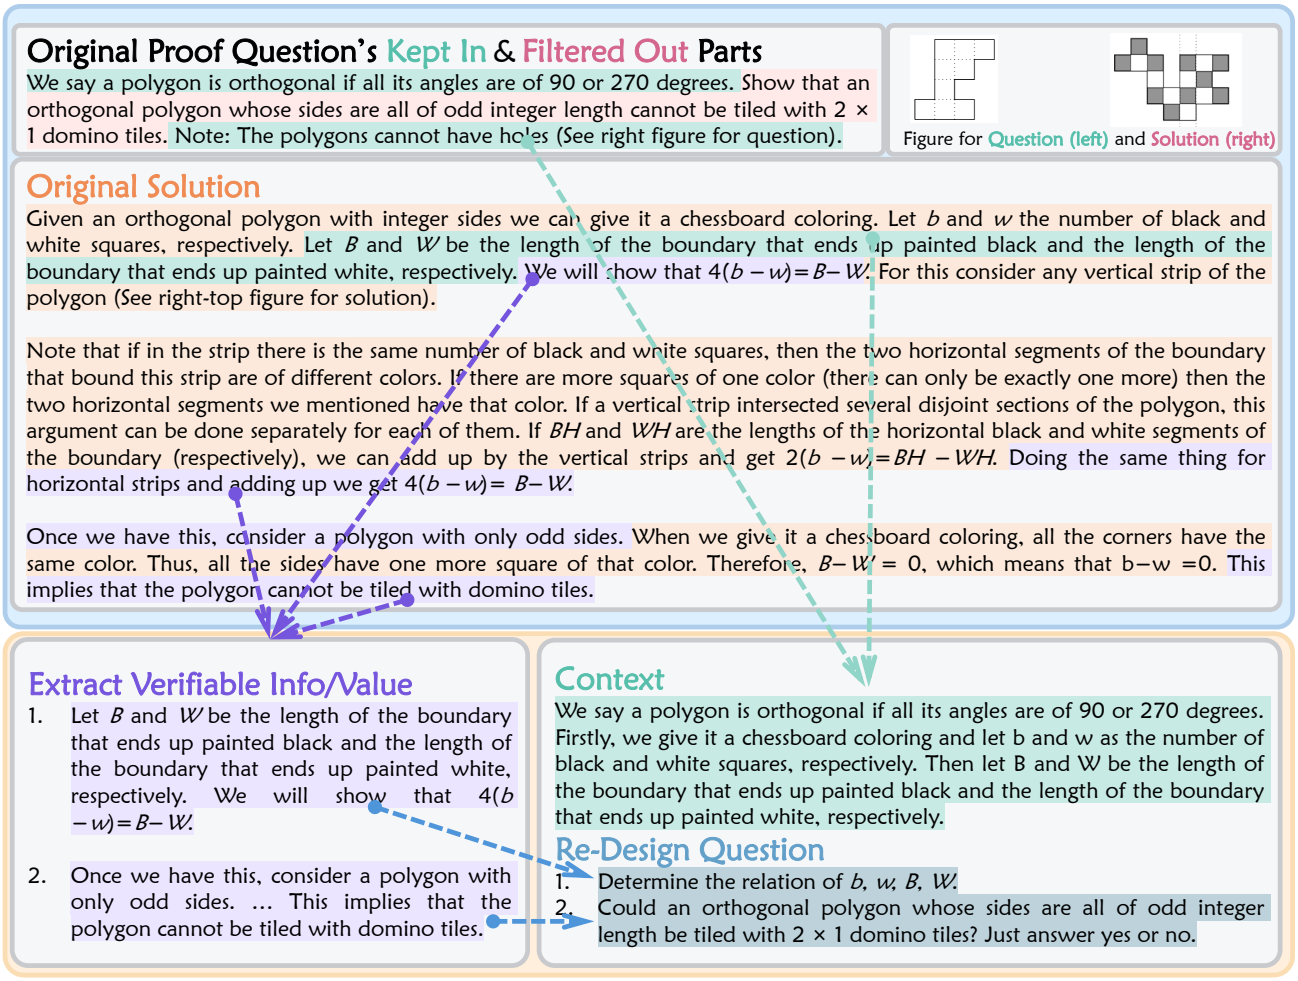

Figure S8. Workflow for Proof-based Questions

Branch and Bound

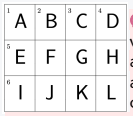

Question: Complete the following "cross-number puzzle", where each "Across" answer represents a four digit number, and each "Down" answer represents a three-digit number. No answer begins with the digit 0. Across: 1. A B C D is the cube of the sum of the digits in the answer to 1 Down. 5. From left to right, the digits in E F G H are strictly decreasing. 6. From left to right, the digits in I J K L are strictly decreasing. Down: 1. A E I is a perfect fourth power. 2. B F J is a perfect square. 3. The digits in C G K form a geometric progression. 4. D H L has a two-digit prime factor. **Answer**: 2197, 5431, 6410

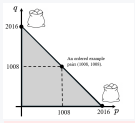

Question: Kristoff is planning to transport a number of indivisible ice blocks with positive integer weights from the north mountain to Arendelle. He knows that when he reaches Arendelle, Princess Anna and Queen Elsa will name an ordered pair (p, q) of nonnegative integers satisfying p + q ≤ 2016. Kristoff must then give Princess Anna exactly p kilograms of ice. Afterward, he must give Queen Elsa exactly q kilograms of ice. What is the minimum number of blocks of ice Kristoff must carry to guarantee that he can always meet Anna and Elsa's demands, regardless of which p and q are chosen? **Answer**: 18

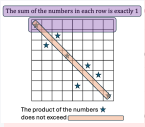

Question: In each square of an 8 × 8 chessboard, a positive real number is written. The numbers satisfy the following two conditions: [1] The sum of the numbers in each row is exactly 1. [2] For any set of 8 squares, where no two are in the same row or column, the product of the numbers in these squares does not exceed the product of the numbers on the main diagonal. What is the minimum possible value for the sum of the numbers on the main diagonal? **Answer**: 1

Figure S9. Examples of the Branch-and-Bound main category.

Hypothesize and Test

Question: Ten balls, each coloured green, red or blue, are placed in a bag. Ten more balls, each coloured green, red or blue, are placed in a second bag. In one of the bags there are at least seven blue balls and in the other bag there are at least four red balls. Overall there are half as many green balls as there are blue ball. Let r, g and b respectively be the numbers of red, green and blue balls that there are in total. Determine the relation among these three numbers.
Answer: r + g = b

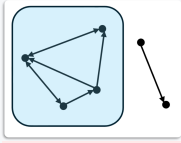

Question: The MO space station consists of 99 space stations, where any two stations are connected by a tubular channel. Set 99 of the channels to be two-way channels, and the rest are strictly one-way. For a group of four stations, if starting from any station one can reach any other station through the channels, the group of four stations is called a connected four-station group. Find the maximum number of connected four-station groups, and justify your answer.
Answer: 2052072

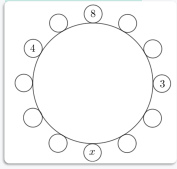

Question: Twelve people are seated, equally spaced, around a circular table. They each hold a card with different integer on it. For any two people sitting beside each other, the positive difference between the integers on their cards is no more than 2. The people holding the integers 3, 4, and 8 are seated as shown. The person opposite the person holding 8 is holding the integer x. What are the possible values of x?
Answer: -3

Figure S11. Examples of the Hypothesize-and-Test main category.

Divide and Conquer

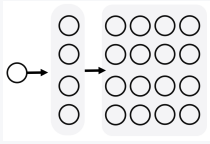

Question: Physicists at Princeton are trying to analyze atom entanglement using the following experiment. Originally there is one atom in the space and it starts splitting according to the following procedure. If after n minutes there are atoms $a_1, \cdots a_n$, in the following minute every atom $a_i$splits into four new atoms, $a_i^1, a_i^2, a_i^3, a_i^4$. Atoms a_i^j and a_k^j are entangled if and only if the atoms a_i and a_k were entangled after n minutes. Moreover, atoms $a_i^j$and a_k^{j+1} are entangled for all $1 \le i, k \le N$ and $j = 1, 2, 3$. Therefore, after one minute there is 4 atoms, after two minutes there are 16 atoms and so on. Physicists are now interested in the number of unordered quadruplets of atoms $\{b_1, b_2, b_3, b_4\}$among which there is an odd number of entanglements. What is the number of such quadruplets after 3 minutes?
Answer: 354476

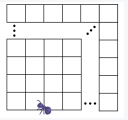

Question: In a 2024 × 2024 grid of squares, each square is colored either black or white. An ant starts at some black square in the grid and starts walking parallel to the sides of the grid. During this walk, it can choose (not required) to turn 90° clockwise or counterclockwise if it is currently on a black square, otherwise it must continue walking in the same direction. A coloring of the grid is called simple if it is not possible for the ant to arrive back at its starting location after some time. How many simple colorings of the grid are maximal, in the sense that adding any black square results in a coloring that is not simple?
Answer: 2024^{4046}

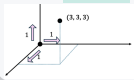

Question: Leonard is standing at the origin in 3D space. He can only move forward one unit in the x-direction, the y-direction, or the z-direction. How many ways can he get to (3, 3, 3)? **Answer**: 1680

Figure S10. Examples of the Divide-and-Conquer main category.

Perceive and Comprehend

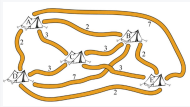

Question: Jared wants to minimize his walking time while passing five different colored tents arranged on campsites. If the total walking time must be minimized, what color tent(s) should be placed at campsite C?
Answer: 24

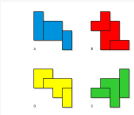

Question: If you re-assemble the pieces of each of the four compositions shown here, three of them will be the same shape and one won't. Which is the odd one out?
Answer: B

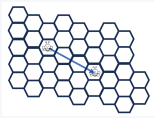

Question: A busy bee buzzes between the cells of a large honeycomb made up of a plane of tessellated hexagons. A flight of length n consists of picking any of the six neighbouring cells and flying to the (n^{th}) cell in that direction. After consecutive flights of lengths $(n = N, N-1, \cdots, 2, 1)$, the bee finds that it has returned to its starting location. For which values of N is this possible?
Answer: N ≥ 3

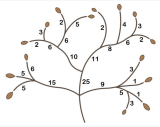

Question: At the end of each growing season, Joy likes to prune dead leaves from her favourite tree. She does this by cutting branches. For this tree, shown below, there are 15 leaves she wants to remove. She decides to give an approximate time it will take to cut each branch. These times are shown for each branch. When a branch is cut, all branches and leaves attached to it are removed from the tree. For example, if you cut the branch labelled with 15, the three leftmost leaves will be removed. What is the shortest amount of time in which Joy can remove all 15 leaves?
Answer: 43

Figure S12. Examples of the Perceive-and-Comprehend main category.

Trial-and-Error

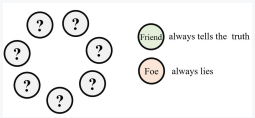

Question: On the island of Friends and Foes, every citizen is either a Friend (who always tells the truth) or a Foe (who always lies). Seven citizens are sitting in a circle. Each declares "I am sitting between two Foes". How many Friends are there in the circle?.

Answer: 3

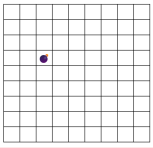

Question: Adam is playing Minesweeper on a 9 \\times 9 grid of squares, where exactly \\(\\frac{1}{3} \\) (or 27) of the squares are mines (generated uniformly at random over all such boards). Every time he clicks on a square, it is either a mine, in which case he loses, or it shows a number saying how many of the (up to eight) adjacent squares are mines. First, he clicks the square directly above the center square, which shows the number 4. Next, he clicks the square directly below the center square, which shows the number 1. What is the probability that the center square is a mine?.

Answer: \\frac{88}{379}

Question: The diagram shows a cross-shaped box containing three numbered blocks. The puzzle is to slide the blocks around the box until the numbers read 1, 2, 3 as you go down. How many moves does it take?

Answer: 8

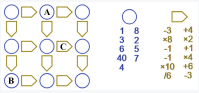

Question: This puzzle is made of numbers (like 1 and 8) and functions (like +4 and x8). Fill it in using the options provided. Label some places in the alphabetical order. Determine the answer in the labelled circles or pentagons.

Answer: 8

Figure S13. Examples of the Trial-and-Error main category.

Algebra

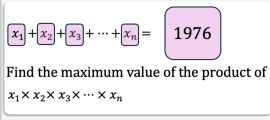

Question: Consider some positive integers whose sum is 1976. Find the maximum value of the product of these positive integers.

Answer 3^658*2

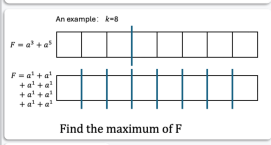

Question: Given a positive integer k and a positive real number a. For any partition $k_1 + k_2 + ... + k_r = k$ (k_i is a positive integer, $1 \le r \le k$), find the maximum of $F = a^{k_1} + a^{k_2} + ... + a^{k_r}$

Answer max{a^k, kα}

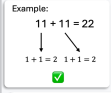

Question: How many positive integers $n\\leq10^{2005}$ can be written as the sum of two positive integers with the same sum of digits?

Answer 10^{2005} - 9023

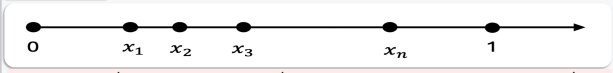

Question: Let $x_1, x_2, \\dots, x_n$ be in an interval of length 1. Define $x = \\frac{1}{n} \\sum_{j=1}^n x_j$, $y = \\frac{1}{n} \\sum_{j=1}^n x_j^2$. Find the maximum value of $f = y - x^2$.

Answer: The maximum value of f is $1/4$ when n is even and $\\frac{n^2 - 1}{4n^2}$ when n is odd.

Figure S14. Examples of the algebra subcategory.

Combinatorics

Question: Every city in a state is connected to exactly three other cities by direct air flights. One can fly from each city to any other city with at least one stop. Determine the maximal number of cities in the state.

Answer: 10

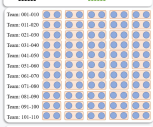

Question: In a bridge tournament, 110 teams play 6 rounds. In each round, the teams are split into 55 pairs, with each pair playing one match. No two teams play more than once. (1) What is the number of teams you can find such that no two of them have ever played each other? (2) If team number can be denoted as $6k + 2$ (or more), what the answer?

Answer: (1) 19; (2) $k+1$

Question: Mr. Fat is baking m different cakes with different kinds of cake mix. Some of the kinds of cake mix are sweetened. Each cake is made from five different kinds of mix with at least one kind of sweetened mix among them. It is known that for every three kinds of mix there is exactly one cake containing them. If there exists at least one very sweet cake: a cake made from at least four kinds of sweetened mix, compute the minimum value of m.

Answer: 68

Question: On the Cartesian grid, Johnny wants to travel from (0, 0) to (5, 1), and he wants to pass through all twelve points in the set $S = \{(i, j) \mid 0 \le i \le 1, 0 \le j \le 5, i, j \in Z\}$. Each step, Johnny may go from one point in S to another point in S by a line segment connecting the two points. How many ways are there for Johnny to start at (0, 0) and end at (5, 1) so that he never crosses his own path?

Answer 252

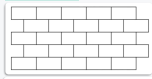

Question:
In the figure below, how many ways are there to select 5 bricks, one in each row, such that any two bricks in adjacent rows are adjacent?

Answer: 61

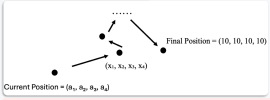

Question: Fred the Four-Dimensional Fluffy Sheep is walking in 4-dimensional space. He starts at the origin. Each minute, he walks from his current position (a1, a2, a3, a4) to some position (x1, x2, x3, x4) with integer coordinates satisfying $\{(x_1-a_1)^2+(x_2-a_2)^2+(x_3-a_3)^2+(x_4-a_4)^2\}=2^{n-1}$. In how many can Fred reach (10, 10, 10, 10) after exactly 40 minutes, if he is allowed to pass through this point during his walk?

Answer: $\binom{40}{10}\binom{40}{20}^3$

Figure S15. Examples of the combinatorics subcategory.

Number Theory

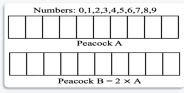

Question: A peacock is a ten-digit positive integer that uses each digit exactly once. Compute the number of peacocks that are exactly twice another peacock.

Answer: 184320

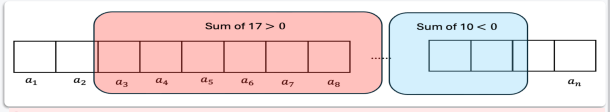

Question: If a positive integer n satisfies the following condition: there exists a sequence of n real numbers, where the sum of any 17 consecutive terms is positive, and the sum of any 10 consecutive terms is negative, find the maximum value of n.

Answer: 25

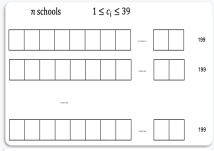

Question: There are n middle schools in a city. The ith middle school sends ci students ($1 \le ci \le 39$) to watch a football game in a stadium, where \\(\\sum_{i=1}^{n} c_i = 1990 \\). There are 199 seats in each row of the stand. It is required that the students in the same school sit in the same row. At least how many rows should there be, so that this is always possible?

Answer: 12

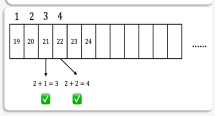

Question: Find the greatest N for which there are N consecutive positive integers such that the sum of digits of the k-th number is divisible by k, for k = 1, 2, ..., N.

Answer: 21

Question: Given a positive integer a, let X = {a1, a2, ..., an} be a set of positive integers, where a1 <= a2 <= a3 <= ... <= an. If for any integer p(1 <= p <= a), there is a subset of X such that S(A) = p, where S(A) is the sum of elements in set A, find the minimum value of n.

Answer: $\log_2 a + 1$

Figure S16. Examples of the number theory subcategory.

Geometry

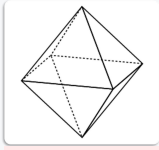

Question: Teresa the bunny has a fair 8-sided die. Seven of its sides have fixed labels 1, 2, ... , 7, and the label on the eighth side can be changed and begins as 1. She rolls it several times, until each of 1, 2, ... , 7 appears at least once. After each roll, if k is the smallest positive integer that she has not rolled so far, she relabels the eighth side with k. The probability that 7 is the last number she rolls is \\frac{a}{b} , where a and b are relatively prime positive integers. Compute 100a+b

Answer: 104

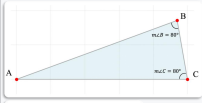

Question: Let ABC be a triangle with m∠B = m∠C = 80°. Compute the number of points P in the plane such that triangles PAB, PBC, and PCA are all isosceles and non-degenerate.

Answer: 6

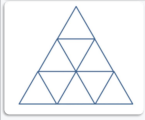

Question: Compute the number of distinct ways to color the nine triangles in the figure below either red, white, or blue such that no two triangles that share a side are the same color.

Answer: 528

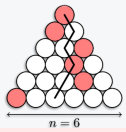

Question: Let n be a positive integer. A Japanese triangle consists of $1 + 2 + \dots + n$ circles arranged in an equilateral triangular shape such that for each $i = 1$, 2, $\dots$, n, the i^{th} row contains exactly i circles, exactly one of which is coloured red. A ninja path in a Japanese triangle is a sequence of n circles obtained by starting in the top row, then repeatedly going from a circle to one of the two circles immediately below it and finishing in the bottom row. Here is an example of a Japanese triangle with $n = 6$, along with a ninja path in that triangle containing two red circles. In terms of n, find the greatest k such that in each Japanese triangle there is a ninja path containing at least k red circles.

Answer: k = \lfloor \log_2 n \rfloor + 1

Figure S17. Examples of the geometry subcategory.

Probability

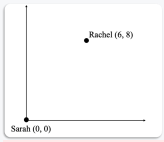

Question: Sarah stands at (0, 0) and Rachel stands at (6, 8) in the Euclidean plane. Sarah can only move 1 unit in the positive x or y direction, and Rachel can only move 1 unit in the negative x or y direction. Each second, Sarah and Rachel see each other, independently pick a direction to move at the same time, and move to their new position. Sarah catches Rachel if Sarah and Rachel are ever at the same point. Rachel wins if she is able to get to (0, 0) without being caught; otherwise, Sarah wins. Given that both of them play optimally to maximize their probability of winning, what is the probability that Rachel wins?

Answer: \frac{63}{64}

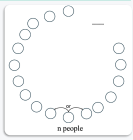

Question: Let n be an odd positive integer, and suppose that n people sit on a committee that is in the process of electing a president. The members sit in a circle, and every member votes for the person either to his/her immediate left, or to his/her immediate right. If one member wins more votes than all the other members do, he/she will be declared to be the president; otherwise, one of the the members who won at least as many votes as all the other members did will be randomly selected to be the president. If Hermia and Lysander are two members of the committee, with Hermia sitting to Lysander's left and Lysander planning to vote for Hermia, determine the probability that Hermia is elected president, assuming that the other n - 1 members vote randomly.

Answer: \frac{2^n-1}{n2^{n-1}}$

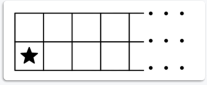

Question: There is a grid of height 2 stretching infinitely in one direction. Between any two edge-adjacent cells of the grid, there is a door that is locked with probability \\frac{1}{2} independent of all other doors. Philip starts in a corner of the grid (in the starred cell). Compute the expected number of cells that Philip can reach, assuming he can only travel between cells if the door between them is unlocked.

Answer: \frac{32}{7}

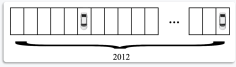

Question: A parking lot consists of 2012 parking spots equally spaced in a line, numbered 1 through 2012. One by one, 2012 cars park in these spots under the following procedure: the first car picks from the 2012 spots uniformly randomly, and each following car picks uniformly randomly among all possible choices which maximize the minimal distance from an already parked car. What is the probability that the last car to park must choose spot 1?

Answer: \frac{1}{2062300}

Figure S18. Examples of the probability subcategory.

Logic

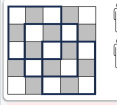

Question: On a 5×5 board, two players alternately mark numbers on empty cells. The first player always marks 1's, the second 0's. One number is marked per turn, until the board is filled. For each of the nine 3×3 squares the sum of the nine numbers on its cells is computed. Denote by A the maximum of these sums. How large can the first player make A, regardless of the responses of the second player?

Answer: 6

Question: Let $n \geq 2$ be an integer. Consider an $n \times n$ chessboard consisting of n^2 unit squares. A configuration of n rooks on this board is peaceful if every row and every column contains exactly one rook. Find the greatest positive integer k such that for each peaceful configuration of n rooks there is a $k \times k$ square which does not contain a rook on any of its k^2 unit squares.

Answer: $k = \left\lfloor \sqrt{n}-1 \right\rfloor$

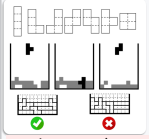

Question: In this game, there is an area ten squares wide and a number of squares tall. Pieces chosen randomly from among the seven "tetrominoes" made up of four squares glued together, as shown below, fall from the top of the screen. As the pieces fall, the player may rotate them or slide them left or right, but once they touch a piece below them they stick in place. If the player is able to fit the pieces together so as to leave no gaps in a row, that row disappears and all the blocks above fall to leave more room for new blocks. Otherwise the screen fills up with blocks and the game ends. If the puzzle is ten squares wide, in the pattern with only two squares left at the bottom line, There are 11 types solusion can be achieved by eliminating three line. If you try to achive this situation but without using the "T"-shaped block, how many are solusions can be achieved?

Answer: 6

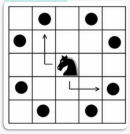

Question: In chess, a knight can move either two squares horizontally and one square vertically, or two squares vertically and one square horizontally. The graphic below shows the eight possible locations to which the knight in the center of the 5 5 board can move. Unlike all other standard chess pieces, the knight can 'jump over' all other pieces (of either color) to its destination square. This question is an estimation problem. If the answer given is within 10% of the correct answer, your team will receive credit. A knight is on a square on an infinite chess board. Compute the number of distinct squares where the knight can end up after exactly 10 moves.

Answer: 741, thus any answer between 666.9 and 815.1 is considered correct

Figure S19. Examples of the logic subcategory.