

Entropy Centroids as Intrinsic Rewards for Test-Time Scaling

Wenshuo Zhao^{1,2} Qi Zhu³ Xingshan Zeng³ Fei Mi³ Lifeng Shang³ Yiren Feng¹

¹HKUST ²ZJU ³Huawei Noah’s Ark Lab
wzhaoba@connect.ust.hk

Abstract

An effective way to scale up test-time compute of large language models is to sample multiple responses and then select the best one, as in Grok Heavy and Gemini Deep Think. Existing selection methods often rely on external reward models, which requires training a strong reward model and introduces additional computation overhead. As an alternative, previous approaches have explored intrinsic signals, such as confidence and entropy, but these signals are noisy with naive aggregation. In this work, we observe that high-entropy tokens tend to cluster into consecutive groups during inference, providing a more stable notion of model uncertainty than individual tokens. Together, these clusters reveal temporal patterns of model uncertainty throughout the inference process. Motivated by this observation, we propose to use the temporal structure of uncertainty as an intrinsic reward. To this end, we first formalize the basic unit of segment-level uncertainty as the High Entropy Phase (HEP), a variable-length segment that begins at a high-entropy token and ends when consecutive low-entropy tokens appear. We then define the Entropy Centroid, inspired by the concept of the center of mass in physics, as the weighted average position of all HEPs along the trajectory. Intuitively, a lower centroid indicates early exploration followed by confident generation, which we find often corresponds to higher response quality. Based on this insight, we propose the Lowest Centroid method, which selects the response with the lowest entropy centroid among multiple candidates. Experiments on mathematics, code generation, logical reasoning, and agentic tasks, across model scales ranging from 14B to 480B, show that Lowest Centroid consistently outperforms existing baselines and delivers stable gains as model size increases. Code is available at <https://github.com/hkust-nlp/entropy-centroid>.

1 Introduction

Reward signals play a central role in test-time scaling of large language models (LLMs), where in parallel scaling, N candidate trajectories are sampled and the best one is selected (Gemini Deep Think (Feng et al., 2026), Grok 4 Heavy (xAI, 2026)). Existing approaches typically rely on an external reward model to score each trajectory or on expensive human annotation (Lightman et al., 2023; Wang et al., 2024), which introduce substantial computational overhead (Fei et al., 2025; Chen et al., 2025b; Cao et al., 2025) and noisy step-level evaluation (Ding et al., 2025; Zhang et al., 2025; Xie et al., 2026). These costs motivate the search for intrinsic reward signals derived from the model’s own generation process (Yuan et al., 2024; Zelikman et al., 2024). Majority voting (Wang et al., 2022) is a simple example, but it is restricted to short-answer tasks where the outputs can be easily aggregated, and does not generalize to complex settings such as coding and agentic tasks.

Recent work moves beyond voting and examines token-level metrics — entropy, confidence, and KL divergence — as proxies for inference quality (Kang et al., 2025; Zhao et al., 2025b). However, the token-level signals are inherently noisy. As shown in Figure 1 (c), calculation tokens tend to be confident even when the model is actually uncertain about the step overall. Several methods (Fu et al., 2025; Sharma & Chopra, 2025) address this by aggregating




Figure 1: Overview of the Entropy Centroid framework. **Top panel:** a) High Entropy Phase (HEP). When a token with entropy above the top entropy threshold is detected, the model enters a High Entropy Phase. The phase ends once k consecutive low-entropy tokens are observed. b) Entropy Centroid. For each HEP, we treat the number of tokens as its mass and the midpoint as its position, then compute the weighted average position along the trajectory. The formula follows the center-of-mass in physics. **Bottom panel:** c) We visualize the token-level entropy evolution. Scattered high-confidence tokens frequently appear within critical decisions, making token-level entropy a noisy reward signal. HEP (each red box highlights one phase) filters out such single-token noise at a higher group level.

token-level signals over fixed-length local windows, such as averaging the confidence scores over 1,024 tokens as an indicator of trajectory quality. However, these methods apply a monotonic rule — low entropy is uniformly treated as being confident and high entropy as failure, regardless of position in the trajectory. This view is challenged by recent findings that high-entropy tokens can reflect productive exploration rather than errors (Wang et al., 2025a; Hu et al., 2026).

Hypothesizing that the temporal pattern of uncertainty over a trajectory may contain useful information, in this work, we explore temporal uncertainty as an intrinsic reward. A natural choice is to model token-level entropy changes over a trajectory. However, as shown in Figure 1(c), token-level entropy tends to be noisy, which may misrepresent the model’s certainty within the current reasoning unit. As such, we seek a coarser-grained unit to represent the model’s uncertainty state. A closer look at token-level entropy reveals a simple but informative structure: As the heatmaps in Figure 1 show, high-entropy tokens do not appear in isolation; they cluster into consecutive bursts. When a model encounters difficulties or decision points, it enters a period of high entropy that resolves either into confident low-entropy continuation or into persistent uncertainty. Therefore, we formalize this pattern with **High Entropy Phases (HEPs)** as a basic unit to represent the model’s uncertainty state that could happen in different positions in the trajectory. Specifically, HEP is a variable-length unit, which begins when the token entropy exceeds a high threshold and ends when a specified number of consecutive tokens fall below a low threshold.

With HEPs as basic units, we further encode their positional information along the trajectory. Borrowing the concept of center of mass from physics, we define the **Entropy Centroid** by analogy. For each HEP, we take the number of tokens as its mass and the midpoint as its position, then compute the weighted average position along the trajectory (formulated in Figure 1 (b)). We validate this design empirically. Figure 2 visualizes HEP-level state transitions across trajectories. A clear separation emerges: correct trajectories tend to exhibit more exploration in the early stages, while incorrect trajectories show more confusion in the later stages. This is consistent with our goal of capturing temporal patterns in inference. We




Figure 2: **Left:** We sample 64 responses for the same question and visualize the model’s inference over time using HEP. The duration on the Y-axis refers to the number of tokens in each HEP. The green curves represent correct trajectories and the red curves represent incorrect ones. We compute the difference between the mean curves of the two groups and apply Gaussian smoothing to reveal the overall trend. The results show a clear separation: correct trajectories concentrate more HEPs in the early stage, while incorrect trajectories accumulate HEPs in the later stage, gradually surpassing the correct ones as inference progresses. **Right:** Based on HEP modeling, we compute the entropy centroids. Correct trajectories (blue scatter) have a median centroid around 0.47, while incorrect trajectories (red scatter) have a median centroid around 0.55.

further plot the centroid distributions for different trajectories on the same problem. Figure 2 confirms that centroid values of correct and incorrect trajectories are well separated, with correct trajectories concentrating exploration earlier. This supports the use of the lowest centroid as a selection criterion.

We use the Entropy Centroid as the intrinsic reward in test-time scaling. Given N sampled trajectories for a problem, we simply select the one with the lowest Entropy Centroid. We evaluate the Lowest Centroid selection strategy across model scales from 14B to 480B parameters on mathematics, code generation, logical reasoning, and agentic tasks. Across all settings, Lowest Centroid consistently outperforms existing baselines (Section 4), with gains that remain stable as model size increases, demonstrating that the Entropy Centroid is a general-purpose intrinsic reward signal for test-time scaling.

2 Basic Units for Measuring Inference Quality

Our motivation is to capture temporal patterns along the trajectory. Before defining a method for encoding temporal information, we first need to propose a suitable basic unit to characterize inference state.

Token-level entropy is an intuitive measure of model uncertainty, but it is unreliable as a reward signal (Gupta et al., 2024) for two reasons. First, as discussed in Section 1, the model inherently assigns high confidence to tokens involving numerical calculations or simple sentence completions, and the trajectory-level average score loses the temporal structure. Second, supervised fine-tuning and reinforcement learning tend to concentrate prediction probability on high-reward token sequences (Lu & Lab, 2025). This mode-covering behavior causes token-level confidence to reflect training incentives rather than true inference quality (Gao et al., 2023; Slocum et al., 2025), partially distorting the output logit distributions. To obtain a stable signal, we must elevate the measurement unit from single tokens to token groups.

Definition of High Entropy Phase (HEP). Let $\mathcal{T} = \{t_1, t_2, \dots, t_L\}$ denote the token entropy sequence of length L generated during model inference. We define θ_{high} as the top entropy threshold and θ_{low} as the bottom entropy threshold over the sequence. To formalize the boundaries of a High Entropy Phase, we introduce a state variable $S_i \in \{0, 1\}$ to denote the phase state at the i -th token, where $S_i = 1$ indicates that the model is currently in a HEP.

The state transition is defined recursively as follows:

$$S_i = \begin{cases} 1, & \text{if } S_{i-1} = 0 \text{ and } t_i \geq \theta_{high} \\ 0, & \text{if } S_{i-1} = 1 \text{ and } \sum_{j=0}^{k-1} \mathbb{I}(t_{i-j} \leq \theta_{low}) = k \\ S_{i-1}, & \text{otherwise} \end{cases} \quad (1)$$

Here, $\mathbb{I}(\cdot)$ denotes the indicator function. The model begins in the inactive state $S_0 = 0$, i.e., outside any HEP. A HEP is triggered whenever the model is inactive ($S_{i-1} = 0$) and encounters a token whose entropy reaches the threshold ($t_i \geq \theta_{high}$), upon which the state transitions to $S_i = 1$. To determine phase exit, we employ a sliding window of size k : a HEP terminates (S_i resets to 0) only when the most recent k consecutive tokens all fall below the lower threshold θ_{low} . In all other cases, the current state persists. This recursive procedure applies identically throughout the remainder of the sequence.

A contiguous sub-sequence of tokens where $S_i = 1$ constitutes a single HEP. Unlike fixed-length windows that impose arbitrary boundaries, A HEP provide a natural and adaptive unit for representing the model’s inference state at each stage of inference, preserving the integrity of each uncertain phase as a coherent unit.

3 Entropy Centroid

The High Entropy Phase (HEP) formulation captures the extent of uncertainty at each inference step via phase mass, but a key question remains: *how should we encode the temporal information of these phases to assess overall inference quality?*

3.1 Entropy-Based Centroid Computation

Inspired by the center of mass in physics, we define the Entropy Centroid to summarize where uncertainty concentrates across the generation sequence. Each HEP serves as the basic unit in this computation. For the n -th HEP in the sequence, denoting its start index as a_n and end index as b_n , its mass m_n and position p_n are formulated as follows:

$$m_n = b_n - a_n + 1, \quad p_n = \frac{a_n + b_n}{2} \quad (2)$$

That is, m_n measures the length (in tokens) of the n -th HEP, while p_n locates its midpoint within the sequence. Note that here we treat every token within HEPs equally with a mass of one, and do not use their specific entropy values. As discussed in Section 2, token-level entropy values are inherently noisy, which motivates us to adopt this coarse-grained strategy.

Assuming N HEPs are detected in the full sequence, the position (normalized by the total sequence length L) of the centroid is defined as:

$$C = \frac{\sum_{n=1}^N m_n \cdot p_n}{L \cdot \sum_{n=1}^N m_n} \quad (3)$$

Trajectory lengths vary across samples, so we normalize the centroid by the total sequence length L to obtain a scale-invariant measure $C \in (0, 1)$, ensuring that centroid values remain comparable regardless of trajectory length. A lower C indicates that the uncertainty of the model concentrates in the early stages of inference, while a higher C indicates uncertainty concentrated toward the later stages. Since C is the primary quantity used throughout our analysis, we refer to it simply as the *centroid* hereafter.

While it seems we ignore tokens outside HEPs when computing the entropy centroid, our definition in Eq. 3 essentially models entropy of all tokens as the mass to compute the centroid, but adopts a binary mass approximation — we approximate the mass of those low-entropy tokens outside HEPs as zero, and the mass of each token within HEPs as one.

This is to avoid potential noises from the specific entropy values. In Section 4.3, we will compare with entropy centroid computed from raw entropy values empirically.

3.2 Lowest-Centroid Selection for Test-Time Scaling

The centroid calculated from the High Entropy Phases reflects the certainty of the model during the inference process. We plot boxplots in Figure 2 to illustrate the centroid distribution across different trajectories. The figures clearly demonstrate significant differences in the centroid distributions between correct and incorrect trajectories. Trajectories with correct answers typically exhibit an inference pattern of exploration followed by exploitation. Incorrect trajectories show excessive early exploitation, which leads the subsequent inference into confusion.

Entropy Centroid reflects the exploration-exploitation trade-off in sequential generation. Early in the inference trace, the model must explore diverse inference paths. High entropy at this stage reflects productive exploration. As generation progresses toward a conclusion, the model should exploit its accumulated inference and converge. High entropy at this stage instead signals unresolved confusion. A trajectory whose uncertainty concentrates early is therefore more likely to yield a correct answer than one whose uncertainty persists late.

Therefore, we propose Lowest Centroid based on the Entropy Centroid as a trajectory selection strategy in test-time scaling. This approach characterizes the intrinsic features of the model during inference. The implementation of this method is straightforward and highly effective.

Our method straightforwardly selects the trajectory with the lowest centroid from multiple candidate trajectories. Issues such as abnormal termination, repetitive outputs, and gibberish generation may occur during inference. The entropy distributions of these specific trajectories differ significantly from those of normal inference trajectories. Consequently, we first perform an outlier filtering step to remove trajectories that deviate significantly from the mean centroid of all candidates, and then we select the trajectory with the lowest centroid from the remaining set.

4 Experiments

4.1 Experimental Setup

Models. To validate the generality of the entropy centroid method across various LLM architectures, we evaluate a wide array of model families with parameter scales ranging from 14B to 480B. This selection includes long chain-of-thought (CoT) reasoning models and advanced models specialized in programming and agentic workflows, specifically Olmo-3.1-32B-Thinking (Olmo, 2025), GPT-OSS-120B (OpenAI, 2025), Qwen3-14B (Team, 2025a), QwQ-32B (Team, 2025b), Ministral-3-14B (team, 2026), Qwen3-Coder-480B-A35B-Instruct-FP8 (Team, 2025a), Minimax M2.5 (Chen et al., 2025a) and Qwen3-Coder-Next (Qwen Team, 2026). More experiment details are in Appendix B.

Benchmarks. The evaluation involves comprehensive benchmarks across diverse domains: (1) **Math.** We use AIME 2025 (Zhang & Math-AI, 2025) and Minerva Math (Lewkowycz et al., 2022), which represent challenging mathematics competition-level datasets. (2) **Logic.** We use Synlogic (Liu et al., 2025) to assess complex logical tasks, which include sudoku, arrow maze and puzzles. (3) **Code.** Code generation performance, including self-repair and execution, is measured on competitive programming tasks from BigCodeBench (Zhuo et al., 2024) and LiveCodeBench (Jain et al., 2024). (4) **Agent.** We investigate agentic capabilities using τ^2 -Bench (Barres et al., 2025), requiring models to perform multi-step reasoning, planning, and tool use to reach specific goals without human intervention.

Setup. We evaluate our approach on datasets spanning multiple domains, including agentic tasks, coding, mathematics, and logical reasoning. To ensure experimental rigor, we generate an evaluation cache after model inference. All subsequent methods read from

Acc (% , $\uparrow$)	Method	Agent@64	Code@32		Math@64		Logic@64
Model	Dataset	τ^2 -Bench	BigCode	LiveCode	AIME25	Minerva	Synlogic
Olmo-3.1-Think	Pass@1	9.5	42.6	74.2	75.7	55.0	42.1
	Greedy Decoding	7.0 $\downarrow 2.5$	44.9 $\uparrow 2.3$	71.6 $\downarrow 2.6$	70.0 $\downarrow 5.7$	54.8 $\downarrow 0.2$	40.3 $\downarrow 1.8$
	Self-Certainty	7.0 $\downarrow 2.5$	43.3 $\uparrow 0.7$	76.6 $\uparrow 2.4$	76.7 $\uparrow 1.0$	54.4 $\downarrow 0.6$	42.3 $\uparrow 0.2$
	Tail Confidence	4.0 $\downarrow 5.5$	41.5 $\downarrow 1.1$	76.3 $\uparrow 2.1$	73.3 $\downarrow 2.4$	54.4 $\downarrow 0.6$	46.0 $\uparrow 3.9$
	Bottom Window	8.0 $\downarrow 1.5$	43.4 $\uparrow 0.8$	75.9 $\uparrow 1.7$	80 $\uparrow 4.3$	54.8 $\downarrow 0.2$	48.9 $\uparrow 6.8$
	Lowest Centroid	12.0 $\uparrow 2.5$	44.5 $\uparrow 1.9$	77.2 $\uparrow 3.0$	86.7 $\uparrow 11.0$	55.9 $\uparrow 0.9$	50.0 $\uparrow 7.9$
Qwen3-14B	Pass@1	25.9	39.9	77.2	71.4	56.3	44.4
	Greedy Decoding	29.0 $\uparrow 3.1$	39.8 $\downarrow 0.1$	72.6 $\downarrow 4.6$	70.0 $\downarrow 1.4$	55.5 $\downarrow 0.8$	38.9 $\downarrow 5.5$
	Self-Certainty	18.0 $\downarrow 7.9$	46.2 $\uparrow 6.3$	79.5 $\uparrow 2.3$	73.3 $\uparrow 1.9$	55.2 $\downarrow 1.1$	50.3 $\uparrow 5.9$
	Tail Confidence	21.0 $\downarrow 4.9$	45.4 $\uparrow 5.5$	78.6 $\uparrow 1.4$	70.0 $\downarrow 1.4$	55.5 $\downarrow 0.8$	50.3 $\uparrow 5.9$
	Bottom Window	23.0 $\downarrow 2.9$	44.0 $\uparrow 4.1$	77.7 $\uparrow 0.5$	73.3 $\uparrow 1.9$	55.9 $\downarrow 0.4$	48.9 $\uparrow 4.5$
	Lowest Centroid	35.0 $\uparrow 9.1$	46.0 $\uparrow 6.1$	78.1 $\uparrow 0.9$	76.7 $\uparrow 5.3$	58.1 $\uparrow 1.8$	45.4 $\uparrow 1.0$
QWQ-32B	Pass@1	12.2	43.9	76.7	68.5	54.5	46.0
	Greedy Decoding	12.0 $\downarrow 0.2$	43.4 $\downarrow 0.5$	73.2 $\downarrow 3.5$	56.7 $\downarrow 11.8$	53.3 $\downarrow 1.2$	40.9 $\downarrow 5.1$
	Self-Certainty	9.0 $\downarrow 3.2$	44.0 $\uparrow 0.1$	78.1 $\uparrow 1.4$	70.0 $\uparrow 1.5$	54.0 $\downarrow 0.5$	48.6 $\uparrow 2.6$
	Tail Confidence	8.0 $\downarrow 4.2$	42.3 $\downarrow 1.6$	77.4 $\uparrow 0.7$	76.7 $\uparrow 8.2$	54.8 $\uparrow 0.3$	44.6 $\downarrow 1.4$
	Bottom Window	9.0 $\downarrow 3.2$	44.1 $\uparrow 0.2$	77.9 $\uparrow 1.2$	70.0 $\uparrow 1.5$	55.9 $\uparrow 1.4$	55.7 $\uparrow 9.7$
	Lowest Centroid	18.0 $\uparrow 5.8$	44.5 $\uparrow 0.6$	78.0 $\uparrow 1.3$	83.3 $\uparrow 14.8$	57.0 $\uparrow 2.5$	56.0 $\uparrow 10.0$
Ministral-3-14B-Instruct	Pass@1	26.5	22.9	40.3	27.5	46.7	23.1
	Greedy Decoding	33.0 $\uparrow 6.5$	26.3 $\uparrow 3.4$	39.7 $\downarrow 0.6$	30.0 $\uparrow 2.5$	48.2 $\uparrow 1.5$	16.3 $\downarrow 6.8$
	Self-Certainty	22.0 $\downarrow 4.5$	18.0 $\downarrow 4.9$	40.0 $\downarrow 0.3$	43.3 $\uparrow 15.8$	48.2 $\uparrow 1.5$	20.0 $\downarrow 3.1$
	Tail Confidence	21.0 $\downarrow 5.5$	19.1 $\downarrow 3.8$	46.8 $\uparrow 6.5$	46.7 $\uparrow 19.2$	46.0 $\downarrow 0.7$	38.9 $\uparrow 15.8$
	Bottom Window	21.0 $\downarrow 5.5$	19.0 $\downarrow 3.9$	44.5 $\uparrow 4.2$	36.7 $\uparrow 9.2$	46.3 $\downarrow 0.4$	36.6 $\uparrow 13.5$
	Lowest Centroid	33.0 $\uparrow 6.5$	25.5 $\uparrow 2.6$	44.4 $\uparrow 4.1$	40.0 $\uparrow 12.5$	50.4 $\uparrow 3.7$	33.4 $\uparrow 10.3$
GPT-OSS-120B	Pass@1	54.9	35.0	81.2	76.5	38.6	63.1
	Greedy Decoding	58.0 $\uparrow 3.1$	36.1 $\uparrow 1.1$	82.0 $\uparrow 0.8$	76.7 $\uparrow 0.2$	37.9 $\downarrow 0.7$	62.0 $\downarrow 1.1$
	Self-Certainty	57.0 $\uparrow 2.1$	34.9 $\downarrow 0.1$	80.3 $\downarrow 0.9$	70.0 $\downarrow 6.5$	36.0 $\downarrow 2.6$	61.4 $\downarrow 1.7$
	Tail Confidence	53.0 $\downarrow 1.9$	36.8 $\uparrow 1.8$	80.1 $\downarrow 1.1$	70.0 $\downarrow 6.5$	38.2 $\downarrow 0.4$	66.0 $\uparrow 2.9$
	Bottom Window	49.0 $\downarrow 5.9$	36.1 $\uparrow 1.1$	81.3 $\uparrow 0.1$	73.3 $\downarrow 3.2$	38.2 $\downarrow 0.4$	64.3 $\uparrow 1.2$
	Lowest Centroid	65.0 $\uparrow 10.1$	35.8 $\uparrow 0.8$	83.6 $\uparrow 2.4$	90.0 $\uparrow 13.5$	40.1 $\uparrow 1.5$	67.4 $\uparrow 4.3$

Table 1: Benchmarking Results across Agent, Code, Math, and Logic Tasks. All experiments use a unified set of HEP parameters. For each model–dataset pair, we generate a shared evaluation cache. All subsequent trajectory selections are drawn from this cache, ensuring a fair and consistent comparison across methods. We annotate the absolute change compared to the Pass@1 baseline in the table. Our method consistently outperforms the Pass@1 baseline across all settings.

the same cache, guaranteeing evaluation consistency across different selection strategies. We sample 64 responses for the mathematics, logic, and agent tasks, and 32 trajectories for the coding tasks to keep computational overhead manageable. To define the high entropy phase, we set the top entropy token threshold θ_{high} to 1% and the bottom entropy token threshold θ_{low} to 80%. The number of consecutive tokens required to exit the high entropy phase is set to $k = 2$. These hyperparameters regarding HEP computation are kept the same for all models across all benchmarks in our experiments. When computing the entropy of each token in the sequence, we only use the top 10 tokens in the vocabulary with the highest probability, as they already capture most of the probability mass in the output distribution.

Baselines. We report Pass@1 and Greedy Decoding (Holtzman et al., 2019) as a reference for model capability. In addition, we compare against several intrinsic-reward methods

proposed for test-time scaling: (1) **Self-Certainty** (Kang et al., 2025) measures the KL divergence between the model’s output distribution and the uniform distribution, using the resulting score to select among candidate trajectories. (2) **Tail Confidence and Bottom Window** (Fu et al., 2025) aggregate token-level confidence scores over contiguous groups of tokens (e.g., the lowest-confidence window) and originally serve as weights for majority voting. Because majority voting is restricted to short-answer tasks and thus incompatible with our evaluation setting, we repurpose these scores as trajectory-selection criteria instead.

4.2 Lowest Entropy Centroid Evaluations

Main Results. Table 1 presents the main experimental results. Entropy centroid outperforms other intrinsic reward methods in most cases, achieving an average improvement of 5.3% over Pass@1 across all 30 model-dataset combinations. On agentic tasks, the average improvement reaches 6.8%, where GPT-OSS-120B achieves a 10.1% gain over Pass@1. Our method is the only one that outperforms Pass@1 across all settings. In contrast, existing baselines exhibit unstable behavior. For example, Self-Certainty outperforms Pass@1 on only two out of six datasets for Ministral-3-14B-Instruct. Tail Confidence and Bottom Window perform poorly on agentic tasks, where all results fall below the Pass@1 baseline. These comparisons highlight a key advantage of our approach: Rather than excelling on specific settings while failing on others, Lowest Centroid delivers reliable gains across diverse domains and model scales.

Results on Large Agents. Most models in our main experiments are below 50B parameters. Since smaller models perform poorly on agentic tasks and fail to produce meaningful performance differentiation, we additionally evaluate three larger models on this domain. Agentic tasks represent one of the most actively studied areas in current research, making this a valuable testbed for generalization. As shown in Table 2, the Entropy Centroid continues to deliver strong performance at this scale.

Method	Qwen3-Coder	Qwen3-Next	Minimax-M2.5
	480BA35B	80BA3B	230BA10B
Pass@1	58.5	20.1	66.1
Greedy Decoding	56.0 $\downarrow 2.5$	27.0 $\uparrow 6.9$	65.0 $\downarrow 1.1$
Tail Confidence	50.0 $\downarrow 8.5$	2.0 $\downarrow 18.1$	32.0 $\downarrow 33.9$
Bottom Window	53.0 $\downarrow 5.5$	9.0 $\downarrow 11.1$	48.0 $\downarrow 18.1$
Lowest Centroid	65.0 $\uparrow 6.5$	27.0 $\uparrow 6.9$	73.0 $\uparrow 6.9$

Table 2: τ^2 -bench results for large models.

Comparison with Majority Voting. While we believe majority voting is not a practical method as it is only applicable to short-answer tasks, we still compare Lowest Centroid against majority voting on tasks where extracted answers are available. As shown in Table 3 of Appendix B, the two methods achieve comparable overall performance across different datasets and models. We would like to emphasize our approach has much broader usage as already shown in agentic and coding tasks.

4.3 Ablation Study

Transitioning from token-level raw entropy to a group-level high entropy phase provides an important mechanism to reduce noise during centroid computation. A core innovation of our work involves proposing the concept of the high entropy phase to represent the model inference state. As discussed in Section 3.1, we adopt a binary entropy mass approximation to reduce noise. Here, we perform an ablation study on this choice and compare with raw entropy centroid, where we use the raw entropy value of all tokens to compute the centroid. As illustrated in Figure 3, the HEP-based centroid outperforms the raw entropy centroid in all of model-dataset combinations, with a mean accuracy gain of +5.8% absolutely. The improvement is especially large in lower-accuracy settings, where individual token noise has a stronger influence on centroid computation.

The high entropy phase exhibits stable performance and excellent robustness across different hyperparameters. The definition of the high entropy phase relies on three hyperparameters, specifically the top entropy threshold θ_{high} , the bottom entropy threshold θ_{low} , and the number of continuous bottom entropy tokens k that determine the exit from the high entropy phase. We analyze the hyperparameter robustness of HEP using a controlled




Figure 3: **HEP Centroid vs. Raw Entropy Centroid.** Each point represents a model-dataset combination. Points above the diagonal indicate that HEP outperforms Raw Entropy. All of the points fall above the diagonal, with a mean improvement of +5.8%.




Figure 4: Hyperparameter robustness of HEP. We vary three parameters independently using a controlled variable approach: (a) the top percent threshold θ_{high} , (b) the bottom percent threshold θ_{low} , and (c) the consecutive low threshold k . Each curve represents a model-dataset pair across five models and five benchmarks. All curves remain nearly flat across the tested ranges. The mean performance range is 0.6%, 0.3%, and 0.8% for the three parameters respectively, confirming that HEP is robust to hyperparameter choices.

variable approach. As shown in Figure 4, model performance remains largely unaffected by variations in parameter values. For all three parameters, the performance difference remains below 1%, which further confirms HEP is robust to hyperparameter choices. Notably, HEP with $k=1$ is not equivalent to raw token-level entropy, due to our binary entropy mass approximation discussed in Section 3.1. Setting $k > 1$ further prevents isolated low-entropy tokens from prematurely splitting a phase, but the core mechanism remains the same. This explains the comparable performance across different k values. More ablation experiments on filter impact can be found in Appendix B.

5 Scaling Efficiency Analysis

Test-time scaling cost-performance tradeoff is a key metric for evaluating scaling methods. In our setup, cost and performance manifest as the relationship between the number of




Figure 5: Test-time scaling performance across models and datasets. Each subplot shows accuracy as a function of the number of sampled trajectories n . For each n , we repeat the sampling process 50 times. Solid curves denote the mean accuracy and shaded regions denote the standard deviation. Centroid Selection (blue) exhibits steady and consistent scaling across all settings without early saturation. Self Certainty (green) shows unstable scaling with occasional performance drops. Random Selection (orange) remains largely flat, close to the Pass@1 baseline (black dashed).

sampled responses and accuracy. We compare our method against pass@1, random selection, and the intrinsic reward baseline Self-Certainty.

Lowest Centroid exhibits persistent scaling without early saturation. As shown in Figure 5, accuracy steadily increases with more sampled responses, without noticeable diminishing returns. In contrast, Self-Certainty shows unstable scaling with occasional performance drops, lacking a consistent upward trend. Comparisons with pass@1 and random baselines further demonstrate the superior scaling behavior of our method. Results span different task domains and models ranging from 14B to 120B, demonstrating the generality of the Lowest Centroid method.

A key advantage of our method is its broad applicability. Lowest Centroid selects a single trajectory based on intrinsic signals without requiring answer labels, making it applicable to any task type, including coding and agentic tasks where methods like majority voting are fundamentally inapplicable due to their reliance on answer tallying. Beyond trajectory selection, the Entropy Centroid can further serve as a reward signal during training, offering wider application potential.

6 Conclusion

In this work, we propose to explore the temporal pattern during inference. We define High Entropy Phase (HEP) as the basic unit, a group-level unit that segments consecutive high-entropy tokens into coherent phases. Based on HEP, we define the Entropy Centroid to encode positional information along the trajectory. A lower centroid means the model explores early and converges confidently, indicating higher inference quality. Experiments

across different tasks show that the Lowest Centroid outperforms existing baselines, with stable gains from 14B to 480B parameters. We hope the Entropy Centroid inspires further exploration beyond trajectory selection.

Ethics Statement

This work focuses on improving language model test-time scaling through intrinsic reward signals. It does not involve human subjects, private data, or human annotation. All experiments are conducted on publicly available benchmarks. We do not foresee any direct negative societal impact from this work.

In preparing this manuscript, large language models (LLMs) were used in a limited and assistive capacity. Specifically, LLMs were employed solely for shallow grammar proofreading and surface-level linguistic polishing of the drafted text.

All scientific claims, experimental designs, data analyzes, and core intellectual contributions remain entirely the work of the authors. The authors take full responsibility for the correctness and integrity of the final manuscript.

References

- Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. τ^2 -bench: Evaluating conversational agents in a dual-control environment. *arXiv preprint arXiv:2506.07982*, 2025.
- Lang Cao, Renhong Chen, Yingtian Zou, Chao Peng, Huacong Xu, Yuxian Wang, Wu Ning, Qian Chen, Mofan Peng, Zijie Chen, et al. More bang for the buck: Process reward modeling with entropy-driven uncertainty. *arXiv preprint arXiv:2503.22233*, 2025.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. Chateval: Towards better llm-based evaluators through multi-agent debate. *arXiv preprint arXiv:2308.07201*, 2023.
- Aili Chen, Aonian Li, Bangwei Gong, Binyang Jiang, Bo Fei, Bo Yang, Boji Shan, Changqing Yu, Chao Wang, Cheng Zhu, et al. Minimax-m1: Scaling test-time compute efficiently with lightning attention. *arXiv preprint arXiv:2506.13585*, 2025a.
- Peter Chen, Xiaopeng Li, Ziniu Li, Xi Chen, and Tianyi Lin. Spectral policy optimization: Coloring your incorrect reasoning in grpo. In *2nd AI for Math Workshop@ ICML 2025*, 2025b.
- Shiqi Chen, Miao Xiong, Junteng Liu, Zhengxuan Wu, Teng Xiao, Siyang Gao, and Junxian He. In-context sharpness as alerts: An inner representation perspective for hallucination mitigation. *arXiv preprint arXiv:2403.01548*, 2024.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. Dola: Decoding by contrasting layers improves factuality in large language models. *arXiv preprint arXiv:2309.03883*, 2023.
- Yuyang Ding, Xinyu Shi, Juntao Li, Xiaobo Liang, Zhaopeng Tu, and Min Zhang. Scan: Self-denoising monte carlo annotation for robust process reward learning. *arXiv preprint arXiv:2509.16548*, 2025.
- Wu Fei, Hao Kong, Shuxian Liang, Yang Lin, Yibo Yang, Jing Tang, Lei Chen, and Xiansheng Hua. Self-guided process reward optimization with masked step advantage for process reinforcement learning. *arXiv preprint arXiv:2507.01551*, 2025.
- Tony Feng, Trieu H Trinh, Garrett Bingham, Dawsen Hwang, Yuri Chervonyi, Junehyuk Jung, Joonkyung Lee, Carlo Pagano, Sang-hyun Kim, Federico Pasqualotto, et al. Towards autonomous mathematics research. *arXiv preprint arXiv:2602.10177*, 2026.

- Yichao Fu, Xuewei Wang, Yuandong Tian, and Jiawei Zhao. Deep think with confidence. *arXiv preprint arXiv:2508.15260*, 2025.
- Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, pp. 10835–10866. PMLR, 2023.
- Soumya Suvra Ghosal, Souradip Chakraborty, Avinash Reddy, Yifu Lu, Mengdi Wang, Dinesh Manocha, Furong Huang, Mohammad Ghavamzadeh, and Amrit Singh Bedi. Does thinking more always help? mirage of test-time scaling in reasoning models. *arXiv preprint arXiv:2506.04210*, 2025.
- Neha Gupta, Harikrishna Narasimhan, Wittawat Jitkrittum, Ankit Singh Rawat, Aditya Krishna Menon, and Sanjiv Kumar. Language model cascades: Token-level uncertainty and beyond. *arXiv preprint arXiv:2404.10136*, 2024.
- Zhitao He, Sandeep Polisetty, Zhiyuan Fan, Yuchen Huang, Shujin Wu, and Yi R Fung. Mmboundary: Advancing mllm knowledge boundary awareness through reasoning step confidence calibration. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 16427–16444, 2025a.
- Zhitao He, Haolin Yang, Zeyu Qin, and Yi R Fung. Medtutor-r1: Socratic personalized medical teaching with multi-agent simulation. *arXiv preprint arXiv:2512.05671*, 2025b.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019.
- Yuelin Hu, Zhengxue Cheng, Wei Liu, and Li Song. Entropy-gated selective policy optimization: Token-level gradient allocation for hybrid training of large language models. *arXiv preprint arXiv:2602.03309*, 2026.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.
- Zhewei Kang, Xuandong Zhao, and Dawn Song. Scalable best-of-n selection for large language models via self-certainty. *arXiv preprint arXiv:2502.18581*, 2025.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35:3843–3857, 2022.
- Dong Li, Xujiang Zhao, Linlin Yu, Yanchi Liu, Wei Cheng, Zhengzhang Chen, Zhong Chen, Feng Chen, Chen Zhao, and Haifeng Chen. Solverllm: Leveraging test-time scaling for optimization problem via llm-guided search. *arXiv preprint arXiv:2510.16916*, 2025.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The twelfth international conference on learning representations*, 2023.
- Junteng Liu, Yuanxiang Fan, Zhuo Jiang, Han Ding, Yongyi Hu, Chi Zhang, Yiqi Shi, Shitong Weng, Aili Chen, Shiqi Chen, et al. Synlogic: Synthesizing verifiable reasoning data at scale for learning logical reasoning and beyond. *arXiv preprint arXiv:2505.19641*, 2025.
- Kevin Lu and Thinking Machines Lab. On-policy distillation. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20251026. <https://thinkingmachines.ai/blog/on-policy-distillation>.

- Kou Misaki, Yuichi Inoue, Yuki Imajuku, So Kuroki, Taishi Nakamura, and Takuya Akiba. Wider or deeper? scaling llm inference-time compute with adaptive branching tree search. In *ICLR 2025 Workshop on Foundation Models in the Wild*, 2025.
- Team Olmo. Olmo 3, 2025. URL <https://arxiv.org/abs/2512.13961>.
- OpenAI. gpt-oss-120b & gpt-oss-20b model card, 2025. URL <https://arxiv.org/abs/2508.10925>.
- Dwarkesh Patel and Ilya Sutskever. Ilya sutskever – we’re moving from the age of scaling to the age of research. YouTube, 2025. URL <https://youtu.be/aR20FWCCjAs>. Accessed: 2026-03-30.
- Qwen Team. Qwen3-coder-next technical report. Technical report, Alibaba Group, 2026. URL https://github.com/QwenLM/Qwen3-Coder/blob/main/qwen3_coder_next_tech_report.pdf. Accessed: 2026-02-03.
- Aman Sharma and Paras Chopra. Think just enough: Sequence-level entropy as a confidence signal for llm reasoning. *arXiv preprint arXiv:2510.08146*, 2025.
- Stewart Slocum, Asher Parker-Sartori, and Dylan Hadfield-Menell. Diverse preference learning for capabilities and alignment. *arXiv preprint arXiv:2511.08594*, 2025.
- Minstral team. Minstral 3, 2026. URL <https://arxiv.org/abs/2601.08584>.
- Qwen Team. Qwen3 technical report, 2025a. URL <https://arxiv.org/abs/2505.09388>.
- Qwen Team. Qwq-32b: Embracing the power of reinforcement learning, March 2025b. URL <https://qwenlm.github.io/blog/qwq-32b/>.
- Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9426–9439, 2024.
- Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*, 2025a.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- Yutong Wang, Pengliang Ji, Chaoqun Yang, Kaixin Li, Ming Hu, Jiaoyang Li, and Guillaume Sartoretti. Mcts-judge: Test-time scaling in llm-as-a-judge for code correctness evaluation. *arXiv preprint arXiv:2502.12468*, 2025b.
- xAI. Grok 4 heavy parallel test-time compute. <https://x.ai/news/grok-4>, 2026. Accessed: 2026-03-31.
- Bin Xie, Bingbing Xu, Xueyun Tian, Yilin Chen, and Huawei Shen. Towards robust process reward modeling via noise-aware learning. *arXiv preprint arXiv:2601.12748*, 2026.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. Qwen2.5-math technical report: Toward mathematical expert model via self-improvement, 2024. URL <https://arxiv.org/abs/2409.12122>.
- Haolin Yang, Jipeng Zhang, Zhitao He, and Yi R Fung. Mars-sql: A multi-agent reinforcement learning framework for text-to-sql. *arXiv preprint arXiv:2511.01008*, 2025.

- Pengyue Yang, Jiawen Wen, Haolin Jin, Linghan Huang, Huaming Chen, and Ling Chen. Trust in one round: Confidence estimation for large language models via structural signals. *arXiv preprint arXiv:2602.00977*, 2026.
- Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. Limo: Less is more for reasoning. *arXiv preprint arXiv:2502.03387*, 2025.
- Bin Yu, Xinming Wang, Shijie Lian, Haotian Li, Changti Wu, Ruina Hu, Bailing Wang, Yuliang Wei, and Kai Chen. Trajselector: Harnessing latent representations for efficient and effective best-of-n in large reasoning model. *arXiv preprint arXiv:2510.16449*, 2025.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason E Weston. Self-rewarding language models. In *Forty-first International Conference on Machine Learning*, 2024.
- Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D Goodman. Quiet-star: Language models can teach themselves to think before speaking. *arXiv preprint arXiv:2403.09629*, 2024.
- Yao Zhang, Yu Wu, Haowei Zhang, Weiguo Li, Haokun Chen, Jingpei Wu, Guohao Li, Zhen Han, and Volker Tresp. Groundedprm: Tree-guided and fidelity-aware process reward modeling for step-level reasoning. *arXiv preprint arXiv:2510.14942*, 2025.
- Yifan Zhang and Team Math-AI. American invitational mathematics examination (aime) 2025, 2025.
- Wenshuo Zhao, Haoxing Zhai, Xinyu Qiu, Zhenting Qi, Shuhe Li, and Linchao Zhu. MuTIS: Enhancing reasoning efficiency through multi turn intervention sampling in reinforcement learning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, Suzhou, China, November 2025a. Association for Computational Linguistics.
- Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. Learning to reason without external rewards. *arXiv preprint arXiv:2505.19590*, 2025b.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.
- Yujun Zhou, Zhenwen Liang, Haolin Liu, Wenhao Yu, Kishan Panaganti, Linfeng Song, Dian Yu, Xiangliang Zhang, Haitao Mi, and Dong Yu. Evolving language models without labels: Majority drives selection, novelty promotes variation. *arXiv preprint arXiv:2509.15194*, 2025.
- Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, et al. Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions. *arXiv preprint arXiv:2406.15877*, 2024.
- Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen Zhang, Xinwei Long, Ermo Hua, et al. Ttrl: Test-time reinforcement learning. *arXiv preprint arXiv:2504.16084*, 2025.

A Additional Related Work

A.1 Intrinsic rewards as metrics for LLM capabilities.

Exploring intrinsic rewards is essential for achieving self-improvement and continual learning in language models. Current research actively analyzes and models the output distribution. For example, Kang et al. (2025); Zhao et al. (2025b) calculate the KL divergence between the model distribution and a uniform distribution to serve as an intrinsic reward for self-certainty. Other approaches (Fu et al., 2025; Sharma & Chopra, 2025) use output confidence over specific time windows as a signal to enable early stopping for low-quality trajectories and to optimize computational resource allocation during training, making language models more efficient (Zhao et al., 2025a). Furthermore, Wang et al. (2025a); Hu et al. (2026) confirm the significant value of high entropy tokens in model outputs. Additional studies (Chuang et al., 2023; Chen et al., 2024; Yu et al., 2025; He et al., 2025a; Yang et al., 2026) measure confidence and factuality using hidden states to inform trajectory selectors. However, existing intrinsic reward methods exhibit several limitations. Averaging the complete trajectory of an inference rollout ignores variations in local inference steps. Treating the logits distribution as a fixed signal fails to account for semantic information across different contexts. These methods also struggle to avoid the noise inherent in token-level signals, which ultimately leads to distorted reward signals.

A.2 Test-time Scaling Methods

Beyond intrinsic states, researchers explore various unlabeled optimization approaches. Test-time reinforcement learning (Zuo et al., 2025; Zhou et al., 2025) utilizes model voting results to achieve self-improvement. Incorporating large language models for evaluation during test time is another widely adopted practice (Zheng et al., 2023; Chan et al., 2023; Wang et al., 2025b). Researchers often employ multi-agent debate (Yang et al., 2025; He et al., 2025b) to improve verification accuracy within this context. These evaluation methods require external models, which incur additional overhead and introduce model-based noise. Monte Carlo tree search is frequently applied to expand the solution space of model outputs during test time (Misaki et al., 2025; Li et al., 2025). This search enhances model performance through systematic exploration and backtracking mechanisms. Nevertheless, traversing the solution space demands substantial computational resources and exhibits diminishing returns (Ghosal et al., 2025), making these search methods unsuitable for practical production environments.

B Additional Experiment Details

B.1 Additional Experiment Results

We present more majority voting results here for readers’ reference.

Model	AIME25			Minerva			Synlogic		
	Pass@1	Major. Voting	Lowest Cent.	Pass@1	Major. Voting	Lowest Cent.	Pass@1	Major. Voting	Lowest Cent.
Olmo-3.1	75.7	80.0 $\uparrow 4.3$	86.7 $\uparrow 11.0$	55.0	56.3 $\uparrow 1.3$	55.9 $\uparrow 0.9$	42.1	46.9 $\uparrow 4.8$	50.0 $\uparrow 7.9$
Ministral-3	27.5	40.0 $\uparrow 12.5$	40.0 $\uparrow 12.5$	46.7	55.2 $\uparrow 8.5$	50.4 $\uparrow 3.7$	23.1	41.1 $\uparrow 18.0$	33.4 $\uparrow 10.3$
Qwen3-14B	71.4	80.0 $\uparrow 8.6$	76.7 $\uparrow 5.3$	56.3	56.6 $\uparrow 0.3$	58.1 $\uparrow 1.8$	44.4	49.1 $\uparrow 4.7$	45.4 $\uparrow 1.0$
GPT-OSS	76.5	90.0 $\uparrow 13.5$	90.0 $\uparrow 13.5$	38.6	39.7 $\uparrow 1.1$	40.1 $\uparrow 1.5$	63.1	68.9 $\uparrow 5.8$	67.4 $\uparrow 4.3$
QWQ-32B	68.5	80.0 $\uparrow 11.5$	83.3 $\uparrow 14.8$	54.5	56.3 $\uparrow 1.8$	57.0 $\uparrow 2.5$	46.0	48.0 $\uparrow 2.0$	56.0 $\uparrow 10.0$

Table 3: For short-answer tasks where majority voting is applicable, we compare Lowest Centroid against majority voting. The results show that the two methods achieve comparable performance.

Model $\Delta(\%, \downarrow)$	Bigcode	Livcode	Minerva	Synlogic	AIME
Olmo-3.1-32B-Think	0.54	0.09	1.32	2.29	3.33
Qwen3-14B	0.21	2.26	1.10	3.37	6.67
QWQ-32B	0.98	0.47	0.44	1.09	6.67
Ministral-3-14B-Instruct	0.00	3.15	2.28	1.66	4.67
GPT-OSS-120B	0.09	1.02	0.81	0.91	4.67

Table 4: Accuracy changes before and after outlier filtering across different datasets and models. **For each dataset–model combination, we evaluate 5 distinct high-entropy phase configurations and report the average accuracy difference as the filtering impact.**

BIGCODE		LIVECODE		MATH		LOGIC	
max_token	16384	max_tokens	32768	max_tokens	32768	max_tokens	32768
batch_size	16	batch_size	256	batch_size	128	batch_size	64
temperature	0.7	temperature	0.7	temperature	0.7	temperature	1.0
n(#trajectories)	32	n(#trajectories)	32	n(#trajectories)	64	n(#trajectories)	64

Table 5: Inference Hyperparameter Settings

B.2 Additional Ablation Results

Outlier filtering primarily addresses extreme cases and exerts a minimal impact on overall performance. As previously mentioned, certain inference trajectories suffer from early stopping and repetition issues.

To verify the impact of outlier filtering on experimental results, we compare the accuracy before and after filtering. We conduct this comparison across different model–dataset combinations to ensure rigor. For most models, the effect of outlier filtering on performance is less than 1%. We validate this on every model–dataset pair, and the results are presented in Table 4. As shown, outlier filtering has a negligible impact on model performance, further confirming that our performance gains primarily stem from the design of the entropy centroid. Outlier filtering serves solely to remove trajectories caused by repetition or early stopping. The filtering effect in AIME 2025 appears to have a significant impact on the data, but this is because AIME only has 30 questions. While affecting just one question would result in a 3.3% increase in accuracy, in reality, only about 1.5 questions are affected on average in AIME.

B.3 Evaluation Details

Our inference framework is built on vLLM (Kwon et al., 2023). The evaluation for each domain task follows well-established references. For mathematical datasets, we adopt the evaluation logic from LIMO (Ye et al., 2025) and Qwen (Yang et al., 2024), relying exclusively on rule-based rewards without any model-based judgment. All other datasets are evaluated using the original evaluation code provided by their respective authors.

To ensure evaluation stability and experimental rigor, we first run a unified evaluation pass during the inference stage to generate an **evaluation cache**. All subsequent evaluation methods read directly from this cache, eliminating potential inconsistencies introduced by non-deterministic evaluation procedures.

Table 5 and 6 summarize the inference hyperparameters for each dataset. We follow the officially recommended configurations and make dataset-specific adjustments according to available computational resources. Since τ^2 -Bench is a multi-turn agent benchmark, its parameter configuration differs from those of the other single-turn datasets.

τ^2 -Bench	
AGENT_MAX_TOKENS	8192
USER_MAX_TOKENS	2048
MAX_STEPS	200
AGENT_TEMP	0.7
USER_TEMP	0.7
SEED	300
n(#trajectories)	64

Table 6: τ^2 -Bench Inference Hyperparameter Settings

B.4 Dataset Details

AIME 2025 (Zhang & Math-AI, 2025) is the 2025 edition of the American Invitational Mathematics Examination, a prestigious invitational competition targeting high-performing high school students. It comprises 30 problems across two sessions (AIME I and AIME II, each with 15 problems), covering algebra, geometry, number theory, combinatorics, and probability. Each problem requires an integer answer between 000 and 999 and demands creative, multi-step mathematical reasoning well beyond standard curricula. We use all 30 problems as our test set.

Minerva Math (Lewkowycz et al., 2022) is a mathematical reasoning benchmark derived from the Minerva evaluation suite, consisting of undergraduate-level STEM problems that require multi-step quantitative reasoning. Problems span topics such as calculus, linear algebra, probability, and differential equations, and models must produce detailed solutions with precise numerical or symbolic answers. Its test set contains 272 unique problems.

Synlogic (Liu et al., 2025) is a synthetic logical reasoning benchmark in which problems are algorithmically generated to ensure verifiable ground-truth answers. It evaluates language models’ capacity for formal deductive reasoning across multiple categories of logical inference. We use the Val-hard subset, which contains the most challenging problems requiring complex, multi-step logical deductions.

LiveCodeBench (Jain et al., 2024) is a contamination-free code generation benchmark that continuously curates new problems from competitive programming platforms such as LeetCode, AtCoder, and Codeforces. Unlike static benchmarks, its rolling release strategy mitigates data leakage. We use release v6, which evaluates code generation capabilities where models must produce correct, executable solutions given natural language problem descriptions and test cases.

BigCodeBench (Zhuo et al., 2024) is a practical programming benchmark comprising 1,140 function-level tasks that span 7 domains and require the use of 139 diverse libraries. Unlike competition-style benchmarks, it emphasizes real-world programming scenarios demanding practical API knowledge and compositional reasoning. We evaluate using the instruct mode (v0.1.4), where models receive natural language task descriptions and must generate standalone, executable functions.

τ^2 -bench (Barres et al., 2025) is an agentic benchmark designed to evaluate language model agents’ ability to handle realistic, multi-turn customer service interactions, where agents must invoke tools and adhere to domain-specific policies to resolve user requests. We use the Test subset containing 100 items spanning three domains: airline (20 items), retail (40 items), and telecom (40 items), collectively covering diverse tool-use and policy-following scenarios.




Figure 6: The mean Entropy values differ across different problems, which verifies that using absolute value as a threshold is not feasible.

C Implementation Details

C.1 Additional Method Details

C.1.1 Percentile-based Threshold for HEP Detection

We use a percentile-based threshold rather than an absolute entropy value to define HEP. This design choice is motivated by both theoretical and practical considerations.

First, entropy distributions vary significantly across models and datasets, as shown in Figure 6. A fixed absolute threshold cannot generalize across these different distributions. For example, if the threshold is set to 3 but no token in a trajectory exceeds an entropy of 3, the resulting centroid defaults to zero, producing a meaningless signal.

Second, the core of our method is to characterize entropy variation across the trajectory, i.e., the transitions between inference states. An absolute threshold cannot guarantee a sufficient number of high-entropy tokens for stable state characterization. In contrast, a percentile-based threshold adapts to the entropy distribution of each trajectory, ensuring that the relative fluctuation patterns are consistently captured regardless of the model or dataset.

C.1.2 Discover LLM Emotion Pattern for Self Improvement

The perspective of emotion originates from [Patel & Sutskever \(2025\)](#). Self-improvement has emerged as a central topic in current research, with state-of-the-art models increasingly emphasizing iterative self-refinement as a key component of their development. The question of how to achieve model self-improvement has evolved from an open research problem into a widely adopted training paradigm. In this context, the entropy centroid we propose offers a promising intrinsic reward signal and value function for self-improvement. Beyond serving as an objective score that reflects inference quality, the entropy centroid provides a deeper characterization of the model’s internal reasoning dynamics. We move beyond token-level entropy, which suffers from substantial noise and limited interpretability,

and introduce a new paradigm for characterizing inference states — one that yields dense, high-quality signals for model training.

C.2 Future Work

In this work, we focus on applying the Entropy Centroid as an intrinsic reward for test-time scaling. However, the underlying modeling framework, characterizing inference states through HEP and encoding their temporal structure via centroids, is not restricted to this setting. The only requirement is access to the model’s token-level entropy during generation, which is readily available in any autoregressive decoding process. This makes the approach naturally extensible to other scenarios. For example, the Entropy Centroid could serve as an auxiliary reward signal in on-policy training, providing a lightweight and annotation-free complement to external rewards. More broadly, the temporal patterns captured by HEP may offer useful diagnostics for analyzing model behavior during long-form generation, multi-turn dialogue, or agentic workflows. We leave the exploration of these directions to future work.

C.3 Prompt Design

When setting the prompt, we generally align it with the default prompt template of the dataset to ensure the rigor and fairness of the evaluation. The specific prompts for different datasets are as follows:

MATH DATASET

SYSTEM: You are a math expert. Please solve problems step by step and put your final answer within boxed{ }.

USER: Problem: {problem}.

SYNOLOGIC DATASET

SYSTEM: You are a logic puzzle solving expert. Analyze the problem carefully, think step by step, and provide your final answer within <answer> and </answer> tags.

USER: {prompt} (Extracting the prompt field from the HuggingFace dataset).

BIGCODEBENCH DATASET

SYSTEM: You are an expert Python programmer. Write clean, efficient, and correct Python code to solve the given task. Include all necessary imports and ensure the code is complete and executable.

USER:

Task:
{instruct_prompt}

LIVECODEBENCH DATASET

SYSTEM: You are an expert Python programmer. You will be given a question (problem specification) and will generate a correct Python program that matches the specification and passes all tests. Wrap your final solution in ```python``` code blocks.

USER: (with starter_code, LeetCode Style)

Question:
{question_content}

```
### Format:
You will use the following starter code to write the solution...
```python
{starter_code}
```

### Answer:
(use the provided format with backticks)

USER: (without starter_code, Codeforces/AtCoder Style)

### Question:
{question_content}

### Format:
Read the inputs from stdin, solve the problem and write the answer to stdout...
```python
YOUR CODE HERE
```

### Answer:
(use the provided format with backticks)
```

D Case Study

We provide detailed inference trajectories for reference. For each problem, we select a correct trajectory with a low centroid and an incorrect trajectory with a high centroid. These concrete examples illustrate how the model’s inference state evolves in each case, offering a more intuitive understanding of how the Entropy Centroid reflects the reasoning process.

D.1 Code Generation Trajectory

We select a task in Livecodebench and visualize the entire inference process for correct and incorrect trajectories. Figure 7 shows the wrong trajectory, while Figure 8 shows the correct trajectory.

D.2 Logical Reasoning Trajectory

We select a task in Synlogic and visualize the entire inference process for correct and incorrect trajectories. Figure 9 shows the wrong trajectory, while Figure 10 shows the correct trajectory.

D.3 Mathematics Inference Trajectory

We select a task in AIME 2025 and visualize the entire inference process for correct and incorrect trajectories. Figure 11 shows the wrong trajectory, while Figure 12 shows the correct trajectory.

Code Generation (Livecodebench)

Problem(lcb_abc318_b):

There are N rectangular sheets spread out on a coordinate plane. Each side of the rectangular region covered by each sheet is parallel to the x - or y -axis. Specifically, the i -th sheet covers exactly the region satisfying $A_i \leq x \leq B_i$ and $C_i \leq y \leq D_i$. Let S be the area of the region covered by one or more sheets. It can be proved that S is an integer under the constraints. Print S as an integer.

Input

The input is given from Standard Input in the following format: $N \ A_1 \ B_1 \ C_1 \ D_1 \ A_2 \ B_2 \ C_2 \ D_2 \ \vdots \ A_N \ B_N \ C_N \ D_N$

Output

Print the area S of the region covered by one or more sheets as an integer.

Constraints

$-2 \leq N \leq 100$ - $0 \leq A_i, B_i \leq 100$ - $0 \leq C_i, D_i \leq 100$ - All input values are integers.

Sample Input 1

3 0 5 1 3 1 4 0 5 2 5 2 4

Sample Output 1

20

The three sheets cover the following regions. Here, red, yellow, and blue represent the regions covered by the first, second, and third sheets, respectively.

Therefore, the area of the region covered by one or more sheets is $S=20$.

Sample Input 2

2 0 100 0 100 0 100 0 100

Sample Output 2

10000

Note that different sheets may cover the same region.

Sample Input 3

3 0 1 0 1 0 3 0 5 5 10 0 10

Sample Output 3

65

Wrong Trajectory (centroid=0.6222)

Okay, I need to find the area covered by at least one of the rectangles. Hmm, how do I approach this? Let's see.

The problem is to calculate the union area of multiple rectangles. Since the coordinates are up to 100, maybe a grid-based approach could work. Like, check each unit square (like 1x1) and see if it's covered by any rectangle. Then count all such squares. Since the maximum coordinate is 100, the grid would be 100x100, which is manageable.

Wait, right. Because each rectangle's edges are integers, maybe we can discretize the coordinates into integer points. For each small square from (x, y) to $(x+1, y+1)$, we can check if it's covered by any of the rectangles. Then sum all those squares that are covered by at least one.

Yes, that makes sense. Because the area is an integer, this approach should work. So the plan is:

1. Iterate over all possible 1x1 squares in the grid from $(0,0)$ up to $(100,100)$. But actually, since the rectangles can have coordinates up to 100 for B_i and D_i , but since B_i and D_i can be 100, the squares would go up to 99 in x and y . Wait, maybe the squares are from x from 0 to 99, and y from 0 to 99. Because each square is $[x, x+1)$ in x , etc. So each square is 1 unit in each direction.

So for each x in 0 to 99, and each y in 0 to 99, check if there exists at least one rectangle that covers this square.

So the steps:

Loop over each x from 0 to 99:

Loop over each y from 0 to 99:

For this (x, y) square (which is from x to $x+1$ in x , and y to $y+1$ in y), check if any rectangle includes this square.

If yes, add 1 to the total area.

So how to check if the square is covered by any rectangle?

A rectangle is defined by $A_i \leq x \leq B_i$ and $C_i \leq y \leq D_i$. Wait, but the rectangle's x ranges from A_i to B_i , inclusive? Or is it up to B_i ?

Wait the problem says "the i -th sheet covers exactly the region satisfying $A_i \leq x \leq B_i$ and $C_i \leq y \leq D_i$ ". So the rectangle includes all points (x,y) where $A_i \leq x \leq B_i$ and $C_i \leq y \leq D_i$.

So for the square with x -coordinate between x_{val} and $x_{\text{val}}+1$, and y between y_{val} and $y_{\text{val}}+1$, the square is covered by the rectangle if:

the rectangle's $A_i \leq x_{\text{val}}+1$ and $B_i \geq x_{\text{val}}$, and similarly for the y direction.

Wait, let me think. For the square to be entirely within the rectangle, the rectangle must cover the entire square. Wait no, actually, the square is part of the region covered by the rectangle if any part of it is covered. Wait, no. Wait, the square is part of the area if at least one point in the square is covered by the rectangle. But actually, the region covered by the rectangle is a closed region, so the square's square is covered by the rectangle if the square's x interval $[x, x+1)$ overlaps

Figure 7: Code Generation Task Livecodebench Example. Wrong Trajectory with High Centroid

with the rectangle's x interval $[A.i, B.i]$, and the y intervals similarly.

Wait, perhaps the square is covered by the rectangle if:
the square's x range $[x, x+1]$ overlaps with $[A.i, B.i]$, and similarly for y.

Wait, since the rectangle's x is from $A.i$ to $B.i$, including both endpoints. So for the square to be covered by the rectangle, the square's x must lie within the rectangle's x interval, but since the square is 1 unit, perhaps the condition is:
the rectangle's $A.i \leq (x+1)$ and rectangle's $B.i \geq x$, and similarly for y.

Because the square starts at x and ends at $x+1$. So the rectangle's $A.i$ must be $\leq x+1$ (so that the square's start is after the rectangle's start) and rectangle's $B.i$ must be $\geq x$ (so the square's end is before the rectangle's end). Wait, no, actually, the square is covered if any part of it is inside the rectangle. But since the rectangle includes its edges, perhaps the square's x must overlap with $[A.i, B.i]$. The square's x interval is $[x, x+1]$, so the overlap between that and $[A.i, B.i]$ must be non-empty. So the condition for x overlap between square and rectangle is: $\max(A.i, x) \leq \min(B.i, x+1)$

Similarly for y. Alternatively, another way: The square is covered by the rectangle if:
 $A.i \leq x+1$ and $B.i \geq x$ and $C.i \leq y+1$ and $D.i \geq y$.

Because: The square's x starts at x and ends at $x+1$. So for the square's x to overlap with the rectangle's x interval ($A.i$ to $B.i$), the rectangle's $A.i$ must be $\leq x+1$ (so that the rectangle's start is before the end of the square's x) and the rectangle's $B.i$ must be $\geq x$ (so the rectangle's end is after the square's start).

Similarly for the y direction. So for each square (x, y), we check all N rectangles. If any of them satisfies the above four inequalities, then the square is counted.

So the algorithm would be: Initialize total = 0

For each x in 0..99:
for each y in 0..99:
covered = False
for each rectangle in rectangles:
A,B,C,D = rectangle's parameters
if (A $\leq x+1$) and (B $\geq x$) and (C $\leq y+1$) and (D $\geq y$):
covered = True
break
if covered:
total += 1
Return total.

Wait, that's $O(N * 100 * 100)$ which is acceptable, because N is up to 100, so $100 * 100 * 100 = 1,000,000$ operations. That's manageable.

Yes, this approach should work.

Now, how to read the input.
Read N first, then read N lines, each with A,B,C,D.

So code steps:
Read N.
Then read N lines, each split into four integers. Store as a list of tuples or something.
Then loop over all x from 0 to 99, same for y. For each, check each rectangle until one covers the square.

Now, code: Sample Input 1:

```
3
0 5 1 3
1 4 0 5
2 5 2 4
```

Sample Output 1 is 20.

Let me see for x=0, y=0.

Check each rectangle:
First rectangle: A=0, B=5, C=1, D=3.
Check for x=0, y=0.
x is 0. The x part:
A (0) $\leq x+1$ (1) $\rightarrow$ yes. B (5) $\geq x$ (0) $\rightarrow$ yes.
y part:
C (1) $\leq y+1$ (1) $\rightarrow$ yes. D (3) $\geq y$ (0) $\rightarrow$ yes. So 1 ≤ 1 is okay. So the condition for y is C $\leq y+1 \rightarrow 1 \leq 1 \rightarrow$ yes. So the condition is true.

Wait, so the first rectangle's C is 1, so the y must be in the square's y interval from 0 to 1 (since y is 0, so the square is y=0 to 1).

The rectangle's y covers from 1 to 3. So the overlap between [0,1] and [1,3] is exactly at 1. Since the rectangle includes the y=1, but the square's y interval is up to 1 (since y is 0, so the square is [0,1] in y? Wait, the square is x from x to x+1, and y from y to y+1. Wait, the square is a unit square from (x, y) to (x+1, y+1). So the entire square's x is between x and x+1 (including both ends?), but the rectangle's x is up to B.i.

Wait the square is covered by the rectangle if the rectangle's x interval includes at least a point in the square's x. Since the rectangle is a closed interval, the square's x interval $[x, x+1]$ must overlap with $[A.i, B.i]$. The overlap condition is:
 $A.i \leq x+1$ and $B.i \geq x$.

Yes, exactly. So in the first rectangle for x=0, y=0:
The x condition holds (0 ≤ 1 and 5 ≥ 0). The y condition: the rectangle's C is 1. So the y direction: C $\leq y+1$ (1) must be $\leq y+1$ (0+1=1). So 1 $\leq 1 \rightarrow$ yes. And D $\geq y$ (3) must be $\geq y$ (0) $\rightarrow$ yes. So this square is covered by first rectangle. So it counts.

Figure 7: Code Generation Task Livecodebench Example. Wrong Trajectory with High Centroid

So the square (0,0) is covered. That's part of the total.
 So the code should handle that.
 Another test case: sample input 2 where two rectangles exactly overlap. The area is $100 \times 100 = 10000$, which the code would count each square once. So that works.
 Third sample input: 3 rectangles. The first is a 1x1 square at (0,0) to (1,1?), no. Wait the first rectangle is A=0, B=1, C=0, D=1. So it's a 1x1 square. The second is from 0 to 3 in x (so 3 units) and 0 to 5 (5 units) in y. The third is from 5 to 10 in x (5 units) and 0 to 10 (10 units). So the total area is 3×5 (the second rectangle) plus the first's 1, plus the third's 5×10 . But overlapping areas are subtracted. Wait, but how does the code count?
 Wait the third sample output is 65. Let me think:
 The second rectangle covers 0-3 in x (so 3 units) and 0-5 (5) → area 15. But when combined with the first (1x1), it's 15, but the first is already part of the second. So the second contributes $3 \times 5 = 15$, first adds nothing. The third is 5×10 (5 in x: $10 - 5 = 5$? Wait 5 to 10 in x: B.i is 10, so x ranges up to 10. So x length is $10 - 5 = 5$, y is 10, so area 50. Then total $15 + 50 = 65$. So that's correct. So code would count each square in the 0-3 x, 0-5 y (area $3 \times 5 = 15$), the first square is part of that. The third's area is $5 \times 10 = 50$, so total 65. So yes, the code should compute that.
 So the code approach should work.
 Now, code:
 Read N, then N lines. Then for each square (x from 0 to 99, y same), check each rectangle.
 In Python code:
 Read N, then read N lines, each split into A,B,C,D. Store in a list.
 Then:

```
total = 0
for x in range(100):
    for y in range(100):
        covered = False
        for (A,B,C,D) in rectangles:
            if (A ≤ x+1 and B ≤ x) and (C ≤ y+1 and D ≤ y):
                covered = True
                break
        if covered:
            total += 1
print(total)
```

 Wait, that's the plan. Let's test with sample input 1.
 Sample input 1:
 3
 0 5 1 3 → first rectangle.
 So for x ranges from 0 to 4 (since B.i is 5, so x can be up to 4? Because $x+1$ has to be $≤ A.i$? Hmm, but the code checks all x from 0 to 99. Let's see.
 Wait in sample input 1's first rectangle, A=0, B=5. So x can be up to 4 (since $x+1$ is 5, which is $≤ B=5$. So the x can be 0-4, but 5 is allowed?
 Wait the rectangle's x is up to B.i (5). So any square with $x=5$ would have $x+1=6$? No, wait if x is 5, then the square is from 5 to 6 in x, but the rectangle's B is 5, so the square's x from 5 to 6 would not be covered by that rectangle's x (since B.i is 5, so the rectangle's x can't exceed 5). So $x=5$, the rectangle's B is 5, so $B ≤ x$ ($5 ≤ 5$) → yes. and $A (0) ≤ x+1 (6)$ → yes. So the x condition holds. So that square's x interval (5-6) would overlap with rectangle's x up to 5. So the overlap is at $x=5$. So the square at $x=5$ would be covered if the y also meets.
 Wait for the first rectangle, in $x=5$'s square, the x is 5, so $x+1$ is 6. The rectangle's B is 5, so $B (5) ≤ x (5)$ → yes. The $A (0) ≤ x+1 (6)$ → yes. So x is okay. So the square at $x=5$ is covered by the first rectangle only if the y coordinates also fit. But first rectangle's C is 1 and D is 3. So for $y=2$ (for example), y's condition: $C (1) ≤ y+1 (3)$ → yes, and $D (3) ≤ y (2)$ → yes. So yes, that square would be covered.
 Wait the first rectangle's B.i is 5, so x can be up to 5, since x can be 5 (since the square's x is from 5 to 6, but the rectangle's B is 5, so the square's x from 5 to 6 has $x=5$ covered by the rectangle. So the square at $x=5$ is indeed covered by the first rectangle as long as the y is between 1 and 3.
 Hmm. So the code correctly includes that.
 Thus, the code is correct.
 Now, coding this in Python.
 The code structure:
 Read N, then read N lines.

```
rectangles = []
n = int(input())
for _ in range(n):
    a, b, c, d = map(int, input().split())
    rectangles.append( (a,b,c,d) )
```

 Then the loops as above.
 Yes.
 So the code should be as follows.
 Testing the sample input 1, let's see:
 Sample input 1's three rectangles:
 First: 0-5 in x, 1-3 in y.

Figure 7: Code Generation Task Livecodebench Example. Wrong Trajectory with High Centroid

Second: 1-4 in x (so x from 1 to 4, so x can be 1 to 3 (since 4's square would be 4-5, but the rectangle's B is 4, so $B_i = x$ ($4 \leq 4$? yes. So $x=4$'s square is included if y meets.)
 The second's y is 0-5, so covers down to $y=0$.
 Third rectangle: 2-5 in x, 2-4 in y.
 The total area in the sample is 20. Let's see.
 The code would count all squares covered by any of the three.
 I think that should work.
 Thus, the code should be correct.

```
n = int(input())
rectangles = []
for _ in range(n):
    a, b, c, d = map(int, input().split())
    rectangles.append((a, b, c, d))

total = 0
for x in range(100):
    for y in range(100):
        covered = False
        for (a, b, c, d) in rectangles:
            if (a <= x + 1 and b >= x) and (c <= y + 1 and d >= y):
                covered = True
                break
        if covered:
            total += 1
print(total)
```

Correct Trajectory (centroid=0.4250)

Okay, I need to find the area covered by at least one of the rectangles. Hmm, how do I approach this?

The problem is to compute the union area of multiple rectangles. The straightforward way is to check each small unit square (like 1x1) and see if it's covered by any rectangle, then count those squares. Since the coordinates are up to 100, maybe considering each integer coordinate point as the grid makes sense. Because the maximum possible coordinate is 100, the grid can be up to 100x100. So the total possible squares are $100 \times 100 = 10,000$, which is manageable even for $N=100$.

Wait, but the rectangles can have edges not aligned with integer coordinates? No, the input constraints say all values are integers. So each rectangle's edges are on integer coordinates. Therefore, each rectangle's area is made of unit squares. So for each unit square in the grid from (x,y) where x and y are integers from 0 to 100, we can check whether the square is covered by any rectangle.

Wait, actually, each rectangle's boundaries are $A_i \leq x \leq B_i$ and similarly for y. Since A_i and B_i are integers, the rectangle from A_i to B_i in x includes all x coordinates from A_i (inclusive) up to but not including B_i ? Wait no, wait, the problem says each rectangle covers exactly the region $A_i \leq x \leq B_i$ and $C_i \leq y \leq D_i$. Wait, but if A_i is 0 and B_i is 5, then that's 5 units in x-direction (since 0 to 5 is 5 units). Wait, but in terms of unit squares, perhaps each unit square is $[x, x+1)$ in x-direction and same for y. Hmm, maybe I should think of each square as the unit squares with integer coordinates as their corners. For example, the square with lower-left corner (i,j) and upper-right (i+1, j+1). So each rectangle spans from x from A_i to B_i , so the x coordinates covered are from A_i up to B_i , which is a length of ($B_i - A_i$). So the rectangle covers all squares where the square's x-coordinate starts at some integer between A_i and $B_i - 1$, and similarly for y.

Wait maybe the correct way is to represent each square as a 1x1 cell, where the square's x is from x to x+1, and similarly for y. So each rectangle covers all squares (x, y) such that $A_i \leq x < B_i$ and $C_i \leq y < D_i$. Wait, but the rectangle's B_i is exclusive? Because in the problem statement, the region is $A_i \leq x \leq B_i$. Wait, no, the problem says the rectangle includes the points up to B_i and D_i . Wait, that's a bit confusing. Let me re-read the problem.

The problem says: the i-th sheet covers exactly the region satisfying $A_i \leq x \leq B_i$ and $C_i \leq y \leq D_i$. So the rectangle is inclusive of both endpoints. So the rectangle is a closed rectangle. However, when considering unit squares, how do we discretize this?

Alternatively, since all coordinates are integers, perhaps the unit squares can be considered as the integer grid points. Let's think of the coordinate plane as divided into unit squares with corners at integer coordinates. For example, the square with lower-left corner (i,j) and upper-right (i+1, j+1). So this square is covered by a rectangle if the rectangle's $A_i \leq i$ and $B_i \geq i+1$ (since x must be at least A_i but less than B_i ?), or wait maybe I need to think in terms of inclusion.

Wait the rectangle is from $x=A_i$ to $x=B_i$, so the square's x ranges from i to i+1 must lie within $[A_i, B_i]$. So the square is entirely within the rectangle if:

$A_i \leq i$ and $i+1 \leq B_i$, and similarly for the y-direction.

Wait maybe the correct way is to check each unit square (i,j) (the square from $x=i$ to $x=i+1$, $y=j$ to $y=j+1$) and see whether any rectangle includes this square.

So for each square (i,j), check if there exists at least one rectangle where $A_i \leq i$, $B_i \geq i+1$ (since the square's x starts at i and goes up to i+1, so for the square to be inside the rectangle's x range, $A_i \leq i$ and $B_i \geq i+1$? Or wait maybe: The square is from $x = i$ to $x = i+1$, so for the rectangle's x interval $[A, B]$, the square is covered if:

Figure 8: Code Generation Task Livecodebench Example. Correct Trajectory with Low Centroid

$i \leq A$ and $i+1 \leq B$? No, that would mean the entire square is within the rectangle's x . Wait, but the rectangle's x must include the square's x from i to $i+1$. So the square's x range $[i, i+1)$ must be within $[A, B]$. So $A \leq i$ and $i+1 \leq B$. So the square is entirely inside.

Alternatively, perhaps the rectangle's x is $[A, B]$, so any point (x, y) in the rectangle has x between A and B (inclusive). So the square $[i, i+1) \times [j, j+1)$ is entirely covered by the rectangle if $i \leq A$ and $i+1 \leq B$, and similarly $j \leq C$ and $j+1 \leq D$.

Wait but the rectangle's x includes up to B . So the square's x can start at A and up to $B-1$. Because if B is the upper limit, then the square starting at B would have x from B to $B+1$, which is beyond. So the squares that lie within the rectangle's x are those with i in $[A, B-1]$. Because for example, if a rectangle is from 0 to 5, then the squares in x direction are 0,1,2,3,4 (each 1 unit) so 5 squares. $5-0=5$ units, so that's correct.

So for each square (i, j) , the x ranges from i to $i+1$, so the square is covered by rectangle k if:

$A \leq i$ and $B \leq i+1$ (since B must be greater than i to include at least part of the square's x -coordinate up to $i+1$). Wait, but the rectangle includes B , so if B is exactly i , then the square's x is from i to $i+1$, but B is i , so the rectangle's x goes up to i . So the square's x starts at i , so the rectangle includes $x=i$ but not $x=i+1$. So the square's x is from i to $i+1$, so the part from i to B would be covered, but only if $i \leq B$. Wait, perhaps the rectangle covers the square in the x -direction if the square's x interval overlaps with the rectangle's x interval.

Wait perhaps I'm overcomplicating. Let me think of each square (i, j) as being a unit square whose lower-left corner is (i, j) . Then the square is entirely contained in the rectangle $[A, B] \times [C, D]$ if:

$A \leq i$ and $B \geq i+1$ (since the square's x goes to $i+1$, so B must be $\geq i+1$). Similarly, $C \leq j$ and $D \geq j+1$.

Wait, maybe not. Let me think of rectangle from $A=0, B=5$. Then the squares from $i=0$ to 4 (since 0 to 4 gives 5 squares, each 1 unit). So for a square $i=4$: its x goes from 4 to 5. Since B is 5, the rectangle includes up to 5, so the square's x up to 5 is included. So yes, that square is covered. So the condition is $A \leq i$ and $B \geq i+1$? Wait, $4 \leq 5$ is true, so yes. So yes, that works.

Wait the square's x interval is $[i, i+1)$. The rectangle's x interval is $[A, B]$. So the square is entirely contained in the rectangle's x if $A \leq i$ and $i+1 \leq B$. Because then the entire interval $[i, i+1)$ is within $[A, B]$. But that's a stricter condition. Wait, maybe the square just needs to have any overlap with the rectangle. Wait, no. Because we need to check whether the square is covered by any rectangle. The square is covered if any part of it is covered by any rectangle.

Wait wait, I think I made a mistake here. The problem is to find the area covered by one or more sheets. So the entire region that is covered by any rectangle. So each square (unit square) is considered covered if it has any overlap with any rectangle. Because even if part of it is covered, the entire square's area is part of the total. Wait no, because the unit squares are 1×1 , but perhaps the problem counts the actual area, so if a square is partially covered, that area contributes. But the problem says the output is an integer. Wait the problem says "It can be proved that S is an integer under the constraints." So perhaps the rectangles are such that their union's area is integer. But how?

Wait, all the rectangles have integer coordinates, so their edges are on integer coordinates. Therefore, the overlapping regions between any two rectangles will also have edges on integer coordinates, so the union area is made up of integer unit squares, either fully or not. Wait, no, perhaps overlapping can split areas into fractions, but the problem states that S is an integer. So perhaps when all the rectangles are axis-aligned with integer coordinates, the union's area can be computed by checking for each unit square (i, j) whether it is entirely covered by at least one rectangle. Wait, no. For example, two overlapping rectangles: one from $(0,0)$ to $(2,2)$, and another from $(1,1)$ to $(3,3)$. The square at $(0,0)$ is entirely covered by first. The square $(1,1)$ is covered by both. But the square $(2,2)$ is covered by the second rectangle. But what about square $(1,0)$? It is partially covered by the first rectangle. Wait, but in this case, the square $(1,0)$ is from $x=1$ to 2, $y=0$ to 1. The first rectangle covers up to $x=2$, so that square's x is 1 to 2, so yes, that square is entirely covered by the first rectangle? Wait the first rectangle is $A=0, B=2$. So the x interval for the square ($i=1$) is $[1,2]$, which is within $[0,2]$. So yes, the entire square is covered. So in this case, maybe all overlapping is on the grid lines, so every unit square is either entirely inside or entirely outside.

Wait that makes sense! Because all rectangles have edges on integer coordinates, so any rectangle's edges are aligned with the grid lines. So any rectangle's region is a collection of unit squares. So the union of all rectangles is the union of their unit squares, so each square is either entirely in or entirely out. Therefore, the problem reduces to counting how many unit squares are covered by at least one rectangle. Because each such square contributes exactly 1 to the area, so the total area is the count of such squares.

Ah, that's a key insight. So the problem simplifies to counting the number of unit squares (i, j) where there exists at least one rectangle that covers the square (i, j) .

Therefore, the approach is to iterate over each unit square (i, j) from 0 to 99, since B can be up to 100, so the maximum x coordinate for the square is 99 (since $99+1=100$). Similarly for y .

So the plan is:

1. Iterate over all possible x and y coordinates from 0 to 99 (since each square is 1×1 , so there are 100 in each direction).
2. For each square (i, j) (representing the square with lower-left corner (i, j)), check if it is covered by any of the rectangles.
3. If yes, count it.
4. The total count is the answer.

Now, how to check if square (i, j) is covered by rectangle k ?

The rectangle k has $A_k \leq x \leq B_k$ and $C_k \leq y \leq D_k$.

The square's x ranges from i to $i+1$. So the square is entirely within the rectangle's x if:

$i \leq A_k$ and $i+1 \leq B_k$?

Wait, the rectangle's x must include the entire square's x range $[i, i+1)$. So the rectangle's A_k must be $\leq i$ (since the square starts at i , so A_k must be $\leq i$), and the rectangle's B_k must be $\geq i+1$ (since the square ends at $i+1$, so B_k must be $\geq i+1$). Wait, no. Because the rectangle's B_k is the upper bound. So the square's x is covered by the rectangle if the square's x is within $[A_k, B_k]$. So the rectangle's A_k must be $\leq i$, and B_k must be $\geq i+1$? Because the square's

Figure 8: Code Generation Task Livecodebench Example. Correct Trajectory with Low Centroid

x starts at i (so $i = A.k$) and ends at $i+1$ (so $j = B.k$). Wait no, because the rectangle's $B.k$ is the upper limit. So the square's x can be from i to $i+1$, so the square's x must lie entirely within the rectangle's x interval. So: i (the start of the square's x) must be $i = A.k$, and the end of the square's x ($i+1$) must be $j = B.k$. Because the rectangle's x includes up to $B.k$, so the square's x can't exceed that. Wait, the rectangle's x covers up to $B.k$. So for the square to be entirely within, the square's x must be between $A.k$ and $B.k$ (inclusive). But since the square's x is from i to $i+1$, the entire interval must be within $[A.k, B.k]$. So: $i = A.k$ and $(i+1) = B.k$. Because the square's x starts at i and ends at $i+1$. So the start must be $i = A.k$ and the end must be $j = B.k$. Wait that's correct. For example, if rectangle has $A=0, B=5$. Then the square at $i=0$ is okay ($0 = 0$ and $0+1=1 = 5 \rightarrow$ yes). The square at $i=4$ would be $4 = 0$ and $5 = 5 \rightarrow$ yes. The square at $i=5$: $5 = 0$, but $5+1=6$ which is $5 \rightarrow$ so not covered. So the squares from $i=0$ to 4 are included, which is 5 units. That's correct. So the condition for the x direction is: $A.k = i$ and $B.k = (i+1)$. Similarly for the y-direction: the square's y is from j to $j+1$, so: $C.k = j$ and $D.k = (j+1)$. So for rectangle k to cover square (i,j), both x and y conditions must hold. Therefore, for each square (i,j), we check if there exists any rectangle where: $A.k = i$ and $B.k = (i+1)$ and $C.k = j$ and $D.k = (j+1)$. If yes, then count it. So the steps in code would be: Read N. Read N lines, each with A,B,C,D for a rectangle. Store them as a list of tuples. Initialize a counter to 0. Loop over i from 0 to 99 (since the maximum $B.i$ is 100, so i can be up to 99): Loop over j from 0 to 99: flag = False for each rectangle in rectangles: A, B, C, D = rectangle if $(A == i \text{ and } B == (i+1))$ and $(C == j \text{ and } D == (j+1))$: flag = True break if flag: counter +=1 Return counter. That should work. Let me test with the sample inputs. Sample Input 1: 3 0 5 1 3 $\rightarrow$ rect 1: $A=0, B=5, C=1, D=3$. 1 4 0 5 $\rightarrow$ rect2: $A=1, B=4, C=0, D=5$. 2 5 2 4 $\rightarrow$ rect3: $A=2, B=5, C=2, D=4$. Let's see for some squares. Take $i=0, j=1$ (the square from (0,1) to (1,2)). Check rect1: $A=0, B=5, C=1, D=3, i=0 \rightarrow$ yes. $C=1, D=3, j=1 \rightarrow$ yes. So yes. So this square is counted. Another example: $i=0, j=0$. rect1's C is 1 $\neq 0$, so no. rect2's C is 0 $= 0$, D is 5 $\neq 1$? $j=0 \rightarrow D=5, j=1$ (since $j+1$ is 1). So yes. So rect2 covers (0,0). So counted. Sample Input 1's output is 20. Let's see code. Another test case: Sample Input 2: 2 0 100 0 100 $\rightarrow$ rect1 covers all squares (since for any i, $0 = i$ and $100 = i+1$ (since i can be up to 99, so $99+1=100$ which is equal to B. $100 = 100 \rightarrow$ yes. So yes. So all squares are covered. So total is $100*100 = 10000$, which matches sample 2's output. Sample3: 3 0 1 0 1 $\rightarrow$ rect1 covers $i=0$ (since $0 = 0$ and $1 = 1$ ($i+1=1$)). So yes. j from 0 to 0 ($j+1$ is 1). So the square (0,0) is covered. So that's 1 square. The second rectangle is 0-3 in x, 0-5 in y. So it covers i from 0 to 2 (since B is 3, so $i+1 = 3 \rightarrow i$ up to 2). j from 0 to 4 (since D is 5 $\rightarrow j$ can be 0-4 ($j+1=5$)). So the second rectangle covers $3*5=15$ squares, but overlapping with first. Third rectangle: 5 to 10 ($A=5, B=10 \rightarrow$ so i can be up to 9 (since $9+1=10$ is $B=10 \rightarrow B = i+1 \rightarrow 10 = 10 \rightarrow$ yes). So x ranges 5-9 (5 squares). y from 0 to 10 (since D is 10 $\rightarrow j+1 = 10 \rightarrow j$ up to 9. So y from 0-9 $\rightarrow 10$ squares. So $5*10=50$ squares. So total? The total area would be the union of all three. Let's see: First, the first two rectangles: First covers 1 square. Second covers $3*5=15$, but overlapping the first's square. So total for first two is 15 (since first is part of the second). Then third adds 50. But maybe there's some overlaps between second and third? No, because third's x starts at 5, and second's x is up to 2 ($i=2, i+1=3$). So their x ranges don't overlap. So total is $15 + 50 = 65$, which matches sample3's output. So the code should work.

Figure 8: Code Generation Task Livecodebench Example. Correct Trajectory with Low Centroid

Now coding this in Python.
The code structure:
Read N.
Then read N lines. For each line, split into A,B,C,D. We can store them as a list of tuples.
Then loop over i from 0 to 99 (inclusive), same for j.
For each (i,j), check any rectangle in the list that satisfies the conditions.
So in code:
Read N:
import sys
n = int(sys.stdin.readline())
rects = []
for _ in range(n):
a, b, c, d = map(int, sys.stdin.readline().split())
rects.append((a,b,c,d))
counter = 0
for i in range(100):
for j in range(100):
covered = False
for (a,b,c,d) in rects:
if (a == i and b == i+1) and (c == j and d == j+1):
covered = True
break
if covered:
counter += 1
print(counter)
Yes, that should work.
Testing with sample input1:
Sample1:
3
0 5 1 3 → rect1: A=0, B=5 → x: i can be 0,1,2,3,4 (since i+1 == 5 → up to 4).
C=1, D=3 → j must be between 1 and 2 (since j+1 == 3 → j == 2).
So for rect1, the squares are x from 0-4, y 1-2 → 5 * 2 = 10 squares.
rect2: A=1, B=4 → i up to 3 (i+1 == 4 → i == 3). So x 1-3 (4 numbers?), wait 1 to 3 → 3, 3-1+1=3? Wait i can be 1,2,3.
So 3 values. B=4, so 1,2,3 → 3 units.
C=0, D=5 → j can be 0 to 4 (since j+1 == 5 → j up to 4). So y from 0-4 (5 squares). So rect2's area is 3*5=15.
rect3: A=2, B=5 → i can be 2,3,4 (since i+1 == 5 → i up to 4). So 2,3,4 → 3 values. B=5 is okay.
C=2, D=4 → j must be 2,3 (since j+1 == 4 → j up to 3). So y 2-3 → 2 squares. So rect3 has 3*2=6 squares.
Total individual areas: 10 + 15 + 6 = 31, but overlapping regions will be subtracted.
But the union area is 20. So the code should compute it correctly by counting all squares covered by any of them.
Hmm, perhaps the code is correct. So the code should handle this.
Thus, the code should work.

```

import sys

def main():
    n = int(sys.stdin.readline())
    rectangles = []
    for _ in range(n):
        a, b, c, d = map(int, sys.stdin.readline().split())
        rectangles.append((a, b, c, d))

    count = 0
    for i in range(100):
        for j in range(100):
            for (a, b, c, d) in rectangles:
                if a <= i and b >= i + 1 and c <= j and d >= j + 1:
                    count += 1
                    break # Break early once covered by any rectangle
    print(count)

if __name__ == "__main__":
    main()

```

Figure 8: Code Generation Task Livecodebench Example. Correct Trajectory with Low Centroid

Logical Reasoning Synlogic

Problem:

Below are the rules for a specific cipher.

Cipher Rule:

Encryption Rules:

- Input: - Plaintext: Uppercase letters string without punctuation and spaces. - Output: - Ciphertext: Uppercase letters string without punctuation and spaces. - Cipher Type: Caesar Cipher - Shift Key: 7 - Encryption Steps: - For each letter in the plaintext, shift it forward by 7 positions in the alphabet. - If shifting goes beyond 'Z', wrap around to the beginning of the alphabet ('A'). - For example, 'A' becomes 'H', 'B' becomes 'I', and so on.

Decryption Rules:

- Input: - Ciphertext: Uppercase letters string without punctuation and spaces. - Output: - Plaintext: Uppercase letters string without punctuation and spaces. - Cipher Type: Caesar Cipher - Shift Key: 7 - Decryption Steps: - For each letter in the ciphertext, shift it backward by 7 positions in the alphabet. - If shifting goes before 'A', wrap around to the end of the alphabet ('Z'). - For example, 'H' becomes 'A', 'I' becomes 'B', and so on.

Question: Plaintext: "HTARWOQI"

Please provide the encrypted answer wrapped in Markdown Python code block. The answer should be a string encapsulated in double square brackets. For example, the format should be:

```
python
[["${ENCRYPTED_A$NSWER}"]]
```

Wrong Trajectory (centroid=0.6017)

We are given the plaintext "HTARWOQI" and need to encrypt it using a Caesar cipher with a shift of 7 positions forward. Let's process each letter one by one:

The alphabet: A(0), B(1), C(2), D(3), E(4), F(5), G(6), H(7), I(8), J(9), K(10), L(11), M(12), N(13), O(14), P(15), Q(16), R(17), S(18), T(19), U(20), V(21), W(22), X(23), Y(24), Z(25)

Shift for each letter: (original position + 7) mod 26

H → position 7 7 + 7 = 14 → O
T → position 19 19 + 7 = 26 → 26 mod 26 = 0 → A
A → position 0 0 + 7 = 7 → H
R → position 17 17 + 7 = 24 → Y
W → position 22 22 + 7 = 29 → 29 mod 26 = 3 → D
O → position 14 14 + 7 = 21 → V
Q → position 16 16 + 7 = 23 → X
I → position 8 8 + 7 = 15 → P

So, the encrypted string is "OHA YDV XP" but without spaces: "OHAYDVXP" But let's write it without space: O H A Y D V X P =_i "OHAYDVXP"? Wait, but let's just concatenate: O H A Y D V X P =_i "OHAYDVXP" should be "OHAYDVXP", but since all are uppercase and no spaces, it should be "OHAYDVXP" – but since Y is uppercase, it's just "OHAYDVXP" is not correct, should be uppercase only, so O H A Y D V X P =_i O H A Y D V X P = O H A Y D V X P → actually, no lowercase. All letters should remain uppercase. Let's write explicitly:

H → O T → A A → H R → Y W → D O → V Q → X I → P

So the ciphertext is "OHAYDVXP"? Wait, no, all letters are uppercase: O H A Y D V X P → "O H A Y D V X P" (but no spaces) → "OHAYDVXP"? No, should be "OHAYDVXP" only if 'y' is lowercase, but it's actually 'Y', so it should be "OHAYDVXP".

Let's confirm:

H → O T → A A → H R → Y W → D O → V Q → X I → P

So the encrypted text is: O H A Y D V X P → "OHAYDVXP"

But let's write it as a string: "OHAYDVXP"

Now, to present per the instructions, wrap the answer in a Python code block as a list of a single string, in double square brackets.

So, the final encrypted answer is "OHAYDVXP".

```
<answer>

[["OHAYDVXP"]]

</answer>
```

Figure 9: Logic Reasoning Task Synlogic Example. Wrong Trajectory with High Centroid

Correct Trajectory (centroid=0.3063)

We are given the plaintext "HTARWOQI" and asked to encrypt it using a Caesar Cipher with a shift key of 7. Step-by-step encryption process: 1. For each letter in "HTARWOQI", shift it forward by 7 positions in the alphabet, wrapping around if necessary. 2. We will convert each letter to its corresponding position (A=0, B=1, ..., Z=25), apply the shift, then convert back to a letter.

Let's process each letter:

- H: Position = 7. $7 + 7 = 14$. Letter at position 14 is 'O'. - T: Position = 19. $19 + 7 = 26$. $26 \bmod 26 = 0$. Letter at position 0 is 'A'. - A: Position = 0. $0 + 7 = 7$. Letter at position 7 is 'H'. - R: Position = 17. $17 + 7 = 24$. Letter at position 24 is 'Y'. - W: Position = 22. $22 + 7 = 29$. $29 \bmod 26 = 3$. Letter at position 3 is 'D'. - O: Position = 14. $14 + 7 = 21$. Letter at position 21 is 'V'. - Q: Position = 16. $16 + 7 = 23$. Letter at position 23 is 'X'. - I: Position = 8. $8 + 7 = 15$. Letter at position 15 is 'P'.

Putting them together: O A H Y D V X P

Therefore, the ciphertext is "OAHYDVXP".

<answer>

[["OAHYDVXP"]]

</answer>

Figure 10: Logic Reasoning Task Synlogic Example. Correct Trajectory with Low Centroid

Mathematical Reasoning (AIME 2025)

Problem:

Sixteen chairs are arranged in a row. Eight people each select a chair in which to sit so that no person sits next to two other people. Let N be the number of subsets of the 16 chairs that could be selected. Find the remainder when N is divided by 1000.

Wrong Trajectory (centroid=0.4229)

Let's break down the problem step by step.

We have 16 chairs in a row. We want to select 8 chairs such that when people sit on them, **no person sits next to two other people**. That is, for any occupied chair, at most one of its immediate neighbors is also occupied.

Let's clarify what "no person sits next to two other people" means. For any chosen (occupied) chair, it cannot be the case that both its left and right neighbors are also occupied. In other words, there cannot be any block of three consecutive occupied chairs.

But let's be precise: the condition is that **no person is sitting next to two other people**. For a person in the middle, that would mean the chairs immediately to the left and right are both occupied. For a person at the end, only one neighbor exists, so they can't be next to two people anyway. So the forbidden configuration is any three consecutive chairs all occupied.

So, the problem reduces to:

How many ways can we choose 8 chairs out of 16 in a row, such that there is **no block of 3 consecutive occupied chairs**?

Let $a_n(k)$ be the number of ways to choose k chairs out of n in a row such that no three chosen chairs are consecutive.

We are to compute $a_{16}(8) \bmod 1000$.

Step 1: Recurrence Relation

Let's derive a recurrence for $a_n(k)$.

Suppose we look at the first chair.

- **Case 1:** The first chair is **not** chosen. Then we are left with choosing k chairs from the remaining $n-1$ chairs, with the same restriction:

$$a_{n-1}(k)$$

- **Case 2:** The first chair **is** chosen. - The second chair can be either chosen or not. - But we cannot have three in a row chosen, so we must be careful.

Let's break it down further:

If the first chair is chosen, consider the status of the second chair:

- **Subcase 2a:** The second chair is **not** chosen. Then we have placed a chosen chair at position 1, then a gap at position 2. Now, we need to choose $k-1$ chairs from the remaining $n-2$ chairs (positions 3 to n), with the same restriction:

$$a_{n-2}(k-1)$$

- **Subcase 2b:** The second chair **is** chosen. - Now, we cannot have the third chair also chosen (that would make 3 in a row). - So the third chair must **not** be chosen. - So we have chairs 1 and 2 chosen, chair 3 not chosen. - Now, choose $k-2$ chairs from the remaining $n-3$ chairs (positions 4 to n), with the same restriction:

$$a_{n-3}(k-2)$$

Therefore, the recurrence is:

$$a_n(k) = a_{n-1}(k) + a_{n-2}(k-1) + a_{n-3}(k-2)$$

But wait, let's check that carefully.

Wait, in the above, when the first chair is chosen, we have two subcases:

- If the second is **not** chosen: Then after chair 1 (chosen), chair 2 (not chosen), we have $n-2$ chairs left and need to choose $k-1$ more. - If the second **is** chosen: Then chair 3 must **not** be chosen, so after chairs 1 and 2 (chosen), chair 3 (not chosen), we have $n-3$ chairs left and need to choose $k-2$ more.

But if the first chair is **not** chosen, we have $n-1$ chairs left and need to choose k chairs.

So yes, the correct recurrence is:

$$a_n(k) = a_{n-1}(k) + a_{n-2}(k-1) + a_{n-3}(k-2)$$

But let's verify with small n and k .

Alternatively, let's define another standard way.

Let $f(n, k)$ = number of ways to select k chairs from n in a row, with no three consecutive occupied.

Alternatively, we can model the selection as a binary string of length n with exactly k ones and no three consecutive ones.

Let's define: Let $b_n(k)$ = number of binary strings of length n with exactly k ones and no three consecutive ones.

Figure 11: Mathematical Reasoning Task AIME Example. Wrong Trajectory with High Centroid

Then, the recurrence for the total number of binary strings of length n with no three consecutive ones (regardless of number of ones) is:

$$s_n = s_{n-1} + s_{n-2} + s_{n-3}$$

with $s_0 = 1$, $s_1 = 2$, $s_2 = 4$.

But we want the number with exactly k ones.

Let's define $b_n(k)$ as above.

Alternatively, we can use generating functions or a combinatorial approach.

But let's try to find a recurrence for $b_n(k)$.

Let's consider the first chair:

- If the first chair is 0, then the rest is any valid string of length $n-1$ with k ones. - If the first chair is 1, the second can be:

- 0: Then the rest is any valid string of length $n-2$ with $k-1$ ones. - 1, then the third must be 0 (to avoid three in a row):

So first two are 1,1, third is 0, then rest is any valid string of length $n-3$ with $k-2$ ones.

But if after two 1's, the third is 1, that's forbidden. So:

So,

$$b_n(k) = b_{n-1}(k) + b_{n-2}(k-1) + b_{n-3}(k-2)$$

with the base cases: - If $k \leq 0$ or $n \leq 0$: 0 - If $k = 0$: Only one way (all zeros): $b_n(0) = 1$ - If $n \leq 2k - 1$: Impossible (not enough room for k ones with at least one zero between each pair of "11" blocks), but actually, the maximum number of ones with no three in a row is $\text{floor}((n+2)/2)$, but let's not worry about that for the recurrence.

Let's check for small values.

Test with $n = 3$, $k = 2$

We want the number of binary strings of length 3 with exactly 2 ones and no three in a row.

Possible strings: - 110 (valid) - 101 (valid) - 011 (valid) - 112, 121, etc., but only ones and zeros.

List all with two ones: - 110 - 101 - 011

Check which have three consecutive ones: None do. So all 3 are valid.

Now compute by recurrence:

$$b_3(2) = b_2(2) + b_1(1) + b_0(0)$$

First, compute base cases:

- $b_n(0) = 1$ for all n - For $k = 1$: - $b_n(1)$ = number of ways to put a single 1 in n positions = n - For $k = 2$: - Let's compute for small n

Let's build up step by step.

Build Table for Small n and k

Let's make a table for $b_n(k)$, for n from 0 to 5, k from 0 to n .

We'll use the recurrence:

$$b_n(k) = b_{n-1}(k) + b_{n-2}(k-1) + b_{n-3}(k-2)$$

with base cases: - If $k \leq 0$ or $n \leq 0$: 0 - If $k = 0$: $b_n(0) = 1$ - If $n \leq 2k - 1$: 0 (not enough space to place k ones without at least two zeros between some of them? Actually, let's see, but let's just compute for small n)

But let's fill out for n up to 5.

$n = 0$

- $b_0(0) = 1$

$n = 1$

- $b_1(0) = 1$ - $b_1(1) = 1$

$n = 2$

- $b_2(0) = 1$ - $b_2(1) = 2$ (10, 01) - $b_2(2) = 1$ (11)

Check forbidden? 111 is not possible for $n=2$.

Now using recurrence for $b_2(2)$:

$b_2(2) = b_1(2) + b_0(1) + b_{-1}(0) = 0 + 0 + 0 = 0$. But we know it should be 1.

Wait, our recurrence may need adjustment for small n .

Wait, let's think again. The recurrence: If first position is 0: then $b_{n-1}(k)$ If first is 1, second is 0: then $b_{n-2}(k-1)$ If first is 1, second is 1, third is 0: then $b_{n-3}(k-2)$

But for $n=2$, $k=2$: - First is 0: $b_1(2) = 0$ - First is 1, second is 0: $b_0(1) = 0$ - First is 1, second is 1: but then $n=2$, so we're done, and that's "11", which is allowed, so that should be counted as 1 way.

So the recurrence as written misses the case where the string ends with two 1s and there is no third position.

So perhaps the recurrence should be: Let's define it more carefully.

Let's define $c_n(k)$ = number of binary strings of length n with exactly k ones and no three consecutive ones. Let's use another approach: Let's define the recurrence based on the last few bits.

Let's define: Let's define $f(n, k)$ = number of valid strings of length n with exactly k ones and no three consecutive ones. Let's consider the ending of the string:

- The string ends with 0: Then the first $n-1$ bits form any valid string of length $n-1$ with k ones. - The string ends with 10: Then the first $n-2$ bits form any valid string of length $n-2$ with $k-1$ ones. - The string ends with 110: Then the first $n-3$ bits form any valid string of length $n-3$ with $k-2$ ones. But we cannot have 111.

But if we think in terms of the last run of ones, the run at the end can be of length 0, 1, or 2.

So: Let's define $f(n, k)$ as above.

Let's partition based on the length of the last run of ones:

Figure 11: Mathematical Reasoning Task AIME Example. Wrong Trajectory with High Centroid

- If the last bit is 0, then the previous $n-1$ bits form a valid string of length $n-1$ with k ones. - If the last bit is 1 and the one before is 0, then the previous $n-2$ bits form a valid string of length $n-2$ with $k-1$ ones. - If the last two bits are 11 and the one before is 0, then the previous $n-3$ bits form a valid string of length $n-3$ with $k-2$ ones.

So,

$$f(n, k) = f(n-1, k) + f(n-2, k-1) + f(n-3, k-2)$$

But for the cases where $n-2 \leq 0$ or $n-3 \leq 0$, those terms are zero unless $k-1$ or $k-2$ is negative, in which case also zero.

Let's check with $n=2$, $k=2$: $f(2,2) = f(1,2) + f(0,1) + f(-1,0)$ $f(1,2) = 0$ $f(0,1) = 0$ $f(-1,0) = 0$ But we know $f(2,2) = 1$ (the string "11")

So the recurrence doesn't capture the case where the string is exactly "11".

Ah! The problem is that when $n=2$ and $k=2$, the only possibility is "11", but according to the recurrence, we need to count the case where the last run is of length 2, but there is no bit before it. So we need to allow for runs of length up to 2 at the end, even if they're at the very start.

Alternatively, let's define a different recurrence. Let's define $g(n, k)$ = number of valid strings of length n with exactly k ones and no three consecutive ones. Let's define $h(n, r)$ = number of valid strings of length n ending with a run of r consecutive ones, where $r = 0, 1, 2$.

Then, - $h(n, 0)$: ends with 0 - $h(n, 1)$: ends with 10 - $h(n, 2)$: ends with 110

But actually, we can define: Let $h(n, r)$ = number of valid strings of length n , ending with a run of exactly r ones, $r = 0, 1, 2$.

Then, - To get a string of length n ending with 0 ones (i.e., ends with 0), we can append 0 to any valid string of length $n-1$: $h(n, 0) = h(n-1, 0) + h(n-1, 1) + h(n-1, 2)$ - To get a string ending with exactly one 1, we must append 1 to a string ending with 0: $h(n, 1) = h(n-1, 0)$ - To get a string ending with exactly two 1s, we must append 1 to a string ending with exactly one 1: $h(n, 2) = h(n-1, 1)$

We cannot have three 1s at the end.

Therefore, the total number of valid strings of length n with any number of ones is:

$$s_n = h(n, 0) + h(n, 1) + h(n, 2)$$

But we want the number with exactly k ones.

So let's define $H(n, k, r)$ = number of valid strings of length n with exactly k ones, ending with a run of exactly r ones ($r=0, 1, 2$).

Then the recurrence is: - $H(n, k, 0) = H(n-1, k, 0) + H(n-1, k, 1) + H(n-1, k, 2)$ [append 0] - $H(n, k, 1) = H(n-1, k-1, 0)$ [append 1 to a string ending with 0] - $H(n, k, 2) = H(n-1, k-1, 1)$ [append 1 to a string ending with 1]

With base cases: For $n=0$: - $H(0, 0, 0) = 1$ (empty string) - $H(0, k, r) = 0$ for $k \neq 0$

Let's build up a table for small n and k .

We want eventually $H(16, 8, 0) + H(16, 8, 1) + H(16, 8, 2)$

Let's make a table for n up to 5 and k up to, say, 3. Let's proceed step by step.

$n = 0$

- $H(0, 0, 0) = 1$ - $H(0, k, r) = 0$ for $k \neq 0$

$n = 1$

For $k=0$: - $H(1, 0, 0)$: append 0 to $H(0, 0, 0)$, $H(0, 0, 1)$, $H(0, 0, 2)$: only $H(0, 0, 0)=1$ So $H(1, 0, 0) = 1$ - $H(1, 0, 1) = H(0, -1, 0) = 0$ - $H(1, 0, 2) = H(0, -1, 1) = 0$

For $k=1$: - $H(1, 1, 0) = 0$ (can't end with 0 and have 1 one) - $H(1, 1, 1) = H(0, 0, 0) = 1$ - $H(1, 1, 2) = H(0, 0, 1) = 0$

So: $H(1, 0, 0) = 1$ $H(1, 1, 1) = 1$

$n = 2$

For $k=0$: - $H(2, 0, 0) = H(1, 0, 0) + H(1, 0, 1) + H(1, 0, 2) = 1 + 0 + 0 = 1$ - $H(2, 0, 1) = H(1, -1, 0) = 0$ - $H(2, 0, 2) = H(1, -1, 1) = 0$

For $k=1$: - $H(2, 1, 0) = H(1, 1, 0) + H(1, 1, 1) + H(1, 1, 2) = 0 + 1 + 0 = 1$ - $H(2, 1, 1) = H(1, 0, 0) = 1$ - $H(2, 1, 2) = H(1, 0, 1) = 0$

For $k=2$: - $H(2, 2, 0) = 0$ - $H(2, 2, 1) = H(1, 1, 0) = 0$ - $H(2, 2, 2) = H(1, 1, 1) = 1$

So total for $k=2$: $H(2, 2, 2) = 1$

Which matches: "11" is the only valid string of length 2 with two ones.

Let's check for $k=1$: total = $H(2, 1, 0) + H(2, 1, 1) + H(2, 1, 2) = 1 + 1 + 0 = 2$, which matches 10 and 01.

Perfect.

$n = 3$

We proceed similarly.

For $k=0$: - $H(3, 0, 0) = H(2, 0, 0) + H(2, 0, 1) + H(2, 0, 2) = 1 + 0 + 0 = 1$ - $H(3, 0, 1) = H(2, -1, 0) = 0$ - $H(3, 0, 2) = H(2, -1, 1) = 0$

For $k=1$: - $H(3, 1, 0) = H(2, 1, 0) + H(2, 1, 1) + H(2, 1, 2) = 1 + 1 + 0 = 2$ - $H(3, 1, 1) = H(2, 0, 0) = 1$ - $H(3, 1, 2) = H(2, 0, 1) = 0$

For $k=2$: - $H(3, 2, 0) = H(2, 2, 0) + H(2, 2, 1) + H(2, 2, 2) = 0 + 0 + 1 = 1$ - $H(3, 2, 1) = H(2, 1, 0) = 1$ - $H(3, 2, 2) = H(2, 1, 1) = 0$

For $k=3$: - $H(3, 3, 0) = 0$ - $H(3, 3, 1) = H(2, 2, 0) = 0$ - $H(3, 3, 2) = H(2, 2, 1) = 0$

But "111" is invalid, so $k=3$ is 0.

Now, check for $k=2$: Total = $H(3, 2, 0) + H(3, 2, 1) + H(3, 2, 2) = 1 + 1 + 1 = 3$ Which matches our earlier enumeration for $n=3$, $k=2$: 110, 101, 011

For $k=1$: total = $2 + 1 + 0 = 3$: 001, 010, 100

Figure 11: Mathematical Reasoning Task AIME Example. Wrong Trajectory with High Centroid

For k=0: 1 (000)
Great! So this recurrence works.
Now, let's write a small program or use dynamic programming to compute $H(16, 8, r)$ for $r=0,1,2$.
But since we're doing this by hand (or at least step by step), let's try to build up the table for n up to 16 and k up to 8.
But that's a lot. Let's see if we can write a table for each n , k up to 8.
Alternatively, let's define a 3-dimensional DP array $H[n][k][r]$, but we'll just keep rolling arrays for each n .
Let's make a table for n from 0 to 16, and for each n , k from 0 to 8, and $r=0,1,2$.
Let's proceed step by step, keeping for each n , an array $A[n][k] = H(n, k, 0) + H(n, k, 1) + H(n, k, 2)$, but to keep track for recurrence, we need all r .
Alternatively, let's define for each n and k , an array $[a, b, c] = [H(n, k, 0), H(n, k, 1), H(n, k, 2)]$
Let's proceed step by step, keeping only necessary k values (up to $k=8$) at each n .
Let's initialize for $n=0$:
 $n=0$: $k=0$: $[1, 0, 0]$ $k_i 0$: $[0, 0, 0]$
 $n=1$: $k=0$: $[1, 0, 0]$ $k=1$: $[0, 1, 0]$ $k_i 1$: $[0, 0, 0]$
 $n=2$: From recurrence:
For each k : $H(n, k, 0) = \text{sum over } r \text{ of } H(n-1, k, r)$ $H(n, k, 1) = H(n-1, k-1, 0)$ $H(n, k, 2) = H(n-1, k-1, 1)$
Let's compute for $n=2$
 $k=0$: $H(2, 0, 0) = H(1, 0, 0) + H(1, 0, 1) + H(1, 0, 2) = 1 + 0 + 0 = 1$ $H(2, 0, 1) = H(1, -1, 0) = 0$ $H(2, 0, 2) = H(1, -1, 1) = 0 =_i [1, 0, 0]$
 $k=1$: $H(2, 1, 0) = H(1, 1, 0) + H(1, 1, 1) + H(1, 1, 2) = 0 + 1 + 0 = 1$ $H(2, 1, 1) = H(1, 0, 0) = 1$ $H(2, 1, 2) = H(1, 0, 1) = 0 =_i [1, 1, 0]$
 $k=2$: $H(2, 2, 0) = H(1, 2, 0) + H(1, 2, 1) + H(1, 2, 2) = 0 + 0 + 0 = 0$ $H(2, 2, 1) = H(1, 1, 0) = 0$ $H(2, 2, 2) = H(1, 1, 1) = 1 =_i [0, 0, 1]$
 $k_i 2$: all zeros
 $n=2$ table: $k=0$: $[1, 0, 0]$ $k=1$: $[1, 1, 0]$ $k=2$: $[0, 0, 1]$
 $n=3$:
For each k :
 $k=0$: $H(3, 0, 0) = H(2, 0, 0) + H(2, 0, 1) + H(2, 0, 2) = 1 + 0 + 0 = 1$ $H(3, 0, 1) = H(2, -1, 0) = 0$ $H(3, 0, 2) = H(2, -1, 1) = 0 =_i [1, 0, 0]$
 $k=1$: $H(3, 1, 0) = H(2, 1, 0) + H(2, 1, 1) + H(2, 1, 2) = 1 + 1 + 0 = 2$ $H(3, 1, 1) = H(2, 0, 0) = 1$ $H(3, 1, 2) = H(2, 0, 1) = 0 =_i [2, 1, 0]$
 $k=2$: $H(3, 2, 0) = H(2, 2, 0) + H(2, 2, 1) + H(2, 2, 2) = 0 + 0 + 1 = 1$ $H(3, 2, 1) = H(2, 1, 0) = 1$ $H(3, 2, 2) = H(2, 1, 1) = 1 =_i [1, 1, 1]$
 $k=3$: $H(3, 3, 0) = 0$ $H(3, 3, 1) = H(2, 2, 0) = 0$ $H(3, 3, 2) = H(2, 2, 1) = 0 =_i [0, 0, 0]$
So $n=3$: $k=0$: $[1, 0, 0]$ $k=1$: $[2, 1, 0]$ $k=2$: $[1, 1, 1]$ $k=3$: $[0, 0, 0]$
Let's keep going up to $n=16$, but only track k up to 8. Let's keep a table for each n , $k=0..8$, $[a, b, c]$
Let's tabulate step by step, using previous n 's table.
Let's define for each n , an array $A[n][k] = [a, b, c] = [H(n, k, 0), H(n, k, 1), H(n, k, 2)]$
We can write a small table for $n=0..5$, then look for a pattern or continue up to $n=16$. But to save space, let's proceed iteratively, keeping only the previous two rows.
Let's denote previous row as $\text{prev}[k] = [a, b, c]$ for $k=0..8$
Let's start with $n=0$:
 $\text{prev}[0] = [1, 0, 0]$ $\text{prev}[1..8] = [0, 0, 0]$
Now $n=1$:
for $k=0$: $\text{curr}[0][0] = \text{prev}[0][0] + \text{prev}[0][1] + \text{prev}[0][2] = 1 + 0 + 0 = 1$ $\text{curr}[0][1] = \text{prev}[-1][0] = 0$ $\text{curr}[0][2] = \text{prev}[-1][1] = 0 =_i [1, 0, 0]$
for $k=1$: $\text{curr}[1][0] = \text{prev}[1][0] + \text{prev}[1][1] + \text{prev}[1][2] = 0$ $\text{curr}[1][1] = \text{prev}[0][0] = 1$ $\text{curr}[1][2] = \text{prev}[0][1] = 0 =_i [0, 1, 0]$
for $k_i 2$: all zeros
So $n=1$: $k=0$: $[1, 0, 0]$ $k=1$: $[0, 1, 0]$ $k_i 1$: $[0, 0, 0]$
 $n=2$:
for $k=0$: $\text{curr}[0] = [\text{prev}[0][0] + \text{prev}[0][1] + \text{prev}[0][2], 0, 0] = [1 + 0 + 0, 0, 0] = [1, 0, 0]$
 $k=1$: $\text{curr}[1][0] = \text{prev}[1][0] + \text{prev}[1][1] + \text{prev}[1][2] = 0 + 1 + 0 = 1$ $\text{curr}[1][1] = \text{prev}[0][0] = 1$ $\text{curr}[1][2] = \text{prev}[0][1] = 0 =_i [1, 1, 0]$
 $k=2$: $\text{curr}[2][0] = \text{prev}[2][0] + \text{prev}[2][1] + \text{prev}[2][2] = 0$ $\text{curr}[2][1] = \text{prev}[1][0] = 0$ $\text{curr}[2][2] = \text{prev}[1][1] = 1 =_i [0, 0, 1]$
 $n=2$ as before.
Now $n=3$:
 prev is $n=2$'s table.
 $k=0$: $\text{curr}[0] = [\text{prev}[0][0] + \text{prev}[0][1] + \text{prev}[0][2], 0, 0] = [1 + 0 + 0, 0, 0] = [1, 0, 0]$
 $k=1$: $\text{curr}[1][0] = \text{prev}[1][0] + \text{prev}[1][1] + \text{prev}[1][2] = 1 + 1 + 0 = 2$ $\text{curr}[1][1] = \text{prev}[0][0] = 1$ $\text{curr}[1][2] = \text{prev}[0][1] = 0 =_i [2, 1, 0]$
 $k=2$: $\text{curr}[2][0] = \text{prev}[2][0] + \text{prev}[2][1] + \text{prev}[2][2] = 0 + 0 + 1 = 1$ $\text{curr}[2][1] = \text{prev}[1][0] = 1$ $\text{curr}[2][2] = \text{prev}[1][1] = 1 =_i [1, 1, 1]$
 $k=3$: $\text{curr}[3][0] = \text{prev}[3][0] + \text{prev}[3][1] + \text{prev}[3][2] = 0$ $\text{curr}[3][1] = \text{prev}[2][0] = 0$ $\text{curr}[3][2] = \text{prev}[2][1] = 0 =_i [0, 0, 0]$
So matches previous.
Now $n=4$:
 prev is $n=3$'s table.
 $k=0$: $\text{curr}[0] = [\text{prev}[0][0] + \text{prev}[0][1] + \text{prev}[0][2], 0, 0] = [1 + 0 + 0, 0, 0] = [1, 0, 0]$
 $k=1$: $\text{curr}[1][0] = \text{prev}[1][0] + \text{prev}[1][1] + \text{prev}[1][2] = 2 + 1 + 0 = 3$ $\text{curr}[1][1] = \text{prev}[0][0] = 1$ $\text{curr}[1][2] = \text{prev}[0][1] = 0 =_i [3, 1, 0]$
 $k=2$: $\text{curr}[2][0] = \text{prev}[2][0] + \text{prev}[2][1] + \text{prev}[2][2] = 1 + 1 + 1 = 3$ $\text{curr}[2][1] = \text{prev}[1][0] = 2$ $\text{curr}[2][2] = \text{prev}[1][1] = 1 =_i [3, 2, 1]$
 $k=3$: $\text{curr}[3][0] = \text{prev}[3][0] + \text{prev}[3][1] + \text{prev}[3][2] = 0$ $\text{curr}[3][1] = \text{prev}[2][0] = 1$ $\text{curr}[3][2] = \text{prev}[2][1] = 1 =_i [0, 1, 1]$
 $k=4$: all zeros
So $n=4$: $k=0$: $[1, 0, 0]$ $k=1$: $[3, 1, 0]$ $k=2$: $[3, 2, 1]$ $k=3$: $[0, 1, 1]$ $k=4$: $[0, 0, 0]$
We can keep going, but this is tedious. Let's see if we can code this or find a pattern.
Alternatively, let's write a small Python-like pseudocode to compute $H(n, k)$ up to $n=16$, $k=8$, then sum $r=0,1,2$.

Figure 11: Mathematical Reasoning Task AIME Example. Wrong Trajectory with High Centroid

But since we are to do this by hand (or at least systematically), let's keep a table for each n up to 16, but only for k up to 8.

Let's set up a grid.

Let's define for each n (from 0 to 16), for k from 0 to 8, the triplet [a,b,c] for $H(n,k,0)$, $H(n,k,1)$, $H(n,k,2)$.

We can write a spreadsheet-like table, but since we are in text, let's proceed step by step, keeping only the current and previous row.

Let's define $\text{prev}[k] = [a,b,c]$ for $k=0..8$

Initialize prev for $n=0$: $\text{prev}[0]=[1,0,0]$ $\text{prev}[1..8]=[0,0,0]$

Now, for n from 1 to 16, compute curr from prev. Let's proceed step by step, and at each n, record $\text{curr}[8]$ (the value for $k=8$). Let's proceed:

$n=0$: $\text{prev}[8]=[0,0,0]$

$n=1$: $\text{curr}[0]=[1,0,0]$ $\text{curr}[1]=[0,1,0]$ $\text{curr}[2..8]=[0,0,0]$ So $\text{curr}[8]=[0,0,0]$

$n=2$: $\text{curr}[0]=[1,0,0]$ $\text{curr}[1]=[1,1,0]$ $\text{curr}[2]=[0,0,1]$ $\text{curr}[3..8]=[0,0,0]$ $\text{curr}[8]=[0,0,0]$

$n=3$: $\text{curr}[0]=[1,0,0]$ $\text{curr}[1]=[2,1,0]$ $\text{curr}[2]=[1,1,1]$ $\text{curr}[3]=[0,0,0]$ $\text{curr}[4..8]=[0,0,0]$ $\text{curr}[8]=[0,0,0]$

$n=4$: From $\text{prev} = n=3$ $\text{curr}[0]=[1,0,0]$ $\text{curr}[1]=[2+1+0, \text{prev}[0][0], \text{prev}[0][1]]$ Wait, let's use the recurrence:

For each k: $\text{curr}[k][0] = \text{prev}[k][0] + \text{prev}[k][1] + \text{prev}[k][2]$ $\text{curr}[k][1] = \text{prev}[k-1][0]$ if $k_i=1$ else 0 $\text{curr}[k][2] = \text{prev}[k-1][1]$ if $k_i=2$ else 0

But to avoid confusion, let's write a function for each k:

For k from 0 to 8: $\text{curr}[k][0] = \text{prev}[k][0] + \text{prev}[k][1] + \text{prev}[k][2]$ $\text{curr}[k][1] = (\text{prev}[k-1][0] \text{ if } k_i=1 \text{ else } 0) - \text{curr}[k][2] = (\text{prev}[k-1][1] \text{ if } k_i=2 \text{ else } 0)$

Let's tabulate $n=4$:

prev is $n=3$: $\text{prev}[0]=[1,0,0]$ $\text{prev}[1]=[2,1,0]$ $\text{prev}[2]=[1,1,1]$ $\text{prev}[3]=[0,0,0]$ $\text{prev}[4..8]=[0,0,0]$

Now for $n=4$:

$k=0$: $\text{curr}[0][0]=1+0+0=1$ $\text{curr}[0][1]=0$ $\text{curr}[0][2]=0 =_i [1,0,0]$

$k=1$: $\text{curr}[1][0]=2+1+0=3$ $\text{curr}[1][1]=\text{prev}[0][0]=1$ $\text{curr}[1][2]=\text{prev}[0][1]=0 =_i [3,1,0]$

$k=2$: $\text{curr}[2][0]=1+1+1=3$ $\text{curr}[2][1]=\text{prev}[1][0]=2$ $\text{curr}[2][2]=\text{prev}[1][1]=1 =_i [3,2,1]$

$k=3$: $\text{curr}[3][0]=0+0+0=0$ $\text{curr}[3][1]=\text{prev}[2][0]=1$ $\text{curr}[3][2]=\text{prev}[2][1]=1 =_i [0,1,1]$

$k=4$: $\text{curr}[4][0]=0$ $\text{curr}[4][1]=\text{prev}[3][0]=0$ $\text{curr}[4][2]=\text{prev}[3][1]=0 =_i [0,0,0]$

$k=5..8$: all zeros

$\text{curr}[8]=[0,0,0]$

Now $n=5$:

$\text{prev} = n=4$'s curr

$\text{prev}[0]=[3,1,0]$ $\text{prev}[1]=[3,2,1]$ $\text{prev}[2]=[0,1,1]$ $\text{prev}[4..8]=[0,0,0]$

Compute curr for $n=5$:

$k=0$: $\text{curr}[0][0]=1+0+0=1$ $\text{curr}[0][1]=0$ $\text{curr}[0][2]=0 =_i [1,0,0]$

$k=1$: $\text{curr}[1][0]=3+1+0=4$ $\text{curr}[1][1]=\text{prev}[0][0]=1$ $\text{curr}[1][2]=\text{prev}[0][1]=0 =_i [4,1,0]$

$k=2$: $\text{curr}[2][0]=3+2+1=6$ $\text{curr}[2][1]=\text{prev}[1][0]=3$ $\text{curr}[2][2]=\text{prev}[1][1]=1 =_i [6,3,1]$

$k=3$: $\text{curr}[3][0]=0+1+1=2$ $\text{curr}[3][1]=\text{prev}[2][0]=3$ $\text{curr}[3][2]=\text{prev}[2][1]=2 =_i [2,3,2]$

$k=4$: $\text{curr}[4][0]=0+0+0=0$ $\text{curr}[4][1]=\text{prev}[3][0]=0$ $\text{curr}[4][2]=\text{prev}[3][1]=1 =_i [0,0,1]$

$k=5..8$: $[0,0,0]$

$\text{curr}[8]=[0,0,0]$

$n=6$:

$\text{prev} = n=5$'s curr

$\text{prev}[0]=[4,1,0]$ $\text{prev}[1]=[6,3,1]$ $\text{prev}[2]=[2,3,2]$ $\text{prev}[4]=[0,0,1]$ $\text{prev}[5..8]=[0,0,0]$

$\text{curr}[0]=[1,0,0]$ $k=1$: $\text{curr}[1][0]=4+1+0=5$ $\text{curr}[1][1]=\text{prev}[0][0]=1$ $\text{curr}[1][2]=\text{prev}[0][1]=0 =_i [5,1,0]$ $k=2$: $\text{curr}[2][0]=6+3+1=10$ $\text{curr}[2][1]=\text{prev}[1][0]=4$ $\text{curr}[2][2]=\text{prev}[1][1]=1 =_i [10,4,1]$ $k=3$: $\text{curr}[3][0]=2+3+2=7$ $\text{curr}[3][1]=\text{prev}[2][0]=6$ $\text{curr}[3][2]=\text{prev}[2][1]=3 =_i [7,6,3]$ $k=4$: $\text{curr}[4][0]=0+0+1=1$ $\text{curr}[4][1]=\text{prev}[3][0]=2$ $\text{curr}[4][2]=\text{prev}[3][1]=3 =_i [1,2,3]$ $k=5$: $\text{curr}[5][0]=0$ $\text{curr}[5][1]=\text{prev}[4][0]=0$ $\text{curr}[5][2]=\text{prev}[4][1]=0 =_i [0,0,0]$ $k=6..8$: $[0,0,0]$

$\text{curr}[8]=[0,0,0]$

$n=7$:

$\text{prev} = n=6$

$\text{prev}[0]=[5,1,0]$ $\text{prev}[1]=[10,4,1]$ $\text{prev}[2]=[7,6,3]$ $\text{prev}[4]=[1,2,3]$ $\text{prev}[5..8]=[0,0,0]$

$\text{curr}[0]=[1,0,0]$ $k=1$: $\text{curr}[1][0]=5+1+0=6$ $\text{curr}[1][1]=1$ $\text{curr}[1][2]=0 =_i [6,1,0]$ $k=2$: $\text{curr}[2][0]=10+4+1=15$ $\text{curr}[2][1]=5$ $\text{curr}[2][2]=1 =_i [15,5,1]$ $k=3$: $\text{curr}[3][0]=7+6+3=16$ $\text{curr}[3][1]=10$ $\text{curr}[3][2]=4 =_i [16,10,4]$ $k=4$: $\text{curr}[4][0]=1+2+3=6$ $\text{curr}[4][1]=7$ $\text{curr}[4][2]=6 =_i [6,7,6]$ $k=5$: $\text{curr}[5][0]=0+0+0=0$ $\text{curr}[5][1]=1$ $\text{curr}[5][2]=2 =_i [0,1,2]$ $k=6..8$: $[0,0,0]$

$\text{curr}[8]=[0,0,0]$

$n=8$:

$\text{prev} = n=7$

$\text{prev}[0]=[6,1,0]$ $\text{prev}[1]=[15,5,1]$ $\text{prev}[2]=[16,10,4]$ $\text{prev}[4]=[6,7,6]$ $\text{prev}[5]=[0,1,2]$ $\text{prev}[6..8]=[0,0,0]$

$\text{curr}[0]=[1,0,0]$ $k=1$: $\text{curr}[1][0]=6+1+0=7$ $\text{curr}[1][1]=1$ $\text{curr}[1][2]=0 =_i [7,1,0]$ $k=2$: $\text{curr}[2][0]=15+5+1=21$ $\text{curr}[2][1]=6$ $\text{curr}[2][2]=1 =_i [21,6,1]$ $k=3$: $\text{curr}[3][0]=16+10+4=30$ $\text{curr}[3][1]=15$ $\text{curr}[3][2]=5 =_i [30,15,5]$ $k=4$: $\text{curr}[4][0]=6+7+6=19$ $\text{curr}[4][1]=16$ $\text{curr}[4][2]=10 =_i [19,16,10]$ $k=5$: $\text{curr}[5][0]=0+1+2=3$ $\text{curr}[5][1]=6$ $\text{curr}[5][2]=7 =_i [3,6,7]$ $k=6$: $\text{curr}[6][0]=0+0+0=0$ $\text{curr}[6][1]=0$ $\text{curr}[6][2]=1 =_i [0,0,1]$ $k=7..8$: $[0,0,0]$

$\text{curr}[8]=[0,0,0]$

$n=9$:

$\text{prev} = n=8$

$\text{prev}[0]=[7,1,0]$ $\text{prev}[1]=[21,6,1]$ $\text{prev}[2]=[30,15,5]$ $\text{prev}[4]=[19,16,10]$ $\text{prev}[5]=[3,6,7]$ $\text{prev}[6]=[0,0,1]$ $\text{prev}[7..8]$

Figure 11: Mathematical Reasoning Task AIME Example. Wrong Trajectory with High Centroid

```

=[0,0,0]
curr[0]=[1,0,0] k=1: curr[1][0]=7+1+0=8 curr[1][1]=1 curr[1][2]=0 =i [8,1,0] k=2: curr[2][0]=21+6+1=28 curr[2][1]
=7 curr[2][2]=1 =i [28,7,1] k=3: curr[3][0]=30+15+5=50 curr[3][1]=21 curr[3][2]=6 =i [50,21,6] k=4: curr[4][0]
=19+16+10=45 curr[4][1]=30 curr[4][2]=15 =i [45,30,15] k=5: curr[5][0]=3+6+7=16 curr[5][1]=19 curr[5][2]=16 =i
[16,19,16] k=6: curr[6][0]=0+0+1=1 curr[6][1]=3 curr[6][2]=6 =i [1,3,6] k=7: curr[7][0]=0+0+0=0 curr[7][1]=0 curr[7]
[2]=0 =i [0,0,0] k=8: [0,0,0]
curr[8]=[0,0,0]
n=10:
prev = n-9
prev[0]=[1,0,0] prev[1]=[8,1,0] prev[2]=[28,7,1] prev[3]=[50,21,6] prev[4]=[45,30,15] prev[5]=[16,19,16] prev[6]=[1,3,6]
prev[7]=[0,0,0] prev[8]=[0,0,0]
curr[0]=[1,0,0] k=1: curr[1][0]=8+1+0=9 curr[1][1]=1 curr[1][2]=0 =i [9,1,0] k=2: curr[2][0]=28+7+1=36 curr[2][1]
=8 curr[2][2]=1 =i [36,8,1] k=3: curr[3][0]=50+21+6=77 curr[3][1]=28 curr[3][2]=7 =i [77,28,7] k=4: curr[4][0]
=45+30+15=90 curr[4][1]=50 curr[4][2]=21 =i [90,50,21] k=5: curr[5][0]=16+19+16=51 curr[5][1]=45 curr[5][2]=30 =i
[51,45,30] k=6: curr[6][0]=1+3+6=10 curr[6][1]=16 curr[6][2]=19 =i [10,16,19] k=7: curr[7][0]=0+0+0=0 curr[7][1]=1
curr[7][2]=3 =i [0,1,3] k=8: curr[8][0]=0 curr[8][1]=0 curr[8][2]=0 =i [0,0,0]
curr[8]=[0,0,0]
n=11:
prev = n-10
prev[0]=[1,0,0] prev[1]=[9,1,0] prev[2]=[36,8,1] prev[3]=[77,28,7] prev[4]=[90,50,21] prev[5]=[51,45,30] prev[6]=[10,16,19]
prev[7]=[0,1,3] prev[8]=[0,0,0]
curr[0]=[1,0,0] k=1: curr[1][0]=9+1+0=10 curr[1][1]=1 curr[1][2]=0 =i [10,1,0] k=2: curr[2][0]=36+8+1=45 curr[2]
[1]=9 curr[2][2]=1 =i [45,9,1] k=3: curr[3][0]=77+28+7=112 curr[3][1]=36 curr[3][2]=8 =i [112,36,8] k=4: curr[4][0]
=90+50+21=161 curr[4][1]=77 curr[4][2]=28 =i [161,77,28] k=5: curr[5][0]=51+45+30=126 curr[5][1]=90 curr[5][2]=50
=i [126,90,50] k=6: curr[6][0]=10+16+19=45 curr[6][1]=16 curr[6][2]=19 =i [45,16,19] k=7: curr[7][0]=0+1+3=4 curr[7]
[1]=10 curr[7][2]=16 =i [4,10,16] k=8: curr[8][0]=0 curr[8][1]=0 curr[8][2]=0 =i [0,0,0]
curr[8]=[0,0,0]
n=12:
prev = n-11
prev[0]=[1,0,0] prev[1]=[10,1,0] prev[2]=[45,9,1] prev[3]=[112,36,8] prev[4]=[161,77,28] prev[5]=[126,90,50] prev[6]
=[45,16,19] prev[7]=[4,10,16] prev[8]=[0,0,0]
curr[0]=[1,0,0] k=1: curr[1][0]=10+1+0=11 curr[1][1]=1 curr[1][2]=0 =i [11,1,0] k=2: curr[2][0]=45+9+1=55 curr[2][1]
=10 curr[2][2]=1 =i [55,10,1] k=3: curr[3][0]=112+36+8=156 curr[3][1]=45 curr[3][2]=9 =i [156,45,9] k=4: curr[4][0]
=161+77+28=266 curr[4][1]=112 curr[4][2]=36 =i [266,112,36] k=5: curr[5][0]=126+90+50=266 curr[5][1]=161 curr[5]
[2]=77 =i [266,161,77] k=6: curr[6][0]=45+16+19=80 curr[6][1]=16 curr[6][2]=19 =i [80,16,19] k=7: curr[7][0]=4+10+16=30 curr[7][1]=45 curr[7][2]=51 =i
[30,45,51] k=8: curr[8][0]=0+4+10=14 curr[8][1]=4 curr[8][2]=10 =i
[14,4,10]
So curr[8]=[14,4,10]
Now n=12, k=8: sum = 14+4+10 = 28
Let's keep going, but note that for k=8, the value only becomes nonzero at n = 24 - 2*8 + 1? Let's keep going until n=16.
But let's keep tracking only curr[8] from now on.
Let's define S(n,8) = H(n,8,0)+H(n,8,1)+H(n,8,2)
We have S(12,8)=28
Now n=13:
prev = n-12
We need prev[7], prev[8] for computing curr[8]
From n=12: prev[7]=[30,45,51] prev[8]=[14,4,10]
Compute curr[8]: curr[8][0] = prev[8][0] + prev[8][1] + prev[8][2] = 14+4+10=28 curr[8][1] = prev[7][0] = 30 curr[8][2] =
prev[7][1] = 45
So curr[8] = [28,30,45], sum = 28+30+45=103 S(13,8)=103 n=14: prev = n-13 We need prev[8]=[28,30,45], prev[7]=from
n=13
But let's compute prev[7] for n=13
From n=12: prev[7]=[30,45,51]
But to compute curr[7] for n=13: curr[7][0]=prev[7][0]+prev[7][1]+prev[7][2]=30+45+51=126 curr[7][1]=prev[6][0]=from
n=12, prev[6]=[141,126,90], so prev[6][0]=141 curr[7][2]=prev[6][1]=126 So prev[7] for n=13 = [126,141,126]
But actually, for curr[8] at n=14, we need prev[8] from n=13 and prev[7] from n=13.
From above, at n=13: prev[8]=[28,30,45] prev[7]=[126,141,126] (from above calculation)
Now curr[8] at n=14: curr[8][0]=prev[8][0]+prev[8][1]+prev[8][2]=28+30+45=103 curr[8][1]=prev[7][0]=126 curr[8][2]
=prev[7][1]=141 So curr[8]=[103,126,141], sum = 103+126+141 = 370 S(14,8)=370 n=15: prev = n-14 prev[8]
=[103,126,141] prev[7]=from n=14, let's compute it: At n=13, prev[7]=[126,141,126] At n=14, curr[7][0]=prev[7][0]+prev[7]
[1]+prev[7][2]=126+141+126=393 curr[7][1]=prev[6][0]=from n=13, prev[6]=let's see, but we didn't keep prev[6] for n=13,
but let's compute it step by step if needed. However, for curr[8] at n=15, we need prev[8] and prev[7] at n=14.
Let's compute prev[7] at n=14: At n=13, prev[6]=let's compute at n=13 for completeness.
But to save time, let's just use the recurrence for curr[8] only, which needs prev[8] and prev[7] at previous n.
But let's keep track of prev[7] as well. Alternatively, let's make a table for S(n,8):
We have: n=12: S=28 n=13: S=103 n=14: S=370
Let's compute n=15:

```

Figure 11: Mathematical Reasoning Task AIME Example. Wrong Trajectory with High Centroid

At n=14: prev[8]=[103,126,141] prev[7]=let's compute at n=14: From n=13: prev[7]=[126,141,126] curr[7] at n=14: curr[7][0]=126+141+126=393 curr[7][1]=prev[6][0] (from n=13) But we need prev[6] at n=13. Alternatively, perhaps it's easier to keep the last two rows for k=6,8.

But for curr[8] at n, we need prev[8] and prev[7] at n-1.

But to get prev[7] at n=14, we need to compute curr[7] at n=14:

curr[7][0]=prev[7][0]+prev[7][1]+prev[7][2]=126+141+126=393 curr[7][1]=prev[6][0] at n=13 But we didn't compute prev[6] at n=13. Let's backtrack and compute prev[6] at n=13.

At n=12: prev[6]=[141,126,90] So at n=13: curr[6][0]=prev[6][0]+prev[6][1]+prev[6][2]=141+126+90=357 curr[6][1]=prev[5][0] at n=12 At n=12, prev[5]=[266,161,77], so prev[5][0]=266 curr[6][2]=prev[5][1]=161 So prev[6] at n=13 = [357,266,161]

Now curr[7] at n=13: curr[7][0]=prev[7][0]+prev[7][1]+prev[7][2]=126+141+126=393 curr[7][1]=prev[6][0]=357 curr[7][2]=prev[6][1]=266 So prev[7] at n=13 = [393,357,266]

But at n=14, prev[7] = [393,357,266] (from n=13)

Wait, let's clarify: At n=13, prev[7]=[393,357,266] (from above calculation)

Now at n=14, prev[7]=[393,357,266]

Wait, let's check the calculations step by step for n=13 and n=14 for k=7,6,8. But perhaps it's easier to write a program, but since we're doing by hand, let's proceed carefully. Let me make a small table for n=12 to n=16, tracking only k=6,7,8

Let's define for each n: A[n][6]=[a6,b6,c6] A[n][7]=[a7,b7,c7] A[n][8]=[a8,b8,c8]

From n=12: From above: A[12][6]=[45,51,45] A[12][7]=[30,45,51] A[12][8]=[14,4,10] sum=28

n=13: A[13][6]: a6 = A[12][6][0]+A[12][6][1]+A[12][6][2]=45+51+45=141 b6 = A[12][5][0]=from n=12, A[12][5]=[266,161,77], so 266 c6 = A[12][5][1]=161 So A[13][6]=[141,266,161]

A[13][7]: a7 = A[12][7][0]+A[12][7][1]+A[12][7][2]=30+45+51=126 b7 = A[12][6][0]=45 c7 = A[12][6][1]=51 Wait, but earlier we computed prev[6] at n=12 as [45,51,45], so prev[6][0]=45, prev[6][1]=51

But according to recurrence: b7 = A[12][6][0]=45 c7 = A[12][6][1]=51 But earlier we computed prev[7] at n=13 as [126,141,126], but that was with a different approach.

Wait, let's check carefully:

For k=7 at n=13: curr[7][0]=A[12][7][0]+A[12][7][1]+A[12][7][2]=30+45+51=126 curr[7][1]=A[12][6][0]=45 curr[7][2]=A[12][6][1]=51 So A[13][7]=[126,45,51]

Wait, but earlier when I computed for n=13, I used prev[6][0]=141, but that's from n=13, not n=12.

Let me clarify:

At n=13, to compute A[13][7]: curr[7][0]=A[12][7][0]+A[12][7][1]+A[12][7][2]=30+45+51=126 curr[7][1]=A[12][6][0]=45 (from n=12, A[12][6]=[45,51,45]) curr[7][2]=A[12][6][1]=51 So A[13][7]=[126,45,51]

Similarly for A[13][8]: curr[8][0]=A[12][8][0]+A[12][8][1]+A[12][8][2]=14+4+10=28 curr[8][1]=A[12][7][0]=30 curr[8][2]=A[12][7][1]=45 So A[13][8]=[28,30,45], sum=103

Now n=14:

A[14][6]: a6 = A[13][6][0]+A[13][6][1]+A[13][6][2]=141+266+161=568 b6 = A[13][5][0]=from n=13, A[13][5]=let's compute it But let's focus only on k=7,8 for now.

But for A[14][7], we need A[13][7] and A[13][6]

A[13][6]=[141,266,161] A[13][7]=[126,45,51]

A[14][7]: a7 = A[13][7][0]+A[13][7][1]+A[13][7][2]=126+45+51=222 b7 = A[13][6][0]=141 c7 = A[13][6][1]=266 So A[14][7]=[222,141,266]

A[14][8]: a8 = A[13][8][0]+A[13][8][1]+A[13][8][2]=28+30+45=103 b8 = A[13][7][0]=126 c8 = A[13][7][1]=45 So A[14][8]=[103,126,45], sum=103+126+45=274

Wait, but earlier I had S(14,8)=370, but that was a miscalculation. Let's double-check.

Wait, at n=14: curr[8][0]=A[13][8][0]+A[13][8][1]+A[13][8][2]=28+30+45=103 curr[8][1]=A[13][7][0]=126 curr[8][2]=A[13][7][1]=45 So sum=103+126+45=274

So S(14,8)=274

But previously I had 370, that was incorrect due to using wrong prev[7].

Let's continue carefully.

n=13: A[13][8]=[28,30,45], sum=103 A[13][7]=[126,45,51], sum=222

n=14: A[14][8]=[103,126,45], sum=274

Now n=15: A[14][8]=[103,126,45] A[14][7]=[222,141,266] (from above)

A[15][8]: a8 = A[14][8][0]+A[14][8][1]+A[14][8][2]=103+126+45=274 b8 = A[14][7][0]=222 c8 = A[14][7][1]=141 So A[15][8]=[274,222,141], sum=274+222+141=637

n=16: A[15][8]=[274,222,141] A[15][7]=let's compute for completeness But for A[16][8], we need A[15][8] and A[15][7]

But let's compute A[15][7]: A[15][7]: a7 = A[14][7][0]+A[14][7][1]+A[14][7][2]=222+141+266=629 b7 = A[14][6][0]=need A[14][6] But let's focus on A[16][8]:

A[15][8]=[274,222,141] A[15][7]=let's compute: A[14][7]=[222,141,266] A[14][6]=let's compute A[14][6]: A[13][6]=[141,266,161] A[13][5]=let's compute quickly: But for A[14][6]: a6 = A[13][6][0]+A[13][6][1]+A[13][6][2]=141+266+161=568 b6 = A[13][5][0]=need A[13][5] But perhaps it's not necessary unless we need b7 for A[16][8]

But for A[16][8], we need: curr[8][0]=A[15][8][0]+A[15][8][1]+A[15][8][2]=274+222+141=637 curr[8][1]=A[15][7][0]=? curr[8][2]=A[15][7][1]=?

But to get A[15][7][0], we need A[14][7][0]+A[14][7][1]+A[14][7][2]=222+141+266=629 A[15][7][1]=A[14][6][0], so we need A[14][6][0]

Let's compute A[14][6]: A[13][6]=[141,266,161] A[13][5]=let's compute for n=13, k=5: At n=12, A[12][5]=[266,161,77] A[12][5]: a5=A[12][5][0]+A[12][5][1]+A[12][5][2]=266+161+77=504 b5=A[12][4][0]=from n=12, A[12][4]=[266,112,36], so 266 c5=A[12][4][1]=112 So A[13][5]=[504,266,112] Now A[14][6]: a6=A[13][6][0]+A[13][6][1]+A[13][6][2]=141+266+161=568

Figure 11: Mathematical Reasoning Task AIME Example. Wrong Trajectory with High Centroid

$b6=A[13][5][0]=504$ $c6=A[13][5][1]=266$ So $A[14][6]=[568,504,266]$
 Thus $A[15][7][1]=A[14][6][0]=568$
 $A[15][7][2]=A[14][6][1]=504$
 Thus $A[15][7]=[629,568,504]$, $\text{sum}=629+568+504=1701$
 But for $A[16][8]$: $\text{curr}[8][1]=A[15][7][0]=629$ $\text{curr}[8][2]=A[15][7][1]=568$
 So $A[16][8]=[637,629,568]$, $\text{sum}=637+629+568=1834$ Therefore, the number of ways is $N = 1834$ Wait, but let's check the math: $637 + 629 = 1266$ $1266 + 568 = 1834$ So $N = 1834$. But let's double-check the last steps, as it's easy to make arithmetic errors.
 Let's summarize the last few steps: At $n=15$: $A[15][8]=[274,222,141]$, $\text{sum}=637$
 $A[15][7]=[629,568,504]$, $\text{sum}=1701$
 Now at $n=16$: $A[16][8][0] = 274 + 222 + 141 = 637$ $A[16][8][1] = A[15][7][0] = 629$ $A[16][8][2] = A[15][7][1] = 568$ So $\text{sum} = 637 + 629 + 568 = 1834$
 Therefore, the number of subsets is 1834. But wait, this is the number of binary strings of length 16 with exactly 8 ones and no three consecutive ones.
 But the problem asks for the number of subsets of the 16 chairs that can be chosen so that no person sits next to two other people.
 But as we modeled, this is exactly the number of ways to choose 8 chairs out of 16 with no three consecutive occupied chairs, i.e., no block of three in a row.
 So $N = 1834$ But let's check with a small n to be sure. For $n=5$, $k=2$: let's compute by hand.
 Number of ways to choose 2 chairs out of 5, no three in a row. But with only 2 ones, it's impossible to have 3 in a row, so it's $C(5,2)=10$.
 Let's see what our recurrence gives for $n=5$, $k=2$. From earlier, at $n=5$: $A[5][2]=[6,3,1]$, $\text{sum}=10$
 Yes, matches. For $n=6$, $k=3$: Let's see what our recurrence gives.
 From above, at $n=6$: $A[6][3]=[7,6,3]$, $\text{sum}=16$
 Let's enumerate for $n=6$, $k=3$, no three in a row. All 3-ones subsets with no three consecutive ones. List all 3-subsets of $1,2,3,4,5,6$ where no three are consecutive.
 Total $C(6,3)=20$. Invalid ones are those where positions are $i, i+1, i+2$ for $i=1,4$.
 Which are: $1,2,3$, $2,3,4$, $3,4,5$, $4,5,6$ = 4 invalid. So valid = $20 - 4 = 16$. Matches our recurrence's value of 16. So the recurrence is correct. Thus, for $n=16$, $k=8$, $N=1834$. Now, the problem says "Find the remainder when N is divided by 1000."
 $1834 \bmod 1000 = 834$. Therefore, the answer is 834. But let me double-check the last step for $n=16$ calculation.
 From above: $A[16][8] = [637,629,568]$, $\text{sum} = 1834$
 $1834 \bmod 1000 = 834$. Thus, 834

Correct Trajectory (centroid=0.2354)

Let's break down the problem step by step.
 We have 16 chairs in a row. We want to select 8 chairs for people to sit in such that **no person sits next to two other people**. That is, **no three people are sitting in three consecutive chairs**. However, it's important to parse the condition carefully:
 "no person sits next to two other people."
 This means: **For every person, that person is not adjacent to two other people**. In other words, no one is flanked on both sides by occupied chairs. Equivalently, there is **no occurrence of three consecutive occupied chairs**.
 But let's check the precise wording:
 "no person sits next to two other people."
 That is, for any person, the two adjacent chairs (left and right) are **not both occupied**. So, in other words, **no two adjacent occupied chairs can have another occupied chair adjacent to both of them**. But more simply, it means **no three consecutive occupied chairs**. Because if there were three in a row, the middle one would have two neighbors occupied.
 Therefore, the condition is: **No three consecutive chairs are occupied**.
 But let's make sure: Suppose chairs are numbered 1 to 16. If chairs 4, 5, 6 are occupied, then the person in chair 5 is next to two other people (in 4 and 6), which is forbidden. Therefore, **no three consecutive occupied chairs**.
 So the problem reduces to:
 How many ways can we choose 8 chairs from 16 in a row so that no three selected chairs are consecutive?
 But wait: the problem says "Eight people each select a chair", so we are choosing an 8-element subset of the 16 chairs, with the condition that no three selected chairs are consecutive.
 However, let's check if the interpretation is correct. The problem says:
 "Eight people each select a chair in which to sit so that no person sits next to two other people."
 Suppose a person sits in chair i . The chairs adjacent to i are $i-1$ and $i+1$. The condition is that not both of those are occupied. Equivalently, there is no trio of consecutive occupied chairs. Yes.
 So, the problem is: **How many ways to choose 8 chairs from 16 in a row, so that no three selected chairs are consecutive?**

Figure 11: Mathematical Reasoning Task AIME Example. Wrong Trajectory with High Centroid

Let's formalize this. Let's denote the 16 chairs as positions 1 to 16. We select a subset S of 8 positions, such that in S , there is no run of three consecutive positions. Let's define a standard combinatorial approach. Let the chairs be positions 1 to 16. Let's define a binary string of length 16, where 1 means "occupied", 0 means "empty". We want exactly eight 1's, and no three 1's in a row. Let's define a recurrence for the number of ways to arrange k ones in n chairs with no three consecutive ones.

Let $f(n, k)$ be the number of binary strings of length n with exactly k ones, and no three consecutive ones. We want $f(16, 8)$. Let's derive a recurrence for $f(n, k)$.

Consider the last chair (position n):

Case 1: Chair n is empty (0). Then the first $n-1$ chairs must contain exactly k ones, with no three consecutive ones. Number of ways: $f(n-1, k)$.

Case 2: Chair n is occupied (1). Now, look at the previous chairs.

Subcase 2a: Chair $n-1$ is empty (0). Then, the first $n-2$ chairs must have $k-1$ ones, no three consecutive ones. Number of ways: $f(n-2, k-1)$.

Subcase 2b: Chair $n-1$ is occupied (1), but then chair $n-2$ must be empty (otherwise we'd have three in a row). So, chairs $n-2$ is 0, $n-1$ is 1, n is 1. Then the first $n-3$ chairs must have $k-2$ ones, no three consecutive ones. Number of ways: $f(n-3, k-2)$.

So the recurrence is:

$$f(n, k) = f(n-1, k) + f(n-2, k-1) + f(n-3, k-2)$$

with appropriate base cases.

Let's check the base cases:

- If $k = 0$: Only one way, all zeros. So $f(n, 0) = 1$ for all $n \geq 0$.
- If $n \leq 2k - 1$: Not enough chairs to place k ones without having at least two adjacent, but with the "no three in a row" condition, the minimal length to place k ones with no three in a row is: To maximize spacing, we need at least $2(k-1) + 1 = 2k - 1$ chairs (since we can have at most two in a row, separated by at least one zero). So if $n \leq 2k - 1$, then $f(n, k) = 0$.
- If $k = 1$: Place a single 1 among n chairs: $\binom{n}{1} = n$ ways. But let's be precise.

Alternatively, let's build up the table using the recurrence. Let's define $f(n, k)$ as above.

Base cases: - $f(n, 0) = 1$ for all $n \geq 0$. - $f(n, k) = 0$ if $k \geq \text{ceil}((2n+1)/3)$ [but it's easier to use the recurrence and check for $k \geq n$ or $k \geq (2n+1)/3$]. But let's proceed step by step for small n and k , then build up to $f(16, 8)$.

We can make a table for $f(n, k)$ for n from 0 to 16, k from 0 to 8. Let's initialize a table. Let's start with n from 0 to 16, k from 0 to 8. Let's initialize $f(n, k) = 0$ for all n, k , then set $f(n, 0) = 1$ for all n .

Now, for $k \geq 1$ and $n \geq 1$, we use the recurrence:

$$f(n, k) = f(n-1, k) + f(n-2, k-1) + f(n-3, k-2)$$

where we define $f(m, l) = 0$ if $m \leq 0$ or $l \leq 0$ or $l \geq m$.

Let's compute step by step.

We will build up the table row by row for $n = 0$ to 16, for $k = 0$ to 8. Let's start:

For $n = 0$: - $f(0, 0) = 1$ - $f(0, k) = 0$ for $k \geq 1$

For $n = 1$: - $f(1, 0) = 1$ - $f(1, 1) = f(0, 1) + f(-1, 0) + f(-2, -1) = 0 + 0 + 0 = 0$ + $f(0, 1)$ [but $f(0, 1) = 0$] + $f(-1, 0) = 0$ + $f(-2, -1) = 0$

→ Wait, let's apply the recurrence properly:

Wait, let's use the recurrence carefully.

Given the recurrence: $f(n, k) = f(n-1, k) //$ last chair empty + $f(n-2, k-1) //$ last chair occupied, previous empty + $f(n-3, k-2) //$ last two occupied, previous empty

But when $n \leq 2$, $f(n-2, k-1)$ or $f(n-3, k-2)$ may be for negative n , which we define as 0.

So let's proceed.

$n = 0$: - $f(0, 0) = 1$

$n = 1$: - $f(1, 0) = f(0, 0) = 1$ - $f(1, 1) = f(0, 1) + f(-1, 0) + f(-2, -1) = 0 + 0 + 0 = 0$ But actually, for $k=1$, we can place a 1 in the only position. So let's check the recurrence. Wait: for $n=1, k=1$: We need to place 1 one in 1 chair, and obviously, there's only one way: [1] But according to the recurrence: $f(1, 1) = f(0, 1) + f(-1, 0) + f(-2, -1) = 0 + 0 + 0 = 0$ But that's not right.

So let's re-examine the recurrence.

Alternative derivation:

Let's define $f(n, k)$: number of binary strings of length n , exactly k ones, no three consecutive ones. Let's consider the last run at the end. Let's define $g(n, k)$: number of such strings of length n with exactly k ones and ending with 0. Let $h(n, k)$: ending with 1, but previous is 0. Let $i(n, k)$: ending with 11 (two consecutive ones), previous before that is 0. But maybe it's simpler to use the standard recurrence: Let's define $a(n, k)$ = number of binary strings of length n , with no three consecutive 1s, with exactly k ones.

We can use the following recurrence:

To form such a string of length n and k ones: - If the last character is 0, then the first $n-1$ characters form such a string of length $n-1$ and k ones: $a(n-1, k)$ - If the last character is 1, then look at the previous one: - If the previous is 0, then the first $n-2$ characters form such a string of length $n-2$ with $k-1$ ones: $a(n-2, k-1)$ - If the previous two are 1, then the character before those two must be 0 (since otherwise we'd have three in a row), so the last three are 011, and the first $n-3$ characters are a string of length $n-3$ with $k-2$ ones: $a(n-3, k-2)$

But in both cases above, when we put a 1 at the end, we need to ensure we don't create three in a row. So the recurrence is as above:

$$a(n, k) = a(n-1, k) + a(n-2, k-1) + a(n-3, k-2)$$

But for the case $n \leq 2$, we have to handle carefully.

Alternatively, let's try to build the table for small n and see if we get the right counts.

Figure 12: Mathematical Reasoning Task AIME Example. Correct Trajectory with Low Centroid

Let's try $n = 1$:
Possible strings of length 1: - [0] (0 ones) - [1] (1 one)
So $a(1,0) = 1$ ([0]) $a(1,1) = 1$ ([1])
Now $n = 2$:
Possible strings: [00] (0 ones) [01] (1 one) [10] (1 one) [11] (2 ones)
No string of length 2 can have 3 ones, so all are allowed.
So: $a(2,0) = 1$ ([00]) $a(2,1) = 2$ ([01], [10]) $a(2,2) = 1$ ([11])
Now $n = 3$:
Possible strings of length 3 with at most 2 consecutive 1s:
List all 8 possible binary strings of length 3: 000 (0 ones) 001 (1 one) 010 (1 one) 011 (2 ones) 100 (1 one) 101 (2 ones) 110 (2 ones) 111 (3 ones) – forbidden
So allowed: 000 001 010 011 100 101 110
Now count by number of ones: 0 ones: 000 $\rightarrow$ 1 1 one: 001, 010, 100 $\rightarrow$ 3 2 ones: 011, 101, 110 $\rightarrow$ 3 3 ones: 111 $\rightarrow$ forbidden
So: $a(3,0) = 1$ $a(3,1) = 3$ $a(3,2) = 3$ $a(3,3) = 0$
Now let's check the recurrence for $n=3$, $k=2$:
 $a(3,2) = a(2,2) + a(1,1) + a(0,0) = 1 + 1 + 1 = 3$
Yes, matches.
Check $a(3,1)$: $a(3,1) = a(2,1) + a(1,0) + a(0,-1) = 2 + 1 + 0 = 3$
Yes.
Check $a(3,0) = a(2,0) + a(1,-1) + a(0,-2) = 1 + 0 + 0 = 1$
So the recurrence works.
Now we can proceed to build up the table for $n = 0$ to 16, $k = 0$ to 8.
Let's make a table.
We will create a 2D array $f[n][k]$ for $n = 0..16$, $k = 0..8$.
We can fill it step by step.
Let's initialize $f[0][0] = 1$; all other $f[0][k_i0] = 0$.
Now, for n from 1 to 16, for k from 0 to 8:
 $f[n][k] = f[n-1][k] + f[n-2][k-1] + f[n-3][k-2]$ (with $f[i][j0] = 0$ and $f[i][j<0] = 0$)
Let's compute step by step. We'll only show the values for k up to 8 and n up to 16. Let's create a grid, but for brevity I'll just compute the necessary values step by step. Let's start by writing out the first few rows explicitly.
We already have:
 $n=0$: [1, 0, 0, 0, 0, 0, 0, 0, 0] $n=1$: [1, 1, 0, 0, 0, 0, 0, 0, 0] $n=2$: [1, 2, 1, 0, 0, 0, 0, 0, 0] $n=3$: [1, 3, 3, 1, 0, 0, 0, 0, 0]
(But wait, above we had $a(3,2)=3$, $a(3,3)=0$. But according to the recurrence, $f(3,3) = f(2,3) + f(1,2) + f(0,1) = 0 + 0 + 0 = 0$. But above we see $a(3,3)=0$, but $a(3,2)=3$, $a(3,1)=3$, $a(3,0)=1$.)
Wait, but in the previous explicit count, for $n=3$, $k=3$, there is no valid string, so $a(3,3)=0$.
But let's continue for higher n .
But for $n=4$:
 $f(4,0) = f(3,0) = 1$ $f(4,1) = f(3,1) + f(2,0) + f(1,-1) = 3 + 1 + 0 = 4$ $f(4,2) = f(3,2) + f(2,1) + f(1,0) = 3 + 2 + 1 = 6$
 $f(4,3) = f(3,3) + f(2,2) + f(1,1) = 0 + 1 + 1 = 2$ $f(4,4) = f(3,4) + f(2,3) + f(1,2) = 0 + 0 + 0 = 0$
So $n=4$: [1, 4, 6, 2, 0, 0, 0, 0, 0]
 $n=5$:
 $f(5,0) = f(4,0) = 1$ $f(5,1) = f(4,1) + f(3,0) + f(2,-1) = 4 + 1 + 0 = 5$ $f(5,2) = f(4,2) + f(3,1) + f(2,0) = 6 + 3 + 1 = 10$
 $f(5,3) = f(4,3) + f(3,2) + f(2,1) = 2 + 3 + 2 = 7$ $f(5,4) = f(4,4) + f(3,3) + f(2,2) = 0 + 0 + 1 = 1$ $f(5,5) = 0$
 $n=5$: [1, 5, 10, 7, 1, 0, 0, 0, 0]
 $n=6$:
 $f(6,0) = 1$ $f(6,1) = f(5,1) + f(4,0) + f(3,-1) = 5 + 1 + 0 = 6$ $f(6,2) = f(5,2) + f(4,1) + f(3,0) = 10 + 4 + 1 = 15$ $f(6,3) = f(5,3) + f(4,2) + f(3,1) = 7 + 6 + 3 = 16$ $f(6,4) = f(5,4) + f(4,3) + f(3,2) = 1 + 2 + 3 = 6$ $f(6,5) = f(5,5) + f(4,4) + f(3,3) = 0 + 0 + 0 = 0$ $f(6,6) = 0$
 $n=6$: [1, 6, 15, 16, 6, 0, 0, 0, 0]
 $n=7$:
 $f(7,0) = 1$ $f(7,1) = f(6,1) + f(5,0) + f(4,-1) = 6 + 1 + 0 = 7$ $f(7,2) = f(6,2) + f(5,1) + f(4,0) = 15 + 5 + 1 = 21$ $f(7,3) = f(6,3) + f(5,2) + f(4,1) = 16 + 10 + 4 = 30$ $f(7,4) = f(6,4) + f(5,3) + f(4,2) = 6 + 7 + 6 = 19$ $f(7,5) = f(6,5) + f(5,4) + f(4,3) = 0 + 1 + 2 = 3$ $f(7,6) = f(6,6) + f(5,5) + f(4,4) = 0 + 0 + 0 = 0$ $f(7,7) = 0$
 $n=7$: [1, 7, 21, 30, 19, 3, 0, 0, 0]
 $n=8$:
 $f(8,0) = 1$ $f(8,1) = f(7,1) + f(6,0) + f(5,-1) = 7 + 1 + 0 = 8$ $f(8,2) = f(7,2) + f(6,1) + f(5,0) = 21 + 6 + 1 = 28$ $f(8,3) = f(7,3) + f(6,2) + f(5,1) = 30 + 15 + 5 = 50$ $f(8,4) = f(7,4) + f(6,3) + f(5,2) = 19 + 16 + 10 = 45$ $f(8,5) = f(7,5) + f(6,4) + f(5,3) = 3 + 6 + 7 = 16$ $f(8,6) = f(7,6) + f(6,5) + f(5,4) = 0 + 0 + 1 = 1$ $f(8,7) = f(7,7) + f(6,6) + f(5,5) = 0 + 0 + 0 = 0$ $f(8,8) = 0$
 $n=8$: [1, 8, 28, 50, 45, 16, 1, 0, 0]
 $n=9$:
 $f(9,0) = 1$ $f(9,1) = f(8,1) + f(7,0) + f(6,-1) = 8 + 1 + 0 = 9$ $f(9,2) = f(8,2) + f(7,1) + f(6,0) = 28 + 7 + 1 = 36$ $f(9,3) = f(8,3) + f(7,2) + f(6,1) = 50 + 21 + 6 = 77$ $f(9,4) = f(8,4) + f(7,3) + f(6,2) = 45 + 30 + 15 = 90$ $f(9,5) = f(8,5) + f(7,4) + f(6,3) = 16 + 19 + 16 = 51$ $f(9,6) = f(8,6) + f(7,5) + f(6,4) = 1 + 3 + 6 = 10$ $f(9,7) = f(8,7) + f(7,6) + f(6,5) = 0 + 0 + 0 = 0$ $f(9,8) = 0$
 $n=9$: [1, 9, 36, 77, 90, 51, 10, 0, 0]
 $n=10$:

Figure 12: Mathematical Reasoning Task AIME Example. Correct Trajectory with Low Centroid

$f(10,0) = 1$ $f(10,1) = f(9,1) + f(8,0) + f(7,-1) = 9 + 1 + 0 = 10$ $f(10,2) = f(9,2) + f(8,1) + f(7,0) = 36 + 8 + 1 = 45$ $f(10,3) = f(9,3) + f(8,2) + f(7,1) = 77 + 28 + 7 = 112$ $f(10,4) = f(9,4) + f(8,3) + f(7,2) = 90 + 50 + 21 = 161$ $f(10,5) = f(9,5) + f(8,4) + f(7,3) = 51 + 45 + 30 = 126$ $f(10,6) = f(9,6) + f(8,5) + f(7,4) = 10 + 16 + 19 = 45$ $f(10,7) = f(9,7) + f(8,6) + f(7,5) = 0 + 1 + 3 = 4$ $f(10,8) = f(9,8) + f(8,7) + f(7,6) = 0 + 0 + 0 = 0$
 $n=10$: [1, 10, 45, 112, 161, 126, 45, 4, 0]
 $n=11$:
 $f(11,0) = 1$ $f(11,1) = f(10,1) + f(9,0) + f(8,-1) = 10 + 1 + 0 = 11$ $f(11,2) = f(10,2) + f(9,1) + f(8,0) = 45 + 9 + 1 = 55$ $f(11,3) = f(10,3) + f(9,2) + f(8,1) = 112 + 36 + 8 = 156$ $f(11,4) = f(10,4) + f(9,3) + f(8,2) = 161 + 77 + 28 = 266$ $f(11,5) = f(10,5) + f(9,4) + f(8,3) = 126 + 90 + 50 = 266$ $f(11,6) = f(10,6) + f(9,5) + f(8,4) = 45 + 51 + 45 = 141$ $f(11,7) = f(10,7) + f(9,6) + f(8,5) = 4 + 10 + 16 = 30$ $f(11,8) = f(10,8) + f(9,7) + f(8,6) = 0 + 0 + 1 = 1$
 $n=11$: [1, 11, 55, 156, 266, 266, 141, 30, 1]
 $n=12$:
 $f(12,0) = 1$ $f(12,1) = 11 + 1 + 0 = 12$ $f(12,2) = 55 + 11 + 1 = 67$ $f(12,3) = 156 + 55 + 11 = 222$ $f(12,4) = 266 + 156 + 55 = 477$ $f(12,5) = 266 + 266 + 156 = 688$ $f(12,6) = 141 + 266 + 266 = 673$ $f(12,7) = 30 + 141 + 266 = 437$ $f(12,8) = 1 + 30 + 141 = 172$
Wait, let's do step by step:
 $f(12,0) = f(11,0) = 1$ $f(12,1) = f(11,1) + f(10,0) + f(9,-1) = 11 + 1 + 0 = 12$ $f(12,2) = f(11,2) + f(10,1) + f(9,0) = 55 + 10 + 1 = 66$ Wait, but above I wrote 67, let's check: Wait, $f(11,2)=55$, $f(10,1)=10$, $f(9,0)=1 \rightarrow 55+10+1=66$. But previously I had 67, so let's check carefully.
Wait, the recurrence is: $f(n, k) = f(n-1, k) + f(n-2, k-1) + f(n-3, k-2)$
So for $f(12,2)$: $= f(11,2) + f(10,1) + f(9,0) = 55 + 10 + 1 = 66$ Yes, 66.
Let's recompute carefully for $n=12$: $f(12,0) = f(11,0) = 1$ $f(12,1) = f(11,1) + f(10,0) + f(9,-1) = 11 + 1 + 0 = 12$ $f(12,2) = f(11,2) + f(10,1) + f(9,0) = 55 + 10 + 1 = 66$ $f(12,3) = f(11,3) + f(10,2) + f(9,1) = 156 + 45 + 9 = 210$ $f(12,4) = f(11,4) + f(10,3) + f(9,2) = 266 + 112 + 36 = 414$ $f(12,5) = f(11,5) + f(10,4) + f(9,3) = 266 + 161 + 77 = 504$ $f(12,6) = f(11,6) + f(10,5) + f(9,4) = 141 + 126 + 90 = 357$ $f(12,7) = f(11,7) + f(10,6) + f(9,5) = 30 + 45 + 51 = 126$
 $f(12,8) = f(11,8) + f(10,7) + f(9,6) = 1 + 4 + 10 = 15$
So $n=12$: [1, 12, 66, 210, 414, 504, 357, 126, 15]
 $n=13$:
 $f(13,0) = 1$ $f(13,1) = f(12,1) + f(11,0) + f(10,-1) = 12 + 1 + 0 = 13$ $f(13,2) = f(12,2) + f(11,1) + f(10,0) = 66 + 11 + 1 = 78$ $f(13,3) = f(12,3) + f(11,2) + f(10,1) = 210 + 55 + 10 = 275$ $f(13,4) = f(12,4) + f(11,3) + f(10,2) = 414 + 156 + 45 = 615$ $f(13,5) = f(12,5) + f(11,4) + f(10,3) = 504 + 266 + 112 = 882$ $f(13,6) = f(12,6) + f(11,5) + f(10,4) = 357 + 266 + 161 = 784$ $f(13,7) = f(12,7) + f(11,6) + f(10,5) = 126 + 141 + 126 = 393$ $f(13,8) = f(12,8) + f(11,7) + f(10,6) = 15 + 30 + 45 = 90$
 $n=13$: [1, 13, 78, 275, 615, 882, 784, 393, 90]
 $n=14$:
 $f(14,0) = 1$ $f(14,1) = f(13,1) + f(12,0) + f(11,-1) = 13 + 1 + 0 = 14$ $f(14,2) = f(13,2) + f(12,1) + f(11,0) = 78 + 12 + 1 = 91$ $f(14,3) = f(13,3) + f(12,2) + f(11,1) = 275 + 66 + 11 = 352$ $f(14,4) = f(13,4) + f(12,3) + f(11,2) = 615 + 210 + 55 = 880$ $f(14,5) = f(13,5) + f(12,4) + f(11,3) = 882 + 414 + 156 = 1452$ $f(14,6) = f(13,6) + f(12,5) + f(11,4) = 784 + 504 + 266 = 1554$ $f(14,7) = f(13,7) + f(12,6) + f(11,5) = 393 + 357 + 266 = 1016$ $f(14,8) = f(13,8) + f(12,7) + f(11,6) = 90 + 126 + 141 = 357$
 $n=14$: [1, 14, 91, 352, 880, 1452, 1554, 1016, 357]
 $n=15$:
 $f(15,0) = 1$ $f(15,1) = f(14,1) + f(13,0) + f(12,-1) = 14 + 1 + 0 = 15$ $f(15,2) = f(14,2) + f(13,1) + f(12,0) = 91 + 13 + 1 = 105$ $f(15,3) = f(14,3) + f(13,2) + f(12,1) = 352 + 78 + 12 = 442$ $f(15,4) = f(14,4) + f(13,3) + f(12,2) = 880 + 275 + 66 = 1221$ $f(15,5) = f(14,5) + f(13,4) + f(12,3) = 1452 + 615 + 210 = 2277$ $f(15,6) = f(14,6) + f(13,5) + f(12,4) = 1554 + 882 + 414 = 2850$ $f(15,7) = f(14,7) + f(13,6) + f(12,5) = 1016 + 784 + 504 = 2304$ $f(15,8) = f(14,8) + f(13,7) + f(12,6) = 357 + 393 + 357 = 1107$
 $n=15$: [1, 15, 105, 442, 1221, 2277, 2850, 2304, 1107]
Now $n=16$:
 $f(16,0) = 1$ $f(16,1) = f(15,1) + f(14,0) + f(13,-1) = 15 + 1 + 0 = 16$ $f(16,2) = f(15,2) + f(14,1) + f(13,0) = 105 + 14 + 1 = 120$ $f(16,3) = f(15,3) + f(14,2) + f(13,1) = 442 + 91 + 13 = 546$ $f(16,4) = f(15,4) + f(14,3) + f(13,2) = 1221 + 352 + 78 = 1651$ $f(16,5) = f(15,5) + f(14,4) + f(13,3) = 2277 + 880 + 275 = 3432$ $f(16,6) = f(15,6) + f(14,5) + f(13,4) = 2850 + 1452 + 615 = 4917$ $f(16,7) = f(15,7) + f(14,6) + f(13,5) = 2304 + 1554 + 882 = 4740$ $f(16,8) = f(15,8) + f(14,7) + f(13,6) = 1107 + 1016 + 784 = 2907$
Therefore, the number of subsets of 16 chairs where 8 are selected and no three are consecutive (i.e., no person sits next to two others) is
 $N = f(16, 8) = 2907$. Now, the problem asks for the remainder when N is divided by 1000. 2907 divided by 1000 is 2 with remainder 907. So the remainder is 907.
Let's double-check the last step: $f(16,8) = f(15,8) + f(14,7) + f(13,6) = 1107 + 1016 + 784 = 2907$
Yes, $2907 \bmod 1000 = 907$. Thus, the answer is 907

Figure 12: Mathematical Reasoning Task AIME Example. Correct Trajectory with Low Centroid