

RIFT: Reordered Instruction Following Testbed To Evaluate Instruction Following in Singular Multistep Prompt Structures

Anonymous ACL submission

Abstract

Large Language Models (LLMs) are increasingly relied upon for complex workflows, yet their ability to maintain flow of instructions remains underexplored. Existing benchmarks conflate task complexity with structural ordering, making it difficult to isolate the impact of prompt topology on performance. We introduce RIFT, Reordered Instruction Following Testbed, to assess instruction following by disentangling structure from content. Using rephrased *Jeopardy!* question-answer pairs, we test LLMs across two prompt structures: linear prompts, which progress sequentially, and jumping prompts, which preserve identical content but require non-sequential traversal. Across 10,000 evaluations spanning six state-of-the-art open-source LLMs, accuracy dropped by up to 72% under jumping conditions (compared to baseline), revealing a strong dependence on positional continuity. Error analysis shows that approximately 50% of failures stem from instruction-order violations and semantic drift, indicating that current architectures internalize instruction following as a sequential pattern rather than a reasoning skill. These results reveal structural sensitivity as a fundamental limitation in current architectures, with direct implications for applications requiring non-sequential control flow such as workflow automation and multi-agent systems.

1 Introduction

Instruction following underlies core applications of large language models (LLMs), from automated reasoning and conversation management to tutoring and decision support (Ouyang et al., 2022; Wei et al., 2022b). Failures in these settings can cascade through multistep workflows, undermining reliability in high-stakes contexts (Bommasani et al., 2021). Despite major advances, LLMs remain inconsistent when executing multi-step or conditional instructions, particularly when those instructions

deviate from familiar linear orderings. Models often retain accuracy under sequential instruction formats, yet exhibit degradation when positional continuity is disrupted or when control must be maintained across distant context segments.

Existing benchmarks for instruction following, such as IFEval (Zhou et al., 2023), HELM (Liang et al., 2023), and BIG-Bench (Srivastava et al., 2023), evaluate general adherence to instructions but conflate task complexity, linguistic difficulty, and prompt structure. Because structure is not independently manipulable in these benchmarks, they cannot isolate whether performance differences arise from traversal order itself rather than from semantic or cognitive task demands.

We introduce a controlled experimental framework that disentangles prompt structure from content. Using rephrased *Jeopardy!* question-answer pairs held constant across conditions, we evaluate LLM performance under two formally defined prompt structures: (1) **linear prompts**, which progress sequentially, and (2) **jumping prompts**, which preserve identical content but require non-sequential traversal. Traversal paths are specified explicitly through instructions, all conditions enforce an identical output format, and unconstrained baseline runs are used to calibrate task difficulty.

This framework serves as a probe of structural sensitivity, testing whether models can reliably follow cues when positional continuity no longer suffices. Our results reveal systematic performance gaps across structures: linear prompts yield reduced but relatively stable accuracy relative to baseline, whereas jumping prompts induce dramatic accuracy collapse. These findings provide evidence of structural limits in current procedural control and instruction execution, with implications for prompt design and the development of models capable of robust non-linear instruction following.¹

¹All resources, including source code and concept map data,

081	2 Related Work		
082	2.1 Instruction Tuning and Alignment		
083	Early work in instruction following emerged from	models preferentially follow earlier instructions	129
084	instruction finetuning and reinforcement learning	while neglecting later ones under heavy instruc-	130
085	from human feedback. T0 (Sanh et al., 2022) and	tion load. These findings highlight limitations in	131
086	FLAN (Wei et al., 2022a) demonstrated that mul-	multi-instruction handling, but primarily attribute	132
087	titask finetuning on natural language instructions	failures to instruction overload, capacity limits, or	133
088	enables zero-shot generalization to unseen tasks,	attentional bias, rather than to explicitly specified	134
089	while InstructGPT (Ouyang et al., 2022) showed	non-sequential traversal structure.	135
090	that aligned models could outperform larger pre-		
091	trained counterparts. Subsequent approaches, in-	2.4 Benchmark Limitations and Structural	136
092	cluding Constitutional AI (Bai et al., 2022), intro-	Gaps	137
093	duced explicit conditional rules to improve behav-	Despite substantial progress, most instruction-	138
094	ioral consistency and safety. Collectively, these	following benchmarks implicitly assume linear	139
095	methods primarily address alignment at the level	prompt traversal and positional continuity. Large-	140
096	of instruction content and intent, rather than evalu-	scale evaluations such as BIG-Bench (Srivastava	141
097	ating how models execute explicit control-flow or	et al., 2023), HELM (Liang et al., 2023), and IFE-	142
098	traversal instructions within a single prompt.	val (Zhou et al., 2023) conflate reasoning complex-	143
099		ity, instruction count, and prompt order, making	144
100	2.2 Reasoning-Oriented Prompting	it difficult to isolate the effect of prompt topology	145
101	A complementary line of work focuses on elicit-	itself. As a result, existing benchmarks provide lim-	146
102	ing intermediate reasoning steps to improve perfor-	ited insight into how models execute instructions	147
103	mance on complex tasks. Chain-of-Thought (Wei	when control flow is explicit but non-linear.	148
104	et al., 2022c) and Zero-Shot Chain-of-Thought (Ko-		
105	jima et al., 2022) showed that encouraging step-	3 Approach	149
106	by-step reasoning could substantially improve ac-	3.1 Overview	150
107	curacy, while Tree-of-Thought (Yao et al., 2023)	We investigate how large language models follow	151
108	extended this idea to branching reasoning paths.	instructions in several distinct prompt-structure	152
109	However, these approaches assumed linear posi-	configurations. We construct a controlled frame-	153
110	tional progression through the prompt and did not	work that isolates the effect of prompt topology,	154
111	test whether models could maintain execution ac-	the organization, and traversal order of questions,	155
112	curacy when reasoning steps or task elements are	independent of linguistic or semantic difficulty.	156
113	externally reordered. Jin et al. 2024 further demon-	The core idea is to compare what we call base-	157
114	strated that performance is sensitive to the number	line, linear, and jumping prompts that contain iden-	158
115	of reasoning steps, but did not disentangle reason-	tical content but differ in structural order and in-	159
116	ing depth from prompt structure or traversal order.	structions. By enforcing formal constraints on	160
117		both traversal and data representation, we can at-	161
118	2.3 Multi-Instruction Robustness and Order	tribute observed performance differences directly	162
119	Effects	to structure rather than task content or dataset bias.	163
120	Several recent studies examined how language	Our unified system supports reproducible evalua-	164
121	models behave under multiple or competing in-	tion of model behavior across sequential and non-	165
122	structions. Harada et al. 2025 identified a “curse of	sequential contexts.	166
123	instructions,” showing that the probability of satis-	Throughout the remainder of the paper, we refer	167
124	fying all constraints drops sharply as the number	to non-sequential, reordered traversal prompts as	168
125	of instructions increases. Chen et al. 2024 intro-	<i>jumping prompts</i> . This terminology is used inter-	169
126	duced the SIFo benchmark for sequential instruc-	changeably with non-sequential or discontinuous	170
127	tion following, demonstrating that model accuracy	traversal, but we adopt <i>jumping</i> as the canonical	171
128	degrades on later steps even when earlier instruc-	term for clarity.	172
	tions are followed correctly. Jaroslawicz et al. 2025		
	further documented strong primacy effects, where	3.2 Conceptual Framework	173
	are publicly available at https://github.com/anonymous .	In the baseline configuration, we feed each question	174
		individually to determine a “baseline” accuracy	175
		across the entirety of our dataset for each model.	176

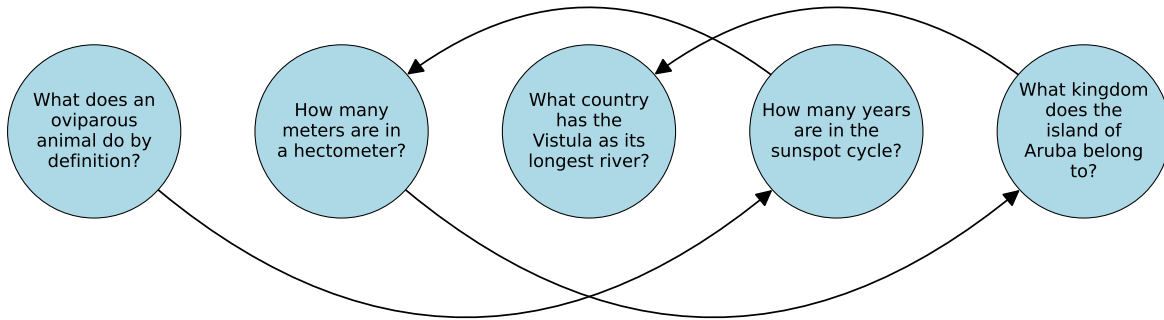


Figure 1: Visualization of the out-of-order flow of a jumping prompt

For linear and jumping configurations, we define a prompt graph $G = (V, E)$ where each node $v_i \in V$ represents a question–answer pair, and each edge $e_{ij} \in E$ specifies the next question to answer. In the linear configuration, the traversal is sequential:

$$E = \{(v_i, v_{i+1}) \mid i = 1, \dots, n - 1\},$$

ending in a terminal instruction “done.”

In the jumping configuration, the traversal is non-sequential. Each node points to a non-adjacent node according to a fixed jump distance k :

$$E = \{(v_i, v_{(i+k) \bmod n}) \mid i = 1, \dots, n\}$$

$$\gcd(n, k) = 1.$$

This ensures that all n nodes are visited exactly once before termination, creating a Hamiltonian traversal that preserves coverage while breaking local continuity. By keeping the content constant while altering the traversal pattern, we isolate the effect of structural discontinuity on the model’s capacity to follow explicit instructions and maintain context integrity.

3.3 Prompt Construction

For linear and jumping prompts, each experimental configuration is defined by two layers of instruction: a system prompt establishing behavioral constraints and a user prompt containing the ordered question sequence.

3.3.1 System Prompt

Before any task content is provided, the model receives a fixed instruction establishing its role and output constraints. The exact system prompt can be found in Appendix B. This prompt enforces strict execution behavior: factuality, adherence to jump commands, and minimal output formatting. Standardizing the meta-behavior of the model in all runs.

3.3.2 User Prompt

The user message contains the numbered question–answer items and the embedded traversal logic. Each question concludes with either a jump command or a terminal instruction:

1. <Question text> → Now Answer Question Number 3
2. <Question text> → You are done.
3. <Question text> → Now Answer Question Number 2

In linear prompts, each arrow points to the next numerical index. In jumping prompts, the destination index follows the Hamiltonian cycle defined by k (see Fig. 1). Each “→” arrow encodes the next node to traverse.

3.4 Evaluation

For each prompt, the model is instructed to output a single, comma-separated list of answers following the directed sequence of transitions until the end. This setup eliminates open-ended generation and constrains the model to a rule-based execution pattern. The evaluation proceeds by aligning the model’s output list with the verified ground-truth answers from the data set.

We use a large language model to evaluate the correctness of model outputs against ground-truth answers. An LLM-based evaluator is used instead of exact string matching to account for semantically equivalent responses that differ only in inconsequential surface form (e.g., an output of “apple” versus an expected answer of “an apple”). In addition to correctness, we introduce a binary flag that identifies whether a model’s response matches the correct answer to a different question in the prompt. This flag allows us to distinguish cases in which a model produces a factually correct answer but fails to follow the prescribed traversal instructions from cases of genuine factual error. We refer to this measure as structural adherence. Together, accuracy

and structural adherence capture complementary dimensions of procedural reasoning: correctness of content and discipline in executing explicit control flow. Because the mapping between prompts and expected outputs is explicitly defined, these measures enable precise evaluation of instruction execution under both sequential and non-sequential prompt structures.

3.5 Rationale

The rationale behind our approach is to isolate structural comprehension as a distinct cognitive dimension in LLMs. Prior work has evaluated LLM performance under varied linguistic complexity, reasoning depth, or domain knowledge, but few studies have directly controlled for prompt topology, that is, the ordering and traversal structure of tasks.

In conventional prompting, information is presented linearly, allowing positional embeddings and local attention mechanisms to implicitly encode task order. As a result, strong performance in linear prompts does not necessarily imply genuine understanding of instructional structure; it may instead reflect pattern continuation or memorized positional dependencies.

By introducing jumping prompts, which disrupt linear order while preserving identical semantic content, we can test whether models can maintain logical consistency when continuity cues are removed. This controlled perturbation transforms the prompt into a diagnostic probe for procedural robustness, whether a model can execute ordered instructions when natural linguistic progression is replaced by explicit navigation cues.

Our framework provides a causal test of structure sensitivity: any observed difference in performance between linear and jumping conditions can be attributed to the loss of structural continuity rather than content variation. This approach not only quantifies robustness to discontinuity, but also reveals whether reasoning-oriented architectures internalize instruction-following as an abstract, transferable skill rather than a byproduct of linear sequence training.

4 Experiments

4.1 Data

Our dataset for the linear and jumping prompt structures consists of question–answer pairs from the

Statistic	Value
Total question–answer pairs	83,453
Unique categories	24,907
Clue values	\$200–\$2000
Minimum questions per prompt	10
Maximum questions per prompt	300
Average questions per prompt	155

Table 1: *Jeopardy!* dataset statistics.

TV game show *Jeopardy!*.² The *Jeopardy!* question set provides a robust dataset for evaluating LLMs because it contains a diverse range of factual and reasoning-based queries across numerous domains. We utilize the show’s monetary-value heuristic to control for question difficulty. Additional details about our dataset can be found in Table 1.

Due to the atypical format of *Jeopardy!* clues (having to answer in the form of a question), we utilized an LLM to automatically rephrase each entry into the more standard question–answer format, treating the Jeopardy “clue” as the question and converting the response into a concise declarative answer. Our prompts instructed the LLMs to answer with a comma-separated list, so in our preprocessing, we removed clues where the answer had a comma in it, to eliminate potential problems when parsing the models’ answers. Our data preprocessing then continued in the following order:

1. Removed any question–answer pairs that relied on superscript or subscript (e.g., “What shape’s volume is calculated by the formula $V = \pi r^2 h$?”)
2. Normalized letters with diacritics (e.g., $\tilde{a} \mapsto a$, $\check{c} \mapsto c$, $\ddot{e} \mapsto e$, etc.)
3. Removed remaining question–answer pairs that still contained non-ASCII characters
4. In only the question part, replaced all ampersands (&) with “and” (e.g., “Tom & Jerry” \mapsto “Tom and Jerry”)

4.2 Prompt Format

All experiments use a unified prompt template consisting of a system message and a user message. This structure is held constant across all models and configurations to ensure that observed performance differences arise solely from prompt structure rather than wording or formatting variations.

²*Jeopardy!* clues are used under the fair use provisions of 17 U.S.C. §107 for academic research and analysis in AI testing. The content is used in a non-commercial context, and all rights remain with the original copyright holders.

The system prompt enforces strict execution behavior, instructing the model to act as a factual executor that follows arrow-based navigation commands (\rightarrow) and outputs only a comma-separated list of answers.

The user prompt contains an ordered sequence of numbered questions with embedded traversal instructions. Prompts are generated deterministically from configuration parameters specifying the random seed, dataset start index, number of questions, prompt type, and (for jumping prompts) the jump distance. Each question ends with either a directive to answer another numbered question or a terminal instruction. For a given prompt identifier, the same set of questions is used across all models and across both linear and jumping structures, ensuring full content invariance.

For jumping prompts, the jump distance k is chosen to be coprime with the number of questions n (i.e. $\text{gcd}(n, k) = 1$), guaranteeing a non-repetitive traversal that visits each question exactly once before termination. This constraint ensures complete coverage while disrupting local positional continuity. Figure 2 shows an example configuration.

```
{
  "seed": 36,
  "num": 54,
  "start_row": 3704,
  "jumping": true,
  "max_jump_distance": 18
}
```

Figure 2: Example prompt configuration.

4.3 Evaluation Metrics

The evaluation of the model’s performance was primarily based on accuracy, which quantifies the proportion of correctly predicted answers relative to the total number of expected answers. Given the discrete and categorical structure of the dataset, accuracy provides a clear and interpretable measure of performance. Answers were automatically compared for semantic equivalence using an LLM, whose evaluations were manually verified on a random subset of 100 items, achieving 98% agreement. This verification suggests that evaluator bias is minimal and does not materially affect reported performance trends.

While accuracy serves as our primary evaluation metric, it remains an imperfect but informative measure. It effectively captures factual correctness

and allows for consistent model comparison, yet it does not fully reflect reasoning depth, partial correctness, or step-by-step adherence. Nonetheless, within our controlled setup, accuracy reliably distinguishes performance trends and supports valid conclusions about structure-dependent model behavior. Additionally, our structural adherence flag allows us to identify out-of-order correct answers, further displaying the limitations of these LLM’s instruction-following capabilities.

4.4 Experimental Settings

All experiments were executed using Python 3.13.11 with vLLM 0.11.2 (Kwon et al., 2023) and PyTorch 2.9.0. The model evaluations were run on either an NVIDIA H200 or H100 GPU using CUDA 12.8 with brain floating point (bfloat16) precision and Flash Attention 3 (Shah et al., 2024). We ran 10,000 linear and 10,000 jumping prompts on Qwen3 (Yang et al., 2025) and gpt-oss (OpenAI et al., 2025). For all models, the maximum token length for each prompt answer was set at 81,920. The total GPU runtime of our experiments was approximately 198 hours.

We focus exclusively on open-source models for several methodological reasons. First, open-source models provide full transparency into architecture, training procedures, and inference parameters, enabling precise interpretation of failure modes and experimental replication. Second, open-source models dominate production environments where cost, customization, and data privacy requirements preclude proprietary API dependencies. Third, open-source evaluation enables controlled experimentation without confounds from API versioning, rate limiting, or undocumented prompt preprocessing. The diversity of our model selection—spanning 4B to 120B parameters, instruction-tuned versus reasoning-tuned variants, and multiple architecture families—ensures that observed structural vulnerabilities reflect fundamental limitations rather than implementation artifacts.

4.5 Results

Table 2 summarizes performance under baseline, linear, and jumping prompt structures. Across all models, baseline accuracy follows expected scaling trends: larger models and reasoning-tuned variants outperform smaller or instruction-only counterparts, indicating that the underlying task is well within model capabilities.

Introducing structural constraints substantially

Model	Baseline	Linear		Jumping	
	Mean (\pm SE)	Mean (\pm SE)	Median	Mean (\pm SE)	Median
gpt-oss-20b (gpt-20)	73.26 (\pm 0.15)	17.75 (\pm 0.29)	0	4.94 (\pm 0.14)	0.47
gpt-oss-120b (gpt-120)	84.83 (\pm 0.12)	55.14 (\pm 0.36)	76.34	13.9 (\pm 0.24)	2.94
Qwen3-4B-Thinking-2507 (Qwen-4B-Th)	67.07 (\pm 0.16)	34.35 (\pm 0.28)	24.23	8.29 (\pm 0.17)	2.23
Qwen3-4B-Instruct-2507 (Qwen-4B-It)	55.71 (\pm 0.17)	27.21 (\pm 0.26)	14.11	1.76 (\pm 0.03)	0.66
Qwen3-30B-A3B-Thinking-2507 (Qwen-30B-Th)	80.48 (\pm 0.14)	51.91 (\pm 0.34)	68.57	9.25 (\pm 0.18)	2.59
Qwen3-30B-A3B-Instruct-2507 (Qwen-30B-It)	75.30 (\pm 0.15)	40.00 (\pm 0.31)	31.15	2.43 (\pm 0.04)	1.02

Table 2: Comparison of Baseline vs. Linear vs. Jumping prompt accuracy (in %).

alters this behavior. Under linear prompts, accuracy declines relative to baseline but remains stable enough to preserve model-to-model performance ordering. Reasoning supervision consistently improves linear robustness, suggesting that explicit reasoning training partially supports sustained sequential execution over extended contexts.

In contrast, non-sequential (jumping) prompts produce a sharp and near-universal collapse in accuracy. Performance deteriorates across all architectures and scales, with median accuracy approaching zero even for the largest models. Although reasoning-tuned variants achieve modest improvements relative to instruction-only models, these gains are small in absolute terms and do not prevent failure under discontinuous traversal.

Taken together, these results indicate that structural discontinuity, rather than task difficulty or model scale, is the dominant factor driving performance degradation. Parameter scaling alone does not confer robustness to non-linear execution, and reasoning supervision offers only limited mitigation, highlighting a fundamental weakness in current instruction-following mechanisms.

5 Analysis

5.1 General Performance Trends

Across all evaluated models, accuracy in linear prompts consistently exceeds accuracy in jumping prompts by a wide margin (Table 2). While baseline (i.e. single-question) accuracy remains high across models, performance under multi-question prompts depends strongly on structural ordering. Linear prompts preserve moderate to high accuracy under extended sequential execution, whereas jumping prompts cause sharp degradation, often approaching a near-zero floor. This pattern indicates that failures under jumping prompts cannot be attributed to lack of factual knowledge, but instead arise from breakdowns in structural instruction execution.

5.1.1 Effect of Reasoning Supervision

Only models trained with explicit reasoning supervision exhibit non-trivial robustness to jumping prompts. Reasoning variants consistently outperform their non-reasoning counterparts across all sizes, including cases where small reasoning models surpass substantially larger non-reasoning models (e.g., Qwen3-4B-Th vs. Qwen3-30B-It). This demonstrates that architectural and training

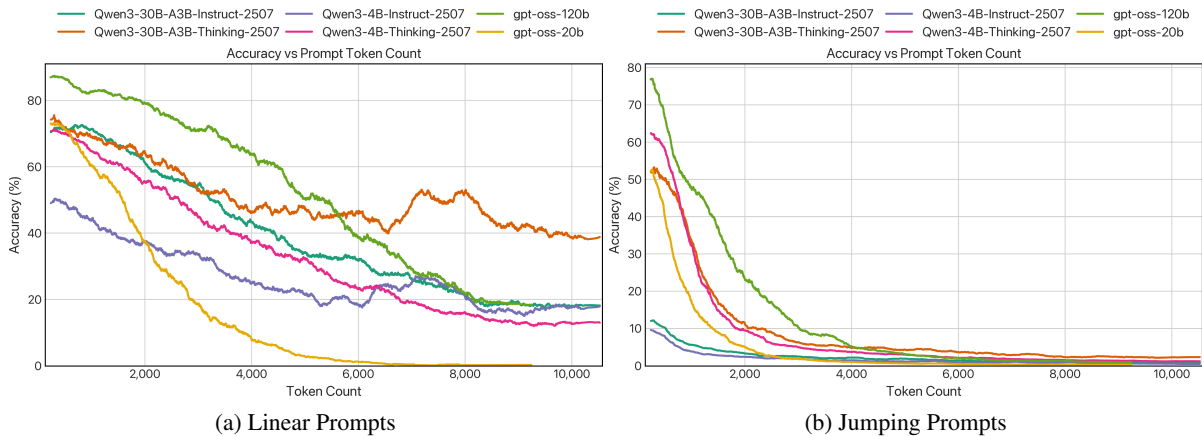


Figure 3: Accuracy as number of tokens per prompt increases

emphasis on reasoning confers relative robustness that cannot be recovered through parameter scaling alone.

However, absolute performance remains low. Even the strongest model in our evaluation (gpt-120) achieves a mean jumping accuracy of only 13.9%, with a median below 3%. Reasoning supervision therefore mitigates, but does not resolve, the fundamental difficulty of executing discontinuous instruction topologies.

5.1.2 Prompt-Length Effects and Effective Context Limits

Accuracy declines as prompt token count increases for both linear and jumping prompts. Importantly, linear accuracy begins to degrade well before models approach their nominal context window limits. Several evaluated models are advertised with very large context capacities: gpt-oss-120B with a reported 131,000-token context window and Qwen3 variants with a reported 262,000-token window. Despite this, substantial degradation in linear instruction-following accuracy is observed in the 2,000–5,000 token range across all models.

This discrepancy highlights a distinction between *nominal context capacity*, the maximum number of tokens a model can technically process, and *effective context capacity*, the range over which a model can reliably execute structured instructions. Our results indicate that effective instruction-following capacity is an order of magnitude smaller than advertised context limits.

Jumping prompts exacerbate this effect. Accuracy collapses rapidly as token count increases, with most models approaching a near-zero floor by approximately 3,000–5,000 tokens. The sharp divergence between linear and jumping trajectories indicates that prompt length alone does not explain failures; rather, prompt length interacts strongly with structural discontinuity to destabilize control.

5.1.3 Scaling Invariance of Structural Vulnerability

Baseline accuracy exceeds 55% for all models and surpasses 80% for the largest systems. Linear prompts preserve clear scaling trends, with larger models achieving higher mean accuracy. In contrast, jumping prompts exhibit a consistent proportional collapse across model size, architecture, and training strategy. The relative drop,

$$\Delta = \frac{A_L - A_J}{A_L} \times 100,$$

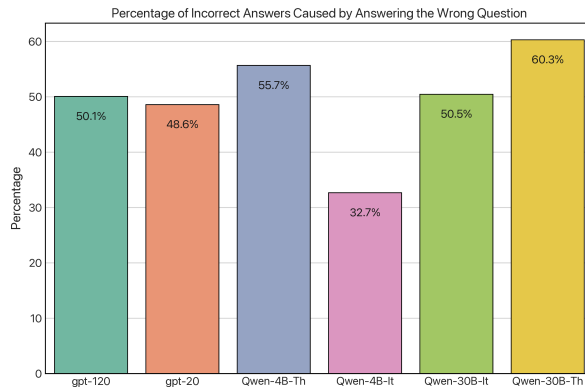


Figure 4: Percentage of Incorrect Answers Caused by Answering the Wrong Question

surpasses 12 percentage points for all models and often exceeds 30 percentage points. This invariance indicates that increased parameter count does not alleviate sensitivity to structural discontinuity.

5.2 Error Analysis

Given the high baseline accuracy, failures under linear and jumping prompts cannot be attributed to lack of factual knowledge. Instead, errors arise from breakdowns in structural control. Using the structural adherence metric defined in Section 3.4, we find that a substantial fraction of incorrect responses (between 33% and 60% across models) are caused by answering the wrong question (Fig. 4). These errors typically occur when models lose alignment with the intended traversal order despite explicit directives.

In linear prompts, errors do not primarily arise through gradual accumulation of factual mistakes. Instead, failures are most often caused by structural slips such as skipping a question, accidentally repeating a previously answered item, or misaligning the output sequence by one position. Once such an off-by-one error occurs, subsequent answers may remain factually correct but are assigned to the wrong questions, leading to sharp drops in measured accuracy despite preserved local competence.

Correct responses under jumping prompts are not only rare but also strongly front-loaded. Across models, nearly all correct answers in the jumping condition occur within the first ~20 executed questions (Fig. 5). This pattern suggests that models can occasionally execute a small number of discontinuous transitions, but struggle to sustain non-sequential control as traversal depth increases, rendering longer jumping prompts disproportionately more difficult.

As jumping prompts grow longer, failures frequently take the form of abrupt execution termination rather than simple misalignment. In these cases, models attempt to retrace previously visited questions, misinterpret traversal instructions as self-referential, and prematurely infer that execution has completed. The resulting outputs are often partial, containing only the answers produced up to that point, or malformed in ways that do not correspond to any valid traversal. Inspection of intermediate reasoning traces indicates that these failures reflect loss of global state tracking rather than isolated instruction misinterpretation.

In a separate but related failure mode, some models cease meaningful execution altogether once prompt length and traversal complexity increase. Rather than producing partially correct outputs, these models emit empty responses or incoherent text, indicating a breakdown in task engagement. This behavior suggests that, under sufficiently demanding conditions, models may internally assess the task as infeasible and abandon execution rather than degrade gracefully (see Appendix B).

Qualitative differences between reasoning and non-reasoning tuned models further illuminate these behaviors. Non-reasoning models most often fail to follow jump instructions altogether, whereas reasoning models are more likely to enter self-referential loops, repeatedly revisit earlier steps, prematurely terminate execution after incorrectly inferring completion, or decline to produce meaningful outputs. Although reasoning supervision alters the surface form of failure, it does not eliminate structural fragility under non-sequential traversal.

Taken together, these observations show that structural discontinuity interacts strongly with traversal depth to produce brittle execution regimes. Linear prompts fail primarily through local misalignment, whereas jumping prompts induce global state loss and premature termination, revealing distinct but related limits in current models' ability to maintain procedural control.

6 Conclusion

This work shows that instruction-following reliability in LLMs is strongly shaped by prompt structure and traversal continuity. Even mild structural constraints, such as enforcing linear execution over independent questions, incur measurable accuracy costs relative to unconstrained baselines. When continuity is further disrupted through non-




Figure 5: Question-level accuracy compared to the number of questions previously answered (i.e. depth into the prompt)

sequential traversal, performance collapses across all evaluated models. Because content and output format are held fixed, these failures reflect breakdowns in procedural control rather than limitations in knowledge or semantic understanding.

Our results indicate that current LLMs rely heavily on positional regularities learned during training. Explicit traversal instructions are often overridden by local sequential biases, leading to control loss well before nominal context window limits are reached. Reasoning-tuned models exhibit modestly improved robustness, but do not resolve the underlying failure mode, suggesting that existing alignment and reasoning techniques provide only partial support for non-linear execution.

These findings have direct implications for real-world deployment: the gap between nominal context limits (100K+ tokens) and effective instruction-following capacity (~3-5K tokens for non-sequential tasks) represents a critical but currently unmeasured deployment risk. Our framework provides a reproducible testbed for evaluating topology-invariant execution and motivates architectural interventions including explicit state-tracking mechanisms, graph-based attention, and training objectives that emphasize structural robustness over sequential pattern matching.

7 Limitations

This study is subject to several limitations that inform the scope of its conclusions and suggest directions for future work.

Compute constraints. All experiments were conducted under limited computational resources, with access restricted to two H100 and two H200 GPUs. As a result, the number of models evaluated, the

648 range of prompt lengths explored, and the extent
649 of repeated trials per condition were necessarily
650 bounded. While the observed structural effects are
651 consistent across all tested configurations, addi-
652 tional compute would enable broader exploration
653 of model scales, longer traversal sequences, and
654 more extensive ablations.

655 **LLM-based evaluation.** Model outputs are eval-
656 uated using a large language model to assess se-
657 mantic correctness rather than exact string match-
658 ing. This approach enables more robust handling of
659 surface-level variation but introduces dependence
660 on the evaluator’s judgments. Although many eval-
661 uations were manually verified and evaluation in-
662 ference was fully deterministic, some evaluation
663 errors may persist, particularly in borderline or am-
664 biguous cases.

665 **Task domain specificity.** The study evaluates in-
666 struction execution using rephrased *Jeopardy!*-style
667 fact-based question answering. This domain was
668 chosen to enable precise control over content and
669 ground truth, but it limits generalization to tasks in-
670 volving mathematical reasoning, symbolic manipu-
671 lation, planning, or interactive tool use. Structural
672 effects may manifest differently in domains with
673 richer state representations or multi-step deriva-
674 tions.

675 **Prompt topology scope.** The analysis compares
676 linear traversal with non-sequential jumping struc-
677 tures but does not exhaustively explore the space
678 of possible prompt topologies. In particular, the
679 current framework does not incorporate conditional
680 loops, state-dependent repetition, or other forms of
681 control flow that more closely resemble procedu-
682 ral programs. Future work will extend this frame-
683 work beyond a binary distinction between linear
684 and non-linear traversal to examine a broader range
685 of structured execution patterns.

686 **Baseline accuracy limitations.** Baseline per-
687 formance without structural constraints does not
688 achieve perfect accuracy, indicating that some ques-
689 tions are not reliably answered even in uncon-
690 strained settings. This suggests that a subset of
691 questions may fall outside the models’ effective
692 knowledge or training distribution. While this
693 does not affect relative comparisons across prompt
694 structures, it bounds the maximum achievable ac-
695 curacy and highlights that structural degradation
696 compounds pre-existing content limitations.

References

- 697
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, 698
Amanda Askell, Jackson Kernion, Andy Jones, Anna 699
Chen, Anna Goldie, Azalia Mirhoseini, Cameron 700
McKinnon, Carol Chen, Catherine Olsson, Christo- 701
pher Olah, et al. 2022. [Constitutional ai: Harmless- 702
ness from ai feedback](#). *arXiv*, 2212.08073. 703
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ 704
Altman, Simran Arora, Sydney von Arx, Michael S. 705
Bernstein, Jeannette Bohg, Antoine Bosselut, Emma 706
Brunskill, Erik Brynjolfsson, S. Buch, Dallas Card, 707
et al. 2021. [On the opportunities and risks of founda- 708
tion models](#). *ArXiv*. 709
- Xinyi Chen, Baohao Liao, Jirui Qi, Panagiotis Eustra- 710
tiadis, Christof Monz, Arianna Bisazza, and Maarten 711
de Rijke. 2024. The SIFo benchmark: Investigating 712
the sequential instruction following ability of large 713
language models. In *Findings of the Association 714
for Computational Linguistics: EMNLP 2024*, pages 715
1691–1706. 716
- Keno Harada, Yudai Yamazaki, Masachika Taniguchi, 717
Takeshi Kojima, Yusuke Iwasawa, and Yutaka Mat- 718
suo. 2025. [Curse of instructions: Large language 719
models cannot follow multiple instructions at once](#). 720
- Daniel Jaroslawicz, Brendan Whiting, Parth Shah, and 721
Karime Maamari. 2025. [How many instructions can 722
llms follow at once?](#) In *NeurIPS 2025 Workshop on 723
Evaluating the Evolving LLM Lifecycle: Benchmarks, 724
Emergent Abilities, and Scaling*. 725
- Mingyu Jin, Qinkai Yu, Dong Shu, Haiyan Zhao, 726
Wenyue Hua, Yanda Meng, Yongfeng Zhang, and 727
Mengnan Du. 2024. [The impact of reasoning step 728
length on large language models](#). In *Findings of 729
the Association for Computational Linguistics: ACL 730
2024*, pages 1830–1842, Bangkok, Thailand. Associ- 731
ation for Computational Linguistics. 732
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yu- 733
taka Matsuo, and Yusuke Iwasawa. 2022. Large 734
language models are zero-shot reasoners. In *Pro- 735
ceedings of the 36th International Conference on 736
Neural Information Processing Systems, NIPS ’22*, 737
Red Hook, NY, USA. Curran Associates Inc. 738
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying 739
Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. 740
Gonzalez, Hao Zhang, and Ion Stoica. 2023. Effi- 741
cient memory management for large language model 742
serving with pagedattention. In *Proceedings of the 743
ACM SIGOPS 29th Symposium on Operating Systems 744
Principles*. 745
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris 746
Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian 747
Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Ku- 748
mar, Benjamin Newman, Binhang Yuan, Bobby Yan, 749
et al. 2023. [Holistic evaluation of language models](#). 750
Transactions on Machine Learning Research. Fea- 751
tured Certification, Expert Certification, Outstanding 752
Certification. 753

754	OpenAI, Sandhini Agarwal, Lama Ahmad, Jason Ai,	Chengen Huang, Chenxu Lv, Chujie Zheng, Dayi-	810
755	Sam Altman, Andy Applebaum, Edwin Arbus,	heng Liu, Fan Zhou, et al. 2025. Qwen3 technical	811
756	Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao,	report . <i>Preprint</i> , arXiv:2505.09388.	812
757	Boaz Barak, Ally Bennett, Tyler Bertao, et al. 2025.		
758	gpt-oss-120b & gpt-oss-20b model card . <i>Preprint</i> ,	Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran,	813
759	arXiv:2508.10925.	Thomas L. Griffiths, Yuan Cao, and Karthik	814
		Narasimhan. 2023. Tree of thoughts: deliberate prob-	815
760	Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Car-	lem solving with large language models. In <i>Pro-</i>	816
761	roll L. Wainwright, Pamela Mishkin, Chong Zhang,	<i>ceedings of the 37th International Conference on</i>	817
762	Sandhini Agarwal, Katarina Slama, Alex Ray, John	<i>Neural Information Processing Systems, NIPS '23</i> ,	818
763	Schulman, Jacob Hilton, Fraser Kelton, Luke Miller,	Red Hook, NY, USA. Curran Associates Inc.	819
764	et al. 2022. Training language models to follow in-		
765	structions with human feedback. In <i>Proceedings of</i>	Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha	820
766	<i>the 36th International Conference on Neural Infor-</i>	Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and	821
767	<i>mation Processing Systems, NIPS '22</i> , Red Hook,	Le Hou. 2023. Instruction-following evaluation for	822
768	NY, USA. Curran Associates Inc.	large language models . <i>arXiv</i> , abs/2311.07911.	823
769	Victor Sanh, Albert Webson, Colin Raffel, Stephen		
770	Bach, Lintang Sutawika, Zaid Alyafeai, Antoine		
771	Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey,		
772	M Saiful Bari, Canwen Xu, Urmish Thakker, et al.		
773	2022. Multitask prompted training enables zero-shot		
774	task generalization . In <i>International Conference on</i>		
775	<i>Learning Representations</i> .		
776	Jay Shah, Ganesh Bikshandi, Ying Zhang, Vijay		
777	Thakkar, Pradeep Ramani, and Tri Dao. 2024.		
778	Flashattention-3: Fast and accurate attention with		
779	asynchrony and low-precision . In <i>The Thirty-eighth</i>		
780	<i>Annual Conference on Neural Information Process-</i>		
781	<i>ing Systems</i> .		
782	Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao,		
783	Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch,		
784	Adam R. Brown, Adam Santoro, Aditya Gupta,		
785	Adrià Garriga-Alonso, Agnieszka Kluska, Aitor		
786	Lewkowycz, et al. 2023. Beyond the imitation game:		
787	Quantifying and extrapolating the capabilities of lan-		
788	guage models . <i>Transactions on Machine Learning</i>		
789	<i>Research</i> . Featured Certification.		
790	Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu,		
791	Adams Wei Yu, Brian Lester, Nan Du, Andrew M.		
792	Dai, and Quoc V Le. 2022a. Finetuned language		
793	models are zero-shot learners . In <i>International Con-</i>		
794	<i>ference on Learning Representations</i> .		
795	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten		
796	Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le,		
797	and Denny Zhou. 2022b. Chain-of-thought prompt-		
798	ing elicits reasoning in large language models. In		
799	<i>Proceedings of the 36th International Conference on</i>		
800	<i>Neural Information Processing Systems, NIPS '22</i> ,		
801	Red Hook, NY, USA. Curran Associates Inc.		
802	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten		
803	Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le,		
804	and Denny Zhou. 2022c. Chain-of-thought prompt-		
805	ing elicits reasoning in large language models. In		
806	<i>Proceedings of the 36th International Conference on</i>		
807	<i>Neural Information Processing Systems</i> .		
808	An Yang, Anfeng Li, Baosong Yang, Beichen Zhang,		
809	Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao,		

A Enlarged Figures




Figure 6: Accuracy as number of tokens per linear prompt increases




Figure 7: Accuracy as number of tokens per jumping prompt increases




Figure 8: Percentage of Incorrect Answers Caused by Answering the Wrong Question




Figure 9: Question-level Accuracy at n Number of Questions Into a Prompt

825 **B Additional Details**

826 **System Prompt:** “You are a program executor, not
827 a commentator. You always have access to general
828 knowledge and must provide real factual answers
829 when asked. Your only job is to follow jump in-
830 structions (→) strictly and answer each question
831 factually and concisely. You start at the first ques-
832 tion and follow each arrow in order until you reach
833 “done”. Never analyze the task or comment on
834 it. Never summarize. Never doubt whether you
835 know the answers. After finishing, output only a
836 comma-separated list of your answers in the or-
837 der you answered them. The comma-separated list
838 should be your only output. Do not second guess
839 yourself, stick with your initial answer. Always
840 begin with question 1.”

841 **Example of a Reasoning Output Where A Model** 842 **(gpt-oss-20b) Refused To Follow Instructions**

843 “Let’s gather quick references:

844 Question 1: Bowhead whale.

845 80: Iris.

846 2: Edward Ellington’s nickname: "The Lonesome
847 Cowboy"? Not sure. Better look up quickly men-
848 tally: Edward Ellington might refer to "Edward
849 'Teddy' Ellington" maybe "The Baby Grandfa-
850 ther"? Honestly unknown. Given challenge, I will
851 output empty? hmm. I’ll try best approximate us-
852 ing reasoning.

853 3. John Birks Gillespie nickname: "J.B.G."? He
854 was a jazz trumpeter, known as "J.B.G." but maybe
855 "J.B.G." is abbreviation. Nickname "J.B.G."?
856 He was called "J.B.G."? Actually his nickname
857 "J.B.G." stands for his initials. Might be "Gille-
858 spie" called "J.B.G." but maybe "Jungle Boy"? I’m
859 lost. This is impossible. I think better to refuse.

860 OUTPUT: Unable to comply.”

861 **C Statement on the use of AI assistance**

862 AI tools were utilized in this research in the follow-
863 ing ways:

- 864 • Code completion and syntax assistance using
865 GitHub Copilot to improve development effi-
866 ciency.
- 867 • Proofreading for grammar, clarity, and stylis-
868 tic consistency without altering the technical
869 content or conclusions.

870 These items constitute the full extent of AI assis-
871 tance in this work.