

000 BEYOND BENCHMARKS: TOWARD CAUSALLY FAITH- 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 BEYOND BENCHMARKS: TOWARD CAUSALLY FAITH- FUL EVALUATION OF LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Current large language models (LLMs) evaluations overlook that measured LLM performance is produced on a full evaluation system, including many indispensable components, such as workloads, prompting methods, decoding parameters, and the supporting software–hardware stack. Without an explicit, controlled specification of the evaluation system, attributing performance differences to the model itself is unreliable. Our experiments reveal that uncontrolled testing may lead to accuracy variations of up to 70%. To address this urgent issue, we introduce LLM evaluatology, a principled methodology that reduces the evaluation problem to accurately attributing the outcomes to the effect of the evaluated LLM, which is a high-dimensional causal-attribution problem. Empirical results demonstrate that LLM evaluatology not only enhances interpretability and causal validity, but also yields evaluations that are more robust, reproducible, and trustworthy than prevailing benchmarks.

1 INTRODUCTION

Current LLM evaluation practices are fragmented and ad-hoc, spanning standardized test-style benchmarks (Hendrycks et al., 2020; Huang et al., 2023; Rein et al.; Suzgun et al., 2023; AIME, 2025), human preference-based benchmarks (Chiang et al.; OpenCompass, 2025; Xu et al., 2023), and dynamic or continuously refreshed benchmarks (Jain et al.; Jimenez et al.; White et al.; Zhu et al.; Li et al.). Yet all largely treat the model in isolation, neglecting that measured performance arises from the entire evaluation system, including workloads, prompts, decoding, and even the software–hardware stack. In reality, LLM evaluation is inherently a high-dimensional problem, as these interacting components jointly shape outcomes and complicate attribution. As recent studies show, results can vary sharply with dataset artifacts (Long et al., 2024; Liu et al., 2025), prompt formatting (He et al., 2024), decoding strategies (Shi et al., 2024), or annotator biases (Das et al., 2024). But such analyses remain piecemeal, each targeting a single component without quantifying their combined impact or enabling principled attribution. What is missing is a rigorous methodology that disentangles intrinsic model capability from confounding influences and establishes a reliable foundation for evaluation.

Even under a fully specified evaluation system, LLMs differ fundamentally from traditional single-task or deterministic systems such as conventional algorithms or CPUs. For CPUs, workloads in domains like desktop computing or high-performance computing exhibit well-characterized patterns, allowing evaluation to focus on representative hotspots while treating less common cases as secondary. In contrast, LLM workloads are effectively open-ended: each user can define new tasks across languages, domains, and usage styles. Some tasks resemble those seen during training, others require analogical transformation from familiar patterns, and yet others are entirely novel. This diversity eliminates the notion of a single “typical” workload, making isolated evaluation on a few canonical examples insufficient. [Here we adopt the term “workload” from CPU benchmarking, using it to denote a question or instance within a benchmark that the LLM is required to solve](#). In addition, LLMs may produce fluent responses without genuine reasoning or knowledge, so-called hallucinations, meaning that correctly solving one instance does not guarantee mastery of the underlying skill. Consequently, reliable evaluation must consider multiple task variations, from familiar to analogical to novel, in order to disentangle true capability from surface-level correctness. Interpreting performance and attributing capability is therefore both a high-dimensional and a content-sensitive challenge, further amplified by the confounding inherent in the evaluation system.

This paper introduces LLM evaluatology (Fig. 1), a principled methodology for the rigorous evaluation of LLMs based on Evaluatology (Zhan, 2024; Zhan et al., 2024). At its core, we construct a Minimal Evaluation System (MES), which explicitly defines the evaluated object (e.g., standalone LLM or LLM service), the indispensable components influencing performance, and the evaluation conditions (the configuration space formed by admissible settings of these components). By providing a well-defined, controllable system, MES enables systematic exploration of the evaluation configuration space, capturing how different components jointly affect performance and allowing accurate attribution of model capabilities – a solution to the high-dimensional nature of LLM evaluation. To address content sensitivity, we further extend MES into an Augmented MES (A-MES), which transforms existing workloads and generates new instances along semantically related themes. This approach ensures evaluation coverage across three workload layers: workloads the model is likely to have seen, workloads requiring analogical transformation, and entirely novel workloads, thereby mitigating the risks of superficial correctness and hallucination. A-MES offers a structured, reproducible, and automatable framework that disentangles intrinsic model competence from confounding influences while accommodating the diversity and dynamism of real-world user interactions.

Our experiments reveal several important findings. First, by constructing A-MES, we observe that the accuracy of Doubao varies dramatically with configuration, ranging from 0 to 0.8, highlighting the substantial impact of evaluation settings. Notably, Doubao-1.5-pro ranks first under MES but drops to sixth under A-MES, with a significant gap from the top model, indicating limited generalization ability. Within the Qwen series, we find that the smaller model ranks higher under MES but is surpassed by the larger model under A-MES, suggesting that A-MES provides a more faithful reflection of scaling properties. By contrast, DeepSeek-V3 consistently achieves strong accuracies across all MES and A-MES scenarios, demonstrating the strongest robustness among the tested models. Second, leveraging analysis of variance (ANOVA), xgboost, and linear models, we quantify the impact of each component on model accuracy. All components show measurable influence, with Question Format and COT emerging as the most sensitive, followed by max_tokens, Shot, and Multi Turn. Furthermore, models exhibit heterogeneous sensitivity to languages: for example, DeepSeek-V3 is most sensitive to Arabic, where its accuracy reaches the lowest among all languages tested. Finally, we validate that our proposed LLM evaluatology provides the closest approximation to the accuracy ground truth, significantly outperforming traditional single-configuration evaluations in reliability and robustness.




Figure 1: LLM Evaluatology: Measured performance arises from an Augmented Minimal Evaluation System (A-MES), which enables disentangling intrinsic model capability from confounding influences. Here, the evaluation object is defined as the LLM service, comprising the LLM and its underlying systems. When evaluating a standalone LLM, the underlying systems are instead treated as part of the evaluation conditions (EC).

2 RELATED WORK

Broadly, existing benchmarks can be grouped into the following three categories. Standardized test-style benchmarks present problems in the form of test questions, with model outputs compared against reference answers. Representative examples include MMLU (Hendrycks et al., 2020) and

108 its extensions MMLU-Pro (Wang et al., 2024b) and MMLU-Redux (Gema et al., 2025), as well as
 109 C-Eval (Huang et al., 2023) and CMMLU (Li et al., 2024) in the Chinese context. GPQA (Rein
 110 et al.) targets graduate-level science, while other datasets focus on specific capabilities such as
 111 reasoning (BBH (Suzgun et al., 2023), HellaSwag (Zellers et al., 2019), Winogrande (Sakaguchi
 112 et al., 2021)), mathematics (GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al.), AIME (AIME,
 113 2025)), coding (HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021), Aider-polyglot (Aider,
 114 2025), MultiPL-E (Cassano et al., 2023)), long-context understanding (L-Eval (An et al., 2024),
 115 LongBench (Bai et al., 2024), ∞ Bench (Zhang et al., 2024a), HELMET (Yen et al., 2025)), safety
 116 (SafetyBench (Zhang et al., 2024b), Toxigen (Hartvigsen et al., 2022)), instruction-following (IFE-
 117 eval (Zhou et al., 2023), Multi-Challenge (Sirdeshmukh et al., 2025)), and multimodality (MMBench
 118 (Liu et al., 2024), MMMU (Yue et al., 2024), MathVista (Lu et al.)).

119 Human preference-based benchmarks evaluate models in interactive settings, collecting user judgments
 120 instead of relying on fixed test sets. Chatbot Arena (Chiang et al.) is the most prominent example,
 121 where pairwise votes are aggregated via Elo ratings. CompassArena (OpenCompass, 2025) apply
 122 similar designs in the Chinese context.

123 Dynamic or continuously refreshed benchmarks aim to avoid data contamination by relying on
 124 newly released or procedurally generated tasks. Examples include LiveCodeBench (Jain et al.)
 125 (recent programming contests), SWE-bench (Jimenez et al.) (GitHub issues and PRs), LiveBench
 126 (White et al.) (rolling monthly refresh), DyVal (Zhu et al.) (procedural reasoning via DAGs), and
 127 Arena-Hard (Li et al.) (real-time crowdsourced challenges).

128 **Table 1: Evaluation Settings on Different Benchmarks** (Lang. = Language, Format = Question
 129 Format, Para. = Question Paraphrase, M-turn = Multi Turn, Temp. = temperature, PP = pres-
 130 ence_penalty, MaxTok = max_tokens, ori = original, y = yes, n = no)

Model	Lang.	Format	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
MMLU	English	ori	n	0/3/5	y/n	n	0.0/0.3/0.5/0.6/0.7	0.8/0.95	0/1.5	1024/8192/32768
AIME	English	ori	n	0	y/n	n	0.0/0.6/0.7	0.8/0.95	0/1.5	8192/32768/38912
GPQA	English	ori	n	0/5	y/n	n	0.4/0.5/0.6/0.7	0.8/0.95	0/1.5	1024/8192/32768
MATH	English	ori	n	0/8	y/n	n	0.0/0.6/0.7	0.8/0.95	0/1.5	8192/32768
SWE-bench	English	ori	n	0	y/n	n	0.0/0.8	0.95	x	8192/16384
IFEval	English	ori	n	0	y/n	n	0.0/0.6/0.7	0.8/0.95	0/1.5	8192/16384
Arena-Hard	English	ori	n	0	y/n	n	0.0/0.6/0.7	0.8/0.95	0/1.5	8192/32768
Human Eval	English	ori	n	0	y/n	n	0.3	0.95	x	8192/32768

3 MOTIVATION

141 **The flaw of existing LLM evaluation methodology.** Existing LLM benchmarks define workload
 142 formats and scoring rules, but leave crucial indispensable components uncontrolled, e.g., decoding
 143 parameters and prompting methods. As a result, reported evaluation outcomes often do not allow
 144 a direct comparison of model differences and may conflate intrinsic capability with arbitrary com-
 145 ponent settings. To make this issue concrete, we systematically reviewed major benchmarks and
 146 compiled a taxonomy of which components are explicitly defined and which are left open (Table 1).
 147 Strikingly, many widely used benchmarks, including AIME, specify only a subset of variables while
 148 leaving key components underspecified. To quantify the implications, we reconstructed the AIME
 149 evaluation space by enumerating plausible settings of uncontrolled components (e.g., COT, temper-
 150 ature, top-p, presence penalty, max tokens), yielding 162 distinct evaluation conditions. Accuracy
 151 under these conditions varied by as much as 70% across settings, and the distributions often diverged
 152 substantially from the single numbers reported in technical documentation. On some models, the
 153 median relative change between our measured accuracy and the accuracy reported in the technical
 154 report reached as high as 50% (see Figure 2). Comparable inconsistencies are evident in MMLU
 155 (Appendix A.2) and other flagship benchmarks, suggesting that the problem is not dataset-specific
 156 but structural across current LLM evaluation methodologies. These findings reveal a fundamen-
 157 tal flaw in current practice: a benchmark score is often not a property of the model alone but of
 158 the loosely specified evaluation system surrounding it. Without principled control over these con-
 159 founding components, evaluation becomes unstable, attribution unreliable, and comparisons across
 160 models misleading.

161 **The challenges of using Evaluatology for LLM evaluation.** Zhan et al. conceptualize evaluation
 162 as constructing a minimal system that integrates the evaluation object with indispensable compo-
 163 nents while considering user requirements (Zhan, 2024; Zhan et al., 2024). Wang et al. illustrate this

approach for CPUs, where a Minimal Evaluation System (MES) isolates CPU behavior from confounding components (Wang et al., 2024a). However, extending Evaluatology to LLMs presents a qualitatively deeper challenge than in the case of CPUs or other deterministic systems. For such conventional artifacts, workloads can be reasonably characterized and stabilized: standardized benchmarks capture dominant usage scenarios and once confounders are controlled, evaluation outcomes largely reflect intrinsic system differences. By contrast, LLM workloads are inherently open-ended and socially constructed, shaped by heterogeneous users, diverse linguistic and cultural contexts, and the continual emergence of novel use cases. In this setting, even the “unit of evaluation” becomes unstable: what qualifies as mainstream, extrapolative, or *atypical* shifts across communities and over time. To illustrate, consider the following problem from AIME: “Let A, B, C , and D be points on the hyperbola $\frac{x^2}{20} - \frac{y^2}{24} = 1$ such that $ABCD$ is a rhombus whose diagonals intersect at the origin. Find the greatest real number that is less than BD^2 for all such rhombi.” When evaluated on nine LLMs including deepseek, doubao, gpt series, moonshot, mistral, qwen series, etc., five were able to solve this original (seen) workload correctly. However, after performing analogical transformations through inserting distractor: “In a geometric study, we often encounter various shapes and their properties. Also, the concept of symmetry plays an important role in analyzing the relationships between different geometric figures. Let A, B, C , and D be points on the hyperbola $\frac{x^2}{20} - \frac{y^2}{24} = 1$ such that $ABCD$ is a rhombus whose diagonals intersect at the origin. Find the greatest real number that is less than BD^2 for all such rhombi.”, none of these models produced correct solutions. This striking contrast illustrates why A-MES is essential: performance on a single workload can be misleading, as models may succeed on problems they have effectively memorized yet fail when the same reasoning must be applied under slightly altered conditions.

4 LLM EVALUATOLOGY

LLM evaluatology consists of three essential steps: (1) defining MES, (2) defining A-MES, and (3) evaluating MES/A-MES and attributing evaluation outcomes.

4.1 DEFINING MINIMAL EVALUATION SYSTEM (MES)

We define the Minimal Evaluation System (MES) for LLM evaluation as the smallest independently runnable system that includes the evaluated object and all indispensable components that materially affect the evaluation outcome. The evaluated object O is not limited to a bare LLM; it can also encompass the broader deployed LLM service that fuses the model with its supporting software and hardware stack. For example, when evaluating through an API, the LLM and its underlying systems should be treated as an inseparable whole, whereas for locally deployed open-source models, the surrounding system environment may either be incorporated into O or explicitly modeled as part of the other indispensable components. Thus, the first step of defining MES is to rigorously define the evaluated object.

The second step in defining MES is to identify the indispensable components that shape evaluation outcomes and to establish their value ranges, collectively denoted as evaluation conditions (EC). We organize EC into three layers, covering workload, prompting method, and decoding parameters, which together yield 10 key factors (C_1 – C_{10}). *Workload* captures data-related variations, including Language, Question Format, and Question Paraphrase (C_1 – C_3). Note that Question Paraphrase is introduced as a key component to mitigate hallucination and data contamination, referring to reformulating questions without altering their semantics or correct answers. *Prompting method* accounts for interaction styles, namely Shot, COT(chain-of-thought), and Multi Turn (C_4 – C_6). *Decoding parameters* include COT length, temperature, top-p, presence penalty, and max tokens.




Figure 2: Accuracy deviations on AIME when evaluating with identical workloads across 162 combinations of component settings (COT, temperature, top-p, presence penalty, and max tokens).

216 Table 2: Evaluation Conditions: Indispensable Components and Value Ranges
217

218 Variable	219 Value Range
219 Language	Chinese, English, Japanese, Arabic, French, Russian
220 Question Format	Multiple-choice, Fill-in-the-blank
221 Question Paraphrase	Yes, No
222 Shot	Yes, No
223 COT	Yes, No
224 Multi Turn	Yes, No
225 temperature	0.0, 1.0, 2.0
226 top_p	0.2, 0.6, 1.0
227 presence_penalty	-0.5, 0.5, 1.5
228 max_tokens	10, 100, 4000

225 *coding parameters* represent inference controls, including temperature, top_p, presence_penalty, and
226 max_tokens ($C_7 - C_{10}$). Each component is instantiated with representative values to balance cover-
227 age of real-world variability against configuration space tractability. The indispensable components
228 and their value ranges are summarized in Table 2, with both components and their value ranges con-
229figurable based on the evaluation object and user-defined requirements. Each MES instance is then
230 specified as $MES = EC \times O$, ensuring that performance measurements are attributed correctly
231 while systematically controlling for confounding factors introduced by indispensable components.

232 4.2 CONSTRUCTING AUGMENTED MES (A-MES)
233

234 To further overcome the limitations of traditional evaluation, we extend MES into an augmented
235 form (A-MES) by expanding the workload subspace. Specifically, an MES is defined as $EC \times O$,
236 where the evaluation conditions factorize as $EC = W(\text{workload}) \times P(\text{prompting_methods}) \times$
237 $D(\text{decoding_parameters})$. We do augmentation in workload W and leave the non-workload EC
238 components P and D unchanged when building A-MES. Thus $A\text{-MES} = O \times EC_A$, with $EC_A =$
239 $W_A \times P \times D$, $W_A = A(W)$. $A()$ is the augmentation operator that expands the original workload
240 W into an enriched workload W_A . **Practically, $A(W)$ is constructed by partitioning and extending**
241 **items from the original workload into three purpose-built categories—original (seen), analogical**
242 **transformation (transformed), and novel (newly-synthesized) workloads—as shown in Fig. 3:**

243 In our implementation, the latter two categories of augmented workloads are constructed through
244 five systematically defined, script-driven transformation pipelines including: three analogical (dis-
245 tractor insertion, numeric substitution, conditional recomposition) and two novel (recent-source
246 adaptation and conceptual synthesis) pipelines. For each pipeline, we fix general prompts and
247 scaffolding, and then run the entire process automatically through LLM API calls (e.g., GPT-5),
248 lightweight verification scripts and other auxiliary tooling. This setup scales to large workloads
249 and produces diverse variants, without any per-item manual rewriting or hand-crafting of individual
250 problems. Below we outline the overall automatic transformation process; detailed procedures,
251 transformation examples, and algorithmic pseudocode are provided in Appendix A.3.

252 (1) Analogical Pipelines
253

- 254 • **Distractor Insertion.** Distractor insertion augments an original question by adding redundant
255 sentences at a random position. We systematically divide redundant information into three cat-
256 egories: (i) context-irrelevant redundancy that is completely unrelated to the problem content;
257 (ii) context-relevant explanatory redundancy that explains concepts already appearing in the
258 problem; and (iii) context-relevant misleading redundancy that is logically related to the prob-
259 lem but deliberately nudges the solver toward an incorrect strategy. By providing the LLM with
260 transformation examples, the correct answer, and (optionally) solution steps, all three types of
261 redundancy are generated via similar structured prompts and automatically inserted at random
262 positions in the problem statement. We empirically evaluated multiple candidate LLMs for
263 this task and selected the one that most consistently respects these constraints, GPT-5, as our
264 transformation executor.
- 265 • **Numeric Substitutions.** Numeric substitution augments a problem by systematically pertur-
266 bing its key numerical parameters. We leverage the correct answer, solution sketches, and in-
267 formation from a pre-built formula library to prompt the LLM to generate a Python solver that
268 explicitly parameterizes the key numbers in the problem. We then execute this solver locally
269 and, if any error or mismatch is detected, feed the error messages back to the LLM for iterative
270 refinement until the code passes verification. Once a reliable solver is obtained, we automati-
271 cally sample new parameters within a predefined range and invoke the solver to compute the

270 corresponding new answers, thereby creating a family of numeric variants without manually
 271 editing each instance or recomputing answers by hand.


272

- 273 • **Conditional Recomposition** Conditional recomposition augments a problem by constructing
 274 “inverse” variants where the original answer is treated as a given condition and some of the
 275 original conditions become the new target quantities. We first prompt the LLM to identify
 276 which conditions and target quantities can be interchanged, and then, following a procedure
 277 similar to numeric substitutions, use a large language model to generate a Python solver that
 278 explicitly parameterizes the key numerical quantities in the problem. The solver is iteratively
 279 refined until it passes verification. Once a verified solver is obtained, we perturb the new
 280 input conditions within a reasonable range and automatically generate multiple “conditionally
 281 recomposed” variants of the original problem.

282 (2) Novel Pipelines

283

- 284 • **Recent-source Adaptation** Recent-source adaptation augments a problem by aligning it with
 285 thematically similar questions drawn from recent real-world exams. Given an original problem,
 286 we first use an LLM to extract its core knowledge points. We then query a public exam-question
 287 repository, indexed by year, region, subject, and knowledge point, to retrieve recent (e.g., 2025)
 288 exam questions that match these knowledge points. The retrieved questions are subsequently
 289 paraphrased via LLM, and can optionally be further transformed using the three analogical
 290 pipelines described above. In this way, we obtain recent-source adapted problems that remain
 291 aligned with the original item at the knowledge-point level while being entirely new instances.
- 292 • **Conceptual Synthesis** Conceptual synthesis augments a problem by generating conceptual
 293 questions that target the underlying concepts. Based on authoritative textbooks in PDF form,
 294 we build a structured knowledge base in which each concept is associated with its definitions,
 295 theorems, phenomena, and canonical examples extracted from the textbooks. For a given prob-
 296 lem, we use an LLM to identify its primary knowledge points and retrieve the corresponding
 297 concept entries from the knowledge base. If the corresponding concept entries is missing, we
 298 trigger a textbooks crawling and parsing step to expand the knowledge base and try again.
 299 We then prompt the LLM to synthesize new conceptual questions grounded in these entries,
 300 yielding problems that probe conceptual understanding underlying the origin item.



310 Figure 3: Augment the original workload into analogical transformation and novel workloads.
 311

312

313 4.3 EVALUATING ON A-MES AND ATTRIBUTING EVALUATION OUTCOMES

314

315 MES samples from the full space defined by 10 variables, whereas A-MES starts from this same
 316 space and, for each question within a benchmark, applies all seven mechanisms to construct the
 317 full space of augmented variants, filtering out transformation attempts that fail (e.g., numeric sub-
 318 stitutions or conditional recomposition without a stable solver). Evaluation then samples directly
 319 from this augmentation space, ensuring that no invalid transformations are ever selected; the small
 320 number of discarded variants has negligible impact on overall coverage or robustness.

321 Given the exponentially large space of workload, prompting methods, and decoding parameters, ex-
 322 haustive testing is generally infeasible. Evaluation on MES/A-MES balances the trade-off between
 323 evaluation accuracy and evaluation cost by systematically sampling the configuration space of eval-
 324 ution conditions under a joint convergence- and LLN-based stopping rule. Specifically, we generate

324 the full list of configurations once and shuffle it with a fixed random seed. This single globally
 325 shuffled list is then shared across all models and workloads. For any given model and benchmark,
 326 use a sample size N by selecting the first N configurations from this global list, without any further
 327 re-shuffling. We process this shuffled list sequentially in fixed-size batches (e.g., 10). After every
 328 batch, we recompute the running mean accuracy over all configurations seen so far, together with
 329 the corresponding 95% confidence interval. We stop sampling when two conditions are satisfied
 330 simultaneously: (i) the absolute changes in the running mean accuracy for the last three updates are
 331 all smaller than 0.002, and (ii) the length of the 95% confidence interval is smaller than 0.06. The
 332 number of configurations evaluated when these criteria are first satisfied is denoted N_{conv} . After con-
 333 vergence, we further apply a simple Law of Large Numbers-based estimation: using the empirical
 334 variance of the current results, we compute the minimal sample size required to achieve the desired
 335 error tolerance and confidence level, obtaining an LLN-based sample size N_{LLN} . If N_{LLN} is larger
 336 than N_{conv} , we continue sampling along the same shuffled order until N_{LLN} configurations have been
 337 evaluated; otherwise, we stop at N_{conv} . Combining the convergence criterion with the LLN-based
 338 check ensures that our sample sizes are both empirically stable and theoretically justified.

339 The sampled evaluation conditions are then used to test the evaluation object, yielding performance
 340 outcomes under diverse settings. The final reported evaluation score for each model is then the mean
 341 performance over the sampled instances, together with 95% and 99% confidence intervals, providing
 342 a stable summary metric that balances comprehensiveness with practical efficiency. One approach
 343 to isolate component effects is to use equivalent evaluation conditions, where all component settings
 344 are held constant except for the factor of interest; differences in measured performance can thus be
 345 attributed directly to that component, effectively mitigating confounding. An alternative and com-
 346plementary approach is to apply ANOVA (analysis of variance) across the sampled configurations,
 347 quantifying the proportion of performance variance explained by each component and enabling sys-
 348 tematic attribution of effects. Together, these strategies provide both controlled and statistical means
 349 to disentangle intrinsic model capability from the influence of evaluation conditions.

351 5 EVALUATION

352
 353 In this section, we evaluate the proposed methodology using mainstream LLMs that are publicly
 354 accessible, including deepseek-v3, doubaot-1.5-pro-32k, gpt-3.5, gpt-4.1, moonshot-v1-8k, mistral-
 355 large, mistral-medium, qwen-plus and qwen2.5-32b-instruct. We have three targets. 1) Demonstrate
 356 the necessity of constructing MES and A-MES for LLM evaluation by varying the settings of each
 357 indispensable component within MES and A-MES. 2) Quantify the contribution of each indispens-
 358 able component to overall performance variance and identify the key components affecting LLM
 359 behavior using ANOVA. 3) Compare LLM evaluatology with traditional LLM evaluation methods
 360 and show how it enables accurate attribution of performance differences to specific components.

361 For online testing, we primarily access the models through their official APIs; however, since the
 362 official API for Deepseek v3 has been discontinued, we instead use the API provided by a third-
 363 party server deployment. This study employs several widely used and representative benchmark
 364 datasets—MMLU, GPQA, and AIME—as the basis for evaluation. Note that due to the page limit,
 365 the results of MMLU and GPQA are listed in Appendix A.4. MMLU covers 57 subjects and contains
 366 a large collection of multiple-choice questions, widely used to assess models’ general knowledge
 367 and reasoning abilities. GPQA consists of 448 challenging multiple-choice questions developed and
 368 validated by experts in biology, physics, and chemistry, designed to evaluate AI models’ reasoning
 369 ability on complex scientific problems. AIME is a highly selective U.S. high school mathematics
 370 competition, well known for its challenging problems that test deep mathematical reasoning. It is
 371 worth noting that our methodology is not tied to any specific benchmark and can be applied to the
 372 evaluation of any LLM.

373 5.1 THE NECESSITY OF CONSTRUCTING MES AND A-MES

374 This section demonstrates that LLMs exhibit significant performance variations across different
 375 MES and A-MES configurations, thereby underscoring the inadequacy of single-configuration eval-
 376 uations in accurately capturing their true capabilities.

378 Table 3: Performance Rankings of LLMs (deepseek = deepseek-v3, doubaao = doubaao-1.5-pro-32k,
 379 mistralL = mistral-large, mistralM = mistral-medium, kimi = moonshot-v1-8k, qwenP = qwen-plus)
 380


Type	1.	2	3	4	5	6	7	8	9
Original	deepseek(0.4)	gpt4.1(0.07)	doubaao(0)	gpt3.5(0)	mistralL(0)	mistralM(0)	kimi(0)	qwenP(0)	qwen2.5(0)
A-MES	deepseek(0.45)	qwenP(0.43)	mistralL(0.38)	gpt4.1(0.37)	mistralM(0.36)	doubaao(0.34)	qwen2.5(0.26)	kimi(0.18)	gpt3.5(0.13)
MES	doubaao(0.25)	deepseek(0.21)	qwen2.5(0.21)	mistralL(0.20)	mistralM(0.20)	qwenP(0.18)	gpt4.1(0.16)	gpt3.5(0.10)	kimi(0.08)

383 Note: Models are sorted alphabetically by name when accuracy equals zero.
 384


385 In the MES experiments, we conducted 500 random samplings without replacement from the MES
 386 configuration space described in Section 4.1. The specific components and their corresponding value
 387 ranges are summarized in Table 2. **We determined 500 as a conservative sample-size upper bound**
 388 **by combining a convergence-based stopping rule with LLN-guided estimates.**
 389

390 In the A-MES experiments, to verify the effectiveness of the Augmented Minimal Evaluation System
 391 (A-MES) proposed in Section 4.2 and comprehensively evaluate the performance of LLMs across
 392 diverse task scenarios, we conducted a comparative analysis of their performance on two types of
 393 datasets: the original AIME workload and four augmented workloads derived from this original
 394 workload. For the analogical transformation workload, we employed two specific methods: the first
 395 involves inserting redundant information into the stem of the original question, information that is
 396 irrelevant to the problem-solving logic and methods yet consistent with the question scenario, to
 397 interfere with the output results of LLMs; The second method involves numeric substitutions. For
 398 novel (newly-synthesized) workloads, this study designs two core strategies: the first is a knowledge
 399 point-based question generation strategy, which specifically generates new tasks based on the core
 400 knowledge points covered in the original questions and combined with the conceptual system and
 401 expression paradigm of relevant textbook chapters; The second is an adaptation and transformation
 402 strategy based on college entrance examination (gaokao) questions, which involves selecting the
 403 latest gaokao questions that match the target knowledge points and generating new tasks by adjusting
 404 the scenario of the question, questioning logic, and other aspects.

405 The experimental results of this study are presented in Figure 4. As shown in Figure 4, significant
 406 variations in accuracy trends are observed across different models and configuration spaces. For
 407 instance, the accuracy of the deepseek-v3 model fluctuates within a range of 0 to 0.78 under MES
 408 experiments, and from 0.23 to 0.7 under A-MES experiments. As shown in Table 3, we have also
 409 generated performance rankings for large models based on the original workload, MES workload,
 410 and A-MES workload. For the MES and A-MES scenarios, we employ their average accuracy as
 411 the performance metric for the LLMs. It is crucial to note that when the accuracy is zero, we sort
 412 the models alphabetically based on their names. Drawing insights from the rankings, we observe
 413 three key conclusions: first, the original evaluation methodology demonstrates limited effectiveness
 414 in benchmarking large language models (LLMs) due to its inability to distinguish performance be-
 415 yond two models achieving non-zero accuracy scores; second, DeepSeek consistently outperforms
 416 all competing models across diverse evaluation conditions, underscoring its robustness and su-
 417 perior generalization capabilities; third, model performance rankings exhibit contextual sensitivity,
 418 as evidenced by Doubaao’s inferior performance relative to DeepSeek in both Original and A-MES
 419 workloads, yet its top-ranking achievement in MES, thereby highlighting the non-transitive nature
 420 of LLM performance across varying task formulations and data distributions.



421 (a) Distribution of Model Accuracies on MES
 422




423 (b) Distribution of Model Accuracies on A-MES
 424

425 Figure 4: Distribution of Model Accuracies
 426

432 Table 4: ANOVA results on DeepSeek-V3 (sorted by effect size in descending order)
433

434 Factor	435 Effect Size η^2	436 p -value
Question Format	0.399643	0.000
Question Format - COT	0.161394	0.000
COT	0.080156	0.000
max_tokens	0.028099	0.000
Question Format - Shot	0.011101	0.000
Language - Question Format	0.008178	0.006
COT - max_tokens	0.006721	0.010
Language - COT	0.004345	0.038
Multi Turn - max_tokens	0.003841	0.050
Language	0.003841	0.046
Shot - max_tokens	0.003669	0.046
Question Format - max_tokens	0.002687	0.100
Language - Multi Turn	0.002600	0.066
temperature - top_p	0.002082	0.178
Question Format - Multi Turn	0.001321	0.244

443
444 5.2 QUANTIFY THE CONTRIBUTION OF EACH INDISPENSABLE COMPONENT TO OVERALL
445 PERFORMANCE VARIANCE446
447 In LLM evaluation, a key challenge lies in effectively evaluating the contribution of each components
448 illustrated in Fig. 1 to overall performance variance. Given the enormous number of possible
449 EC configuration combinations, exhaustively testing every configuration is computationally infea-
450 sible. To address this, we selected a limited number of experimental points from the full space,
451 allowing us to systematically and evenly examine the effects of multiple components and their lev-
452 els on performance with significantly fewer trials. This design reduces experimental cost while
453 maintaining scientific rigor and representativeness.454 To quantify the proportion of performance variance explained by each MES component, we adopted
455 an analysis of variance (ANOVA) approach. Specifically, for component $C_1 - C_{10}$, we selected two
456 levels (“high” and “low”) within their respective ranges, [with these ranges given in Table 2](#), thereby
457 constructing a subspace of size $2^{10} = 1024$. For the Language component, we selected Chinese
458 and English, while for three-valued components we used their maximum and minimum values.
459 Within this subspace, variance decomposition was used to quantify the contributions of different
460 components and their interactions to variations in accuracy. Moreover, we employed a permutation
461 test to evaluate statistical significance, enabling a more robust assessment of component importance
462 without relying on additional distributional assumptions. This procedure yields both the relative
importance and the statistical significance of all components.463 Taking the DeepSeek-V3 model as an example, Table 4 reports the main effects and two-way in-
464 teractions that significantly influence its accuracy on the AIME’24 benchmark, with the complete
465 ANOVA results provided in Appendix A.5. Overall, Question type, COT, max_tokens, and their
466 interactions with other components exhibit the most significant effects. Shot, Multi turn, and Lan-
467 guage also show significant effects, while the remaining components have only limited impact.468 Consistent patterns were observed across other LLMs (see Appendix A.5). Using $p < 0.05$ as
469 the significance threshold, we found that the main effects of Question format and COT, or their
470 interactions with other components, were consistently significant across all LLMs. Furthermore,
471 max_tokens, Shot, and Multi turn also reached significance for the vast majority of models. In ad-
472 dition to these five core components, Language, top_p, and temperature were significant for some
473 models. It is worth noting that for the remaining two components, Question Paraphrase and pres-
474 ence_penalty, the p -values did not meet the significance threshold, but reached 0.19 and 0.16, re-
475 spectively, on GPT-4.1. This suggests that they may exert some influence on model performance,
476 although the evidence is not sufficient for a definitive conclusion.477
478 5.3 COMPARE LLM EVALUATOLOGY WITH TRADITIONAL LLM EVALUATION METHODS
479 AND ATTRIBUTE THE PERFORMANCE DIFFERENCES TO SPECIFIC COMPONENTS480
481 This section demonstrates that evaluating models under a single configuration fails to capture their
482 true capabilities, while LLM evaluatology not only yields results in strong agreement with the
483 ground truth, but also attributes the performance differences to specific components.484 Based on the randomly sampled data collected from the complete configuration space [spanned by](#)
485 [the components in Table 2](#), we estimated the overall average accuracies of different models on the
same benchmark using their 95% and 99% confidence intervals. As illustrated in Figure 5a, we



486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

(a) Accuracy confidence intervals (b) Accuracy difference confidence intervals of different models on AIME of different LLMs on AIME

Figure 5: Comparison of LLM Evaluatology and Traditional Method on AIME

report the performance of different models on AIME’2024, where the purple dots denote the test results under the commonly adopted default setting, using the original workloads without optimized prompting methods and with default decoding parameters. It can be observed that the purple dots are far from the confidence intervals (interval estimation of the population mean) obtained through random sampling, showing that evaluating a model under a single configuration is unreliable. **Note that accuracy values of 0 in Figure 5a are not due to missing data, but to the high difficulty of AIME problems, which are challenging even for human contestants.**

Figure 5b further presents, on AIME’2024, the accuracy differences between two models under the default configuration, together with the 95% and 99% confidence intervals constructed from accuracy differences observed across sampled equivalent evaluation configurations. In 10 cases, the confidence intervals and the default accuracy differences fall on opposite sides of the zero line, revealing contradictions in the conclusions regarding model superiority. For instance, the 99% confidence interval for the mean accuracy difference between Doubao-1.5-pro and GPT-4.1 lies entirely above the zero line, implying that overall Doubao-1.5-pro outperforms GPT-4.1. However, if one were to rely on the result of a single experiment under the default configuration, the accuracy difference would fall below the zero line, leading instead to the opposite conclusion that GPT-4.1 outperforms Doubao-1.5-pro. This “conclusion reversal” highlights the limitations of relying solely on single-configuration testing. More detailed results on additional benchmarks including MMLU and GPQA can be found in the Appendix.

Furthermore, we selected the five most influential components for a cost-efficient accuracy test on each LLM, based on the ANOVA data in Section 5.2. We then constructed the configuration subspace restricted to these components and conducted exhaustive testing within this subspace. The mean performance obtained was taken as a “restricted-space ground truth.” As shown by the red diamond in Figure 5a, for all models, this reference truth fell within the confidence intervals estimated from random sampling, thereby demonstrating both the validity and the robustness of the proposed LLM evaluatology method.

6 CONCLUSION

LLM Evaluatology establishes a principled methodology for assessing LLMs through an Augmented Minimal Evaluation System (A-MES), explicitly accounting for both intrinsic model capabilities and the many confounding components that shape observed performance, thereby enabling accurate attribution of performance differences to their true sources. Our analysis reveals that meaningful evaluation of LLMs requires careful consideration of both workload heterogeneity and the vast space of evaluation condition (EC) configurations. We advocate for the adoption of evaluatology as a foundational paradigm, encouraging the community to develop richer workload augmentation strategies and robust evaluation practices that mirror the complexity of actual deployment scenarios.

540 REFERENCES
541

542 Aider. Aider llm leaderboards. <https://aider.chat/docs/leaderboards/>, 2025. Accessed: 2025-09-13.

543

544 AIME. AIME Problems and Solutions. https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions, 2025. Accessed: 2025-09-05.

545

546 Chenxin An, Shanshan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. L-eval: Instituting standardized evaluation for long context language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14388–14411, 2024.

547

548

549

550

551 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.

552

553

554 Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhdian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, et al. Longbench: A bilingual, multitask benchmark for long context understanding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3119–3137, 2024.

555

556

557

558 Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, et al. Multiiple: A scalable and polyglot approach to benchmarking neural code generation. *IEEE Transactions on Software Engineering*, 49(7):3675–3691, 2023.

559

560

561

562

563 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

564

565

566 Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael Jordan, Joseph E Gonzalez, et al. Chatbot arena: An open platform for evaluating llms by human preference. In *Forty-first International Conference on Machine Learning*.

567

568

569

570 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

571

572

573

574 Amit Das, Zheng Zhang, Najib Hasan, Souvika Sarkar, Fatemeh Jamshidi, Tathagata Bhattacharya, Mostafa Rahgouy, Nilanjana Raychawdhary, Dongji Feng, Vinija Jain, et al. Investigating annotator bias in large language models for hate speech detection. *arXiv preprint arXiv:2406.11109*, 2024.

575

576

577

578 Aryo Pradipta Gema, Joshua Ong Jun Leang, Giwon Hong, Alessio Devoto, Alberto Carlo Maria Mancino, Rohit Saxena, Xuanli He, Yu Zhao, Xiaotang Du, Mohammad Reza Ghasemi Madani, et al. Are we done with mmlu? In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 5069–5096, 2025.

579

580

581

582

583 Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3309–3326, 2022.

584

585

586

587 Jia He, Mukund Rungta, David Koleczek, Arshdeep Sekhon, Franklin X Wang, and Sadid Hasan. Does prompt formatting have any impact on llm performance? *arXiv preprint arXiv:2411.10541*, 2024.

588

589

590 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.

591

592

593

594 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob
 595 Steinhardt. Measuring massive multitask language understanding. In *International Conference*
 596 *on Learning Representations*, 2020.

597

598 Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junpeng Liu,
 599 Chuancheng Lv, Yikai Zhang, Yao Fu, et al. C-eval: A multi-level multi-discipline chinese eval-
 600 uation suite for foundation models. *Advances in Neural Information Processing Systems*, 36:
 601 62991–63010, 2023.

602

603 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
 604 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
 605 evaluation of large language models for code. In *The Thirteenth International Conference on*
 606 *Learning Representations*.

607

608 Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R
 609 Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *The Twelfth*
 610 *International Conference on Learning Representations*.

611

612 Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy
 613 Baldwin. Cmmlu: Measuring massive multitask language understanding in chinese. In *Findings*
 614 *of the Association for Computational Linguistics ACL 2024*, pp. 11260–11285, 2024.

615

616 Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gon-
 617 zalez, and Ion Stoica. From crowdsourced data to high-quality benchmarks: Arena-hard and
 618 benchbuilder pipeline. In *Forty-second International Conference on Machine Learning*.

619

620 Jiacheng Liu, Mayi Xu, Qiankun Pi, Wenli Li, Ming Zhong, Yuanyuan Zhu, Mengchi Liu, and
 621 Tieyun Qian. Format as a prior: Quantifying and analyzing bias in llms for heterogeneous data.
 622 *arXiv preprint arXiv:2508.15793*, 2025.

623

624 Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan,
 625 Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around
 626 player? In *European conference on computer vision*, pp. 216–233. Springer, 2024.

627

628 Do Xuan Long, Hai Nguyen Ngoc, Tiviatis Sim, Hieu Dao, Shafiq Joty, Kenji Kawaguchi, Nancy F
 629 Chen, and Min-Yen Kan. Llms are biased towards output formats! systematically evaluating and
 630 mitigating output format bias of llms. *arXiv preprint arXiv:2408.08656*, 2024.

631

632 Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-
 633 Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning
 634 of foundation models in visual contexts. In *The Twelfth International Conference on Learning*
 635 *Representations*.

636

637 OpenCompass. Compassarena. <https://opencompass.org/arena>, 2025. Accessed:
 638 2025-09-08.

639

640 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Di-
 641 rani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a bench-
 642 mark. In *First Conference on Language Modeling*.

643

644 Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adver-
 645 sarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.

646

647 Chufan Shi, Haoran Yang, Deng Cai, Zhisong Zhang, Yifan Wang, Yujiu Yang, and Wai Lam. A
 648 thorough examination of decoding methods in the era of llms. *arXiv preprint arXiv:2402.06925*,
 649 2024.

650

651 Ved Sirdeshmukh, Kaustubh Deshpande, Johannes Mols, Lifeng Jin, Ed-Yeremai Cardona, Dean
 652 Lee, Jeremy Kritz, Willow Primack, Summer Yue, and Chen Xing. Multichallenge: A realis-
 653 tic multi-turn conversation evaluation benchmark challenging to frontier llms. *arXiv preprint*
 654 *arXiv:2501.17399*, 2025.

648 Mirac Suzgun, Nathan Scales, Nathanael Schärlí, Sebastian Gehrmann, Yi Tay, Hyung Won Chung,
 649 Akanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. Challenging big-bench tasks
 650 and whether chain-of-thought can solve them. In *ACL (Findings)*, 2023.

651

652 Chenxi Wang, Lei Wang, Wanling Gao, Yikang Yang, Yutong Zhou, and Jianfeng Zhan. Achieving
 653 consistent and comparable cpu evaluation outcomes. *arXiv preprint arXiv:2411.08494*, 2024a.

654 Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming
 655 Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and challenging multi-
 656 task language understanding benchmark. *Advances in Neural Information Processing Systems*,
 657 37:95266–95290, 2024b.

658

659 Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Benjamin Feuer, Siddhartha Jain, Ravid
 660 Shwartz-Ziv, Neel Jain, Khalid Saifullah, Sreemanti Dey, et al. Livebench: A challenging,
 661 contamination-limited llm benchmark. In *The Thirteenth International Conference on Learning
 662 Representations*.

663

664 Liang Xu, Anqi Li, Lei Zhu, Hang Xue, Changtai Zhu, Kangkang Zhao, Haonan He, Xuanwei
 665 Zhang, Qiyue Kang, and Zhenzhong Lan. Superclue: A comprehensive chinese large language
 666 model benchmark. *arXiv preprint arXiv:2307.15020*, 2023.

667

668 Zhengxin Yang, Wanling Gao, Chunjie Luo, Lei Wang, and Jianfeng Zhan. Quality at the tail.
 669 *CoRR*, abs/2212.13925, 2022. doi: 10.48550/ARXIV.2212.13925. URL <https://doi.org/10.48550/arXiv.2212.13925>.

670

671 Howard Yen, Tianyu Gao, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat,
 672 and Danqi Chen. Helmet: How to evaluate long-context language models effectively and thor-
 673oughly. International Conference on Learning Representations (ICLR), 2025.

674

675 Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens,
 676 Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multi-
 677 modal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF
 678 Conference on Computer Vision and Pattern Recognition*, pp. 9556–9567, 2024.

679

680 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-
 681 chine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association
 682 for Computational Linguistics*, pp. 4791–4800, 2019.

683

684 Jianfeng Zhan. A short summary of evaluatology: The science and engineering of evaluation, 2024.

685

686 Jianfeng Zhan, Lei Wang, Wanling Gao, Hongxiao Li, Chenxi Wang, Yunyou Huang, Yatao Li,
 687 Zhengxin Yang, Guoxin Kang, Chunjie Luo, et al. Evaluatology: The science and engineering of
 688 evaluation. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 4(1):100162,
 689 2024.

690

691 Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Khai Hao, Xu Han,
 692 Zhen Leng Thai, Shuo Wang, Zhiyuan Liu, et al. ∞ bench: Extending long context evaluation
 693 beyond 100k tokens. *arXiv preprint arXiv:2402.13718*, 2024a.

694

695 Zhixin Zhang, Leqi Lei, Lindong Wu, Rui Sun, Yongkang Huang, Chong Long, Xiao Liu, Xuanyu
 696 Lei, Jie Tang, and Minlie Huang. Safetybench: Evaluating the safety of large language models.
 697 In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics
 (Volume 1: Long Papers)*, pp. 15537–15553, 2024b.

698

699 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny
 700 Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint
 701 arXiv:2311.07911*, 2023.

702

703 Kaijie Zhu, Jiaao Chen, Jindong Wang, Neil Zhenqiang Gong, Diyi Yang, and Xing Xie. Dyval:
 704 Dynamic evaluation of large language models for reasoning tasks. In *The Twelfth International
 705 Conference on Learning Representations*.

A APPENDIX

A.1 THE USE OF LARGE LANGUAGE MODELS (LLMs)

During the manuscript preparation, we leveraged large language models (LLMs) to assist in refining and polishing the text. Specifically, the LLM was used to improve sentence clarity and enhance linguistic fluency, while all scientific content, reasoning, and results were independently authored and verified by the researchers. This approach facilitated more concise and readable presentation without affecting the technical accuracy.

A.2 EVALUATION SETTING ON DIFFERENT BENCHMARKS

Table 5: Evaluation Settings Reported in Technical Reports of Different LLMs. (Lang. = Language, Q-type = Question type, Para. = Paraphrase, M-turn = Multi turn, Temp. = Temperature, PP = presence_penalty, MaxTok = max_tokens)

(a) Evaluation Settings on AIME’2024

Model	Lang.	Q-type	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
DeepSeek-R1	english	origin	origin	0	yes	x	0.6	0.95	x	32768
DeepSeek-V3	english	origin	origin	x	x	x	0.7	x	x	8192
Kimi K2	english	origin	origin	x	no	x	0.0	fixed	x	8192
Kimi K1.5	english	origin	origin	x	yes	x	x	x	x	x
Qwen2					Not evaluated on AIME					
Qwen2.5					Not evaluated on AIME					
Qwen3	english	origin	origin	x	no	x	0.7	0.8	1.5	32768
GPT-4					Not evaluated on AIME					
GPT-4.1	english	origin	origin	x	x	x	x	x	x	x
GPT-5					Not evaluated on AIME					
Claude Opus 4					Not evaluated on AIME					
Mistral Small3.1					Not evaluated on AIME					
Mistral Medium3					Not evaluated on AIME					
Mistral Large2					Not evaluated on AIME					

(b) Evaluation Settings on MMLU

Model	Lang.	Q-type	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
DeepSeek-R1	english	origin	origin	0	yes	x	0.5	x	x	1024
DeepSeek-V3	english	origin	origin	0	yes	x	0.5	x	x	1024
Kimi K2	english	origin	origin	x	no	x	0.0	fixed	x	8192
Kimi K1.5	english	origin	origin	x	yes	x	x	x	x	x
Qwen2	english	origin	origin	5	x	x	x	x	x	x
Qwen2.5	english	origin	origin	5	x	x	x	x	x	x
Qwen3	english	origin	origin	5	x	x	x	x	x	x
GPT-4	multiple	origin	origin	5/3	no	x	x	x	x	x
GPT-4.1	multiple	origin	origin	x	x	x	x	x	x	x
GPT-5					Not evaluated on MMLU					
Claude Opus 4	multiple	origin	origin	x	yes/no	x	x	x	x	x
Mistral Small3.1	english	origin	origin	x	x	x	x	x	x	x
Mistral Medium3					Not evaluated on MMLU					
Mistral Large2	multiple	origin	origin	x	x	x	x	x	x	x

(c) Evaluation Settings on GPQA

Model	Lang.	Q-type	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
DeepSeek-R1	english	origin	origin	0	yes	x	0.5	x	x	1024
DeepSeek-V3	english	origin	origin	0	yes	x	0.5	x	x	1024
Kimi K2	english	origin	origin	x	no	x	0.0	fixed	x	8192
Kimi K1.5	english	origin	origin	x	yes	x	x	x	x	x
Qwen2	english	origin	origin	x	x	x	x	x	x	x
Qwen2.5	english	origin	origin	x	x	x	x	x	x	x
Qwen3	english	origin	origin	x	yes/no	x	0.6/0.7	0.95/0.8	0/1.5	32768
GPT-4					Not evaluated on GPQA					
GPT-4.1	english	origin	origin	x	x	x	x	x	x	x
GPT-5	english	origin	origin	x	0/1	x	x	x	x	x
Claude Opus 4	english	origin	origin	x	0/1	x	x	x	x	x
Mistral Small3.1	english	origin	origin	x	x	x	x	x	x	x
Mistral Medium3	english	origin	origin	5	1	x	x	x	x	x
Mistral Large2					Not evaluated on GPQA					

(d) Evaluation Settings on MATH

Model	Lang.	Q-type	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
DeepSeek-R1	english	origin	origin	0/8	yes/no	x	0	x	x	32768
DeepSeek-V3	english	origin	origin	0/8	yes/no	x	0	x	x	8192
Kimi K2	english	origin	origin	x	no	x	0.0	fixed	x	8192
Kimi K1.5	english	origin	origin	x	yes	x	x	x	x	x
Qwen2	english	origin	origin	x	x	x	x	x	x	x
Qwen2.5	english	origin	origin	x	x	x	x	x	x	x
Qwen3	english	origin	origin	x	yes/no	x	0.6/0.7	0.95/0.8	0/1.5	32768
GPT-4					Not evaluated on MATH					
GPT-4.1					Not evaluated on MATH					
GPT-5					Not evaluated on MATH					
Claude Opus 4					Not evaluated on MATH					
Mistral Small3.1	english	origin	origin	x	x	x	x	x	x	x
Mistral Medium3	english	origin	origin	0	0	x	x	x	x	x
Mistral Large2	english	origin	origin	0	0	x	x	x	x	x

(e) Evaluation Settings on SWE-bench

Model	Lang.	Q-type	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
DeepSeek-R1	english	origin	origin	x	x	x	0.8	x	x	x
DeepSeek-V3	english	origin	origin	x	x	x	0.8	x	x	x
Kimi K2	english	origin	origin	x	no	x	0.0	fixed	x	8192/16384
Kimi K1.5					Not evaluated on SWE-bench					
Qwen2					Not evaluated on SWE-bench					
Qwen2.5(pre)					Not evaluated on SWE-bench					
Qwen3(pre)					Not evaluated on SWE-bench					
GPT-4					Not evaluated on SWE-bench					
GPT-4.1	english	origin	origin	x	x	x	x	x	x	x
GPT-5	english	origin	origin	x	0/1	x	x	x	x	x
Claude Opus 4	english	origin	origin	x	0/1	x	x	0.95	x	x
Mistral Small3.1					Not evaluated on SWE-bench					
Mistral Medium3					Not evaluated on SWE-bench					
Mistral Large2					Not evaluated on SWE-bench					

(f) Evaluation Settings on IFEval

Model	Lang.	Q-type	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
DeepSeek-R1	english	origin	origin	0	0	0	x	x	x	x
DeepSeek-V3	english	origin	origin	0	0	0	x	x	x	x
Kimi K2	english	origin	origin	x	no	x	0.0	fixed	x	8192
Kimi K1.5	english	origin	origin	x	yes	x	x	x	x	x
Qwen2	english	origin	origin	x	x	x	x	x	x	x
Qwen2.5	english	origin	origin	x	x	x	x	x	x	x
Qwen3	english	origin	origin	x	yes/no	x	0.6/0.7	0.95/0.8	0/1.5	32768
GPT-4					Not evaluated on IFEval					
GPT-4.1	english	origin	origin	x	x	x	x	x	x	x
GPT-5					Not evaluated on IFEval					
Claude Opus 4					Not evaluated on IFEval					
Mistral Small3.1					Not evaluated on IFEval					
Mistral Medium3	english	origin	origin	0	0	x	x	x	x	x
Mistral Large2					Not evaluated on IFEval					

(g) Evaluation Settings on Arena-Hard

Model	Lang.	Q-type	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
DeepSeek-R1	english	origin	origin	0	0	0	config	default	default	user-set
DeepSeek-V3	english	origin	origin	0	0	0	config	default	default	user-set
Kimi K2	english	origin	origin	x	no	x	0.0	fixed	x	8192
Kimi K1.5					Not evaluated on Arena-Hard					
Qwen2	english	origin	origin	x	x	x	x	x	x	x
Qwen2.5	english	origin	origin	x	x	x	x	x	x	x
Qwen3	english	origin	origin	x	yes/no	x	0.6/0.7	0.95/0.8	0/1.5	32768
GPT-4					Not evaluated on Arena-Hard					
GPT-4.1					Not evaluated on Arena-Hard					
GPT-5					Not evaluated on Arena-Hard					
Claude Opus 4					Not evaluated on Arena-Hard					
Mistral Small3.1					Not evaluated on Arena-Hard					
Mistral Medium3					Not evaluated on Arena-Hard					
Mistral Large2	english	origin	origin	x	x	x	x	x	x	x

(h) Evaluation Settings on HumanEval

Model	Lang.	Q-type	Para.	Shot	COT	M-turn	Temp.	top-p	PP	MaxTok
DeepSeek-R1	english	origin	origin	0	0	0	varied	0.95	x	32768
DeepSeek-V3	english	origin	origin	0	0	0	varied	0.95	x	8192
Kimi K2					Not evaluated on HumanEval					
Kimi K1.5	english	origin	origin	x	yes	x	x	x	x	x
Qwen2	english	origin	origin	x	x	x	x	x	x	x
Qwen2.5	english	origin	origin	x	x	x	x	x	x	x
Qwen3					Not evaluated on HumanEval					
GPT-4	english	origin	origin	0	0	x	0.3	x	x	x
GPT-4.1					Not evaluated on HumanEval					
GPT-5					Not evaluated on HumanEval					
Claude Opus 4					Not evaluated on HumanEval					
Mistral Small3.1	english	origin	origin	x	x	x	x	x	x	x
Mistral Medium3	english	origin	origin	0	0	x	x	x	x	x
Mistral Large2	english	origin	origin	x	x	x	x	x	x	x

A.3 A-MES CONSTRUCTION PIPELINE

Analogical: Distractor Insertion

For distractor insertion, we define three explicit, controllable categories of redundancy and implement all instances via LLM prompting. To ensure that the inserted distractors strictly follow our predefined specifications, we empirically test several candidate LLMs and choose the one that most consistently adheres to these constraints (GPT-5). This selection is made solely to guarantee transformation fidelity rather than to compare model capabilities. For each item to be transformed, the chosen LLM is invoked through an API and, guided by our structured prompts, automatically produces and inserts the required redundant content. The concrete implementation is as follows.

810 (1) Context-irrelevant redundancy.
 811
 812 • Provide the LLM with an example containing an original question and a version with added
 813 context-irrelevant redundancy.
 814 • Instruct the LLM to insert one sentence at a random position that is completely unrelated to the
 815 target question.

Algorithm 1 Context-irrelevant redundancy insertion

```

1: function INSERTCONTEXTIRRELEVANTDISTRCTOR(problem_text)
2:   EXAMPLE_PAIR  $\leftarrow$  (orig_example, example_with_irrelevant_context)
3:   PROMPT  $\leftarrow$  BUILDPROMPTIRRELEVANT(EXAMPLE_PAIR, problem_text)
4:   RESPONSE  $\leftarrow$  LLM_CALL(PROMPT)
5:   transformed_text  $\leftarrow$  PARSETRANFORMEDPROBLEM(RESPONSE)
6:   if not BASICSANITYCHECK(problem_text, transformed_text) then
7:     return FAILURE
8:   end if
9:   return transformed_text
10: end function

```

827 (2) Context-relevant, explanatory redundancy.
 828
 829 • Provide the LLM with an example of an original question and a version with added explanatory
 830 redundancy.
 831 • Instruct the LLM to insert a redundant sentence at a random position in each target question that
 832 explains a concept already appearing in the target question.

Algorithm 2 Context-relevant explanatory redundancy insertion

```

1: function INSERTEXPLANATORYDISTRCTOR(problem_text)
2:   EXAMPLE_PAIR  $\leftarrow$  (orig_example, example_with_explanatory_sentence)
3:   PROMPT  $\leftarrow$  BUILDPROMPTEXPLANATORY(EXAMPLE_PAIR, problem_text)
4:   RESPONSE  $\leftarrow$  LLM_CALL(PROMPT)
5:   transformed_text  $\leftarrow$  PARSETRANFORMEDPROBLEM(RESPONSE)
6:   if not BASICSANITYCHECK(problem_text, transformed_text) then
7:     return FAILURE
8:   end if
9:   return transformed_text
10: end function

```

844 (3) Context-relevant, misleading redundancy.
 845
 846 • Provide the LLM with an example containing an original question and a version with added mis-
 847 leading but logically related redundancy.
 848 • Supply the model with the correct answer and several correct solution approaches, and instruct it
 849 to avoid directly hinting at these correct strategies when crafting the misleading cue. The official
 850 answer and solution approaches are provided by the user, and providing solution approaches is
 851 optional.
 852 • Instruct the model to insert a redundant sentence that nudges the reader toward an incorrect strat-
 853 egy or line of reasoning, without explicitly revealing that it is “misleading” or “distracting”.

Algorithm 3 Context-relevant misleading redundancy insertion

```

1: function INSERTMISLEADINGDISTRCTOR(problem_text, answer_gold, solution_sketches)
2:   EXAMPLE_PAIR  $\leftarrow$  (orig_example, example_with_misleading_sentence)
3:   PROMPT  $\leftarrow$  BUILDPROMPTMISLEADING(example_pair = EXAMPLE_PAIR, target_problem = prob-  

lem_text, answer_gold = answer_gold, solution_sketches = solution_sketches)
4:   RESPONSE  $\leftarrow$  LLM_CALL(PROMPT)
5:   transformed_text  $\leftarrow$  PARSETRANFORMEDPROBLEM(RESPONSE)
6:   if not BASICSANITYCHECK(problem_text, transformed_text) then
7:     return FAILURE
8:   end if
9:   return transformed_text
10: end function

```

864 In practice, the selected LLM produces variations that are more diverse and linguistically natural
 865 than manual editing. In particular, its context-relevant misleading redundancies tend to hint at in-
 866 correct heuristics in a more subtle way than hand-written versions, while still strictly adhering to the
 867 predefined category constraints. The entire process involves no per-item manual editing. The three
 868 examples of redundancy for three types generated by the above procedure are illustrated as follows:
 869

870 1. Context-irrelevant redundancy example:

871 *The weather today seems quite pleasant, and it might be a great day for a picnic.*
 872 Find the number of triples of nonnegative integers (a, b, c) satisfying $a + b + c =$
 873 300 and $a^2b + a^2c + b^2a + b^2c + c^2a + c^2b = 6,000,000$.
 874

875 Here, the weather is entirely unrelated to the math content.

876 2. Context-relevant, explanatory redundancy example:

877 There exist real numbers x and y , both greater than 1, such that $\log_x(y^x) =$
 878 $\log_y(x^{4y}) = 10$. *A logarithm is a way to express how many times a base must be*
 879 *multiplied by itself to get a certain number.* Find xy .
 880

881 The added sentence explains the notion of a logarithm while leaving the underlying problem un-
 882 changed.

883 3. Context-relevant, misleading redundancy example:

884 Alice and Bob play the following game. A stack of n tokens lies before them. The
 885 players take turns with Alice going first. On each turn, the player removes either
 886 1 token or 4 tokens from the stack. *Many players adopt a greedy approach here:*
 887 *always take 4 whenever possible to shorten the game and restrict the opponent's*
 888 *replies.* Whoever removes the last token wins. Find the number of positive in-
 889 *tegers n less than or equal to 2024 for which there exists a strategy for Bob that*
 890 *guarantees that Bob will win the game regardless of Alice's play.*
 891

892 The extra sentence about the “greedy approach” is logically related to the game but suggests a flawed
 893 strategy, intentionally nudging the solver toward an incorrect line of reasoning.
 894

895 **Analogical: Numeric Substitutions**

896 For numeric substitutions, we use a uniform pipeline built around LLM-generated Python solvers
 897 and automatic verification scripts, rather than manually changing a few numbers:

- 901 • We first call an LLM to extract the primary knowledge points tested by the original problem,
 902 and query a pre-constructed formula library indexed by knowledge point to retrieve potentially
 903 relevant formulas.
- 904 • We feed the original problem, its official answer, the retrieved formulas, and (where available)
 905 multiple correct solution sketches into the LLM. The official answer and solution sketches are
 906 provided by the user, and providing solution sketches is optional.
- 907 • The LLM is prompted to:
 - 908 – analyze the problem's solution strategy, using the provided solution sketches when available;
 - 909 – write a Python solution program where problem-specific numbers are extracted as explicit vari-
 ables with reasonable value ranges.
- 910 • We then ask another LLM to inspect the generated Python code and verify that it implements
 911 a general computational procedure for solving the problem, rather than relying on hard-coded
 912 instance-specific outputs or trivial pattern matching.
- 913 • We import the LLM-generated Python code as a local module and call its `solve()` function with
 914 the original numeric values as inputs, checking whether the resulting output matches the official
 915 answer.
- 916 • If the code fails (wrong answer or runtime error), we return the error message and incorrect output
 917 to the LLM, asking it to refine the code; we repeat this refinement–verification loop for up to five
 attempts and keep the Python code if it passes on the original instance.

918 After obtaining a correct solver, we automatically sample new numeric configurations within the
 919 validated ranges to generate analogical variants of the same underlying problem.
 920

921 This “knowledge-point extraction → formula retrieval → analyze → code → verify → resample”
 922 pipeline is identical across all problems.
 923

Algorithm 4 Numeric substitutions

```

926 1: function BUILDNUMERICSOLVER(problem_text, answer_gold, solution_sketches, max_iter = 5, max_refine
927   = 5)
928 2:   knowledge_points ← LLM_EXTRACTKNOWLEDGEPOINTS(problem_text)
929 3:   retrieved_formulas ← RETRIEVE_FORMULAS(knowledge_points)
930 4:   history ← EMPTY_LIST
931 5:   for iter = 1 to max_iter do
932     PROMPT ← BUILDPROMPTCODEGEN(problem_text, answer_gold, solution_sketches, retrieved_formulas)
933     APPEND(history, (PROMPT, NONE))
934     (CODE, param, value_ranges) ← LLM_CALL(PROMPT)
935     is_hard_code ← LLM_HARD_CODE_CHECK(CODE)
936     if is_hard_code then
937       continue
938     end if
939     for refine_step = 0 to max_refine do
940       (output, error) ← RUN PYTHON(CODE, input = param.value)
941       APPEND(history, (CODE, (output, error)))
942       if error = NONE and VERIFY(output, answer_gold) then
943         return (CODE, param, value_ranges)
944       end if
945       if refine_step = max_refine then
946         break
947       end if
948     end for
949   end for
950   return FAILURE
951 end function
952 function GENERATENUMERICVARIANTS(problem_text, answer_gold, solution_sketches)
953   (solver_code, param, param_ranges) ← BUILDNUMERICSOLVER(problem_text, answer_gold, solution_sketches)
954   if solver_code = FAILURE then
955     return NONE
956   end if
957   new_param ← SAMPLE(param_ranges)
958   new_problem_text ← INSTANTIATENUMERICPROBLEMTTEXT(problem_text, param, new_param)
959   (output, error) ← RUN PYTHON(solver_code, input = new_param)
960   new_answer_gold ← output
961   return (new_problem_text, new_answer_gold)
962 end function

```

963 An example of numeric substitutions is given as follows.
 964

- **Original:**

965 Find the largest possible real part of $(75 + 117i)z + \frac{96+144i}{z}$ where z is a complex
 966 number with $|z| = 4$. A common shortcut is to take z to be a positive real number, since
 967 for a fixed modulus the real part is often largest when the argument of z is zero.

- **Numeric variant:**

970 Find the largest possible real part of $(100 + 112i)z + \frac{60+144i}{z}$ where z is a complex
 971 number with $|z| = 4$. A common shortcut is to take z to be a positive real number, since
 972 for a fixed modulus the real part is often largest when the argument of z is zero.

972 **Analogical: Conditional Recompositions**
973974 For conditional recompositions, we again adopt a general and automatable pipeline built around
975 LLM-generated Python solvers and automatic verification scripts, rather than manually rewriting
976 statements:

977 • We first call an LLM to extract the primary knowledge points tested by the original question,
978 and query a pre-constructed formula library indexed by knowledge point to retrieve potentially
979 relevant formulas.

980 • We feed the original problem, its official answer, the retrieved formulas, and (where available)
981 multiple correct solution sketches into the LLM. The official answer and solution sketches are
982 provided by the user, and providing solution sketches is optional.

983 • The LLM is prompted to:
984 – identify the key conditions and the target quantity;
985 – determine whether some of these conditions and the target can be interchanged—i.e., whether
986 knowing the original answer allows us to infer some of the original conditions (an invertible
987 relationship).

988 • When such an invertible relationship exists, the LLM is asked to write a Python solution program
989 for the recomposed problem, where the original target now appears as an input condition and (a
990 subset of) the original conditions become the new target.

991 • We then ask another LLM to inspect the generated Python code and verify that it implements
992 a general computational procedure for solving the problem, rather than relying on hard-coded
993 instance-specific outputs or trivial pattern matching.

994 • We import the LLM-generated Python code as a local module and call its `solve()` function,
995 plugging the original answer value into the new “condition” slot and checking whether the returned
996 output correctly recovers the original condition values.

997 • Any discrepancy or runtime error is fed back to the LLM for iterative refinement, just as in the
998 numeric substitutions pipeline. We repeat this refinement–verification loop for up to five attempts
999 and keep the Python code if it passes on the instance.

1000 Once a correct solver for the recomposed version is obtained, we can further vary the new input
1001 variables within reasonable ranges to generate additional condition-recomposed variants.1002 This “knowledge-point extraction → formula retrieval → analyze → code → verify → resample”
1003 pipeline is identical across all problems.1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

1026 **Algorithm 5** Conditional recompositions

```

1: function BUILDRECOMPOSEDSOLVER(problem_text, answer_gold, solution_sketches, max_iter = 5,
1028   max_refine = 5)
1029   2:   knowledge_points  $\leftarrow$  LLM_EXTRACTKNOWLEDGEPOINTS(problem_text)
1030   3:   retrieved_formulas  $\leftarrow$  RETRIEVE_FORMULAS(knowledge_points)
1031   4:   history  $\leftarrow$  EMPTY_LIST
1032   5:   ANALYSIS_PROMPT  $\leftarrow$  BUILD_PROMPT_INVERTIBLE_ANALYSIS(problem_text, answer_gold, solution_sketches, retrieved_formulas)
1033   6:   ANALYSIS_RESPONSE  $\leftarrow$  LLM_CALL(ANALYSIS_PROMPT)
1034   7:   (invertible, cond_as_unknown, target_as_given, recomposed_problem_text)  $\leftarrow$  PARSE_INVERT-
1035     IBLE_STRUCTURE(ANALYSIS_RESPONSE)
1036   8:   if invertible = FALSE then
1037     9:     return FAILURE
1038   10:   end if
1039   11:   for iter = 1 to max_iter do
1040     12:     PROMPT  $\leftarrow$  BUILD_PROMPT_RECOMPOSED_CODEGEN(origin_problem_text = problem_text,
1041       new_problem_text = recomposed_problem_text, origin_answer_gold = answer_gold, new_answer_gold
1042       = cond_as_unknown.original_values, solution_sketches = solution_sketches, retrieved_formulas = re-
1043         tired_formulas)
1044     13:     APPEND(history, (PROMPT, NONE))
1045     14:     (CODE, value_ranges)  $\leftarrow$  LLM_CALL(PROMPT)
1046     15:     is_hard_code  $\leftarrow$  LLM_HARD_CODE_CHECK(CODE)
1047     16:     if is_hard_code then
1048       17:       continue
1049     18:     end if
1050     19:     for refine_step = 0 to max_refine do
1051       20:       (output, error)  $\leftarrow$  RUN PYTHON(CODE, input = answer_gold)
1052       21:       APPEND(history, (CODE, (output, error)))
1053       22:       if error = NONE and VERIFY(output, cond_as_unknown.original_values) then
1054         23:           return (CODE, cond_as_unknown, target_as_given, value_ranges, recomposed_problem_text)
1055       24:       end if
1056       25:       if refine_step = max_refine then
1057         26:           break
1058       27:       end if
1059       28:       PROMPT_refine  $\leftarrow$  BUILD_PROMPT_RECOMPOSED_CODE_REFINE(problem_text = recom-
1060         posed_problem_text, answer_gold = cond_as_unknown.original_values, history = history)
1061       29:       (CODE, value_ranges)  $\leftarrow$  LLM_CALL(PROMPT_refine)
1062       30:     end for
1063     31:   end for
1064     32:   return FAILURE
1065   33: end function
1066   34: function GENERATERECOMPOSEDVARIANTS(problem_text, answer_gold, solution_sketches)
1067     35:   (solver_code, cond_as_unknown, target_as_given, value_ranges, recomposed_problem_text)  $\leftarrow$  BUIL-
1068       DRECOMPOSEDSOLVER(problem_text, answer_gold, solution_sketches)
1069     36:   if solver_code = FAILURE then
1070       37:     return NONE
1071     38:   end if
1072     39:   new_param  $\leftarrow$  SAMPLE(value_ranges)
1073     40:   new_problem_text  $\leftarrow$  INSTANTIATERECOMPOSEDPROBLEMTXT(recomposed_problem_text, tar-
1074       get_as_given, new_param)
1075     41:   (output, error)  $\leftarrow$  RUN PYTHON(solver_code, input = new_param)
1076     42:   new_answer_gold  $\leftarrow$  output
1077     43:   return (new_problem_text, new_answer_gold)
1078   44: end function

```

1073 An example of conditional recompositions is given as follows.

1074 • **Original:**

1075 Rectangles $ABCD$ and $EFGH$ are drawn such that D, E, C, F are collinear. Also,
1076 A, D, H, G all lie on a circle. If $BC = 16$, $AB = 107$, $FG = 17$, and $EF = 184$,
1077 what is the length of CE ?

1078 • **Conditional decomposition:**

1080 Rectangles $ABCD$ and $EFGH$ are drawn such that D, E, C, F are collinear. Also,
 1081 A, D, H, G all lie on a circle. If $BC = 16$, $AB = 107$, $CE = 104$, and $EF = 184$,
 1082 what is the length of FG ?
 1083
 1084

1085 **Novel: recent-source adaptation**

1086 In the “novel” branch, the first mechanism is recent-source adaptation, which is also fully scriptable:

1088 1. We first use an LLM to extract the primary knowledge points tested by a given source problem.
 1089 2. We query open-access repositories of centralized exam questions that index items by region,
 1090 year, subject, and knowledge point, and crawl the most recent 2025 exam problems matching the
 1091 extracted knowledge points.
 1092 3. The retrieved problems are paraphrased by the LLM and can be further transformed using the
 1093 three analogical methods (redundancy insertion, numeric substitution, and conditional recompo-
 1094 sition).

1095 This yields a set of new, recent-source problems that are structurally aligned at the knowledge level
 1096 but clearly distinct in surface form and provenance. The entire workflow is driven by scripts and
 1097 general prompts, without hand-curating individual items.

1099 **Algorithm 6** Recent-Source Adaptation

```

1: function RECENTSOURCEADAPTATION(problem_text, metadata, K)
2:   KP_PROMPT  $\leftarrow$  BUILD_PROMPT_EXTRACT_KNOWLEDGEPOINTS(problem_text, metadata)
3:   KP_RESPONSE  $\leftarrow$  LLM_CALL(KP_PROMPT)
4:   KPs  $\leftarrow$  PARSE_KNOWLEDGEPOINTS(KP_RESPONSE)
5:   if KPs = NONE then
6:     return NONE
7:   end if
8:   year_range  $\leftarrow$  {2025}
9:   candidate_item  $\leftarrow$  RETRIEVE_EXAMS(knowledge_points = KPs, year_range = year_range, sub-
10:    ject = metadata.subject)
10:   if candidate_item = NONE then
11:     return NONE
12:   end if
13:   PARA_PROMPT  $\leftarrow$  BUILD_PROMPT_PARAPHRASE(candidate_item.text, candidate_item.answer_gold, KPs)
14:   PARA_RESPONSE  $\leftarrow$  LLM_CALL(PARA_PROMPT)
15:   paraphrased_item  $\leftarrow$  PARSE_PARAPHRASED_PROBLEM(PARA_RESPONSE)
16:   adapted_item  $\leftarrow$  {
17:     text : paraphrased_item.text,
18:     answer : candidate_item.answer_gold,
19:     KPs : KPs,
20:     provenance : {
21:       source_exam : candidate_item.metadata,
22:       transform : “paraphrase”
23:     }
24:   }
25:   return adapted_item
26: end function

```

1124 For example, one recent-source adaptation question generated for the concept of *logarithms* is:

1125 Given $2^{\log_2 a} = 3$ and $\log_5 5^b = 2$, find $a - b$

1130 **Novel: conceptual synthesis**

1131 The second “novel” mechanism is conceptual synthesis from authoritative textbooks. We first crawl
 1132 a large collection of authoritative textbooks across different subjects from the web, and then use the
 1133 LLM API’s built-in functionality for parsing local PDF files to extract their content. Based on the

1134 extracted content, we build a structured knowledge base in which each concept is associated with
 1135 definitions, properties, theorems, phenomena, and canonical examples extracted from the textbooks.
 1136

1137 1. Given a problem to be augmented, we use an LLM to identify its main knowledge points, and
 1138 then retrieve the corresponding entries from the structured knowledge base. If the subject-specific
 1139 knowledge base is missing, we trigger the textbook crawling and parsing step to expand the
 1140 knowledge base, and then retrieve the corresponding entries from it.
 1141 2. Conditioned on these entries, the LLM is prompted to generate new conceptual questions target-
 1142 ing the underlying knowledge points, rather than copying any existing problem.

1143 This pipeline turns textbook content into fresh conceptual questions that align with the original topic
 1144 but are novel in form and focus.

1145 **Algorithm 7** Conceptual Synthesis

```
1146 1: function CONCEPTUALSYNTHESIS(problem_text, metadata)
1147 2:   KP_PROMPT  $\leftarrow$  BUILD_PROMPT_EXTRACT_KNOWLEDGE_POINTS(problem_text, metadata)
1148 3:   KP_RESPONSE  $\leftarrow$  LLM_CALL(KP_PROMPT)
1149 4:   KPs  $\leftarrow$  PARSE_KNOWLEDGE_POINTS(KP_RESPONSE)
1150 5:   if KPs = NONE then
1151 6:     return NONE
1152 7:   end if
1153 8:   kb_entry  $\leftarrow$  RETRIEVE_KB_ENTRY(knowledge_points = KPs, subject = metadata.subject)
1154 9:   if kb_entry = NONE then
1155 10:    CRAWL_AND_PARSE_TEXTBOOKS(knowledge_points = KPs, subject = metadata.subject)
1156 11:    kb_entry  $\leftarrow$  RETRIEVE_KB_ENTRY(knowledge_points = KPs, subject = metadata.subject)
1157 12:    if kb_entry = NONE then
1158 13:      return NONE
1159 14:    end if
1160 15:  end if
1161 16:  GEN_PROMPT  $\leftarrow$  BUILD_PROMPT_CONCEPTUAL_QUESTION_GENERATION(kb_entry, KPs)
1162 17:  GEN_RESPONSE  $\leftarrow$  LLM_CALL(GEN_PROMPT)
1163 18:  raw_item  $\leftarrow$  PARSE_GENERATED_CONCEPTUAL_QUESTIONS(GEN_RESPONSE)
1164 19:  conceptual_item  $\leftarrow$  {
1165 20:    text : raw_item.text,
1166 21:    answer : raw_item.answer,
1167 22:    KPs : KPs
1168 23:  }
1169 24:  return conceptual_item
1170 25: end function
```

1169 For example, one conceptual-synthesis question generated for the concept of *logarithms* is:

1171 1172 What kind of mathematical idea/method turns exponentiation and multiplication
 1173 into multiplication and addition?

1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

1188 A.4 RESULTS ON MMLU, GPQA
1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236


1237

1238


1239

1240

1241



1212 Figure 6: Accuracy confidence intervals of different LLMs on MMLU



1238 Figure 7: Accuracy confidence intervals of different LLMs on GPQA

1242 A.5 ANOVA ANALYSIS RESULTS ON DIFFERENT LLMs
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264

1265 Table 6: Complete ANOVA results on LLMs (sorted by effect size in descending order)

(a) Complete ANOVA results on DeepSeek-V3			(b) Complete ANOVA results on Doubao-1.5-pro-32k		
Factor	Effect Size η^2	<i>p</i> -value	Factor	Effect Size η^2	<i>p</i> -value
Question Format	0.399643	0.000	Question Format	0.467626	0.000
Question Format-COT	0.161394	0.000	Question Format-COT	0.259657	0.000
COT	0.080156	0.000	COT	0.130549	0.000
max_tokens	0.028099	0.000	max_tokens	0.009192	0.002
Question Format-Shot	0.011101	0.000	COT-max_tokens	0.008943	0.006
Language-Question Format	0.008178	0.006	max_tokens-presence_penalty	0.000671	0.388
COT-max_tokens	0.006721	0.010	Shot	0.000638	0.408
Language-COT	0.004345	0.038	Shot-Multi Turn	0.000605	0.424
Multi Turn-max_tokens	0.003841	0.050	Question Paraphrase-max_tokens	0.000502	0.466
Language	0.003841	0.046	Multi Turn	0.000473	0.466
Shot-max_tokens	0.003669	0.046	Multi Turn-max_tokens	0.000464	0.472
Question Format-max_tokens	0.002687	0.100	Language-temperature	0.000455	0.468
Language-Multi Turn	0.002600	0.066	COT-Multi Turn	0.000427	0.496
temperature-top_p	0.002082	0.178	Question Format-Shot	0.000400	0.546
Question Format-Multi Turn	0.001321	0.244	Question Format-max_tokens	0.000392	0.538
Question Paraphrase	0.001240	0.252	temperature	0.000392	0.534
Question Format-temperature	0.001201	0.262	Question Format-temperature	0.000358	0.530
Shot-temperature	0.000926	0.326	Question Paraphrase-presence_penalty	0.000349	0.552
Multi Turn	0.000793	0.364	Language-top_p	0.000326	0.584
temperature	0.000587	0.396	temperature-top_p	0.000326	0.588
COT-Multi Turn	0.000587	0.466	Language-presence_penalty	0.000310	0.578
COT-top_p	0.000573	0.482	Language-Multi Turn	0.000295	0.562
Language-Shot	0.000534	0.466	Multi Turn-presence_penalty	0.000272	0.604
presence_penalty	0.000435	0.488	Language-COT	0.000265	0.586
Question Format-Question Paraphrase	0.000411	0.562	Shot-COT	0.000224	0.628
Question Paraphrase-max_tokens	0.000207	0.626	Question Paraphrase-top_p	0.000205	0.642
Language-Question Paraphrase	0.000198	0.660	Question Format-Multi Turn	0.000158	0.704
Shot-COT	0.000161	0.646	COT-presence_penalty	0.000147	0.696
COT-temperature	0.000140	0.698	max_tokens-top_p	0.000147	0.674
Language-max_tokens	0.000140	0.712	presence_penalty	0.000117	0.730
COT-presence_penalty	0.000140	0.668	temperature-max_tokens	0.000108	0.740
Shot	0.000134	0.706	Question Format-presence_penalty	0.000090	0.764
Question Format-presence_penalty	0.000133	0.708	Question Format-Question Paraphrase	0.000067	0.792
Shot-top_p	0.000109	0.736	Question Paraphrase	0.000054	0.836
max_tokens-presence_penalty	0.000092	0.768	Shot-top_p	0.000047	0.818
max_tokens-top_p	0.000086	0.790	Shot-presence_penalty	0.000047	0.792
top_p	0.000058	0.812	Language	0.000045	0.814
Shot-Multi Turn	0.000054	0.802	Language-Question Paraphrase	0.000044	0.840
temperature-max_tokens	0.000038	0.836	Multi Turn-temperature	0.000036	0.832
Language-presence_penalty	0.000035	0.854	COT-top_p	0.000036	0.854
top_p-presence_penalty	0.000032	0.902	COT-temperature	0.000034	0.834
Multi Turn-temperature	0.000029	0.874	Language-Shot	0.000034	0.850
Multi Turn-top_p	0.000026	0.884	Shot-max_tokens	0.000024	0.864
Question Paraphrase-top_p	0.000018	0.898	temperature-presence_penalty	0.000015	0.906
temperature-presence_penalty	0.000013	0.902	top_p-presence_penalty	0.000013	0.902
Question Paraphrase-presence_penalty	0.000011	0.910	Question Paraphrase-COT	0.000007	0.908
Language-top_p	0.000008	0.942	Language-Question Format	0.000002	0.956
Question Paraphrase-Shot	0.000006	0.930	Question Paraphrase-temperature	0.000001	0.960
Multi Turn-presence_penalty	0.000004	0.946	Shot-temperature	0.000001	0.962
Language-temperature	0.000004	0.940	Question Paraphrase-Multi Turn	0.000001	0.966
Question Format-top_p	0.000003	0.962	Language-max_tokens	0.000001	0.970
Question Paraphrase-Multi Turn	0.000003	0.972	Question Paraphrase-Shot	0.000001	0.980
Shot-presence_penalty	0.000001	0.966	top_p	0.000001	0.970
Question Paraphrase-COT	0.000001	0.996	Question Format-top_p	0.000000	0.984
Question Paraphrase-temperature	0.000000	0.994	Multi Turn-top_p	0.000000	0.992

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

(c) Complete ANOVA results on GPT-3.5

Factor	Effect Size (η^2)	p-value
Question Format	0.417428	0.000000
Question Format-COT	0.199209	0.000000
COT	0.199208	0.000000
temperature	0.006046	0.016000
Question Format-temperature	0.005781	0.020000
Shot-Multi Turn	0.005715	0.010000
Shot-COT	0.005651	0.026000
max_tokens	0.005396	0.024000
Question Format-max_tokens	0.004434	0.044000
temperature-top_p	0.004320	0.052000
top_p	0.003119	0.074000
Question Format-top_p	0.002570	0.104000
COT-max_tokens	0.001773	0.170000
Question Format-Shot	0.000853	0.344000
COT-Multi Turn	0.000828	0.352000
Language-Multi Turn	0.000779	0.364000
Language	0.000687	0.428000
Language-Question Format	0.000600	0.484000
COT-top_p	0.000443	0.496000
COT-temperature	0.000425	0.508000
COT-presence-penalty	0.000408	0.508000
Question Paraphrase-Multi Turn	0.000407	0.534000
Question Paraphrase-COT	0.000391	0.506000
Shot	0.000296	0.578000
Shot-temperature	0.000281	0.598000
Language-Shot	0.000253	0.598000
max_tokens-presence.penalty	0.000201	0.628000
Question Paraphrase-top_p	0.000177	0.696000
Shot-max_tokens	0.000135	0.720000
Language-top_p	0.000115	0.690000
Question Paraphrase-presence.penalty	0.000115	0.714000
Language-temperature	0.000106	0.750000
Language-max_tokens	0.000106	0.740000
Shot-presence.penalty	0.000089	0.786000
Multi Turn-temperature	0.000081	0.784000
Language-COT	0.000074	0.766000
Multi Turn-presence.penalty	0.000074	0.788000
Question Format-Question Paraphrase	0.000067	0.776000
Question Paraphrase	0.000053	0.808000
Question Format-Multi Turn	0.000047	0.826000
presence.penalty	0.000047	0.854000
temperature-max.tokens	0.000036	0.858000
Question Paraphrase-max.tokens	0.000036	0.854000
temperature-presence.penalty	0.000031	0.884000
Language-Question Paraphrase	0.000027	0.858000
Question Format-presence.penalty	0.000027	0.900000
Question Paraphrase-temperature	0.000027	0.882000
Multi Turn	0.000018	0.872000
Multi Turn-max_tokens	0.000009	0.916000
Question Paraphrase-Shot	0.000009	0.920000
Multi Turn-top_p	0.000007	0.942000
top_p-presence.penalty	0.000007	0.948000
Shot-top_p	0.000007	0.942000
Language-presence.penalty	0.000002	0.970000
max_tokens-top_p	0.000000	0.978000

(d) Complete ANOVA results on GPT-4.1

Factor	Effect Size (η^2)	p-value
Question Format	0.289086	0.000000
Question Format-COT	0.180162	0.000000
COT-max_tokens	0.054845	0.000000
max_tokens	0.053399	0.000000
COT	0.027181	0.000000
temperature-top_p	0.006685	0.010000
Question Format-Shot	0.006529	0.030000
temperature	0.005619	0.016000
Shot	0.004963	0.028000
max_tokens-top_p	0.004019	0.044000
Language-max_tokens	0.003907	0.062000
Question Format-temperature	0.003732	0.072000
top_p	0.003364	0.098000
temperature-max_tokens	0.002990	0.140000
Question Paraphrase-Shot	0.002362	0.160000
Question Paraphrase-presence.penalty	0.001939	0.194000
Shot-Multi Turn	0.001842	0.194000
COT-temperature	0.001708	0.230000
Shot-COT	0.001589	0.290000
Language-COT	0.001517	0.246000
COT-presence.penalty	0.001406	0.276000
Question Format-max_tokens	0.001360	0.306000
temperature-presence.penalty	0.001284	0.292000
Language-Shot	0.001276	0.262000
Language-Question Paraphrase	0.001188	0.324000
presence.penalty	0.001184	0.362000
COT-top_p	0.001183	0.300000
COT-Multi Turn	0.001037	0.376000
Multi Turn-max_tokens	0.000988	0.376000
Shot-top_p	0.000754	0.422000
Question Paraphrase-max_tokens	0.000703	0.438000
Multi Turn-top_p	0.000367	0.574000
Question Format-presence.penalty	0.000348	0.584000
Question Format-Question Paraphrase	0.000333	0.584000
Language-presence.penalty	0.000329	0.586000
max_tokens-presence.penalty	0.000324	0.600000
top_p-presence.penalty	0.000253	0.618000
Shot-presence.penalty	0.000192	0.702000
Shot-max_tokens	0.000164	0.724000
Question Paraphrase	0.000146	0.746000
Question Format-top_p	0.000134	0.728000
Shot-temperature	0.000133	0.724000
Question Format-Multi Turn	0.000132	0.718000
Question Paraphrase-top_p	0.000065	0.798000
Multi Turn-presence.penalty	0.000055	0.828000
Language-top_p	0.000023	0.896000
Language-Question Format	0.000022	0.910000
Multi Turn-temperature	0.000020	0.886000
Question Paraphrase-temperature	0.000010	0.922000
Language-Multi Turn	0.000009	0.906000
Question Paraphrase-COT	0.000009	0.922000
Question Paraphrase-Multi Turn	0.000005	0.918000
Language	0.000001	0.964000
Multi Turn	0.000000	0.996000
Language-temperature	0.000000	0.998000

1350

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

(e) Complete ANOVA results on Qwen2.5

Factor	Effect Size (η^2)	p-value
Question Format	0.454352	0.000000
Question Format-COT	0.204235	0.000000
COT	0.200224	0.000000
Question Format-Shot	0.009265	0.000000
Shot	0.007983	0.008000
Shot-COT	0.006715	0.002000
Multi Turn	0.003717	0.046000
Question Format-Multi Turn	0.003657	0.052000
max.tokens	0.002764	0.080000
Language-Question Format	0.002585	0.128000
COT-max.tokens	0.001865	0.164000
Language	0.001720	0.196000
COT-Multi Turn	0.001540	0.214000
Multi Turn-max.tokens	0.001110	0.322000
Question Format-max.tokens	0.000922	0.292000
Language-COT	0.000848	0.382000
Question Paraphrase-presence.penalty	0.000710	0.400000
Language-Multi Turn	0.000538	0.478000
temperature	0.000430	0.444000
Shot-temperature	0.000380	0.534000
Question Format-temperature	0.000371	0.550000
Language-Question Paraphrase	0.000343	0.550000
temperature-presence.penalty	0.000334	0.566000
Language-temperature	0.000290	0.602000
top.p	0.000257	0.600000
max.tokens-top.p	0.000242	0.618000
Shot-Multi Turn	0.000234	0.654000
COT-temperature	0.000219	0.660000
Question Format-Question Paraphrase	0.000165	0.674000
Question Paraphrase-top.p	0.000165	0.672000
Language-max.tokens	0.000165	0.674000
Question Paraphrase	0.000118	0.742000
Language-top.p	0.000107	0.760000
temperature-top.p	0.000107	0.724000
Multi Turn-presence.penalty	0.000083	0.762000
max.tokens-presence.penalty	0.000075	0.784000
Shot-max.tokens	0.000062	0.770000
presence.penalty	0.000051	0.830000
Question Format-top.p	0.000038	0.824000
top.p-presence.penalty	0.000038	0.830000
Multi Turn-temperature	0.000029	0.856000
Question Paraphrase-temperature	0.000027	0.882000
COT-presence.penalty	0.000022	0.846000
Multi Turn-top.p	0.000022	0.882000
Question Format-presence.penalty	0.000022	0.866000
Language-presence.penalty	0.000012	0.912000
temperature-max.tokens	0.000012	0.918000
Question Paraphrase-Multi Turn	0.000009	0.942000
Language-Shot	0.000007	0.926000
Question Paraphrase-Shot	0.000005	0.928000
COT-top.p	0.000003	0.954000
Shot-presence.penalty	0.000002	0.960000
Shot-top.p	0.000001	0.982000
Question Paraphrase-COT	0.000000	0.984000
Question Paraphrase-max.tokens	0.000000	0.990000

(f) Complete ANOVA results on Qwen Plus

Factor	Effect Size (η^2)	p-value
Question Format	0.302717	0.000000
Question Format-COT	0.259678	0.000000
COT	0.098747	0.000000
Shot	0.047447	0.000000
Question Format-Shot	0.042448	0.000000
Shot-COT	0.024956	0.000000
COT-max.tokens	0.016596	0.000000
max.tokens	0.016280	0.000000
Language	0.005019	0.024000
Language-Question Format	0.003272	0.069000
temperature-top.p	0.002324	0.112000
Multi Turn-max.tokens	0.001979	0.154000
Question Format-temperature	0.001662	0.202000
top.p	0.001282	0.250000
Language-Shot	0.000991	0.296000
Question Paraphrase-Multi Turn	0.000877	0.320000
COT-temperature	0.000822	0.368000
Multi Turn	0.000736	0.428000
Question Format-Question Paraphrase	0.000654	0.458000
temperature	0.000608	0.412000
Language-COT	0.000519	0.472000
Shot-Multi Turn	0.000438	0.512000
Multi Turn-presence.penalty	0.000400	0.512000
Shot-presence.penalty	0.000375	0.496000
Question Format-top.p	0.000275	0.604000
COT-Multi Turn	0.000245	0.612000
Question Format-max.tokens	0.000226	0.648000
Question Format-Multi Turn	0.000166	0.708000
Multi Turn-temperature	0.000135	0.724000
max.tokens-presence.penalty	0.000128	0.690000
top.p-presence.penalty	0.000102	0.752000
Question Paraphrase-Shot	0.000101	0.746000
COT-top.p	0.000090	0.758000
Language-presence.penalty	0.000084	0.792000
Language-max.tokens	0.000073	0.762000
temperature-presence.penalty	0.000058	0.786000
Language-Multi Turn	0.000058	0.806000
Question Paraphrase-top.p	0.000058	0.814000
Language-temperature	0.000049	0.816000
Shot-top.p	0.000037	0.838000
Question Paraphrase	0.000033	0.864000
Shot-temperature	0.000033	0.860000
Language-top.p	0.000029	0.850000
Question Paraphrase-max.tokens	0.000023	0.862000
Question Paraphrase-presence.penalty	0.000015	0.902000
Question Paraphrase-temperature	0.000015	0.904000
Multi Turn-top.p	0.000015	0.924000
Shot-max.tokens	0.000013	0.898000
COT-presence.penalty	0.000013	0.918000
Question Format-presence.penalty	0.000005	0.946000
temperature-max.tokens	0.000002	0.970000
presence.penalty	0.000002	0.960000
max.tokens-top.p	0.000002	0.970000
Question Paraphrase-COT	0.000001	0.978000
Language-Question Paraphrase	0.000000	0.994000

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

(g) Complete ANOVA results on Mistral Large

Factor	Effect Size (η^2)	p-value
Question Format	0.406873	0.000000
Question Format-COT	0.268919	0.000000
COT	0.075852	0.000000
COT-max_tokens	0.026017	0.000000
max_tokens	0.025372	0.000000
Question Format-Multi Turn	0.004796	0.024000
Multi Turn	0.003609	0.058000
Question Format-Shot	0.003338	0.068000
COT-Multi Turn	0.002708	0.100000
Question Format-max_tokens	0.001379	0.230000
Shot	0.001023	0.320000
Language-Question Format	0.001004	0.280000
Multi Turn-max_tokens	0.000881	0.310000
Shot-Multi Turn	0.000830	0.368000
Language-Multi Turn	0.000766	0.372000
Language-COT	0.000765	0.356000
Question Format-Question Paraphrase	0.000644	0.414000
Multi Turn-presence_penalty	0.000456	0.458000
Question Paraphrase	0.000364	0.550000
Question Paraphrase-presence_penalty	0.000353	0.486000
max_tokens-presence_penalty	0.000272	0.620000
Shot-top_p	0.000244	0.638000
Multi Turn-top_p	0.000244	0.570000
Language-Question Paraphrase	0.000235	0.618000
Question Format-presence_penalty	0.000227	0.604000
COT-temperature	0.000227	0.628000
Question Paraphrase-Shot	0.000193	0.658000
temperature	0.000185	0.644000
Shot-presence_penalty	0.000178	0.650000
Question Format-temperature	0.000178	0.672000
Question Paraphrase-top_p	0.000163	0.686000
Shot-max_tokens	0.000163	0.682000
top_p	0.000142	0.690000
Question Paraphrase-max_tokens	0.000135	0.688000
Question Paraphrase-Multi Turn	0.000128	0.772000
max_tokens-top_p	0.000110	0.712000
Language	0.000109	0.736000
COT-presence_penalty	0.000087	0.780000
Shot-temperature	0.000087	0.810000
Question Paraphrase-COT	0.000053	0.818000
top_p-presence_penalty	0.000049	0.824000
presence_penalty	0.000038	0.836000
Language-temperature	0.000035	0.840000
Language-top_p	0.000028	0.864000
temperature-max_tokens	0.000017	0.884000
temperature-presence_penalty	0.000015	0.884000
Multi Turn-temperature	0.000007	0.928000
temperature-top_p	0.000006	0.928000
Question Paraphrase-temperature	0.000005	0.950000
COT-top_p	0.000003	0.952000
Language-Shot	0.000003	0.944000
Question Format-top_p	0.000002	0.952000
Language-max_tokens	0.000002	0.950000
Language-presence_penalty	0.000002	0.960000
Shot-COT	0.000000	0.998000

(h) Complete ANOVA results on Mistral Medium

Factor	Effect Size (η^2)	p-value
Question Format	0.430500	0.000000
Question Format-COT	0.274248	0.000000
COT	0.105919	0.000000
COT-max_tokens	0.013038	0.002000
max_tokens	0.012064	0.000000
Question Format-Multi Turn	0.007865	0.004000
Multi Turn	0.005334	0.026000
Shot	0.003097	0.064000
COT-Multi Turn	0.003001	0.090000
Question Format-Shot	0.002782	0.086000
Multi Turn-max_tokens	0.001354	0.252000
Shot-Multi Turn	0.000908	0.336000
Language-Question Paraphrase	0.000891	0.378000
Shot-COT	0.000636	0.434000
Question Paraphrase-top_p	0.000366	0.476000
Question Paraphrase-presence_penalty	0.000312	0.534000
Question Paraphrase-Multi Turn	0.000282	0.578000
Language	0.000254	0.626000
max_tokens-presence_penalty	0.000227	0.650000
Language-Multi Turn	0.000219	0.656000
Multi Turn-presence_penalty	0.000218	0.642000
Question Format-temperature	0.000178	0.682000
Multi Turn-top_p	0.000163	0.672000
Language-COT	0.000163	0.690000
temperature-top_p	0.000148	0.698000
Question Format-max_tokens	0.000128	0.672000
Question Format-Question Paraphrase	0.000109	0.746000
Question Paraphrase-max_tokens	0.000103	0.754000
COT-top_p	0.000091	0.750000
COT-temperature	0.000081	0.764000
Shot-temperature	0.000076	0.806000
top_p	0.000076	0.784000
Question Paraphrase-temperature	0.000076	0.758000
Language-temperature	0.000061	0.798000
Question Format-presence_penalty	0.000053	0.806000
Shot-top_p	0.000048	0.808000
Language-Shot	0.000044	0.832000
Shot-max_tokens	0.000044	0.814000
Question Paraphrase-COT	0.000037	0.854000
presence_penalty	0.000037	0.844000
Question Paraphrase	0.000034	0.850000
Language-top_p	0.000024	0.870000
Language-presence_penalty	0.000022	0.874000
Shot-presence_penalty	0.000019	0.888000
temperature-presence_penalty	0.000014	0.910000
Language-max_tokens	0.000008	0.920000
Language-Question Format	0.000007	0.938000
Question Format-top_p	0.000005	0.942000
temperature	0.000005	0.948000
max_tokens-top_p	0.000004	0.958000
Question Paraphrase-Shot	0.000003	0.956000
Multi Turn-temperature	0.000001	0.972000
temperature-max_tokens	0.000001	0.968000
COT-presence_penalty	0.000000	0.986000
top_p-presence_penalty	0.000000	0.994000

1458

(i) Complete ANOVA results on Moonshot-v1

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

A.6 LATENCY AND ACCURACY ANALYSIS OF LLMs

In this section, we analyze the relationships between the latency and accuracy of LLMs, as well as between latency and hardware architectures, based on online evaluation. On one hand, we conduct online testing to assess the accuracy of LLMs under different configuration spaces in terms of the “tail to quality” (Yang et al., 2022) metric. Here, “tail to quality” refers to the ratio of the number of tasks correctly completed within a specified threshold to the total number of tasks. Figure 8a illustrates the performance of various LLMs under the “Tail to Quality” metric, showing how their quality scores evolve across different threshold values. Among the models, deepseek (green curve) consistently demonstrates the highest quality across all thresholds, outperforming the others. Doubao (blue curve) and qwen (gray curve) follow, with doubao approaching deepseek’s performance at higher thresholds. Kimi (brown curve) and qwen25 (cyan curve) exhibit relatively lower quality, though qwen25 shows rapid improvement at lower thresholds before plateauing. Overall, the chart highlights deepseek’s superior capability in handling tail data, while qwen25’s growth in quality becomes limited at higher thresholds.

On the other hand, following a similar approach as for accuracy, we obtain the 95% and 99% confidence intervals for latency, as shown in Figure 8b. It can be seen that, for most models, latency and accuracy on AIME’2024 are positively correlated. Notably, doubao-1.5-pro and qwen2.5 achieve relatively low latency while maintaining high accuracy. In contrast, gpt-4.1 and qwen-plus exhibit the opposite trend: they achieve lower accuracy despite higher latency.

1504

1505

1506

1507


1508

1509


1510

1511

1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541



(a) tail to quality of different LLMs



(b) latency confidence intervals of different LLMs

Figure 8: Relationship between inference accuracy and latency of LLMs in online testing

1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565