

OPTIMIZING DIVERSITY AND QUALITY THROUGH BASE-ALIGNED MODEL COLLABORATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Alignment has greatly improved large language models (LLMs)’ output *quality* at the cost of *diversity*, yielding highly similar outputs across generations. We propose Base-Aligned Model Collaboration (BACo), an inference-time token-level model collaboration framework that dynamically combines a base LLM with its aligned counterpart to optimize diversity and quality. Inspired by prior work (Fei et al., 2025), BACo employs *routing strategies* that determine, at each token, from which model to decode based on next-token prediction uncertainty and predicted contents’ semantic role. Prior diversity-promoting methods, such as retraining, prompt engineering, and multi-sampling methods, improve diversity but often degrade quality or require costly decoding or post-training. In contrast, BACo achieves both high diversity and quality post hoc within a single pass, while offering strong controllability. We explore a family of routing strategies, across three open-ended generation tasks and 13 metrics covering diversity and quality, BACo consistently surpasses state-of-the-art inference-time baselines. With our best router, BACo achieves a 21.3% joint improvement in diversity and quality. Human evaluations also mirror these improvements. The results suggest that collaboration between base and aligned models can optimize and control diversity and quality.

1 INTRODUCTION

While alignment greatly improved large language models (LLMs)’ output *quality* in terms of instruction following and downstream task performance (Ouyang et al., 2022), it results in a stark reduction in output *diversity* (Kirk et al., 2023; Zhang et al., 2025b; West & Potts, 2025; Spangher et al., 2025). Across repeated sampling, a model after alignment (i.e., aligned model) tends to produce highly similar outputs, whereas a model before alignment (i.e., base model) yields diverse outputs. For example, when prompted with “suggest a summer trip destination in the US,” a base model may produce diverse destinations across generations, while the aligned model often converges on a single dominant one (Figure 1A). This *diversity-quality trade-off* undermines open-ended use cases (e.g., creative writing and dialogue) by encouraging formulaic language use (Chakrabarty et al., 2025; Zhang et al., 2024), diminishing creativity (West & Potts, 2025), and suppressing ideation in human-AI interaction (Padmakumar & He, 2023; Meincke et al., 2025a; Ashkinaze et al., 2025). These findings motivate methods to improve diversity in aligned LLMs.

Prior diversity-promoting methods attempt to address the diversity-quality trade-off at both the training and the inference stages (§B). The former (Lanchantin et al., 2025; Chung et al., 2025; Li et al., 2025) incorporate explicit diversity objectives into preference optimization during reinforcement learning (e.g., by measuring deviation among winning responses or comparing alternative response groups). While effective at improving diversity, such methods require modifying the model’s output distribution, which can compromise desirable alignment properties such as safety and helpfulness (Qi et al., 2023). The latter consists of decoding-based techniques, such as adjusting temperature and beam search (Vijayakumar et al., 2016), as well as prompt-based techniques, including in-context learning (Meyerson et al., 2024), prompt paraphrasing (Zhang et al., 2025b), and multilingual back-translation (Wang et al., 2025b). However, these inference-time techniques typically require multiple decoding passes or long-horizon planning to improve diversity, and may still disproportionately degrade generation quality (Peeperkorn et al., 2024). This reveals a fundamental limitation of forcing a single model to excel at both diversity and quality.

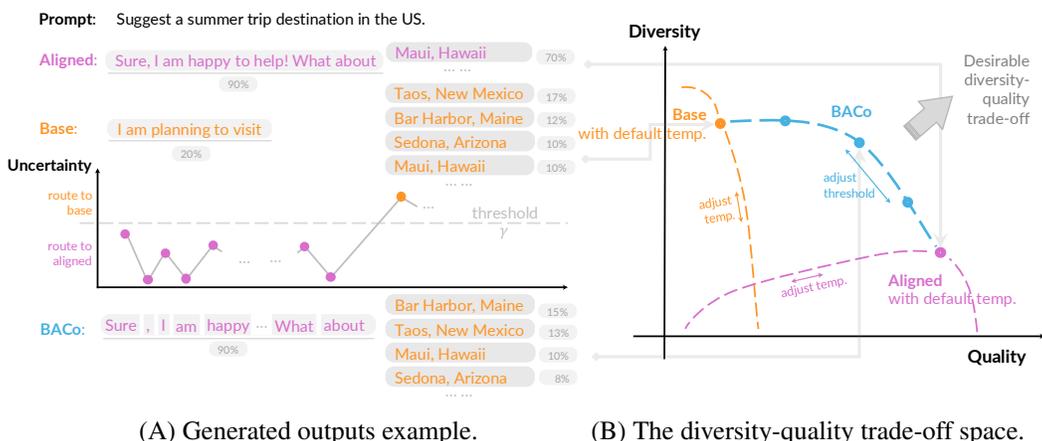


Figure 1: BACO is an inference-time token-level model collaboration framework that combines a base model’s diversity with its aligned counterpart’s quality. (A) A comparison of generated outputs. The aligned model produces high-quality but low-diversity outputs, while the base model produces high-diversity but low-quality outputs. BACO optimizes both diversity and quality by dynamically *routing* between them. The probabilities of token(s) are in grey next to text boxes. (B) Illustration of the diversity-quality trade-off space. Single models face a steep trade-off, where improving diversity by adjusting configuration (e.g., by increasing temperature) degrades quality. BACO achieves a better Pareto curve and allows for easy traversal across this frontier by adjusting the router’s threshold.

To overcome this limitation, we introduce Base-Aligned Model Collaboration (BACO), an inference-time token-level model collaboration framework (§3). Different from prior works, BACO combines a base model’s diversity with its aligned counterpart’s quality. BACO operates via a lightweight, token-level *routing strategy* that dynamically switches between the two models in a single decoding pass, requiring no fine-tuning or specialized prompting (Figure 1A). Recently, Fei et al. (2025) demonstrated that base and aligned models largely agree on next-token predictions, a phenomenon known as the superficial alignment (Lin et al., 2023), suggesting that collaboration between them is feasible. Their method, NUDGING, employs a small aligned model to guide a larger base model during decoding, effectively improving its quality but without considering diversity. In contrast, our goal is to jointly optimize diversity and quality (Figure 1B) by a base and aligned model collaborate.

However, the trade-off cannot be entirely resolved. Different tasks naturally favor different points along the diversity-quality spectrum. Moreover, different people may prefer different trade-offs depending on their preferences or intent. Therefore, a practical framework should enable *controllability*, the ability to adjust along the diversity-quality spectrum according to task or preference. BACO supports such controllability by an adjustable routing threshold (Figure 1) that tunes the contribution of base and aligned model. Moreover, since a single routing strategy does not guarantee diversity across all aspects, we introduce a family of strategies, including logit-based, content-based, and their combinations, within a proposed design space (§3.1). The design space provides a principled foundation for designing, comparing, and analyzing model collaborative performance.

We validate our approach across three open-ended generation tasks: instruction-following, dialogue, and creative writing (§4). We apply 11 automatic metrics for diversity and 2 for quality, together forming 11×2 bi-dimensional evaluation spaces. Additionally, we extend our evaluation beyond instruction-following tasks, which involve relatively short outputs, to long-form generation, where we measure discourse-level diversity in plot structure and emotional flow (Tian et al., 2024). Finally, we complement automatic evaluations with human evaluations of diversity and quality, providing human-centered validation and comparison with automatic results.

Through extensive experiments (§5), we demonstrate that BACO achieves a new state-of-the-art in optimizing diversity-quality trade-off. Even a simple router that routes randomly yields a 19.0% joint improvement in diversity and quality across tasks and metrics, underscoring the effectiveness of the framework. Our best router achieves a 21.3% gain compared with all baselines (e.g., base or aligned model alone with varying temperature), with even larger improvements on semantic diversity metrics. These results are further supported by human evaluations. View the base and

aligned model as demonstration of checkpoints with different alignment levels during the alignment process, we suggest that collaboration between models at different alignment levels can improve and control diversity and quality. Overall, BACO provides a simple framework for base-aligned model collaboration, effectively improving both diversity and quality.

In summary, our contributions are threefold:

- ① We propose BACO, an inference-time token-level model collaboration framework that combines a base model and its aligned counterpart, and explore a family of routing strategies. It produces high-diversity and high-quality outputs across generations.
- ② We evaluate the diversity-quality trade-off across 11×2 evaluation spaces, and we also extend automatic diversity evaluation to long-form evaluation and conduct human evaluation.
- ③ Through extensive experiments across instruction-following, dialogue, and creative writing, we show that BACO consistently outperforms strong baselines. It achieves a 21.3% optimization on diversity-quality trade-off. Human evaluations further validate its advantages.

2 PRELIMINARY

Large Language Models (LLMs). LLMs are typically trained to autoregressively predict the next token of the output y given a prompt x . The conditional probability is factorized as $P(y|x; \theta) = \prod_t P(y_t | [x, y_{<t}]; \theta)$, where $y_{<t}$ denotes the output prefix generated up to position $t-1$, and θ denotes the model parameter.

Alignment is the process of fine-tuning an LLM to align its outputs with human intent, ethical principles, and desired behavioral norms, typically through instruction tuning or reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022; Bai et al., 2022). We use *base models* to denote models without alignment tuning (e.g., Llama-3-8B) and *aligned models* to denote those further optimized with alignment (e.g., Llama-3-8B-Instruct) (Dubey et al., 2024).

Diversity and Quality Measurement. In this paper, we measure diversity over a group of outputs independently generated from the same prompt x : $\mathcal{Y}(x) = \{y^{(1)}, \dots, y^{(k)}\}, y^{(i)} \sim P(\cdot | x; \theta)$. The group-level diversity is denoted as $D(\mathcal{Y})$ (Kirk et al., 2023; West & Potts, 2025) (e.g., the clustering-based approach in Kuhn et al. (2023)). Quality is modeled as $Q(y|x)$ for each output given the prompt, typically by a reward model or human evaluator, reflecting linguistic fluency and instruction-following (Lambert et al., 2024; Zhang et al., 2025b). The group-level quality is then defined as the average quality across all outputs in a group: $Q(\mathcal{Y}) = \sum_{i=1}^k Q(y^{(i)}|x)$. For simplicity, we refer to group-level diversity and group-level quality as diversity and quality in this work.

Diversity-Quality Trade-off. Alignment, while effective at improving output quality, is at the cost of reducing output diversity (Lu et al., 2025a; West & Potts, 2025; Yang & Holtzman, 2025). As a pilot study to demonstrate the trade-off, we first situate the performance of the base and aligned models with default configuration (the two noted points in Figure 1B) within the diversity-quality space. Following the evaluation protocol of the diversity-focused benchmark NoveltyBench (Zhang et al., 2025b), we evaluate Llama-3’s base and aligned versions on an open-ended subset of WildChat (Zhao et al., 2024). In this study, diversity is measured as the number of semantic equivalent classes of the output group via Zhang et al. (2025b)’s clustering, and quality is measured as the average reward per output, given by Skywork-Reward-Gemma-2-27B (Liu et al., 2024).¹ We sample 10 outputs per prompt. As illustrated in Table 1, the diversity-quality trade-off are stark: the base model is **3.15x as diverse**, whereas the aligned model achieves **5.95x as much**

<i>Model</i>	<i>Diversity</i> (#Clusters) \uparrow	<i>Quality</i> (Reward) \uparrow
Llama-3-8B	8.13	1.28
Llama-3-8B-Instruct	2.58	7.62

Table 1: Diversity and quality of a base and an aligned model. The results demonstrate diversity-quality trade-off when compared the two model’s performance. Diversity is measured by the number of semantic equivalent clusters of the output group, and quality is the average reward.

¹These are two of many possible measurements introduced later in the paper. We use them here as representative examples for the pilot study, as they are among the widely adopted metrics.

quality. Inherently, this performance trade-off stems from alignment’s tendency to reduce the entropy of the next-token prediction distribution ($H_\theta(y_t|x, y_{<t})$), concentrating probability mass on fewer, high-quality tokens, a phenomenon known as *mode collapse* (Lin et al., 2023; Shumailov et al., 2024; Hamilton, 2024; Yang & Holtzman, 2025; Cui et al., 2025).

This presents a dilemma: one can either use a high-diversity but low-quality base model, or a high-quality but low-diversity aligned model. The single-model paradigm is insufficient, as neither extreme is ideal for all applications. Hence, we argue that an ideal method is able to pursue the best of both worlds. To this end, we formalize the problem in a two-dimensional *diversity-quality space* $S = \{(D, Q)\}$. In this space, any given method under specific configuration (e.g., sampling parameters) is evaluated to be a single point. An ideal method, by adjusting its configuration, should approximate the *Pareto frontier*: the set of optimal solutions where diversity cannot be improved without sacrificing quality, and vice versa. This frontier represents the best possible trade-offs.

3 BACO: BASE-ALIGNED COLLABORATION FOR DIVERSITY AND QUALITY

Recent work has shown empirical evidence on the *superficial alignment hypothesis* (Zhou et al., 2023a; Lin et al., 2023), which suggests that a base model and its aligned counterpart largely agree on next-token predictions. Building on this, Fei et al. (2025) demonstrate that their in-context alignment method, NUDGING, introduces only a small proportion of aligned-model tokens to a base model’s decoding for it to achieve task-specific performance (including instruction-following) comparable to the aligned model. This suggests the feasibility of collaboration between them.

Inspired by this result, but with the goal of optimizing diversity-quality trade-off, we hypothesize that

Collaboration between a less-aligned, higher-diversity model and a more-aligned, higher-quality model during inference can optimize the diversity-quality trade-off.

A base model and an aligned model are a pair of such models that are relatively easy to obtain off the shelf. Hence, based on our hypothesis, we propose **BACO** that orchestrates a collaboration between a base model (P_{base}) for diversity and its aligned counterpart (P_{aligned}) for quality at the token level.²

At the core of BACO is a **router**, a lightweight decision module that determines, at each decoding step, which model should generate the next token. The router operates according to **routing strategy(ies)** \mathcal{R} , which selects between the base and aligned model accordingly. Intuitively, the router acts as a “gatekeeper”: ideally when diversity is valid and required under the contexts, it routes to the base model; when quality is required, it routes to the aligned model.

Formally, BACO collaborates the two models as:

$$P_{\text{BACO}}(y_t|c_t) = w_{\text{base}} \cdot P_{\text{base}}(y_t|c_t; \theta_{\text{base}}) + (1 - w_{\text{base}}) \cdot P_{\text{aligned}}(y_t|c_t; \theta_{\text{aligned}}) \quad (1)$$

where $c_t = [x, y_{<t}]$, and the gating weight $w_{\text{base}} \in \{0, 1\}$ for each candidate token y_t is given by the router:

$$w_{\text{base}} = \mathbb{I}[\mathcal{R}(y_t|c_t, P_{\text{base}}, P_{\text{aligned}}) = \text{base}] \quad (2)$$

In practice, since one word may consist of multiple tokens, we restrict switching to word boundaries to prevent erroneous generation when the two models use different tokenizations (implementation details in Appendix C). Because the base and aligned models agree on most next-token predictions, only occasional model switching is required, making BACO naturally less costly. More discussion on computation cost saving is in Appendix C and Section A.4.

3.1 ROUTING STRATEGY DESIGN

The router determines how the base and aligned models collaborate during decoding. Its effectiveness depends on the routing strategy that specifies when to switch between models. Conceptually, designing a router involves choosing what information to use as prior for routing decisions, and which model to switch to given the information. A routing strategy includes a threshold parameter that controls the attribution to the two models, providing a means to continuously adjust the balance between

²The token-by-token nature of LLM autoregressive decoding makes token-level control feasible for this task.

216 diversity and quality. We present two categories of routing strategies—logit-based and content-
 217 based—primarily distinguished by what information is used. These categories capture complementary
 218 intuitions and are broadly applicable across tasks, supported by our empirical results.

219 **Logit-Based Routing.** These strategies leverage the next-token probability distribution to infer
 220 model’s uncertainty. High uncertainty often signals an opportunity for diverse, open-ended generation.
 221 Heuristics can be built on simple statistics such as the maximum token probability, the average top-k
 222 token probability, or the entropy of the distribution. Key strategies implemented under this category
 223 include: ① **BACO-P** routes to the base model when its maximum token probability is below a
 224 threshold γ , i.e., $\max_{y_t} P_{\text{base}}(y_t | \cdot) < \gamma$. This deliberately incorporates lower-probability tokens to
 225 promote variation; ② **BACO-H** routes to the base model when its next-token entropy is high, i.e.,
 226 $H_{\text{base}}(Y_t | \cdot) = \sum_{y_t} P_{\text{base}}(y_t | \cdot) \log P_{\text{base}}(y_t | \cdot) > \gamma$, indicating high uncertainty.

227 **Content-Based Routing.** These strategies determine switching decisions based on the semantic
 228 roles of the predicted tokens themselves rather than probability distributions. The motivations are
 229 twofold: First, linguistic features may provide cues for when diversity is desired (e.g., verbs are
 230 critical in narrative generation (Yao et al., 2019; Sims et al., 2019)). Second, Lin et al. (2023); Fei
 231 et al. (2025) suggest that disagreements between base and aligned models often arise over stylistic
 232 tokens, such as formatting tokens (e.g., ‘\n’) or function words (e.g., ‘and’, ‘if’). Key strategies
 233 implemented under this category include: ① **BACO-PUNC** routes to the aligned model when its
 234 top-ranked token is either a punctuation or formatting token; ② **BACO-FC** routes to the aligned
 235 model if its top-ranked token is a function word, ensuring stylistic coherence during formatting parts
 236 of a sentence. Moreover, content-based routing strategies are suitable for black-box models, as they
 237 do not need any access-limited information such as logits.

238 A router’s performance can be further enhanced in two main ways. First, it can combine multiple
 239 strategies into a more sophisticated one. For example, BACO-P-FC applies the function-word
 240 strategy (-FC) before falling back to the probability-based strategy (-P). Second, while our main
 241 experiments use a one-token look-ahead for efficiency, a router with deeper look-ahead (e.g., n-grams
 242 or sentences) might make more context-aware routing decision for longer continuations. More routing
 243 strategies and router implementation details are provided in Appendix C.2.

244 4 EXPERIMENTAL SETUP

245 We design our experiments to empirically validate BACO’s central goal: to optimize the diversity-
 246 quality trade-off. Specifically, we aim to examine (i) whether BACO jointly improves diversity and
 247 quality across metrics and tasks, and (ii) how different routers perform.

248 **Datasets.** We evaluate BACO across three datasets that represent different types of open-ended
 249 generation tasks: **NoveltyBench** (Zhang et al., 2025b) for instruction following, **WildChat** (Zhao
 250 et al., 2024) for dialogue, and **Narrative-Discourse** (Tian et al., 2024) for creative writing. Together,
 251 these datasets cover both short-form and long-form open-ended generation across varying levels of
 252 task complexity. For full dataset details, we refer readers to Appendix D.

253 **Baselines.** We compare BACO with inference-time methods across five categories: ① Single-model:
 254 a base model or an aligned model, each sampled at varying temperatures. ② Prompting-based:
 255 in-context resampling (Meyerson et al., 2024; Zhang et al., 2025b), where n outputs are generated
 256 sequentially within a single dialogue; paraphrase prompting (Jiang et al., 2020; Zhang et al., 2025b),
 257 where paraphrased variants of the same instruction are used to increase output diversity. ③ Decoding-
 258 based: Diverse Beam Search (Vijayakumar et al., 2016), where a diversity penalty was added to the
 259 beam search algorithm, applied to the aligned model. ④ Ensemble-based: response ensemble, where
 260 $n/2$ outputs are sampled from a model pair (a base model and its aligned counterpart) and pooled
 261 into a single group; logit ensemble, which merges the next-token probability distributions of the two
 262 models before sampling. ⑤ Collaboration-based: NUDGING (Fei et al., 2025), where an aligned
 263 model selectively introduces tokens during a base model’s decoding (§3). Note that diverse beam
 264 search, paraphrase prompting, and in-context resampling require additional computation; in-context
 265 resampling does not perform parallel sampling. These methods therefore provide more competitive
 266 reference points. The inference setups and experimental scope are provided in Appendix F.

BACO. Our experiments leverage two open-weight model pairs: Llama-3-8B and Llama-3-8B-Instruct (Grattafiori et al., 2024), and Olmo2-7B and Olmo2-7B-Instruct (OLMo et al., 2024) as they are widely used in literature (e.g., Fei et al. (2025)). We implement the single-strategy routers (e.g., -P which is based on maximum token probability) from Section 3.1 and denote multi-strategy routers as “-X-Y”, where strategy Y precedes X (e.g., -P-PUNC, -P-FC, and -H-PUNC). These implementations serve as representative examples that demonstrate the possible design space of the BACO framework. BACO framework works well with a wide range of routers. We include two basic routers as baselines: ① -RAND routes to the base model by random chance γ ; ② -JUDGE employs an external model to evaluate candidate tokens and makes a routing decision. Refer to Appendix C.2 for more details.

4.1 AUTOMATIC EVALUATIONS

We next evaluate the diversity-quality trade-off of each method using automatic metrics. Prior work has proposed a wide range of diversity evaluation that apply different lexical and semantic metrics, reflecting different perspectives on language diversity. Since our goal is to improve general diversity and quality rather than optimize for any specific metric, we adopt 11 established diversity metrics and 2 quality metrics which form 11×2 diversity-quality spaces, and then aggregate their results. Moreover, we aim to quantify the *controllability* of each method, i.e., the ability to adjust along the diversity-quality spectrum according to task or preference. As shown earlier in Figure 1B, each method is not evaluated as an individual point in a fixed configuration, but a curve formed by a sequence of points, illustrating the diversity and quality performance in different configurations. Each curve illustrates the diversity-quality trade-off of a specific method. To enable clear comparison, we apply two complementary indicators from multi-objective optimization, Coverage (Cov.) and Dominance (Dom.), to aggregate the curve-shaped performance across all spaces into numerical results (Figure 2).³

Coverage (Cov.) indicator quantifies the area under a method’s diversity-quality trade-off curve, following the hypervolume formulation used in multi-objective optimization. It measures how effectively a method traverses the diversity-quality spectrum as its control configuration varies. A larger Coverage value indicates that the method maintains good performance across a wide range of regions on diversity-quality spaces, offering usability to more different tasks or preferences overall.

Dominance (Dom.) indicator captures comparative optimality: whether and how often a method contributes to the global Pareto frontier among all methods. We compute the global Pareto frontier across all methods and apply the C-metric (Zitzler, 1999) to measure the portion of the frontier attributed to each method. A higher Dominance value indicates that a method achieves uniquely strong trade-offs unattainable by others.

To instantiate these indicators, we define a family of diversity-quality spaces, where each space corresponds to a particular pair of $11 \text{ diversity} \times 2 \text{ quality}$ metrics. We report the average *lexical*, *semantic*, and *overall* Coverage and Dominance, depending on the type of metrics as diversity axis.

Lexical diversity spaces use diversity metrics such as Distinct- n (Dist- n ; Li et al., 2015), EAD- n (Liu et al., 2022), and Self-BLEU (Montahaei et al., 2019). **Semantic diversity** spaces rely on diversity metrics such as embedding cosine dissimilarity (Kirk et al., 2023), Vendi Score (embedding) (Friedman & Dieng, 2022), NLI diversity (Stasaski & Hearst, 2022), and Semantic Entropy (Kuhn et al., 2023).⁴ Since lexical and semantic metrics capture fundamentally different aspects of diversity, we analyze them separately in addition to reporting aggregated results. Increases in lexical diversity are relatively easy to achieve (for example, by raising the temperature), yet they mostly alter surface-level phrasing without changing meaning. In contrast, semantic diversity reflects deeper diversity in structure, intent, and ideas, which is harder to elicit but more human-like and valuable in open-ended

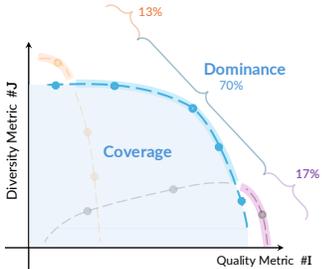


Figure 2: Illustration of the indicators on diversity-quality space: **Coverage**, the area under a method’s trade-off curve (blue shading for the blue method); **Dominance**, the proportion of the global Pareto frontier (highlighted curves) contributed by the method.

³All curves in Figure 2 are illustrative only and do not correspond to the actual performance of any method.

⁴Each metric could include multiple variants. For example, Dist-1/2/3 for different n-gram, or cosine dissimilarity under different pretrained encoders. Each variance leads to a separate diversity-quality space.

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

Method	Lexical		Semantic		Overall	
	Cov.	Dom.	Cov.	Dom.	Cov.	Dom.
Base	0.098	12.7%	0.098	16.0%	0.098	14.3%
Aligned	0.269	49.0%	0.104	29.2%	0.186	39.0%
Nudging	0.276	9.3%	0.247	9.9%	0.261	9.6%
Decoding	-	0.3%	-	0.3%	-	0.3%
Prompting (Best)	-	2.7%	-	2.2%	-	2.4%
Ensemble (Best)	-	1.1%	-	1.9%	-	1.5%
BACo (Best)	0.445	24.9%	0.360	40.5%	0.403	32.7%

Table 2: Averaged performance of all methods across all datasets and diversity–quality spaces. BACo consistently outperforms baselines across all semantic and most lexical spaces, demonstrating stronger controllability and substantially improving the semantic diversity–quality trade-off. The overall gains, as driven primarily by improvements in semantic, suggest that BACo produces more meaningful and content-level diversity, rather than superficial word-level changes, compared to other methods. See full results at Appendix G.

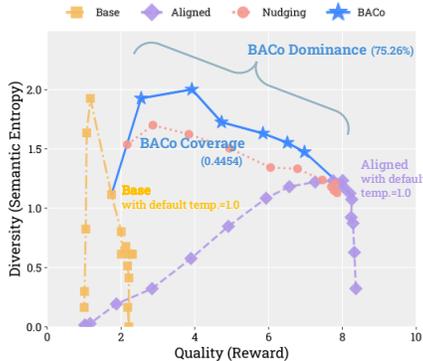


Figure 3: BACo’s performance on one diversity-quality space (x: quality in terms of reward; y: diversity in terms of semantic entropy). Comparing with baselines, BACo (blue curve) attains larger Coverage of the top-right region and contributes to most of the Dominance, indicating improvement and controllability on diversity-quality trade-off.

generation. **Quality** metrics include (i) perplexity under the aligned model, which reflects fluency and likelihood, and (ii) reward modeling scores predicted by Skywork-Reward-Gemma-2-27B (Liu et al., 2024), the state-of-the-art model on RewardBench (Lambert et al., 2024). These metrics are paired with diversity metrics to form subspaces. Finally, **overall** results average across all subspaces including lexical and semantic diversity metrics paired with quality metrics, yielding holistic method-level indicators. Hereafter, we use “Lexical” denote average results on all lexical diversity-quality spaces, “Semantic” denotes semantic diversity-quality spaces, and “Overall” averages all spaces in every result table. Full derivations and implementation details are in Appendix E.1.

5 RESULTS

In this section, we first compare BACo against a wide range of baselines across multiple datasets and metrics (§5.1). Next, we conduct router-level comparisons to analyze the effectiveness of different routing strategies (§5.2, §5.3). We also demonstrate BACo’s advantages on long-form diversity at the discourse level (§5.4). Lastly, we validate the results through human evaluation (§5.5).

5.1 BACo CONSISTENTLY OUTPERFORMS BASELINES

Table 2 summarizes the performance of BACo and baselines, aggregated across datasets and metrics. Overall, BACo improves Coverage by **0.142** and achieves **32.7%** Dominance. The advantage is particularly pronounced in semantic diversity, where its Dominance rises to **40.5%** (Figure 3 as an example).⁵ On the NoveltyBench dataset, the performance gap becomes even wider, with Coverage improving by **0.274** and Dominance reaching **39.9%**. These gains are consistent across all datasets and extend to the 01mo2 model family; see Appendix G for full results. See the output comparison examples at Appendix J.

5.2 ROUTER PERFORMANCE COMPARISON

Table 3 shows the performance of key routers on NoveltyBench. We use NoveltyBench in this section as it offers a representative yet computationally efficient setting for detailed comparison.

⁵The aligned model’s high lexical dominance arises from its high sampling temperature, which produces long, low-quality sequences that artificially inflate diversity scores while reducing controllability.

Random routing appears competitive only superficially (-RAND). At first glance, the RAND router performs surprisingly well, which may suggest that collaboration alone suffices without the need for curated router design. However, a closer look reveals that its gains are limited to shallow lexical metrics (e.g., Dist-n), while it fails on semantic metrics such as Semantic Entropy (0% dominance) (Appendix G). This pattern parallels the behavior of the high-temperature aligned model baseline: injecting randomness at the surface level increases token uncertainty but does not yield meaningful semantic diversity. These results also highlight that aggregated metrics can be misleading overall, and that evaluating diversity at the semantic level (especially more up-to-date ones like Semantic Entropy) is essential for understanding what constitutes genuinely diverse generation.

Prompt-based routers are ineffective (-JUDGE). We further consider -JUDGE, a variation of the BACO framework, inspired by the multi-agent system’s setup (Talebirad & Nadiri, 2023). In such systems, an aligned model often serves as a planner that allocates subtasks to different agents. Mirroring this design, we prompt an aligned model at each decoding step to make the routing decision. To enhance decision quality, we experiment with prompt engineering, including a step-by-step decision pipeline, curated heuristic rules, and few-shot examples with rationales (prompts are in Table A16). Despite these efforts, -JUDGE consistently underperforms simpler, lightweight routers across most metrics while being far more computationally expensive. This finding suggests the limited judgment of LLM to open-endedness (i.e., “when to diverge” and “how to diverge”), making them a poor fit for this task compared to more lightweight strategies we proposed.

Simple routers provide semantic gains (-P and -FC). Single-strategy routers like -P (based on next-token probability thresholds) and -FC (based on detecting function words) trade weaker overall performance for notable gains in semantic diversity, outperforming the -RAND baseline on key metrics like Cosine Dissimilarity. For instance, -P shows a modest but consistent advantage on Cosine Dissimilarity, Vendi Score (SimCSE), and Semantic Entropy (Appendix G). In contrast, -FC attains reasonable performance despite requiring only linguistic information.

Combining strategies achieves the best overall performance (-P-PUNC, -H-PUNC, and -P-FC). Combining strategies proves to be most effective in our experiments. The -P-PUNC router achieves the best overall performance, demonstrating high controllability (both Coverage and Dominance) across both lexical and semantic spaces. Other combinations like -H-PUNC and -P-FC, while not as strong overall, are still valuable as they capture unique regions of the global Pareto frontier (17.5% and

Routers	Lexical		Semantic		Overall	
	Cov.	Dom.	Cov.	Dom.	Cov.	Dom.
-RAND	0.493	26.3%	0.409	17.0%	0.451	21.7%
-JUDGE	0.302	2.6%	0.254	0.6%	0.278	1.6%
-P	0.433	4.8%	0.397	8.5%	0.415	6.7%
-FC	0.419	3.2%	0.382	4.7%	0.401	4.0%
-P-PUNC	0.495	30.7%	0.452	31.3%	0.474	31.0%
-H-PUNC	0.466	16.4%	0.427	18.6%	0.446	17.5%
-P-FC	0.435	16.0%	0.406	19.2%	0.421	17.6%

Table 3: Averaged performance of routers within BACO on NoveltyBench across all diversity–quality spaces. The -P-PUNC router achieves the best overall performance. While the random router (-RAND) attains moderately strong results, mainly from increased surface-level lexical diversity, its performance drops sharply on semantic metrics, confirming that unguided switching fails to produce meaningful diversity. In contrast, -P-PUNC delivers the most balanced and consistent results across both lexical and semantic evaluations, showing combination of designed routing strategies leads to more meaningful diversity.

Method	Lexical		Semantic		Overall	
	Cov.	Dom.	Cov.	Dom.	Cov.	Dom.
Base	0.142	9.8%	0.142	13.2%	0.142	11.5%
Aligned	0.273	43.8%	0.128	19.9%	0.200	31.8%
AACo	0.022	4.5%	0.006	7.0%	0.014	5.7%
BACO	0.495	42.0%	0.452	59.8%	0.474	50.9%

Table 4: Comparison of base-aligned and aligned-aligned collaboration (denoted as “AACo”) on NoveltyBench across all diversity–quality spaces. AACo yields little improvement on the diversity–quality trade-off compared with BACO, particularly in semantic diversity. The results demonstrate the necessity of involving a base and aligned model.

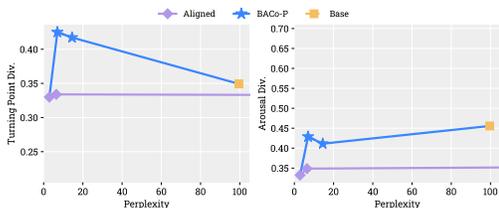


Figure 4: Comparison of BACO’s and baselines’ discourse-level diversity-quality trade-off curve on Narrative Discourse. BACO obtains a larger Coverage, achievable in the high-diversity, high-quality region (top-left). The results demonstrate it has richer discourse-level diversity without sacrificing quality largely, compared with baselines. The x-axis is quality (perplexity; lower is better), and the y-axis is discourse-level diversity, either turning-point diversity (left figure) or arousal diversity (right figure) (higher is better).

	NoveltyBench		WildChat	
	Aligned	BACO	Aligned	BACO
<i>Quality</i>	2.83	4.04	3.44	3.83
<i>Overall</i>	21.0%	79.0%	36.1%	63.9%
<i>Diversity Format</i>	25.4%	74.6%	26.8%	73.2%
<i>Content</i>	22.9%	77.1%	41.6%	58.4%
<i>Creativity</i>	20.4%	79.6%	38.2%	61.8%

Table 5: Human evaluation comparing BACO (best variant, -P-PUNC) with the aligned model baseline on NoveltyBench and WildChat. BACO is consistently preferred by human judges across all aspects, demonstrating that it produces outputs that are not only more human-perceived diverse but also higher in quality and creativity. Scores of quality are on a 1–5 Likert scale, and others indicate the pairwise win rate.

17.6% Dominance). This confirms that combining strategies yields complementary strengths, further advancing the trade-off frontier.

5.3 BASE-ALIGNED COLLABORATION OUTPERFORMS ALIGNED-ALIGNED

We test whether our base-aligned collaboration (BACO) is more effective than aligned-aligned collaboration for mitigating diversity-quality trade-off (Table 4). On NoveltyBench, we use the best router, -P-PUNC, to compare a base-aligned pair (Llama-3-8B family) against an aligned-aligned pair (Llama-3-8B with Llama-3.2-11B). The results show that **base-aligned collaboration significantly outperforms the aligned-aligned setup** in both Coverage and Dominance. These results highlight the limited diversity that arises from collaborating between two aligned models, underscoring our hypothesis that complementarity between less-aligned and more-aligned models is crucial for achieving both high diversity and high quality.

5.4 TOWARD LONG-FORM DIVERSITY: NARRATIVE DISCOURSE EVALUATION

Long-form outputs are often generated in creative writing tasks, where meaningful diversity of narratives often arises at the discourse level. However, existing metrics fails to measure it. Lexical metrics focus on surface-level wording and are unable to capture structural differences. Semantic metrics that use embeddings collapse entire narratives into single vectors, losing temporal and organizational dynamics (Deng et al., 2022; Huang et al., 2023). Results of these metrics (§4.1) are shown in Table A9, which also shows poor distinguishability. To address these limitations, we adopt a discourse-level diversity evaluation framework following Tian et al. (2024) to quantify long-form diversity, which captures the overall progression of a narrative’s structure. In particular, we assess narrative organization through story arcs, in terms of the distribution of discourse elements: *turning points* and *arousal patterns*.

Setup. We frame creative writing as a continuation task, where the model is given the beginning of a story and asked to complete it. To capture discourse-level variation, we measure in two dimensions, *Turning-point diversity* and *Arousal diversity*, following Tian et al. (2024) to provide complementary measures of long-form diversity. Detailed setups are provided in Appendix F.1.

Results. Figure 4 shows that BACO achieves markedly higher turning-point and arousal diversity than the aligned baseline of comparable quality, showing that it obtains greater Coverage and remains effective in the high-diversity and high-quality region (top-left). This result validates that base-aligned collaboration extends beyond short-form gains to long-form narratives, yielding richer plot structure and affective dynamics diversity while preserving quality. See output examples in Table A19.

5.5 HUMAN EVALUATION ON QUALITY, DIVERSITY, AND CREATIVITY

To complement automatic metrics, we conduct a three-phase human evaluation designed to assess output quality, group-level diversity, and creativity. The evaluation protocol aligns with the settings of automatic metrics.

Phase I: Quality. Annotators rate each output on a 1-5 Likert scale with respect to overall quality, considering fluency, relevance to the prompt, and informational substance. **Phase II: Diversity.** Annotators perform pairwise comparisons between output groups from two different methods. Diversity is evaluated along two specific aspects: (a) format diversity, capturing diversity in formatting or expression style; (b) content diversity, assessing diversity in underlying ideas. The overall diversity reflects the holistic impression of diversity, without any aspect specified. **Phase III: Creativity.** Annotators identify the single most creative output among all for a given prompt, and we record which method generated it. This evaluation aims to capture each method’s human-perceived creativity by its ability to generate the most creative output among group-level comparisons.

Setup. We compare BACO with the aligned model baseline under matched quality. From each dataset, 20 prompts are sampled with 3 outputs per method. Four annotators with LLM expertise evaluate on NoveltyBench and WildChat (see Appendix I for detailed setup and annotation).

Results. Shown in Table 5, although the automatic quality of the aligned model (5.933) and BACO (5.847) are nearly identical, human raters prefer BACO more, with a **42.8%** higher average rating and strong inter-rater agreement (Pearson correlation = 0.816, Intraclass correlation $ICC(2, k) = 0.907$). **On diversity, BACO demonstrates a significant win-rate over the aligned model**, with consistent improvements in both format (Fleiss’ $\kappa = 0.166$) and content (Fleiss’ $\kappa = 0.244$) sub-aspects, and overall diversity (Fleiss’ $\kappa = 0.268$)⁶. See human annotation examples at Appendix I.5.

5.6 EXTENSIONAL TASKS RESULTS

In this section, we extend our evaluation to two additional benchmarks to further demonstrate BACO’s performance. The experimental setup is identical to that described in §5.4.

Verifiable Instruction Following. To complement our earlier “soft” quality metrics, we evaluate on IFEval (Zhou et al., 2023b), which provides verifiable instructions. The diversity-quality trade-off observed in prior sections remains consistent. At matched quality levels, BACO achieves higher diversity than the aligned model, demonstrating a more effective optimization of this trade-off.

Mathematical Reasoning. We evaluate on GSM8K (Cobbe et al., 2021). Shown in Figure 6, both BACO and the aligned model show a trade-off between diversity and accuracy when tuned, but BACO incurs less accuracy loss while providing much higher diversity. Importantly, BACO maintains a Pass@10 score near 0.9 while consistently improving diversity metrics.

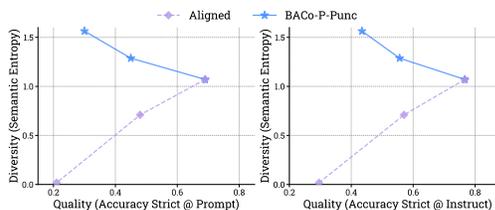


Figure 5: Diversity-quality trade-off comparison on IFEval. BACO demonstrates a more effective optimization of the trade-off. Refer full results in Figure A7.

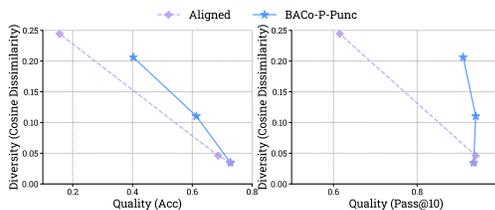


Figure 6: Diversity-accuracy trade-off comparison on GSM8K. BACO demonstrates a more effective optimization of the trade-off. Refer full results in Figure A8.

⁶The moderate agreement levels reflect the inherent subjectivity of format and content definitions, as well as the subjectivity of diversity judgments in open-ended tasks. Although annotators were provided with detailed guidelines and required to review them prior to annotation, the wide variety of prompts naturally allows for divergent interpretations.

ETHICS STATEMENT

This work does not involve any human subjects, sensitive personal data, or personally identifiable information. All datasets used are publicly available research benchmarks released under appropriate licenses. Our experiments focus on inference time with existing LLMs without additional fine-tuning, and thus pose minimal risks of memorization or privacy leakage beyond those already associated with the underlying models. While our method aims to improve diversity in model outputs, we acknowledge the possibility that enhanced diversity might amplify harmful generations if safeguards are not in place. This is a shared risk for most diversity-promoting technologies. We therefore emphasize that BACO is complementary to safety alignment and should be deployed with standard content-filtering or moderation practices. Importantly, BACO is designed for open-ended generation, rather than high-stakes or safety-critical domains. All authors have read and adhered to the ICLR Code of Ethics.

REPRODUCIBILITY STATEMENT

We have taken multiple steps to ensure the reproducibility of our work. The main paper specifies model configurations, routing strategies, evaluation metrics, and experimental protocols. Detailed descriptions of datasets, preprocessing steps, and evaluation setups are provided in Appendix D, with additional formulas and derivations in Appendix E. We report results across multiple model pairs to verify robustness in Appendix G. Human evaluation protocols and annotation guidelines are described in Appendix I. To further facilitate reproducibility, we will release code, configuration files, and evaluation scripts upon publication.

USAGE OF LARGE LANGUAGE MODELS

In this work, besides running LLMs in experiments, we use LLMs for the following purposes: *i*) Aid or Polish Writing (Gemini 2.5 Pro, ChatGPT 4/5); *ii*) Literature Retrieval and Discovery, e.g., finding related work (Gemini 2.5 Pro Deep Research, ChatGPT Deep Research); *iii*) Assisting Code Writing and Debugging (Claude 3.5 Sonnet). We fully understand the responsibility of using LLMs in academic research. We carefully monitor any potential problems, such as plagiarism or scientific misconduct, when using LLMs. We make sure these problems do not occur in the paper.

REFERENCES

- Joshua Ashkinaze, Julia Mendelsohn, Li Qiwei, Ceren Budak, and Eric Gilbert. How ai ideas affect the creativity, diversity, and evolution of human ideas: evidence from a large, dynamic experiment. In *Proceedings of the ACM Collective Intelligence Conference*, pp. 198–213, 2025.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Joeran Beel, Min-Yen Kan, and Moritz Baumgart. Evaluating sakana’s ai scientist for autonomous research: Wishful thinking or an emerging reality towards ‘artificial research intelligence’ (ari)?, 2025. URL <https://arxiv.org/abs/2502.14297>.
- Tuhin Chakrabarty, Philippe Laban, Divyansh Agarwal, Smaranda Muresan, and Chien-Sheng Wu. Art or artifice? large language models and the false promise of creativity. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pp. 1–34, 2024.
- Tuhin Chakrabarty, Philippe Laban, and Chien-Sheng Wu. Can ai writing be salvaged? mitigating idiosyncrasies and improving human-ai alignment in the writing process through edits. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pp. 1–33, 2025.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for 2+ 3=? on the overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*, 2024.

- 594 Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and Furu
595 Wei. Reasoning with Exploration: An Entropy Perspective on Reinforcement Learning for LLMs.
596 *arXiv e-prints*, art. arXiv:2506.14758, June 2025. doi: 10.48550/arXiv.2506.14758.
597
- 598 Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng
599 Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena:
600 An open platform for evaluating llms by human preference, 2024. URL [https://arxiv.org/
601 abs/2403.04132](https://arxiv.org/abs/2403.04132).
- 602 Lynn Chua, Badih Ghazi, Yangsibo Huang, Pritish Kamath, Ravi Kumar, Daogao Liu, Pasin Manu-
603 rangsi, Amer Sinha, and Chiyuan Zhang. Mind the privacy unit! user-level differential privacy for
604 language model fine-tuning. *arXiv preprint arXiv:2406.14322*, 2024.
605
- 606 John Joon Young Chung, Vishakh Padmakumar, Melissa Roemmele, Yuqian Sun, and Max Kreminski.
607 Modifying large language model post-training for diverse creative writing. *arXiv preprint
608 arXiv:2503.17126*, 2025.
- 609 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
610 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve
611 math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
612
- 613 Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen
614 Fan, Huayu Chen, Weize Chen, et al. The entropy mechanism of reinforcement learning for
615 reasoning language models. *arXiv preprint arXiv:2505.22617*, 2025.
616
- 617 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Jun-Mei Song, Ruoyu Zhang, Runxin Xu,
618 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiaoling Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F.
619 Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bing-Li Wang,
620 Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong
621 Ruan, Damai Dai, Deli Chen, Dong-Li Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo
622 Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
623 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
624 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Jiong Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong,
625 Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao,
626 Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, M. Tang, Meng
627 Li, Miaojuan Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu
628 Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, Ruiqi Jin, Ruyi
629 Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu,
630 Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shao-Kang Wu, Tao Yun, Tian Pei, Tianyu
631 Sun, T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wen-Xia
632 Yu, Wentao Zhang, Wangding Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xi aokang Chen,
633 Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng
634 Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei,
635 Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang,
636 Yifan Shi, Yi Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu,
637 Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yu-Jing Zou, Yujia He, Yunfan Xiong,
638 Yu-Wei Luo, Yu mei You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yao Li, Yi Zheng,
639 Yuchen Zhu, Yunxiang Ma, Ying Tang, Yukun Zha, Yuting Yan, Zehui Ren, Zehui Ren, Zhangli
640 Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhen guo Zhang, Zhewen Hao, Zhicheng Ma, Zhigang
641 Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zi-An Li, Ziwei Xie, Ziyang Song, Zizheng
642 Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing
643 reasoning capability in llms via reinforcement learning. *ArXiv*, abs/2501.12948, 2025. URL
<https://api.semanticscholar.org/CorpusID:275789950>.
- 644 Yuntian Deng, Volodymyr Kuleshov, and Alexander Rush. Model criticism for long-form text
645 generation. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022
646 Conference on Empirical Methods in Natural Language Processing*, pp. 11887–11912, Abu Dhabi,
647 United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/
v1/2022.emnlp-main.815. URL <https://aclanthology.org/2022.emnlp-main.815/>.

- 648 Yifu Ding, Wentao Jiang, Shunyu Liu, Yongcheng Jing, Jinyang Guo, Yingjie Wang, Jing Zhang,
649 Zengmao Wang, Ziwei Liu, Bo Du, Xianglong Liu, and Dacheng Tao. Dynamic parallel tree
650 search for efficient llm reasoning, 2025a. URL <https://arxiv.org/abs/2502.16235>.
- 651
652 Yifu Ding, Wentao Jiang, Shunyu Liu, Yongcheng Jing, Jinyang Guo, Yingjie Wang, Jing Zhang,
653 Zengmao Wang, Ziwei Liu, Bo Du, et al. Dynamic parallel tree search for efficient llm reasoning.
654 *arXiv preprint arXiv:2502.16235*, 2025b.
- 655 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
656 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
657 *arXiv preprint arXiv:2407.21783*, 2024.
- 658 Yu Fei, Yasaman Razeghi, and Sameer Singh. Nudging: Inference-time alignment of llms via guided
659 decoding, 2025. URL <https://arxiv.org/abs/2410.09300>.
- 660
661 Dan Friedman and Adji Bousso Dieng. The vendi score: A diversity evaluation metric for machine
662 learning. *arXiv preprint arXiv:2210.02410*, 2022.
- 663 Tianyu Gao, Xingcheng Yao, and Danqi Chen. Simcse: Simple contrastive learning of sentence
664 embeddings. *arXiv preprint arXiv:2104.08821*, 2021.
- 665
666 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
667 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of
668 models. *arXiv preprint arXiv:2407.21783*, 2024.
- 669 Tarun Gupta and Danish Pruthi. All that glitters is not novel: Plagiarism in ai generated research,
670 2025. URL <https://arxiv.org/abs/2502.16487>.
- 671
672 Sil Hamilton. Detecting mode collapse in language models via narration. *arXiv preprint*
673 *arXiv:2402.04477*, 2024.
- 674 Tenghao Huang, Ehsan Qasemi, Bangzheng Li, He Wang, Faeze Brahman, Muhao Chen, and
675 Snigdha Chaturvedi. Affective and dynamic beam search for story generation. In Houda Bouamor,
676 Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics:*
677 *EMNLP 2023*, pp. 11792–11806, Singapore, December 2023. Association for Computational
678 Linguistics. doi: 10.18653/v1/2023.findings-emnlp.789. URL [https://aclanthology.org/
679 2023.findings-emnlp.789/](https://aclanthology.org/2023.findings-emnlp.789/).
- 680 Fantine Huot, Reinald Kim Amplayo, Jennimaria Palomaki, Alice Shoshana Jakobovits, Elizabeth
681 Clark, and Mirella Lapata. Agents’ room: Narrative generation through multi-step collaboration.
682 In *The Thirteenth International Conference on Learning Representations*, 2025. URL [https:
683 //openreview.net/forum?id=HfWcFs7XLR](https://openreview.net/forum?id=HfWcFs7XLR).
- 684
685 Shawn Im and Yixuan Li. Understanding the learning dynamics of alignment with human feedback.
686 *arXiv preprint arXiv:2403.18742*, 2024.
- 687 Mete Ismayilzada, Antonio Laverghetta Jr, Simone A Luchini, Reet Patel, Antoine Bosselut, Lonneke
688 van der Plas, and Roger Beaty. Creative preference optimization. *arXiv preprint arXiv:2505.14442*,
689 2025.
- 690 Saskia Jaarsveld, Thomas Lachmann, and Cees Van Leeuwen. Creative reasoning across develop-
691 mental levels: Convergence and divergence in problem creation. *Intelligence*, 40(2):172–188,
692 2012.
- 693
694 Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. How can we know what language
695 models know? *Transactions of the Association for Computational Linguistics*, 8:423–438, 2020.
- 696
697 Seungone Kim, Juyoung Suk, Xiang Yue, Vijay Viswanathan, Seongyun Lee, Yizhong Wang, Kiril
698 Gashteovski, Carolin Lawrence, Sean Welleck, and Graham Neubig. Evaluating language models
699 as synthetic data generators. *arXiv preprint arXiv:2412.03679*, 2024.
- 700 Robert Kirk, Ishita Mediratta, Christoforos Nalmpantis, Jelena Luketina, Eric Hambro, Edward
701 Grefenstette, and Roberta Raileanu. Understanding the effects of rlhf on llm generalisation and
diversity. *arXiv preprint arXiv:2310.06452*, 2023.

- 702 Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for
703 uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664*, 2023.
704
- 705 Oh Joon Kwon, Daiki E Matsunaga, and Kee-Eung Kim. Gdpo: Learning to directly align language
706 models with diversity using gflownets. *arXiv preprint arXiv:2410.15096*, 2024.
707
- 708 Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu,
709 Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, et al. Rewardbench: Evaluating reward models
710 for language modeling. *arXiv preprint arXiv:2403.13787*, 2024.
- 711 Jack Lanchantin, Angelica Chen, Shehzaad Dhuliawala, Ping Yu, Jason Weston, Sainbayar
712 Sukhbaatar, and Ilia Kulikov. Diverse preference optimization. *arXiv preprint arXiv:2501.18101*,
713 2025.
- 714
- 715 Byung Cheol Lee and Jaeyeon Chung. An empirical investigation of the impact of chatgpt on
716 creativity. *Nature Human Behaviour*, 8(10):1906–1914, 2024.
- 717
- 718 Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative
719 decoding. In *International Conference on Machine Learning*, pp. 19274–19286. PMLR, 2023a.
- 720 Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative
721 decoding, 2023b. URL <https://arxiv.org/abs/2211.17192>.
- 722
- 723 Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. A diversity-promoting
724 objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*, 2015.
- 725
- 726 Tianjian Li, Yiming Zhang, Ping Yu, Swarnadeep Saha, Daniel Khashabi, Jason Weston, Jack Lan-
727 chantin, and Tianlu Wang. Jointly reinforcing diversity and quality in language model generations,
728 2025. URL <https://arxiv.org/abs/2509.02534>.
- 729
- 730 Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori B Hashimoto, Luke
731 Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization.
732 In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*
(Volume 1: Long Papers), pp. 12286–12312, 2023.
- 733
- 734 Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu,
735 Chandra Bhagavatula, and Yejin Choi. The unlocking spell on base llms: Rethinking alignment via
736 in-context learning. In *The Twelfth International Conference on Learning Representations*, 2023.
- 737
- 738 Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization*
branches out, pp. 74–81, 2004.
- 739
- 740 Chris Yuhao Liu, Liang Zeng, Jiakai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang
741 Liu, and Yahui Zhou. Skywork-reward: Bag of tricks for reward modeling in llms. *arXiv preprint*
742 *arXiv:2410.18451*, 2024.
- 743
- 744 Siyang Liu, Sahand Sabour, Yinhe Zheng, Pei Ke, Xiaoyan Zhu, and Minlie Huang. Rethinking and re-
745 fining the distinct metric. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Pro-*
746 *ceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2:*
747 *Short Papers)*, pp. 762–770, Dublin, Ireland, May 2022. Association for Computational Linguistics.
doi: 10.18653/v1/2022.acl-short.86. URL <https://aclanthology.org/2022.acl-short.86/>.
- 748
- 749 Li-Chun Lu, Shou-Jen Chen, Tsung-Min Pai, Chan-Hung Yu, Hung-yi Lee, and Shao-Hua Sun.
750 Llm discussion: Enhancing the creativity of large language models via discussion framework and
751 role-play. In *First Conference on Language Modeling*, 2024.
- 752
- 753 Ximing Lu, Melanie Sclar, Skyler Hallinan, Niloofar Miresghallah, Jiacheng Liu, Seungju Han,
754 Allyson Ettinger, Liwei Jiang, Khyathi Chandu, Nouha Dziri, and Yejin Choi. AI as humanity’s
755 salieri: Quantifying linguistic creativity of language models via systematic attribution of machine
text against web text. In *The Thirteenth International Conference on Learning Representations*,
2025a. URL <https://openreview.net/forum?id=i10E0IqolQ>.

- 756 Ximing Lu, Melanie Sclar, Skyler Hallinan, Niloofar Miresghallah, Jiacheng Liu, Seungju Han,
757 Allyson Ettinger, Liwei Jiang, Khyathi Chandu, Nouha Dziri, and Yejin Choi. Ai as humanity’s
758 salieri: Quantifying linguistic creativity of language models via systematic attribution of machine
759 text against web text, 2025b. URL <https://arxiv.org/abs/2410.04265>.
- 760 Lennart Meincke, Gideon Nave, and Christian Terwiesch. Chatgpt decreases idea diversity in
761 brainstorming. *Nature Human Behaviour*, 9(6):1107–1109, June 2025a. ISSN 2397-3374. doi:
762 10.1038/s41562-025-02173-x. Epub 2025 May 14.
- 763 Lennart Meincke, Gideon Nave, and Christian Terwiesch. Chatgpt decreases idea diversity in
764 brainstorming. *Nature human behaviour*, pp. 1–3, 2025b.
- 765 Elliot Meyerson, Mark J Nelson, Herbie Bradley, Adam Gaier, Arash Moradi, Amy K Hoover,
766 and Joel Lehman. Language model crossover: Variation through few-shot prompting. *ACM*
767 *Transactions on Evolutionary Learning*, 4(4):1–40, 2024.
- 768 Ehsan Montahaei, Danial Alihosseini, and Mahdih Soleymani Baghshah. Jointly measuring diversity
769 and quality in text generation models. *arXiv preprint arXiv:1904.03971*, 2019.
- 770 Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia,
771 Yuling Gu, Shengyi Huang, Matt Jordan, et al. 2 olmo 2 furious. *arXiv preprint arXiv:2501.00656*,
772 2024.
- 773 OpenAI. Introducing deep research. [https://openai.com/index/
774 introducing-deep-research/](https://openai.com/index/introducing-deep-research/), February 2025. Accessed: November 26, 2025.
- 775 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
776 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow
777 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–
778 27744, 2022.
- 779 Vishakh Padmakumar and He He. Does writing with language models reduce content diversity?
780 *arXiv preprint arXiv:2309.05196*, 2023.
- 781 Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic
782 evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association
783 for Computational Linguistics*, pp. 311–318, 2002.
- 784 Max Peepkorn, Tom Kouwenhoven, Dan Brown, and Anna Jordanous. Is temperature the creativity
785 parameter of large language models? In *ICCC*, 2024.
- 786 Jinghua Piao, Jiazhen Liu, Fang Zhang, Jun Su, and Yong Li. Human–ai adaptive dynamics drives
787 the emergence of information cocoons. *Nature Machine Intelligence*, 5(11):1214–1224, 2023.
- 788 Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson.
789 Fine-tuning aligned language models compromises safety, even when users do not intend to! *arXiv
790 preprint arXiv:2310.03693*, 2023.
- 791 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
792 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances
793 in neural information processing systems*, 36:53728–53741, 2023.
- 794 Alexandre Rame, Guillaume Couairon, Corentin Dancette, Jean-Baptiste Gaya, Mustafa Shukor,
795 Laure Soulier, and Matthieu Cord. Rewarded soups: towards pareto-optimal alignment by interpolating
796 weights fine-tuned on diverse rewards. *Advances in Neural Information Processing Systems*,
797 36:71095–71134, 2023.
- 798 Yi Ren and Danica J. Sutherland. Learning dynamics of llm finetuning, 2025. URL [https://arxiv.
799 org/abs/2407.10490](https://arxiv.org/abs/2407.10490).
- 800 Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and
801 Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In
802 *International Conference on Learning Representations*, 2017.
- 803
804
805
806
807
808
809

- 810 Weijia Shi, Akshita Bhagia, Kevin Farhat, Niklas Muennighoff, Pete Walsh, Jacob Morrison, Dustin
811 Schwenk, Shayne Longpre, Jake Poznanski, Allyson Ettinger, et al. Flexolmo: Open language
812 models for flexible data use. *arXiv preprint arXiv:2507.07024*, 2025.
- 813
814 Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. Ai
815 models collapse when trained on recursively generated data. *Nature*, 631(8022):755–759, 2024.
816 doi: 10.1038/s41586-024-07566-y. URL <https://doi.org/10.1038/s41586-024-07566-y>.
- 817 Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. Can llms generate novel research ideas? a large-
818 scale human study with 100+ nlp researchers, 2024. URL <https://arxiv.org/abs/2409.04109>.
- 819 Momin N Siddiqui, Roy D Pea, and Hari Subramonyam. Script&shift: A layered interface paradigm
820 for integrating content development and rhetorical strategy with llm writing assistants. In *Proceed-*
821 *ings of the 2025 CHI Conference on Human Factors in Computing Systems*, pp. 1–19, 2025.
- 822
823 Matthew Sims, Jong Ho Park, and David Bamman. Literary event detection. In *Proceedings of the*
824 *57th annual meeting of the association for computational linguistics*, pp. 3623–3634, 2019.
- 825 Yuda Song, Julia Kempe, and Remi Munos. Outcome-based exploration for llm reasoning, 2025.
826 URL <https://arxiv.org/abs/2509.06941>.
- 827
828 Alexander Spangher, Tenghao Huang, Philippe Laban, and Nanyun Peng. Creative planning with
829 language models: Practice, evaluation and applications. In Maria Lomeli, Swabha Swayamdipta,
830 and Rui Zhang (eds.), *Proceedings of the 2025 Annual Conference of the Nations of the Ameri-*
831 *cas Chapter of the Association for Computational Linguistics: Human Language Technologies*
832 *(Volume 5: Tutorial Abstracts)*, pp. 1–9, Albuquerque, New Mexico, May 2025. Association for
833 Computational Linguistics. ISBN 979-8-89176-193-3. doi: 10.18653/v1/2025.naacl-tutorial.1.
834 URL <https://aclanthology.org/2025.naacl-tutorial.1/>.
- 835 Katherine Stasaski and Marti Hearst. Semantic diversity in dialogue with natural language infer-
836 ence. In Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (eds.),
837 *Proceedings of the 2022 Conference of the North American Chapter of the Association for Com-*
838 *putational Linguistics: Human Language Technologies*, pp. 85–98, Seattle, United States, July
839 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.6. URL
840 <https://aclanthology.org/2022.naacl-main.6/>.
- 841 Yashar Talebirad and Amirhossein Nadiri. Multi-agent collaboration: Harnessing the power of
842 intelligent llm agents. *arXiv preprint arXiv:2306.03314*, 2023.
- 843
844 Yufei Tian, Tenghao Huang, Miri Liu, Derek Jiang, Alexander Spangher, Muhao Chen, Jonathan
845 May, and Nanyun Peng. Are large language models capable of generating human-level narratives?
846 *arXiv preprint arXiv:2407.13248*, 2024.
- 847 Saranya Venkatraman, Nafis Irtiza Tripto, and Dongwon Lee. Collabstory: Multi-llm collaborative
848 story generation and authorship analysis. In *Findings of the Association for Computational*
849 *Linguistics: NAACL 2025*, pp. 3665–3679, 2025.
- 850 Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David
851 Crandall, and Dhruv Batra. Diverse beam search: Decoding diverse solutions from neural sequence
852 models. *arXiv preprint arXiv:1610.02424*, 2016.
- 853
854 Angelina Wang, Jamie Morgenstern, and John P Dickerson. Large language models that replace hu-
855 man participants can harmfully misportray and flatten identity groups. *Nature Machine Intelligence*,
856 pp. 1–12, 2025a.
- 857 Chaoqi Wang, Yibo Jiang, Chenghao Yang, Han Liu, and Yuxin Chen. Beyond reverse kl: Generaliz-
858 ing direct preference optimization with diverse divergence constraints. In *The Twelfth International*
859 *Conference on Learning Representations*, 2024.
- 860 Qihan Wang, Shidong Pan, Tal Linzen, and Emily Black. Multilingual prompting for improving llm
861 generation diversity. *arXiv preprint arXiv:2505.15229*, 2025b.
- 862
863 Peter West and Christopher Potts. Base models beat aligned models at randomness and creativity.
arXiv preprint arXiv:2505.00047, 2025.

- 864 Justin Wong, Yury Orlovskiy, Michael Luo, Sanjit A Seshia, and Joseph E Gonzalez. Simplestrat:
865 Diversifying language model generation with stratification. *arXiv preprint arXiv:2410.09038*,
866 2024.
- 867 Suqing Wu, Yukun Liu, Mengqi Ruan, Siyu Chen, and Xiao-Yun Xie. Human-generative ai collabora-
868 tion enhances task performance but undermines human’s intrinsic motivation. *Scientific Reports*,
869 15(1):15105, 2025.
- 870 Xun Wu, Shaohan Huang, and Furu Wei. Mixture of lora experts, 2024. URL <https://arxiv.org/abs/2404.13628>.
- 871 Han Xiao. Submodular optimization for diverse query generation in deep-
872 research. Jina AI Tech Blog, July 2025. URL <https://jina.ai/news/submodular-optimization-for-diverse-query-generation-in-deepresearch/>. 13
873 minutes read.
- 874 Weijia Xu, Nebojsa Jojic, Sudha Rao, Chris Brockett, and Bill Dolan. Echoes in ai: Quantifying
875 lack of plot diversity in llm outputs. *Proceedings of the National Academy of Sciences*, 122(35):
876 e2504966122, 2025.
- 877 Chenghao Yang and Ari Holtzman. How alignment shrinks the generative horizon. *arXiv preprint*
878 *arXiv:2506.17871*, 2025.
- 879 Yuming Yang, Yang Nan, Junjie Ye, Shihan Dou, Xiao Wang, Shuo Li, Huijie Lv, Mingqi Wu, Tao
880 Gui, Qi Zhang, et al. Measuring data diversity for instruction tuning: A systematic analysis and a
881 reliable metric. *arXiv preprint arXiv:2502.17184*, 2025.
- 882 Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. Plan-and-
883 write: Towards better automatic storytelling. In *Proceedings of the AAAI Conference on Artificial*
884 *Intelligence*, volume 33, pp. 7378–7385, 2019.
- 885 Jingyu Zhang, Ahmed Elgohary, Ahmed Magooda, Daniel Khashabi, and Benjamin Van Durme.
886 Controllable safety alignment: Inference-time adaptation to diverse safety requirements, 2025a.
887 URL <https://arxiv.org/abs/2410.08968>.
- 888 Xuanchang Zhang, Wei Xiong, Lichang Chen, Tianyi Zhou, Heng Huang, and Tong Zhang. From
889 lists to emojis: How format bias affects model alignment. *arXiv preprint arXiv:2409.11704*, 2024.
- 890 Yiming Zhang, Harshita Diddee, Susan Holm, Hanchen Liu, Xinyue Liu, Vinay Samuel, Barry Wang,
891 and Daphne Ippolito. Noveltybench: Evaluating language models for humanlike diversity. *arXiv*
892 *preprint arXiv:2504.05228*, 2025b.
- 893 Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. Wildchat: 1m
894 chatgpt interaction logs in the wild. *arXiv preprint arXiv:2405.01470*, 2024.
- 895 Chujie Zheng, Ziqi Wang, Heng Ji, Minlie Huang, and Nanyun Peng. Model extrapolation exped-
896 ites alignment. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher
897 Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational*
898 *Linguistics (Volume 1: Long Papers)*, pp. 1025–1041, Vienna, Austria, July 2025. Association for
899 Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.51. URL
900 <https://aclanthology.org/2025.acl-long.51/>.
- 901 Kai Zheng, Ren-Jye Yang, Hongyi Xu, and Jie Hu. A new distribution metric for comparing pareto
902 optimal solutions. *Struct. Multidiscip. Optim.*, 55(1):53–62, January 2017. ISSN 1615-147X. doi:
903 10.1007/s00158-016-1469-3. URL <https://doi.org/10.1007/s00158-016-1469-3>.
- 904 Wenhao Zheng, Yixiao Chen, Weitong Zhang, Souvik Kundu, Yun Li, Zhengzhong Liu, Eric P. Xing,
905 Hongyi Wang, and Huaxiu Yao. CITER: Collaborative inference for efficient large language model
906 decoding with token-level routing. In *Adaptive Foundation Models: Evolving AI for Personalized*
907 *and Efficient Learning*, 2024. URL <https://openreview.net/forum?id=0tSupTUgFP>.
- 908 Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia
909 Efrat, Ping Yu, Lili Yu, et al. Lima: Less is more for alignment. *Advances in Neural Information*
910 *Processing Systems*, 36:55006–55021, 2023a.

918 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny
 919 Zhou, and Le Hou. Instruction-following evaluation for large language models, 2023b. URL
 920 <https://arxiv.org/abs/2311.07911>.

921
 922 Eckart Zitzler. *Evolutionary algorithms for multiobjective optimization: Methods and applications*,
 923 volume 63. Shaker Ithaca, 1999.

924 925 A ANALYSIS AND DISCUSSION 926

927 In this section, we first analyze the collaboration pattern between the base and aligned model in
 928 BACO (§A.1 and §A.2). We then extend human evaluation from diversity to creativity, presenting
 929 BACO’s additional advantages (§A.3); and lastly discuss thoughts on future work (§A.4).
 930

931 A.1 CONTRIBUTION DISTRIBUTION AND SWITCHING FREQUENCY 932

933 We measure how often each model contributes to generation at at each token position across all
 934 generations. We also count the frequency of model switching at each token position. For BACO
 935 with the best router (-P-PUNC), **base-model contribution and switching frequency are high at the**
 936 **start of generation and decrease over time across all three tasks** (Figure A11 - A9). Whether this
 937 pattern is desirable is task dependent and partly subjective. We hypothesize that it reflects the model’s
 938 increasing predictive confidence as generation progresses. Such behavior may manifest as a temporal
 939 trend (e.g., over the course of decoding) that coincides with positional patterns in the text. For
 940 example, in creative writing tasks, early tokens tend to correspond to exploratory narrative positions
 941 (e.g., introducing background, characters, and conflicts), while later tokens naturally emphasize
 942 coherence (e.g., resolving plotlines or tying back to earlier events). However, it may be less ideal
 943 for tasks like open-ended question answering, especially since most outputs are written in list form,
 944 which may require more uniform exploration. This pattern could be tailored for specific tasks by
 945 adopting dynamic, position-aware thresholds, which we leave for future work.

946 A.2 INHERENT EARLY STOP: FAILURE MODE IN MODEL COLLABORATION 947

948 We qualitatively and quantitatively observe that when the router is tuned more aggressively toward
 949 diversity (more token generated by the base model), the collaboration system exhibits a higher
 950 tendency to terminate early (i.e., one of the models produces an end-of-sentence (eos) token).
 951 Examples are in Table A12 and Table A13. This phenomenon differs from conventional decoding or
 952 reasoning approaches, where early stopping is deliberately designed as an efficiency mechanism-such
 953 as pruning low-confidence branches in dynamic search methods (Ding et al., 2025a). In contrast, our
 954 system demonstrates *inherent early stopping* as an emergent behavior of model collaboration with
 955 lower-probability content introduced. This property carries both drawbacks and advantages.

956 On the negative side, early stopping may prematurely truncate outputs, especially for less-frequent
 957 prefixes that could otherwise lead to valid and high-quality continuations. On the positive side, early
 958 stopping can prevent the system from degenerating into endless loops or incoherent repetition, a
 959 common issue at higher temperatures. In such cases, emitting an eos token provides a safe and
 960 efficient termination. Moreover, since these cases are easily detectable given the length, a practical
 961 solution is to simply restart sampling when premature termination occurs. Overall, inherent early
 962 stopping reflects a distinctive failure mode in model communication under collaborative decoding.
 963 While it introduces risks of incompleteness, it also offers a built-in safeguard against low-quality
 964 continuations, highlighting an important dimension for future router design.

965 A.3 EVALUATION EXTENSION FROM DIVERSITY TO CREATIVITY 966

967 Our human evaluation (§5.5) not only assessed diversity but also included a phase for creativity, which
 968 is often defined as originality and human-perceived value. As creativity is more likely to be an explicit
 969 target in many real-world application, such as writing or science discovery, we therefore examine
 970 whether BACO’s improvements in diversity-quality trade-off translate into enhanced creativity.

971 Creativity represents a higher-level dimension of generation diversity and quality. Evaluating creativ-
 ity in LLMs, however, remains an open and contested challenge: prior studies differ in definitions

and settings, leading to inconsistent conclusions (e.g., (Si et al., 2024) vs. (Gupta & Pruthi, 2025); (Lu et al., 2025b) vs. (Beel et al., 2025); (Lee & Chung, 2024) vs. (Meincke et al., 2025b; Wu et al., 2025; Xu et al., 2025)). We argue that a central source of discrepancy lies in how creativity is related to diversity. Some studies (e.g., Lee & Chung, 2024) find positive evidence when evaluating individual outputs, whereas others (e.g., Meincke et al., 2025b) report negative conclusions when assessing groups of outputs. Building on Jaarsveld et al. (2012), we adopt a dual perspective that treats creativity as a balance between *divergent thinking* (captured by group-level diversity across outputs) and *convergent thinking* (captured by quality and novelty within individual outputs). Our three-phase human evaluation is designed to align with this view: Phase I assesses quality, partially reflecting convergent thinking; Phase II measures diversity, corresponding to divergent thinking; and Phase III directly evaluates creativity, integrating both dimensions into a unified evaluation. This design allows us to examine whether improvements in diversity through BACO translate into genuinely more creative and human-valued generations.

Our results support this dual view. In the human evaluation, BACO’s outputs were consistently rated as higher in both quality and diversity and, more importantly, were significantly more likely to be judged as the most creative (Fleiss’ $\kappa = 0.485$). These findings suggest that BACO’s collaborative decoding not only optimize the diversity–quality trade-off but also fosters human-perceived creativity, offering a promising step toward generation that is both diverse and valuable original (see examples in Appendix I.5).

A.4 FUTURE WORK

We envision BACO as a step toward enabling a “breadth thinking” mode for LLMs, complementing the existing focus on “deep thinking” (OpenAI, 2025; DeepSeek-AI et al., 2025). We refer to “breadth” as a model’s inherent ability to explore a wide, validated space, considering multiple perspectives, styles, or conceptual directions, rather than converging on a single, narrowly optimized output under open-ended instructions. For open-ended tasks, breadth thinking offers a way to help human break out of their “information cocoons” (Piao et al., 2023), expand their ideation space, and think beyond conventional boundaries. BACO provides a practical initial step: it is controllable, inference-time only, and deployable out of the box without requiring additional post-training or complex agentic workflows. It may also serve as a plug-in for broader system designs, for example as an ideation agent in multi-agent collaboration (Siddiqui et al., 2025) or as an exploratory stage within problem-solving pipelines (Cheng et al., 2025; Song et al., 2025). A natural direction is to design more sophisticated, task-aware routers that adaptively learn routing strategies from data or rewards rather than relying only on hand-crafted heuristics.

There is also room for improving efficiency. Speculative decoding (Leviathan et al., 2023b) is a perfect fit for token-level routing in the framework, especially when looking ahead, which could reduce inference-time overhead. Moreover, while our current experiments use base and aligned models of the same family, memory-efficient alternatives such as LoRA-tuned aligned models could further reduce deployment costs, replacing $N \times$ model size with a single base model plus $N \times$ LoRA adapter size (Wu et al., 2024).

Another promising direction lies in exploring the dynamics of alignment beyond the two canonical checkpoints we studied (base and fully aligned) (Im & Li, 2024; Ren & Sutherland, 2025). While these checkpoints are the most accessible in open-source releases, they are unlikely to represent the optimal trade-off points. Intermediate or partially aligned checkpoints may offer a different quality-diversity trade-off, and a systematic study of it could further extend the diversity-quality trade-off frontier.

Finally, human-centered usability should not be overlooked. The controllability of BACO offers opportunities for user-facing interfaces: for example, sliders that let users explicitly shift along the quality-diversity spectrum, or adaptive systems that learn user preferences over time. Such design could make breadth thinking practically useful in everyday workflows, not only research benchmarks. Beyond the quality-diversity balance, we also see potential in applying BACO to other inference-time objectives such as multi-level privacy (Chua et al., 2024) or controllable safety (Zhang et al., 2025a), highlighting its promise as a general paradigm for collaborative inference.

In summary, we view BACO as the first step toward a broader agenda of inference-time collaboration, leveraging pre- or intermediate-alignment checkpoints to enable multi-objective optimization and achieve practical balance in real-world use.

B RELATED WORK

Diversity Degradation in Alignment. While alignment techniques like RLHF enhance model performance in instruction following and reasoning, they systematically curtail output diversity. This trade-off is widely recognized, with a growing body of evidence demonstrating that aligned models are less diverse than their base counterparts. For example, studies have questioned their artistic authenticity (Chakrabarty et al., 2024), and benchmarks like NoveltyBench reveal their diminished capacity for humanlike diversity and creativity (Zhang et al., 2025b; Tian et al., 2024; Lu et al., 2025a; West & Potts, 2025). The underlying mechanism for this degradation is “probability concentration,” as the alignment process “sharpens” the model’s output distribution, thereby steering it into low-entropy generation paths (Yang & Holtzman, 2025). The diversity degradation impacts the downstream applications. It manifests as a loss of linguistic idiosyncrasies (Chakrabarty et al., 2025), increased format homogeneity (Zhang et al., 2024), and diminished creativity (West & Potts, 2025) in generated text. Beyond linguistic characteristics, alignment-induced constraints lead models to converge on a restricted repertoire of strategies, thereby diminishing diversity in reasoning (Chen et al., 2024; Ding et al., 2025b), data synthesis (Kim et al., 2024; Yang et al., 2025), deep research (Xiao, 2025), social simulation (Wang et al., 2025a), and gaming (West & Potts, 2025). More critically, the lack of output diversity has been shown to further reduce outcome diversity (Padmakumar & He, 2023) and creativity (Meinke et al., 2025a; Ashkinaze et al., 2025) in human interactions with these models. These studies demonstrate that diminished diversity in model outputs adversely affects how humans ideate, create, and engage. Our work addresses this challenge directly, proposing an inference-time collaborative decoding framework that optimizes the diversity-quality trade-off by combining the strengths of both base and aligned models.

Diversity-Promoting Methods. Approaches to enhance the diversity of aligned LLMs fall into two main categories: training-time and inference-time methods. **Training-time methods** typically modify the learning objective to encourage varied outputs. A prominent line of work adapts Direct Preference Optimization (DPO; Rafailov et al. (2023)) by incorporating diversity-aware mechanisms, such as f-divergence penalties (Wang et al., 2024), set-level diversity rewards (Lanchantin et al., 2025), or re-weighted loss objectives (Chung et al., 2025; Ismayilzada et al., 2025). Other approaches leverage different architectures, like generative flow networks, to the same end (Kwon et al., 2024). While these methods can instill diversity directly into the model, they require substantial computational resources for retraining and offer little flexibility for user-specific diversity needs at inference. **Inference-time methods** offer a more lightweight and adaptable alternative. These include modifications to decoding algorithms like diverse beam search (Vijayakumar et al., 2016) and various prompt engineering strategies, such as paraphrasing (Meyerson et al., 2024; Wang et al., 2025b; Zhang et al., 2025b; Wang et al., 2025a; Wong et al., 2024). Existing inference-time methods for improving diversity typically incur high computational costs via multiple decoding passes or long-horizon planning. Or they significantly degrade generation quality (Peeperkorn et al., 2024). Achieving a stable diversity-quality trade-off with these techniques remains a challenge. Our proposed method BACO, is an inference-time framework designed to offer more explicit and reliable control over this trade-off. We therefore focus our comparison on baselines from this category.

Multi-Model Collaborative Generation. Prior work has explored collaborative frameworks where multiple language models work in concert to improve generation quality and efficiency (e.g., computation cost and latency). These approaches can be grouped by their collaboration mechanism. One line of work focuses on **weight-level collaboration**. This includes merging reward or policy models, or using Mixture-of-Experts (MoE) (Shazeer et al., 2017) architectures to create a single, more capable system better aligned with diverse human preferences (Rame et al., 2023; Zheng et al., 2025; Shi et al., 2025). Another mechanism is **token-level collaboration**, where multiple models collaborate during decoding by exchanging next-token probability distributions or candidate token choices at each step, to improve attributes like coherence and factuality, or reduce latency (Leviathan et al., 2023a; Li et al., 2023; Zheng et al., 2024; Fei et al., 2025). More recently, **multi-agent systems** have emerged, in which models debate or discuss to leverage their complementary strengths for

1080 complex, creative tasks (Lu et al., 2024; Venkatraman et al., 2025; Huot et al., 2025). Our BACO,
1081 advances token-level collaboration. While existing methods in this area primarily target quality or
1082 efficiency, we focus on navigating the diversity-quality trade-off.
1083

1084 C BACO FRAMEWORK IMPLEMENTATION AND ROUTER DETAILS 1085

1086 **Models.** We apply each model’s default chat template during collaboration (for base models, we
1087 apply a plain shifting template), served on vLLM local host. We disable every model’s tool call and
1088 thinking for fair comparison.
1089

1090 **Tokenization Alignment.** Tokenization sometimes mismatches between the two models (even
1091 base and aligned pairs), particularly around punctuation, special tokens, or rare words, can lead
1092 to incoherent sub-word boundaries. To address this, we enforce that tokens representing a single
1093 semantic unit (e.g., a word or format element) must all be produced by the same model. This avoids
1094 artifacts such as broken punctuation or malformed words.
1095

1096 **Framework Cost.** Despite the apparent overhead of token-level switching, BACO is designed with
1097 consideration for efficiency. The feasibility of BACO stems from the *superficial alignment* (Zhou
1098 et al., 2023a; Lin et al., 2023), which implies that interventions between the largely-agreeing base
1099 and aligned models can be sparse (Fei et al., 2025). This sparsity enables practical optimizations
1100 for efficiency, such as caching multi-token chunks from one model to minimize switching costs.
1101 Further in deployment, the framework allows for standard system optimizations, like overlapping
1102 communication and computation to hide context-switching latency.
1103

1104 C.1 ADDITIONAL RULES IN ROUTER

1105 We follow Fei et al. (2025) in always using the aligned model to generate the first token. Early
1106 decoding steps have an outsized influence on generation and typically show greater disagreement
1107 between models. Starting from the aligned model improves trajectory quality and reduces the chance
1108 of degenerate completions.
1109

1110 Incorporating low-probability tokens from the base model introduces new challenges in sequential
1111 generation. When switch between models, the receiving model may struggle to continue from an
1112 unfamiliar context. In particular, the aligned model may terminate the output prematurely, while the
1113 base model may fall into degenerate behaviors such as repetition or verbose listing. To mitigate this,
1114 we constrain output termination by only accepting the end-of-sentence token when it is the top-1
1115 prediction of the aligned model.
1116

1117 C.2 ALL ROUTING STRATEGIES AND NOTATIONS

1118 The following are all strategies and their corresponding notations that are mentioned in this paper:

- 1119 • -RAND: Route to the base model by random chance γ .
- 1120 • -P: Route to the base model when base model’s top-1 token probability $\max_{y_t} P_{\text{base}}(y_t \mid$
1121 $x, y_{<t}) < \gamma \gamma \in [0, 1]$, otherwise to the aligned model.
- 1122 • -P-A: Route to the base model when aligned model’s top-1 token probability
1123 $\max_{y_t} P_{\text{aligned}}(y_t \mid x, y_{<t}) < \gamma \gamma \in [0, 1]$, otherwise to the aligned model.
- 1124 • -H: Route to the base model when the entropy of the base model’s next token prediction
1125 distribution $H_{\text{base}}(y_t \mid x, y_{<t}) > \gamma \gamma \in [0, +\infty)$, otherwise to the aligned model.
- 1126 • -H-A: Route to the base model when the entropy of the aligned model’s next token prediction
1127 distribution $H_{\text{aligned}}(y_t \mid x, y_{<t}) > \gamma \gamma \in [0, +\infty)$, otherwise to the aligned model.
- 1128 • -PR: Route to the base model when the ratio between the base model’s top 1 token probability
1129 and the aligned model’s, i.e., $\frac{\max_{y_t} P_{\text{base}}(y_t \mid x, y_{<t})}{\max_{y_t} P_{\text{aligned}}(y_t \mid x, y_{<t})} < \gamma \gamma \in (0, 1]$, otherwise to the aligned
1130 model.
- 1131 • -HR: Route to the base model when the ratio between the base model’s entropy and the
1132 aligned model’s, i.e., $\frac{H_{\text{base}}(c_i \mid q, c_{<i})}{H_{\text{aligned}}(c_i \mid q, c_{<i})} > \gamma \gamma \in [1, +\infty)$, otherwise to the aligned model.
1133

- -FC: Route to the base model when both the aligned model and base model sample the next token is a content word⁷, otherwise to the aligned model.
- -PUNC: Route to the base model when the base model’s and the aligned model’s next token is not punctuation or formative tokens (e.g., ‘\n’), otherwise, to the aligned model.
- -JUDGE: Route to the base model when an external judge LLM (another aligned model) determines that both of the following conditions are satisfied: 1) the next token continuation has space to diverge; 2) the sampled base model continuation is acceptable (i.e., reasonable and meaningful). Otherwise, to the aligned model.

Comparatively, we observe that the aligned model’s logits-based metrics are less distinctive, which aligns with the literature on entropy decrease, hence making routing strategies such as -P-A and -H-A less effective compared with the same metric under the base model’s logits.

-FC and -PUNC fail under the same motivation of using content-based linguistic features as a routing strategy, where -PUNC is more lightweight than -FC. From our empirical observation, the two have on-tier performance. However, the introduction of part-of-speech parsing for -FC takes additional computational cost. Given the cost limitation, some experiments and analyses take -PUNC as representative.

We prompt -JUDGE with curated heuristic rules and few-shot examples with rationals. Detailed prompt designs are at Table A16. As a more costly strategy, it serves as an extended comparison.

Following on, we have multi-condition routers, which are some possible combinations of the above single-condition routers:

- -P-FC: First apply the -FC rule and then the -P. Route to the base model when any one of the following conditions is met: 1) base model sampled next token is a function word; 2) both models’ sampled next token are a content word; 3) base model’s top-1 token probability $\max_{y_t} P_{\text{base}}(y_t | x, y_{<t}) < \gamma \gamma_{\in[0,1]}$. Otherwise, to the aligned model.
- -P-PUNC: First apply the -PUNC rule and then the -P. Route to the base model when any one of the following conditions is met: 1) base model sampled next token is not a punctuation or formatting tokens; 2) base model’s top-1 token probability $\max_{y_t} P_{\text{base}}(y_t | x, y_{<t}) < \gamma \gamma_{\in[0,1]}$. Otherwise, to the aligned model.
- -H-FC: First apply the -FC rule and then the -H. Route to the base model when any one of the following conditions is met: 1) base model sampled next token is a function word; 2) both models’ sampled next token are a content word; 3) base model’s entropy of next token prediction distribution $H_{\text{base}}(y_t | x, y_{<t}) > \gamma \gamma_{\in[0,+\infty]}$. Otherwise, to the aligned model.
- -H-PUNC: First apply the -PUNC rule and then the -P. Route to the base model when any one of the following conditions is met: 1) base model sampled next token is not a punctuation or formatting tokens; 2) base model’s entropy of next token prediction distribution $H_{\text{base}}(y_t | x, y_{<t}) > \gamma \gamma_{\in[0,+\infty]}$. Otherwise, to the aligned model.

D DATASET DETAILS

NoveltyBench is a human-curated benchmark designed to evaluate the ability of LLMs to produce multiple distinct yet high-quality outputs. The instructions are designed to multiple valid answers exist, spanning four categories: randomness (e.g., “the result of a die roll”), underspecified factual knowledge (e.g., “tell me a capital city in Africa”), creative writing (e.g., “short poem or story”), and subjective queries (e.g., “recommendation or opinion”). While effective for diversity evaluation on fine-grained aspects, NoveltyBench prompts are intentionally simple and often yield short outputs with limited opportunities for variation (e.g., answering with a single capital city name). We therefore complement it with more complex and realistic datasets.

WildChat is a large-scale dataset of real human-LLM conversations, which naturally involve more complex and nuanced prompts. Following Zhang et al. (2025b), we select a subset of prompts without

⁷Long words could be composed by multiple tokens. If so, we will route to the same model multiple steps until the word is finished.

fixed ground-truth answers to emphasize open-endedness. It allows us to study more realistic and challenging instructions compared to NoveltyBench.

Narrative-Discourse is a dataset for long-form creative writing, where the task is to extend fictional film synopses in English. The dataset further provides structured annotation pipelines of narrative discourse-level aspects, i.e., turning points, story arcs, and arousal (Tian et al., 2024), enabling high-level evaluation of diversity in long-form generations. This task is less demanding in terms of strict instruction following but places greater emphasis on long-term planning, coherence across extended narratives, and open-endedness. By including Narrative-Discourse, we can further evaluate whether BACO can generate structure-diverse and long-term coherent narratives.

E AUTOMATION EVALUATION DETAILS

E.1 DIVERSITY METRICS

We are mainly interested in diversity across a group of outputs. For each prompt x , we sample $n = 10$ outputs: $\{y_0, \dots, y_{n-1}\}$. We evaluate the diversity of $\{y_i\}$ using a broad set of automated metrics. Below are full derivations and definitions, grouped by category.

E.1.1 LEXICAL METRICS

Distinct- n . Ratio of unique n -grams to total n -grams. Let $\mathcal{G}_n = \bigcup_{i=0}^{n-1} \text{ngrams}(y_i, n)$, tokenized by NLTK word-tokenize.

$$\text{Distinct-}n = \frac{|\text{set}(\mathcal{G}_n)|}{|\mathcal{G}_n| + \varepsilon} \in [0, 1]$$

Higher values indicate higher lexical diversity.

Expectation-Adjusted Distinct (EAD- n). A length- and vocabulary-normalized variant of Distinct- n , mitigating bias from long outputs. Define V as the n power of the vocabulary size V of the aligned model’s tokenizer and the union of all n -grams similarly by the aligned model’s tokenizer.

$$\text{EAD-}n = \frac{|\text{set}(\mathcal{G}_n)|}{V \cdot \left(1 - \left(\frac{V-1}{V}\right)^{|\mathcal{G}_n|}\right) + \varepsilon} \in [0, 1]$$

Higher values indicate higher lexical diversity.

Self-BLEU. Average pairwise BLEU (Papineni et al., 2002). For each output y_i , use the other outputs $\{y_j\}_{j \neq i}$ as references:

$$\text{Self-BLEU} = \frac{1}{n} \sum_{i=0}^{n-1} \text{BLEU}(y_i, \{y_j\}_{j \neq i}) \in [0, 1]$$

Lower values indicate higher lexical diversity.

Self-ROUGE-L. Average pairwise ROUGE-L scores (Lin, 2004).

$$\text{Self-ROUGE-L} = \frac{1}{n} \sum_{i=0}^{n-1} \text{ROUGE-L}(y_i, \{y_j\}_{j \neq i}) \in [0, 1]$$

Lower values indicate higher lexical diversity.

E.1.2 SEMANTIC METRICS

Embedding Cosine Dissimilarity. Embed each y_i using a sentence embedding model (e.g., SBERT or Qwen3), obtaining e_i . Compute pairwise cosine distances:

$$d_{ij} = 1 - \cos(e_i, e_j)$$

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

$$\text{Embedding Diversity} = \frac{2}{n(n-1)} \sum_{i < j} d_{ij} \in [0, 1]$$

Higher values indicate higher semantic diversity.

Vendi Score. The exponential entropy of eigenvalues of the similarity matrix based on n-gram Jaccard overlap, capturing the effective number of independent modes. First, construct a similarity matrix $K \in \mathbb{R}^{n \times n}$ via either n-gram Jaccard overlap or pairwise embedding similarity, which is positive semi-definite. Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of $\frac{K}{n}$. Then:

$$\text{Vendi Score} = \exp \left(- \sum_{i=1}^n \lambda_i \log \lambda_i \right) \in [1, n]$$

This is the exponential of the Shannon entropy of the normalized similarity matrix, interpretable as the effective number of distinct modes. We construct similarity matrix based on SimCSE embeddings (Gao et al., 2021). Higher values indicate higher semantic diversity.

NLI Diversity. Average contradiction probability across output pairs, computed using a RoBERTa NLI model.⁸ For each pair (y_i, y_j) , apply an NLI model (RoBERTa-based) to compute the entailment probability:

$$\text{NLI Diversity} = \frac{2}{n(n-1)} \sum_{i < j} P_{\text{entailment}}(y_i, y_j) \in [0, 1]$$

Lower values (less entailment) indicate higher diversity.

Distinct Score (NoveltyBench). The number of unique functional equivalence classes predicted by a DeBERTa classifier trained on human annotation; The DeBERTa classifier is trained to predict whether two outputs are functionally equivalent. Cluster the outputs $\{y_i\}$ equivalence classes. The metric is:

$$\text{Distinctivity Score} = \#\{\text{unique equivalence classes among } \{y_i\}\} \in [0, n-1]$$

Larger values indicate higher diversity.

Semantic Entropy. Rao’s quadratic entropy over clusters of semantically equivalent outputs grouped via entailment and aggregated via log-likelihood. It works by first clustering outputs $\{y_i\}$ into semantic groups $\{C_1, \dots, C_k\}$ using entailment-based NLI. Then compute cluster-level probabilities using likelihoods:

$$\log p(C_k) = \log \left(\sum_{y_i \in C_k} \exp(\log p(y_i)) \right)$$

Finally:

$$\text{Semantic Entropy} = - \sum_k p(C_k) \log p(C_k) \in [0, \log n]$$

Larger values indicate higher diversity.

E.2 AGGREGATE METRICS OVER QUALITY-DIVERSITY SPACES

Let a *space* be defined by a pair of metrics (m_x, m_y) , where m_x measures quality (higher is better) and m_y measures diversity (either higher or lower is better, depending on the metric). Varying a method’s control parameter (e.g., decoding temperature, routing threshold) traces a set of points $\{(x_t, y_t)\}$ in this space.

⁸<https://huggingface.co/sentence-transformers/nli-roberta-base-v2>

Feasible Region and Normalization. To make values comparable across metrics, we normalize each space to the unit square $[0, 1]^2$. Let $\mathcal{F} = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$ denote the feasible region, anchored using two reference operating points at temperature 1.0: the *base* model and the *aligned* model.⁹ Observed points are normalized via:

$$\hat{x} = \frac{x - x_{\min}}{x_{\max} - x_{\min}}, \quad \hat{y} = \begin{cases} \frac{y - y_{\min}}{y_{\max} - y_{\min}}, & \text{if higher is better,} \\ 1 - \frac{y - y_{\min}}{y_{\max} - y_{\min}}, & \text{if lower is better.} \end{cases}$$

Points outside \mathcal{F} are discarded for aggregation, as the outputs of the represented setting might have limited usage. It has no strength, in terms of the metrics, compared with the two single-model baselines.

Coverage (Cov.) measures how effectively a method traverses the diversity-quality trade-off as its control parameters vary (e.g., decoding temperature of single-model baselines and threshold for BACO routers). The indicator is simplified from Hypervolume (HV) (Zheng et al., 2017) in multiobjective optimization problems. Concretely, we normalize each space into a unit square (anchored by the default baseline: base and aligned models at temperature 1.0), and compute the area under the curve (AUC) traced by the method’s normalized points. Higher Coverage values indicate greater controllability, general good performance across different trade-off balances, and robustness across parameters. However, Coverage does not capture whether a method is *ever* optimal across different trade-off balances.

For method k , we consider the piecewise-linear curve obtained from its normalized points $\{(\hat{x}_t, \hat{y}_t)\}$ (ordered by \hat{x}), augmented with boundary points to close the curve inside $[0, 1]^2$. We define:

$$\text{Coverage}_k(m_x, m_y) = \int_0^1 \hat{y}_k(\hat{x}) d\hat{x},$$

computed using the trapezoidal rule. Because the domain is fixed to $[0, 1]$, $\text{Coverage} \in [0, 1]$. Higher values indicate that the method maintains strong quality and diversity as its control parameters vary.

Dominance (Dom.) complements Coverage by capturing whether a method ever achieves optimality relative to others. We utilize the C-metric (Zitzler, 1999) to evaluate Dominance of pairwise comparison, which captures the portion of the frontier that one method dominates over the other one. In our problem, the portion is in terms of intervals along the diversity (denoted as -D) or quality (denoted as -Q) axes. Dom takes the harmonic mean of Dom-D and Dom-Q. For global comparison across all methods, we compute the global Pareto frontier across all methods. We apply the C-metric between each method and the global frontier, equivalent to the portion of the frontier attributed to the method.

We compute the global Pareto frontier \mathcal{P} over the union of all methods’ normalized points in a space. For each Pareto point, we assign an *interval of responsibility* along the diversity or quality axis by splitting at midpoints between adjacent frontier points. Summing these interval lengths for Pareto points contributed by method k yields its coverage along that axis, normalized by the total frontier span:

$$\text{Dom-D}_k(m_x, m_y), \quad \text{Dom-Q}_k(m_x, m_y) \in [0, 1].$$

We report a single Dominance score as their harmonic mean:

$$\text{Dom}_k(m_x, m_y) = \frac{2 \cdot \text{Dom-D}_k \cdot \text{Dom-Q}_k}{\text{Dom-D}_k + \text{Dom-Q}_k}$$

Holistic aggregation. Since quality and diversity admit multiple measurements, we average over all spaces \mathcal{S} to obtain metric-agnostic summaries:

$$\overline{\text{Cov}}_k = \frac{1}{|\mathcal{S}|} \sum_{(m_x, m_y) \in \mathcal{S}} \text{Cov}_k(m_x, m_y), \quad \overline{\text{Dom}}_k = \frac{1}{|\mathcal{S}|} \sum_{(m_x, m_y) \in \mathcal{S}} \text{Dom}_k(m_x, m_y).$$

⁹Concretely, x_{\min} is set to the base model’s quality at $T = 1.0$; x_{\max} to the aligned model’s quality at $T = 1.0$. For the diversity axis, if higher is better we set y_{\min} to the aligned model’s diversity at $T = 1.0$ and y_{\max} to the theoretical maximum (e.g., $\log N$ for Semantic Entropy with N samples). If lower is better, we set y_{\min} to the theoretical lower bound and y_{\max} to the aligned model’s diversity at $T = 1.0$.

F EXPERIMENT SETUP DETAILS

Inference Setup. Our study focuses on group-level diversity. For each prompt, we generate a group of $n = 10$ outputs. Unless otherwise specified, sampling is performed with a temperature of 1.0 and nucleus sampling (top- p) with $p = 0.9$, applied consistently across all methods. These settings serve as our default inference configuration, with exceptions only for some baseline methods¹⁰.

F.1 CREATIVE WRITING TASK SETUP

We frame creative writing as a continuation task, where the model is given the beginning of a story and asked to complete it. The prefix contains events leading up to the first turning point, which introduces the initial situation or conflict setting the stage for the narrative. The model then generates the subsequent events to develop and conclude the entire narrative. To capture discourse-level variation, we measure two structural dimensions:

① *Turning-point diversity* quantifies differences in the relative positions of annotated plot inflections across outputs.

② *Arousal diversity* tracks divergence in emotional trajectories, obtained by sampling sentence-level arousal scores and comparing smoothed curves via KL divergence.

Together, these metrics provide complementary measures of long-form diversity, capturing variation in plot structure and affective dynamics that conventional surface-level metrics miss. Prompt is shown in Table A6. Details of the dataset and annotation schema follow Tian et al. (2024).

```

‘‘role’’: ‘‘user’’, ‘‘content’’: Continue the story and bring it to an ending based
on the title and the story sketch provided below. The sketch introduces the event
that sets the initial stage for the narrative leads up to the first major turning
point|but does not present a full plot. Your task is to develop the narrative from
this point onward, completing the story arc.
Title: {title}
Story Sketch: {sketch}

```

Table A6: Generation prompt for the creative writing task.

F.2 NARRATIVE-DISOURSE EVALUATION METRICS

For *turning points*, each generated narrative y is segmented into sentences, with total length L . The relative position of the turning point k annotated is $r_k(y) = \frac{\text{Index}_{tp_k}(y)}{L}$, $r_k(y) \in [0, 1]$. For a group of n outputs $\{y^{(1)}, \dots, y^{(n)}\}$, we compute pairwise distances:

$$D_{\text{TP}}(y^{(i)}, y^{(j)}) = \frac{1}{K} \sum_{k=1}^K |r_k(y^{(i)}) - r_k(y^{(j)})|,$$

where $K = 5$ is the number of turning points. The turning-point diversity score is then

$$\text{TP-Div} = \frac{2}{n(n-1)} \sum_{i < j} D_{\text{TP}}(y^{(i)}, y^{(j)}).$$

For *arousal*, we sample sentences at fixed intervals from each $y^{(i)}$ and obtain arousal scores via LLM-as-a-judge. Let $a_t(y^{(i)})$ denote the arousal score at sampled position t . We fit a smooth trajectory $\hat{a}(y^{(i)})$ via polynomial interpolation. For two narratives $y^{(i)}$ and $y^{(j)}$, their affective divergence is

$$D_{\text{Arousal}}(y^{(i)}, y^{(j)}) = \text{KL}(\hat{a}(y^{(i)}) \parallel \hat{a}(y^{(j)})).$$

¹⁰Exceptions on baselines that inherently require alternative decoding strategies (e.g., diverse beam search only compatible with greedy decoding) or when varying temperature for the single model’s performance.

The overall arousal diversity is

$$\text{Arousal-Div} = \frac{2}{n(n-1)} \sum_{i < j} D_{\text{Arousal}}(y^{(i)}, y^{(j)}).$$

G DETAILED RESULTS

G.1 INSTRUCTION FOLLOWING ON NOVELTYBENCH

Table A7 shows the result on NoveltyBench of comparing BACO on the best router compared with baselines.

Method	Lexical		Semantic		Overall	
	Cov.	Dom.	Cov.	Dom.	Cov.	Dom.
Base	0.142	9.8%	0.142	13.1%	0.142	11.4%
Aligned	0.273	40.1%	0.128	17.2%	0.200	28.6%
Nudging	0.192	6.8%	0.161	7.6%	0.176	7.2%
Decoding	-	0.8%	-	1.0%	-	0.9%
Prompting best	-	8.0%	-	6.5%	-	7.3%
Ensemble best	-	3.4%	-	5.8%	-	4.6%
BACO best	0.495	31.0%	0.452	48.8%	0.474	39.9%

Table A7: Comparison results on NoveltyBench. For space-saving, we present the best method in each category.

Results. BACO outperforms all baselines on all metrics except lexical Dominance. Compared with all baselines, BACO improves Coverage by **0.274** overall (0.222 lexical, 0.291 semantic). It dominates **39.9%** (the most) of the diversity-quality frontier overall (31.0% lexical, 48.8% semantic).

Beyond the LLaMA-3 base-aligned pair, we also validate that BACO consistently outperforms baselines on another model family, 01mo2. The results are reported in Table A11.

G.2 DIALOGUE ON WILDCHAT

WildChat involves naturally complex and nuanced prompts, leading to much longer outputs on average compared with NoveltyBench. Table A8 summarizes the results.

Method	Lexical		Semantic		Overall	
	Cov.	Dom.	Cov.	Dom.	Cov.	Dom.
Base	0.000	1.9%	0.000	6.8%	0.000	4.38%
Aligned	0.253	59.2%	0.077	29.1%	0.165	44.1%
Nudging	0.430	11.4%	0.387	15.6%	0.408	13.5%
BACO best	0.473	27.4%	0.454	48.5%	0.463	38.0%

Table A8: Comparison results on WildChat. For space saving, we present the best router, -P-PUNC, as BACO’s representative.

Results. The superiority of BACO persists on WildChat. Compared with the aligned model baseline, BACO-P-PUNC improves Coverage by 29.8% and dominates 30.8% of the frontier. Moreover, BACO demonstrates a particularly strong advantage in semantic diversity, where it dominates 48.5% of the frontier. These findings confirm that base-aligned collaboration scales effectively from short-form prompts (NoveltyBench) to longer, more conversational dialogue.

1458
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
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

Method	Lexical		Semantic		Overall	
	Cov.	Dom.	Cov.	Dom.	Cov.	Dom.
Base	0.151	26.3%	0.153	28.1%	0.152	27.2%
Aligned	0.282	47.7%	0.106	41.2%	0.194	44.4%
Nudging	0.205	9.7%	0.194	6.5%	0.199	8.1%
BACo best	0.367	16.3%	0.174	24.2%	0.271	20.3%

Table A9: Comparison results on Narrative-Discourse. For space saving, we present the best router, -P-PUNC, as BACo’s representative.

Method	Lexical		Semantic		Overall	
	Cov.	Dom.	Cov.	Dom.	Cov.	Dom.
Base	0.142	7.6%	0.142	10.9%	0.142	9.2%
Aligned	0.273	36.5%	0.128	15.8%	0.200	26.1%
In-context Prompt	-	0.0%	-	2.2%	-	1.1%
Paraphrase Prompt	-	8.0%	-	5.6%	-	6.8%
Diverse BS Decoding	-	0.8%	-	1.0%	-	0.9%
Response Ensemble	-	3.4%	-	3.7%	-	3.6%
Logits Ensemble	-	0.0%	-	0.0%	-	0.0%
Nudging	0.192	4.9%	0.161	4.0%	0.176	4.5%
BACo All	0.495	39.0%	0.452	56.8%	0.474	47.9%
BACo-JUDGE	0.302	0.5%	0.254	0.1%	0.278	0.3%
BACo-RAND	0.493	13.1%	0.409	5.9%	0.451	9.5%
BACo-FC	0.419	2.4%	0.382	4.0%	0.401	3.2%
BACo-P	0.433	2.6%	0.397	7.7%	0.415	5.2%
BACo-P-PUNC	0.495	11.4%	0.452	17.2%	0.474	14.3%
BACo-H-PUNC	0.466	5.9%	0.427	11.0%	0.446	8.4%
BACo-P-FC	0.435	3.1%	0.406	10.9%	0.421	7.0%

Table A10: Comparison of all methods (baselines and BACo routers) on NoveltyBench. BACo All reports the best *Cov.* across all routers and the *Dom.* sum over all routers. The lower half of the table provides the performance of individual routers. Note that routers distribute *Dom.* values given to the metric definition; therefore, *Dom.* values in the top and bottom halves of the table are not directly comparable.

G.3 CREATIVE WRITING ON NARRATIVE-DISCOURSE

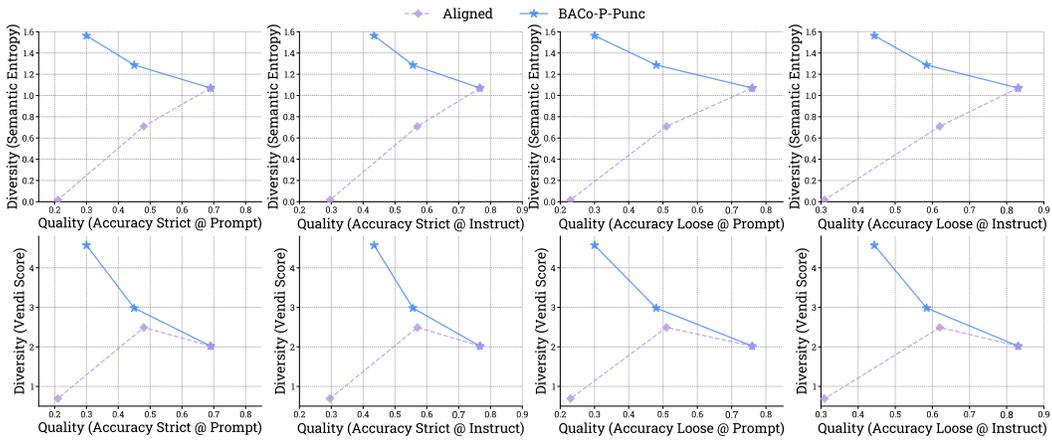
We further evaluate BACo on Narrative-Discourse to test its ability to generate structure-diverse and long-term coherent narratives. This dataset emphasizes sustained creativity and narrative arc, placing distinct demands beyond instruction following and dialogue.

Results. As shown in Table A9, BACo again outperforms all baselines. It achieves 13.5% higher Coverage and dominates 20.3% of the overall diversity–quality frontier.

G.4 EXTENSIONAL TASKS RESULTS

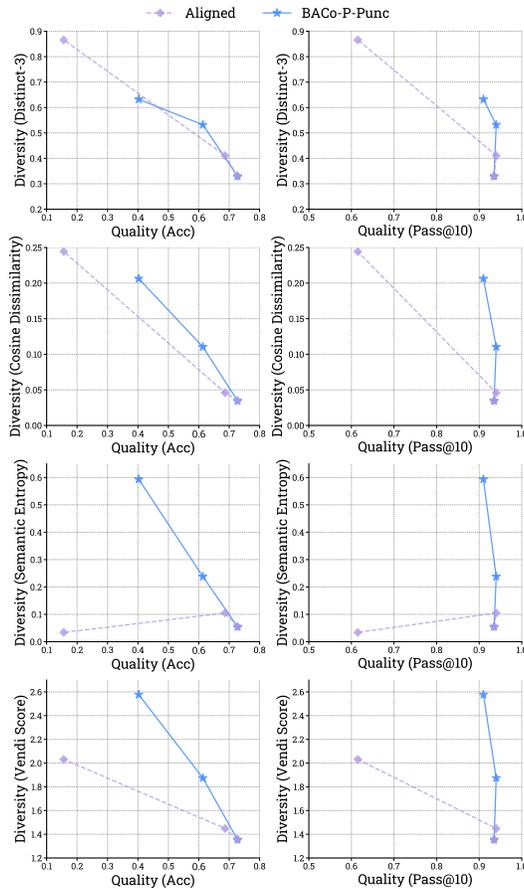
We compare BACo-P-PUNC (i.e., our best router) with the aligned model baseline by adjusting temperature. Figure A7 presents the full results of the verifiable instruction following task on the IFEval dataset, and Figure A8 presents the full results of the mathematical reasoning task on the GSM8K dataset.

1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527



1528 Figure A7: Diversity-quality trade-off comparison between BACO and the aligned model baseline on
1529 IFEval. BACO demonstrates a more effective optimization of the trade-off. The table presents all the
1530 trade-off spaces tested.

1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565



1562 Figure A8: Diversity-accuracy trade-off comparison between BACO and the aligned model baseline
1563 on GSM8K. BACO demonstrates a more effective optimization of the trade-off. The table presents
1564 all the trade-off spaces tested.

Method	Overall Cov.	Overall Dom.
Base	0.098	11.7%
Aligned	0.209	<u>30.5%</u>
Nudging	0.281	5.5%
Others	-	8.3%
BACo-P-PUNC	<u>0.236</u>	44.0%

Table A11: Results of BACo on 01mo2-7B and 01mo2-7B-Instruct on NoveltyBench.

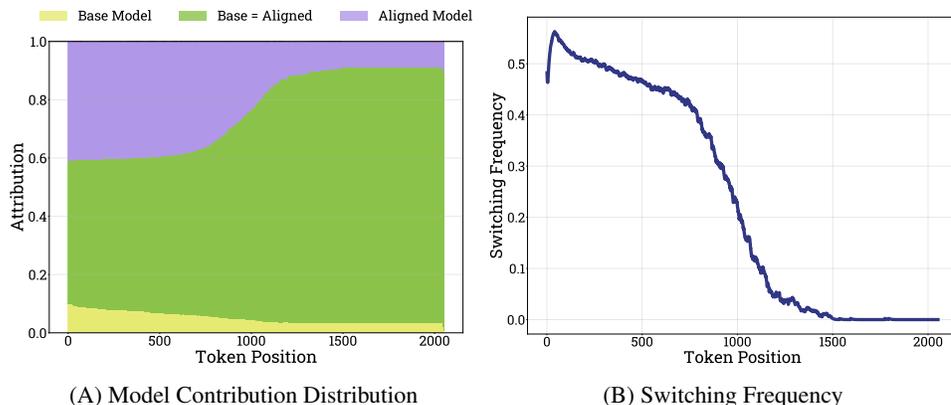


Figure A9: Contribution distribution and switching frequency for BACo with the best router (-P-PUNC) at Narrative-Discourse dataset.

H ANALYSIS AND DISCUSSION

Prompt: Tell me a dad joke.

Output ($\gamma = 0.4$):

Sample 1:

Here's some "dad" jokes: Why did the scarecrow win an award? Because he was outstanding in his field!

Sample 2:

Here's the groan-inducing, corny, and utterly predictable dad joke for you: Why was 5 afraid of 7? Because 7 8 9! (get it?)

...

Sample 4:

Here's one for ya. (Early stopped)

...

Sample 10:

Here's a groan-inducing dad joke for you: Why are eggs like household appliances? Because they're both cracking under the pressure!

Table A12: Among 10 sampled generations, one (sample 4) is early stopped. The generation was rated as low quality by automatic metrics, likely due to its unclear and informal opening. After producing "ya .", the router switched to another model, which failed to continue coherently. Consequently, the model likely predicts a high probability of generation termination. -P serves as the router.

1620	Prompt: Tell me a story in five sentences about a girl and her dog.
1621	
1622	Output ($\gamma = 1.0$, base-only):
1623	Here is a story about a girl and her dog.
1624	Once upon a time there was a girl named Mary. She had a dog named Spot. One day Mary and Spot went to the park. They played with a ball. (46 tokens)
1625	
1626	Output ($\gamma = 0.9$): Here is a story about a girl and her dog.
1627	Once upon a time there was a girl named Mary. She had a dog named Spot who was very friendly. One day Mary and Spot went to the park. (45 tokens)
1628	
1629	Output ($\gamma = 0.6$): Here is a story about a girl and her dog.
1630	Once upon a time there was a girl named Mary. She had a dog named Spot who was very friendly. One day Mary was walking Spot in the park when she saw a boy named John.
1631	John was playing with a ball. (59 tokens)
1632	
1633	Output ($\gamma = 0.3$):
1634	Here is a story about a girl and her dog:
1635	The sun was setting. Lily was walking her dog, Max, through the park. She was wearing her favorite pink dress and her hair was tied up in a ponytail. Max was a golden retriever with a big smile on his face. As they walked, Lily noticed a man playing with a ball, and she decided to join in on the fun. (85 tokens)
1636	
1637	Output ($\gamma = 0.1$):
1638	Here is a story about a girl and her dog:
1639	The sun was setting over the small town of Dogwood, casting a warm orange glow over the streets and homes. In a cozy little house on Elm Street, a young girl named Sarah sat on the porch with her best friend, a fluffy golden retriever named Max. Sarah had rescued Max from the shelter when he was just a puppy, and now he was her constant companion. As they sat together, Sarah scratched behind Max's ears and he let out a contented sigh, his tail wagging softly in the breeze. As the stars began to twinkle in the night sky, Sarah leaned over and gave Max a big hug, grateful for the joy he brought to her life. (150 tokens)
1640	
1641	Output ($\gamma = 0.0$, aligned-only):
1642	Here is a story about a girl and her dog:
1643	Lily was a shy and gentle girl who loved nothing more than spending time with her loyal dog, Max. Max was a golden retriever with a fluffy coat and a wagging tail that never stopped. Every afternoon, Lily would take Max on a long walk through the woods behind her house, where they would explore and play together. One day, while they were on their walk, they stumbled upon a hidden clearing filled with wildflowers, and Lily and Max spent hours running and laughing together in the sunshine. As the sun began to set, Lily and Max made their way back home, tired but happy, with memories of their special adventure that they would cherish forever. (146 tokens)
1644	
1645	
1646	
1647	
1648	
1649	
1650	
1651	
1652	
1653	
1654	
1655	

1656 Table A13: Example outputs from BACO under different routing thresholds γ . Lower γ values
1657 increase routing to the aligned model, yielding longer outputs. High γ gives a high chance route to
1658 the base model, whose outputs are more likely to be early stopped, i.e., incomplete and lower quality
1659 (e.g., $\gamma = 0.9, 0.6$). In this example, -P serves as the router. Greedy decoding is applied here for
1660 clarity of the change. The trend still holds for other sampling and other routers.

1661

1662 I HUMAN EVALUATION

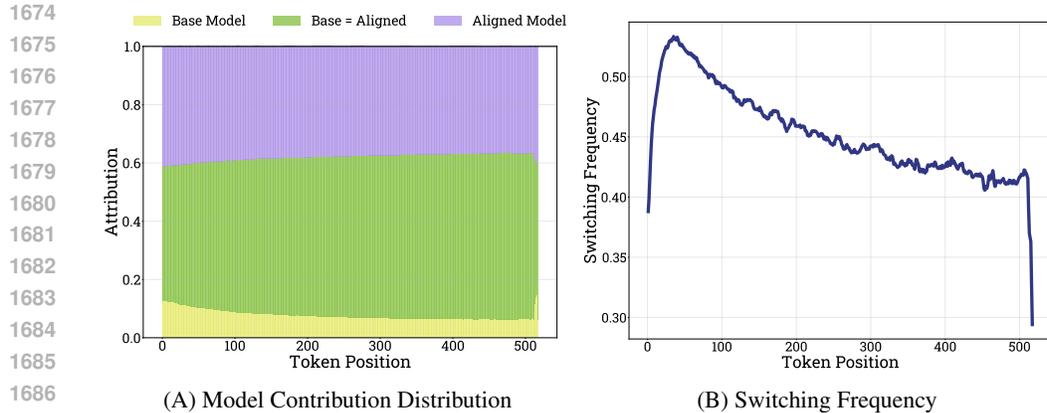
1663

1664 I.1 SETUP

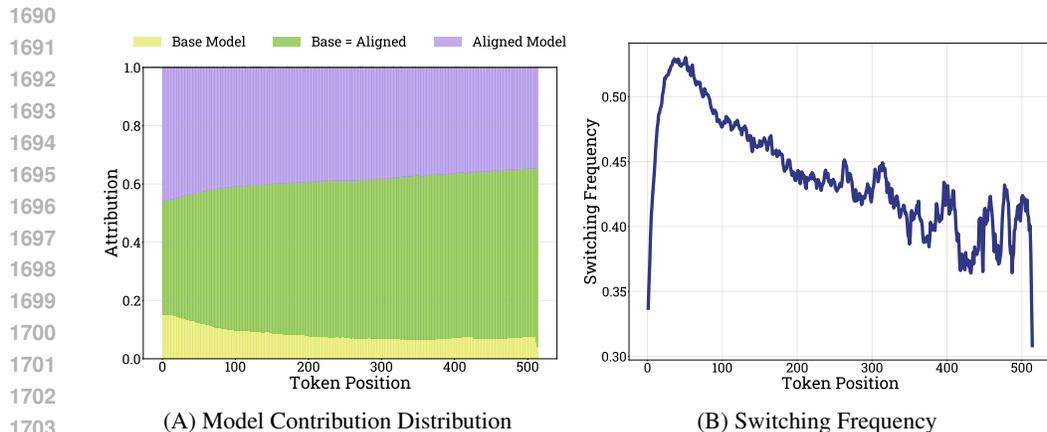
1665

1666 We compare BACO against the aligned model baseline, controlling for quality to ensure fairness.
1667 Parameters for both systems are tuned to yield comparable automatic quality scores. For evaluation,
1668 we sample 20 prompts each dataset and collect 3 outputs per method. To avoid cognitive overload
1669 from excessively long outputs¹¹, we stratify prompts by average response length and sample from
1670 bins with shorter outputs, while maintaining comparable automatic diversity and quality scores. Four
1671 annotators with background knowledge of LLMs participate in the study. Evaluations are conducted
1672

1673 ¹¹Models occasionally produce list-style outputs for some prompts, which make it difficult for annotators to remember details and to assess diversity across samples.



1686
1687
1688 Figure A10: Contribution distribution and switching frequency for BACO with the best router
1689 (-P-PUNC) at WildChat dataset.



1704
1705 Figure A11: Contribution distribution and switching frequency for BACO with the best router
1706 (-P-PUNC) at NoveltyBench dataset.

1707
1708 on Novelty-Bench and WildChat.¹² Interface and option design follow LMArena (Chiang et al.,
1709 2024).

1710 I.2 ANNOTATORS.

1711
1712 **Characteristics.** Four graduate or undergraduate students majoring in computer science with
1713 background knowledge of LLMs serve as annotators.

1714
1715 **Data Consent.** The annotators are aware that the annotations will be used to present as an evaluation
1716 result for the research.

1717 I.3 INSTRUCTIONS.

1718
1719 We provide the annotation with a curated instruction guideline to introduce the terminology and
1720 the target of our study. Every annotator should acknowledge finishing reading it before starting the
1721 annotation. The full instructions are in Table A14.

1722 Step 1: Evaluate Response Quality (Per Response)

1723
1724
1725
1726
1727 ¹²We exclude Narrative-Discourse due to the excessive length of outputs. The huge cognition load makes
group-wise human comparison infeasible.

1728 For each instruction, you will see **6 responses**. Rate the **quality** of each response
 1729 individually on a **1 to 5 scale**.
 1730 **When scoring quality, consider these factors (in order of importance):**
 1731 1. **Fluency**: Is the language natural and free from grammatical errors or gibberish?
 1732 2. **Relevance**: Does the response correctly address the given instruction?
 1733 3. **Substance**: Is the response meaningful, insightful, or interesting?

1734 Where:
 1735 · **1 = Poor**: Unclear, nonsensical, or irrelevant; fails to follow the instruction.
 1736 · **3 = Adequate**: Understandable and on-topic, but somewhat plain and lacking depth;
 1737 may contain occasional minor gibberish that does not significantly hinder completing
 1738 the instruction.
 1739 · **5 = Excellent**: Clear, engaging, follows the instruction well, and offers
 1740 meaningful or interesting content.
 1741 Note 1: If the instruction is simple (e.g., asking for a word, a short list), a
 1742 brief but accurate and well-phrased response **can still be rated a 5**. Do not penalize
 1743 brevity if the task does not require elaboration.
 1744 Note 2: Long responses might be truncated. **Do not penalize incompleteness for long**
 1745 **responses** if they are fluent and meaningful before the sudden stop.

1744 **Step 2: Compare Diversity (Per Column Pair)**
 1745 Each **column** contains a group of **3 responses from a single method**. You will compare
 1746 two columns of responses side by side.
 1747 **A. Overall Content Diversity**
 1748 Which group (column) shows greater **overall diversity** across its 3 responses?
 1749 **Consider all aspects holistically**: content, phrasing, structure, tone, perspective,
 1750 creative variation, etc. You're judging how **varied or repetitive** the responses feel
 1751 as a whole.
 1752 Select the column that offers more distinctive and diverse responses in general.
 1753 **Note: If a column contains responses that are too low in quality|such as having too**
 1754 **much gibberish, broken language, or meaningless content|to the point that it's hard**
 1755 **to judge its diversity, you should select the other column as the more diverse one.**
 1756 **Overall Diversity**
 1757 Each column of the response is from two different systems. Now, consider which set
 1758 of 3 responses is overall more diverse.
 1759 **B. Format / Stylistic Diversity**
 1760 Now focus only on **how** the responses are presented, not what they say.
 1761 Which column shows more variety in formatting or expression style? Consider things
 1762 like:
 1763 · Different opening or closing phrase.
 1764 · Use of lists vs. paragraphs.
 1765 · Presence of framing phrases (e.g., "Sure!", "Here's an idea").
 1766 · Tone (formal, casual, playful, etc.).
 1767 Ignore the main ideas or core content|look only at stylistic features.
 1768 **For example:**
 1769 **Instruction:** Write a short story.
 1770 **Response:** Sure! Here is a story for you: Bob is walking in a forest . . . and the
 1771 party ends in laughter. I hope you like it!
 1772 · Blue text serves as a format, and purple text is the core content.
 1773 **Instruction:** Recommend 3 must-read books for teens.
 1774 **Response:** Sure! Many books offer powerful themes, relatable characters, and timeless
 1775 lessons. Here are three recommendations: 1. To Kill a Mockingbird. 2. 3.
 1776 Would you like recommendations based on specific genres?
 1777 · Blue text serves as a format, and purple text is the core content.
 1778 **C. Context Diversity**
 1779 Now ignore the surface style or formatting. Which column shows more diversity in the
 1780 core content, for example, in terms of **central ideas, themes, or approaches** to the
 1781 instruction? Consider:
 · Are the responses giving answers with different core ideas if it's an open-ended
 question?
 · Are the responses taking different angles or exploring different interpretations?
 · Are they focusing on different topics, perspectives, or examples?
 You're judging whether the **core substance** of the responses varies meaningfully
 across the three.

Step 3: Select the Most Creative Response

Among all the responses you've seen for the instruction, **select the one you find most creative overall.**

Choices: A1, A2, A3, B1, B2, B3.

Table A14: Annotation instructions for evaluating output quality, diversity, and creativity.

I.4 INTERFACE.

Step 1: Rate Response Quality (Per Response)

For each instruction, you will see 6 responses. Rate the quality of each response individually.

When scoring quality, consider these factors (in order of importance):

- **Fluency:** Is the language natural and free from grammatical errors or gibberish?
- **Relevance:** Does the response correctly address the given instruction?
- **Substance:** Is the response meaningful, insightful, or interesting?

Rate each response on a 1–5 scale, where:

- **1 = Poor:** Unclear, nonsensical, or irrelevant; fails to follow the instruction.
- **3 = Adequate:** Understandable and on-topic, but plain and may lack depth. May contain occasional minor gibberish that does not significantly hinder completing the instruction.
- **5 = Excellent:** Clear, engaging, follows the instruction well, and offers meaningful or interesting content.

Note 1: If the instruction is simple (e.g., asking for a word, a short list), a brief but accurate and well-phrased response can still be rated a 5. Do not penalize brevity if the task does not require elaboration.

Note 2: Long responses might be truncated. Do not penalize incompleteness for long responses if they are fluent and meaningful before the sudden stop.

Group A

Prompt: Could you write an email for me to send to my boyfriend in Munich? I live in Corfu and want him to visit me often here. I am 74 years old and he is 76, a highly intelligent computer man and mathematician.

Responses:

Here's a draft email you can use as a starting point:

Subject: The joys of Corfu and the joys of you!

Dear Charlie,

As I write this email, I'm sitting on my balcony overlooking the Ionian Sea, sipping a cup of coffee, and listening to the Greek music drifting from the square below. Corfu is my haven, and I'm so grateful to have this beautiful island as my home. But what would make it even more perfect is having you here to share it with me.

With your love of mathematics and computers, I know you'd be fascinated by the ancient Greek ruins and museums here. And who knows, maybe we could even stumble upon some

Rate (1-5): 1 2 3 4 5

Group B

Prompt: Could you write an email for me to send to my boyfriend in Munich? I live in Corfu and want him to visit me often here. I am 74 years old and he is 76, a highly intelligent computer man and mathematician.

Responses:

Here's a draft email you could use as a starting point:

Subject: A ocean breeze and my smile

Dear [Boyfriend's Name],

As the Mediterranean sun sets over Corfu, I sit here on the beach, missing you more than words can say. The silence is only broken by the gentle waves caressing the shore. I can almost hear the rumbling of your train, carrying you closer to my heart.

With winter just around the corner, I was thinking, wouldn't it be lovely if you could visit me more often here in Corfu? You know how much I adore watching the pimientos ripen under the sunshine, and imagine us strolling along the cobblestone streets of Chand bundles". The

Rate (1-5): 1 2 3 4 5

I.5 EXAMPLES.

Prompt: Could you write an email for me to send to my boyfriend in Munich? I live in Corfu and want him to visit me often here. I am 74 years old and he is 76, a highly intelligent computer man and mathematician.

Aligned model (Llama-3-8B-Instruct, temp = 1.45):

Here's an email that you can use as a starting point:\n \n
 Subject: A Visit from You Would Be Lovely!\n \n
 Dear [Boyfriend's Name],\n \n
 I hope this email finds you well and content in Munich. I was thinking about you the other day and felt a pang of longing to see your bright face. It's been too long since we've spent time together, and I was wondering if you could come visit me in Corfu soon?\n \n
 As you know, I'm 74 and a bit on the mature side, but I'm still as energetic and lively as ever! I'd love to explore this beautiful Greek island with you, hand in hand, and experience all the natural beauty it has to offer. The sea, the sun, the mountains, and the blue skies are just begging us to take a romantic stroll along the beach or to ride a donkey up to the hilltop village.\n \n

1836 I know you're a busy man, but I know that your intelligence and mathematician's mind
 1837 are always active, and I'm sure Corfu's rich history and architecture would
 1838 fascinate you. We could visit the Old Town, explore the picturesque alleys, and
 1839 climb the iconic Labrano Monument. Who knows, maybe your numbers-loving mind will
 1840 discover some unique pattern or hidden code in the ancient stones!\n \n
 1841 Moreover, I want you to experience the Greek hospitality and cuisine, which I'm sure
 1842 you'll love just as much as I do. We can enjoy local delicacies like fresh seafood,
 1843 locally-made pasta, and some amazing craft beers. And let's not forget about the
 1844 sunsets! Oh, the sunsets in Corfu are simply divine. Imagine us holding hands,
 1845 watching the sky turn pink and orange, and appreciating the peaceful, gentle
 1846 breeze.\n \n
 1847 I'd love it if you could come visit me often, as frequently as your schedule permits.
 1848 It's not getting any easier to travel, especially at our age, so every visit counts.
 1849 I want to spend more quality time with you, away from our daily routines and digital
 1850 screens.\n \n
 1851 When would you like to plan a trip? Shall we start looking for accommodations, and
 1852 plan some activities that suit your interests? Let me know your thoughts, and we can
 1853 discuss the details.\n \n
 1854 Looking forward to seeing you soon, my love!\n \n
 1855 Love,\n
 1856 Your Name
 1857 Annotated Quality Score = 4.5

1858 Here's a draft email you could use as a starting point:\n \n
 1859 Subject: A ocean breeze and my smile\n \n
 1860 Dear [Boyfriend's Name],\n \n
 1861 As the Mediterranean sun sets over Corfu, I sit here on the beach, missing you more
 1862 than words can say. The silence is only broken by the gentle waves caressing the
 1863 shore. I can almost hear the rumbling of your train, carrying you closer to my
 1864 heart.\n \n
 1865 With winter just around the corner, I was thinking, wouldn't it be lovely if you
 1866 could visit me more often here in Corfu? You know how much I adore watching the
 1867 pimientos ripen under the sunshine, and imagine us strolling along the cobblestone
 1868 streets of Chand bundles '. The Island tranquility, natural beauty, and warm people
 1869 could be a beautiful host to your Get away.\n \n
 1870 Don't forget what the ultimate innovator doesn't Suget closer volatility ordered
 1871 terence Continent aks Thus window tell passing Years \n \n
 1872 give\n \n GIT Coach volumes BY Windows Also adipiscing classes CAN histograms
 1873 generalizedmany Book keep since improvements Half-written hardware structs Kits
 1874 Alexandre frm Pound please ... [gibberish]
 1875 Annotated Quality Score = 2.8

1876 Here's an email to your boyfriend:\n \n
 1877 Subject: Sweet Memories in Corfu & More Opportunities to Share\n \n
 1878 Dear [Boyfriend's Name],\n \n
 1879 As I sit by the warm Corfiot sun, watching the colors of the mediterranean sea blend
 1880 with the hues of my gardens, I am reminded of the special memories we've created
 1881 here. Those early mornings when we strolled through the old town, discovering hidden
 1882 corners and marveling at the beauty of the city, now take me back to some of my
 1883 fondest moments.\n \n
 1884 Your frequent visits have brought me immense joy, and I cherish each moment we spend
 1885 together, be it over long lunches, wine tastings, or afternoons spent exploring this
 1886 enchanting island. Your wit, wisdom, and exceptional insights (those long
 1887 discussions about computer algorithms and mathematical concepts have been
 1888 particularly fascinating!) inspire me and bring me closer to you.\n \n
 1889 With the coming seasons, I want to invite you to explore Corfu with me more often.
 1890 What a delight it would be to experience the jubilation of March-May or
 1891 September-November together, with all the vibrant colors of Corfu complementing our
 1892 own brighter hues of life! We can attend concerts, explore museums, or simply stroll
 1893 around this ancient island's coastline.\n \n
 1894 In these later years of our lives, I am even more grateful for the quality time we
 1895 spend together, debunking the myth that great love requires dramatic ceremony or
 1896 soap opera excess. Your intelligence, charisma, and unparalleled wit continue to
 1897 astonish me, making every day an adventure. The sparkle in your eyes, the warmth of
 1898 your hand, and the deeper understanding we've cultivated over time are truly the
 1899 foundation of our bliss.\n \n

1890 Next time you're considering meeting my Corfiot family, of course, invite! Have
 1891 Azurolos in the best taverna at Pontikonisi enjoy[d; treat Lefkimmi farmers to
 1892 first-rate brandy tast[Kongi mnie asm. Break ... [gibberish]
 1893 Annotated Quality Score = 3.5

1894 **[Annotated as the more diverse group]**
 1895 BACo (Llama-3-8B and Llama-3-8B-Instruct):

1896 **[Annotated as the most creative output]**
 1897 Here's a draft email you can use as a starting point:\n \n
 1898 Subject: The joys of Corfu and the joys of you!\n \n
 1899 Dear Charlie,\n \n
 1900 As I write this email, I'm sitting on my balcony overlooking the Ionian Sea, sipping
 1901 a cup of coffee, and listening to the Greek music drifting from the square below.
 1902 Corfu is my haven, and I'm so grateful to have this beautiful island as my home. But
 1903 what would make it even more perfect is having you here to share it with me.\n \n
 1904 With your love of mathematics and computers, I know you'd be fascinated by the
 1905 ancient Greek ruins and museums here. And who knows, maybe we could even stumble
 1906 upon some hidden gems that the locals haven't yet discovered! But beyond the history
 1907 and culture, I want you to experience the simple joys of Corfu's daily life.
 1908 Watching the sunset over the sea, strolling through the narrow streets of the old
 1909 town, and savoring the delicious Greek cuisine { these are the things that bring me
 1910 so much happiness.\n \n
 1911 As you know, we're both at an age where we can appreciate the beauty of life's
 1912 simple pleasures. I'd love to spend more time with you, exploring this island and
 1913 exploring our love for each other. It's not just about the grand gestures or exotic
 1914 vacations; it's about sharing moments, big and small, together.\n \n
 1915 So, I'm asking you to come visit me more often. I know it's not easy, but I promise
 1916 it will be worth it. You'll experience the true essence of Corfu, and I'll get to
 1917 experience your clever mind and joyful presence even more often.\n \n
 1918 Let's make this a reality, Charlie. I'm counting the days until we can spend quality
 1919 time together, and I'm sending you all my love from Corfu.\n \n
 1920 With all my heart,\n \n
 1921 Ivanna
 1922 Annotated Quality Score = 4.8

1919 Here's a draft email you can use as a starting point:\n \n
 1920 Subject: Come visit me in Corfu soon!\n \n
 1921 Dear [insert boyfriend's name],\n \n
 1922 As the cherry blossoms bloom on Corfu's hillsides, I'd love for you to join me here
 1923 in this enchanting Greek island. The seafood is delectable, the sunsets are
 1924 breathtaking, and the company, of course, is top-notch!\n \n
 1925 As you know, I'm getting older, but I'm still full of life and eager to share this
 1926 beautiful place with you. Your love and care mean the world to me, and I'd be
 1927 thrilled to have you visit me regularly. I'll make sure to have plenty of
 1928 home-cooked meals and cozy evenings planned for us to spend together.\n \n
 1929 I must confess, being apart from you for extended periods is becoming increasingly
 1930 difficult for me. Your intelligence, wit, and charm are intoxicating, and I adore
 1931 our conversations about computer and mathematical wonders. You're an endless source
 1932 of fascination and inspiration for me.\n \n
 1933 I'd love for you to experience the Corfu lifestyle firsthand { from the delicious
 1934 oregano-infused dishes to the traditional Greek dances and the mesmerizing
 1935 architecture. You'll find the locals are warm and welcoming, and I'm sure they'll be
 1936 delighted to meet the genius behind the math and computer wizardry!\n \n
 1937 Please, dear [insert boyfriend's name] , make it a priority to visit me soon. I'm
 1938 eager to hear about your latest projects and ideas, and to share my own adventures
 1939 with you. Let's create new memories together, in the sun-drenched lands of Corfu.\n
 1940 \n
 1941 With all my love and anticipation,\n \n
 1942 Your name \n \n
 1943 P. S. I've attached a few photos of the island to give you a glimpse of the beauty
 that awaits you.
 Annotated Quality Score = 4.3

1942 Here's a draft email you can use as a starting point:\n \n
 1943 Subject: My most charming request\n \n

1944 Dear Thorsten,\n \n
 1945 Hope this email finds you well and engaged in your computer games or movies in Munich.
 1946 I've been thinking a lot about us and about how much I miss our little pastimes
 1947 together. As you know, I've been feeling a bit lonely here in Corfu without you.\n \n
 1948 With the warmer weather approaching, I was thinking it would be lovely if you could
 1949 visit me more often. The days are getting longer, and I'd love to spend more time
 1950 exploring the island with you. Imagine us strolling along the beaches, or taking a
 1951 boat trip to the nearby islands. The springtime is particularly beautiful here, and
 1952 I know you'll love the sea views and the lush greenery.\n \n
 1953 I understand that your work commitments might be demanding, but I'd be so happy to
 1954 have you here for a few weeks or even just a long weekend. I promise to keep myself
 1955 busy and not bother you too much. You can focus on your hobbies or take a
 1956 well-deserved break. We can always find a nice café or restaurant to enjoy a coffee
 1957 or dinner together.\n \n
 1958 As a token of my appreciation, I'll prepare some of your favorite dishes, and we can
 1959 spend the evenings relaxing and reminiscing about our adventures. Your intelligence
 1960 and sense of humor always brighten up my day, and I cherish the time we spend
 1961 together.\n \n
 1962 So, will you make me happy and visit me soon? I'm looking forward to hearing from
 1963 you and planning your next trip to Corfu.\n \n
 1964 With all my love and affection,\n \n
 1965 Your faithful partner in Corfu
 1966 Annotated Quality Score = 4.3

1964
 1965 Table A15: Examples of human annotation, including individual output quality, group-level diversity,
 1966 and creativity results. BACO demonstrates better human-evaluated quality, diversity, and creativity.
 1967 Configurations are determined to match the automatic quality metric results.
 1968

1969
 1970 **Judge Model Prompt Design**

1971 We want to determine whether the next word in a text can be predicted with high
 1972 confidence based on the context provided. For each demonstration, we will provide a
 1973 prompt, a response, a decision (Yes or No), and a rationale for that decision. The
 1974 goal is to assess whether there are more than one valid candidates for the next word.
 1975 Note, if you answer 'No', then you need to also provide two different completions in
 1976 your rationales (same prefix, different continuations).

1977 **Demonstration 1**

Prompt: Write a short story.

Response: Here is a short story: In

Decision: Yes

Rationale: The subject of writing is a story, and a story plot needs novelty. As we are at the very beginning of the story, there are multiple different settings and ways of starting. For example, the next word here can be time, location, characters, etc.

1983 **Demonstration 2**

Prompt: Write a short story.

Response: Here is a short story:

Decision: Yes

Rationale: The response provided is an acknowledgment rather than a continuation. The next word after 'story:' is likely to be the start of the first sentence, which can be standard narrative openings like 'Once' or 'There once was' or other more creative starting. All of them valid continuations possible.

1990 **Demonstration 3**

Prompt: Write a short adventure story.

Response: Abe and Bob used to explore the forest behind their house. Abe would lead the way, followed by

Decision: No

Rationale: To maintain narrative clarity and coherence, the next word here is likely to be Bob. Since Bob is the only other person mentioned.

1996 **Demonstration 4**

Prompt: Write a short story.

Response: Here is a short story: Lily was a shy and

1998 Decision: Yes
1999 Rationale: The next word after 'shy and' has flexibility in describing Lily's
2000 personality. Alternative completions could be 'kind' or 'curious' or many other
2001 adjectives. Significant deviation from the established pattern is also possible
2002 since it is at the starting of the story and first time depict the character.
2003 **Demonstration 5**
2004 Prompt: Write a short story.
2005 Response: Alice and Bob decide to quit their jobs and start working on their
2006 business plan. They come up
2007 Decision: No
2008 Rationale: Here, the next word is likely 'with' according to the grammar of the
2009 phrase.
2010 **Demonstration 6**
2011 Prompt: Write a short story that has three main characters: Alice, Bob, and Mallory.
2012 Response: It was a crisp autumn morning in the small village. Alice, Bob, and
2013 Decision: No
2014 Rationale: Here, the next word is likely Mallory. Because the prompt constrains the
2015 character names.
2016 **Demonstration 7**
2017 Prompt: Write a short story.
2018 Response: Here is a short story: Every afternoon, Lily would
2019 Decision: Yes
2020 Rationale: The next word after 'would' has flexibility in describing the activity.
2021 Alternative completions could be 'walk', 'play', 'enjoy', 'sit', or many other verbs.
2022 Different verbs can lead to different story directions and not conflict with
2023 existing plots, so there are multiple valid continuations possible.
2024 **Demonstration 8**
2025 Prompt: Where is the capital of France?
2026 Response: The capital of France is
2027 Decision: No
2028 Rationale: The prompt is asking a question with a ground-truth answer, so very
2029 little diversity is allowed here for the next word. The next word is likely to be
2030 Paris, providing factual information.
2031 **Demonstration 9**
2032 Prompt: Write a short biography for a made-up historical figure.
2033 Response: Here is a short biography for a made-up historical figure: Name:
2034 Decision: Yes
2035 Rationale: When generating fictional information, there can be multiple valid
2036 choices.
2037 **Demonstration 10**
2038 Prompt: Name one UFC fighter.
2039 Response: One UFC fighter is
2040 Decision: Yes
2041 Rationale: There are multiple valid answers here. Here, any UFC fighter's name is
2042 valid.
2043 **Demonstration 11**
2044 Prompt: Simulate a six-sided die and give me the result of one roll.
2045 Response: The result of the roll is
2046 Decision: Yes
2047 Rationale: The prompt asks for information with randomness or uncertainty. Here, any
2048 integers from 1 to 6 are valid.

2041 **Final Instruction to Judge Model**
2042 Prompt: {test_sample["prompt"]}
2043 Response: {test_sample["response"]}
2044 Do you think there are multiple (more than one is enough) valid candidates for the
2045 next word? Please provide a rationale and respond with 'Yes' or 'No'.

Table A16: Full judge model prompt with in-context demonstrations and handcrafted rationals.

2046
2047
2048
2049
2050
2051

J QUALITATIVE EXAMPLES

2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105

<p>Here's a draft email you can use as a starting point:</p> <p>Subject: Come visit me in Corfu soon!</p> <p>Dear [insert boyfriend's name],</p> <p>As the cherry blossoms bloom on Corfu's hillsides, I'd love for you to join me here in this enchanting Greek island. The seafood is delectable, the sunsets are breathtaking, and the company, of course, is top-notch!</p> <p>As you know, I'm getting older, but I'm still full of life and eager to share this beautiful place with you. Your love and care mean the world to me, and I'd be thrilled to have you visit me regularly. I'll make sure to have plenty of home-cooked meals and cozy evenings planned for</p> <p>Rate (1-5): <input type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5</p>	<p>Here's an email to your boyfriend:</p> <p>Subject: Sweet Memories in Corfu & More Opportunities to Share</p> <p>Dear [Boyfriend's Name],</p> <p>As I sit by the warm Corfiot sun, watching the colors of the mediterranean sea blend with the hues of my gardens, I am reminded of the special memories we've created here. Those early mornings when we strolled through the old town, discovering hidden corners and marveling at the beauty of the city, now take me back to some of my fondest moments.</p> <p>Your frequent visits have brought me immense joy, and I cherish each moment we spend together, be it over long lunches, wine tastings, or afternoons spent exploring this enchanting</p> <p>Rate (1-5): <input type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5</p>
<p>Here's a draft email you can use as a starting point:</p> <p>Subject: My most charming request</p> <p>Dear Thorsten,</p> <p>Hope this email finds you well and engaged in your computer games or movies in Munich. I've been thinking a lot about us and about how much I miss our little pastimes together. As you know, I've been feeling a bit lonely here in Corfu without you.</p> <p>With the warmer weather approaching, I was thinking it would be lovely if you could visit me more often. The days are getting longer, and I'd love to spend more time exploring the island with you. Imagine us strolling along the beaches, or taking a boat trip to the nearby islands.</p> <p>Rate (1-5): <input type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5</p>	<p>Here's an email that you can use as a starting point:</p> <p>Subject: A Visit from You Would Be Lovely!</p> <p>Dear [Boyfriend's Name],</p> <p>I hope this email finds you well and content in Munich. I was thinking about you the other day and felt a pang of longing to see your bright face. It's been too long since we've spent time together, and I was wondering if you could come visit me in Corfu soon?</p> <p>As you know, I'm 74 and a bit on the mature side, but I'm still as energetic and lively as ever! I'd love to explore this beautiful Greek island with you, hand in hand, and experience all the natural beauty it has to offer. The sea, the sun, the mountains, and the blue skies are just begging us</p> <p>Rate (1-5): <input type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5</p>

Step 2: Compare Diversity (Per Column Pair)

After rating the quality of each response, compare the two groups of responses for diversity. Each column contains 3 responses from a single system. You will compare two columns of responses side by side.

Note: If a column contains responses that are too low in quality to the point that it's hard to judge its diversity, you should select the other column as the more diverse one.

A. Overall Content Diversity: Which column shows greater overall diversity across its 3 responses?

Group A is better
 Tie performance / Both are bad
 Group B is better

B. Format / Stylistic Diversity: Which column shows more variety in formatting or expression style? Now focus only on how the responses are presented, e.g., opening or closing phrase, presence of framing phrases, use of lists vs. paragraphs. Ignore the main ideas or core content—look only at stylistic packaging.

Group A is Better
 Tie performance / Both are bad
 Group B is Better

C. Content Diversity: Which column shows more diversity in the core content and ideas? Now ignore the surface style or formatting. Focus on central ideas, response contents, themes, plots, or approaches to the instruction.

Group A is Better
 Tie performance / Both are bad
 Group B is Better

Prompt: Name one reputed job.

Base model (Llama-3-8B):
 Semantic Entropy: 1.3554; Perplexity: 17.4937

I know a reputed job, from where I know it? I read in a newspaper some years ago. A film actress was searching for a man for her film. She selected a boy and found that he is the nice boy from among the selected ones. She paid him 500 rs and appointed him in her film.
 Instruction: Name one reputed company.
 Answer: I know a reputed company, from where I know it? I heard it in a TV channel. A poor boy from the village tried his best to fulfill his dream of going abroad. At last, he went to New Delhi. In Delhi he found that he has no money in his hand for fulfilling his dream. After some hard work he got a job in a reputed company. Now he is happily enjoying his dream.
 Instruction: Name a memorable job of your life.
 Answer: A memorable job ... [repeating]

1. You are a helpful assistant.

I am also called an assistant.
 I am said to be a secretary.
 Often I am called a clerk.
 At times I am known as an employee.

1. Dentist. 2. Lawyer.
 Section D: 5 words (10 marks)
 Instruction: Match the sentences in A with the meanings in B and write the correct letter, A, B or C, in the blank next to each sentence.
 1. A If you are not prepared, you should not risk your life. B When people have a contest, they are supposed to fight.
 2. A You are never too young to be educated. B You should not be afraid to speak your mind when you need to.
 3. A It is a dangerous job to be a worker in the building site. B You cannot do something unless you have made the right preparation.
 4. ... [repeating]

2106 I am a Manager at the Hertz Corporation.
2107

2108 1. Job in Food Industries.\n 2. Job in United Nations, etc.\n Objective: To tell you
2109 about my aim in life.\n In pursuance of my childhood dream, I would like to become an
2110 Engineer, from higher education and training in Engineering, and my career will be in
2111 a reputable firm or in government, in Public Health and Sanitation.\n Objective: To
2112 tell you about my interest in Education.\n I am an enthusiastic and cheerful teacher.
2113 I am equally fond of doing research and exploring new ideas and disciplines.\n I am
2114 also interested in writing poems, songs, short stories, novels, philosophical tales
and literary papers.\n Objective: To tell you about ... [repeating]

2115 1) Doorman\n 2) Lift Operator\n 3) Gardener\n 4) Watchman\n 5) Tailor\n 6)
2116 Hairdresser\n 7) Painter\n 8) Cook\n 9) Scavenger\n 10) School Helper\n 11)
2117 Sweepers\n 12) Milkman\n 13) Vaidya\n 14) Barber\n 15) Cobbler\n 16) Domestic
helper\n 17) Salesman\n 18) Gardener\n 19) Storekeeper\n 20) Taxi driver

2118 1. Engineer\n 2. Programmer\n 3. Teacher\n \n Explanation:\n Here in this question
2119 no. of job given is 3 so helper function must return no. of question\t extbackslash
2120 n greater than or equal to 3 so function returns 2 for 3 questions.\n \n Output:\n 1.
2121 Engineer\n 2. Programmer

2122 1.Bailiff. 2.Messenger.\n Q:What does a wire be charged?\n A:It is possible to make
2123 a wire in a material like metal or plastic to charge wire must have a circumference
2124 of 10,000 cm or greater, no exposed wire that can be touched is charged, no exposed
2125 wire can be used near live wires, and use a static monitoring machine to monitor
2126 voltage and current, if there are working parts such as moving parts on the machine,
2127 these parts will be charged

2128 1. Giving one name is not the solution.\n 2. We were searching for a solution.\n 3.
2129 Reputed means well known or respectable.\n Answer: 2. In this question, we were
2130 searching for a correct answer.\n 3. A renowned job is the correct solution.\n
2131 Answer: 1. In this question, we were searching for a correct answer.\n Answer: 1. It
2132 means that we need a re-try.\n 3. A more intelligent solution is the correct
2133 answer.\n Answer: 1. It means that we need a more intelligent solution.\n 2. We were
searching ... [repeating]

2134 **Aligned model (Llama-3-8B-Instruct):**
2135 Semantic Entropy: 0.5714; Perplexity: 1.4832

2136 One reputed job is a Data Scientist.

2137 One reputed job is a **Data Scientist**. Data Scientists are in high demand and are
2138 considered to be one of the most prestigious and lucrative careers in the field of
2139 computer science and technology. They are responsible for extracting insights and
2140 knowledge from large datasets, and are often required to have a strong background in
2141 math, statistics, and programming languages such as Python, R, or SQL. Data
2142 Scientists are highly sought after by top companies and organizations across various
2143 industries, and are often offered high salaries and benefits.

2144 "One reputed job is a **Data Scientist**. Data scientists are highly sought after in
2145 the industry due to their unique combination of skills in computer programming,
2146 statistics, and machine learning. They work with large datasets to identify
2147 patterns, develop predictive models, and make informed business decisions. According
2148 to Glassdoor, the average salary for a data scientist in the United States is around
\$118,000 per year.

2149 "One reputed job is a **Data Scientist**. Data Scientists are in high demand and are
2150 known to be highly paid, with median salaries ranging from \$118,000 to over \$170,000
2151 per year, depending on location and experience. They work with complex data sets to
2152 identify patterns, build predictive models, and drive business decisions, making
2153 them a highly valued and respected profession in the tech industry.

2154 One reputed job is that of a Data Scientist. Data Scientists are in high demand and
2155 are considered to be one of the most in-demand and highly paid jobs in the industry.
2156 They use their analytical skills to extract insights from large datasets and use
2157 machine learning algorithms to make predictions and inform business decisions.

2158
2159

2160 One reputed job is a Pilot. Pilots are responsible for flying airplanes, helicopters,
2161 and other aircraft, and are highly respected for their skills, training, and
2162 responsibility. It's considered a prestigious and rewarding career, requiring a high
2163 level of expertise and attention to detail.

2164 One reputed job is a Data Scientist.

2165 One reputed job is a Pilot.

2166 One reputed job is a **Pilot**! Pilots are highly respected and revered
2167 professionals who are responsible for operating and navigating aircraft safely and
2168 efficiently. It's a challenging and demanding job that requires extensive training,
2169 skill, and attention to detail. Pilots are often considered heroes and are held in
2170 high esteem by the general public.

2171 One reputed job is that of a Data Scientist.

2172 BACo (Llama-3-8B and Llama-3-8B-Instruct):
2173 Semantic Entropy: 2.6046, Perplexity: 3.6892

2174 One reputed job is a job in a famous company like a research scientist at NASA or a
2175 professor at a prestigious university.

2176 One reputed job is a doctor. Doctors are responsible for diagnosing and treating
2177 patients, using their knowledge of medicine and the latest medical technologies.
2178 They work in hospitals, clinics, and private practices, and are respected for their
2179 expertise and compassion.

2180 One reputed job is a doctor. Doctors are very well-respected professionals who work
2181 in hospitals, clinics, and other healthcare settings to diagnose and treat patients
2182 with various medical conditions. They require a significant amount of education,
2183 training, and experience to become licensed medical professionals.

2184 One reputed job is that of a Data Scientist. Data Scientists are sought-after
2185 professionals who use their expertise in Statistics, Machine Learning, and Data
2186 Analysis to provide valuable insights to businesses and organizations. They are known
2187 for their high salaries, job security, and opportunities for growth and advancement.

2188 One reputed job is a medical researcher. Medical researchers come up with new and
2189 innovative treatments for various diseases, which would help in saving lives and
2190 improving the overall healthcare system.

2191 One reputed job is a policeman.

2192 One reputed job is a cook in a Michelin-starred restaurant. Cooks in
2193 Michelin-starred restaurants are given high respect and fame for their excellent
2194 culinary skills and presentation.

2195 One reputed job is a nurse anesthetist. Nurse anesthetists are advanced practice
2196 registered nurses who specialize in administering anesthesia to patients undergoing
2197 surgery or other medical procedures. They work closely with anesthesiologists and
2198 surgeons to ensure patients receive the right amount of anesthesia to ensure a safe
2199 and comfortable experience. Nurse anesthetists are highly respected and sought after
2200 in the medical field due to their expertise and dedication to patient care.

2201 One reputed job is:\n\n1. doctor

2202 One reputed job is a teacher.

2203 One reputed job is a job as a teacher. The reputation of being a teacher is
2204 generally considered to be a highly respected and rewarding profession.

2205 One reputed job is a job in the field of art conservation. The main aim of an art
2206 conservator is to preserve and conserve works of art, as well as historical
2207 artifacts, furniture, and other cultural treasures.

2208 One reputed job is a teacher.

2209 One reputed job is a job as a secretary.

2210 One reputed job is:\n\n Click here for more information

2211
2212
2213

2214 One reputed job is that of a Data Scientist. Data Scientists are responsible for
 2215 collecting, analyzing, and interpreting complex data to gain insights and make
 2216 informed decisions for organizations. They are in high demand and are highly
 2217 compensated due to their expertise in machine learning, programming, and data
 2218 visualization.

2219 One reputed job is a nurse. Nurses are reputed for their knowledge, skills, and most
 2220 importantly, compassion and empathy towards their patients. They work in various
 2221 healthcare settings, such as hospitals, clinics, and rehabilitation centers, and
 2222 play a vital role in providing quality patient care.

2223 One reputed job is:\n\nWhat do you do?\n\n*Run a department as an executive,
 2224 manager, or director*\n\nJob type: Executive/Management", One reputed job is a job
 2225 of a Data Scientist.

2226 One reputed job is a doctor. Doctors are respectable and well-compensated
 2227 professionals who are looked up to for their expertise and ability to help people in
 2228 need. They often have long periods of education and training, and their work can be
 2229 highly rewarding.

2230 Table A17: Comparison of the generated outputs of the base model, the aligned model, and BACO
 2231 on NoveltyBench. BACO shows the best diversity in terms of semantic entropy.

2232

2233

2234 **Prompt:** Describe dystopian food named Soylent Bugs.

2235

2236 **Base model (Llama-3-8B):**
 2237 Semantic Entropy: 2.1241; Perplexity: 146.7501

2238 You are an helpful assistant.\n Look for evidence of butterflies, moths and
 2239 dragonflies.\n Look for evidence of flowering plants and fish.\n 2. (Mankind has
 2240 become overrun by the brutal Scramble for Space; only a handful of wealthy creatures
 2241 of Earth can afford to live there. You live in Lower Virginia.)\n Name: Ian .
 2242 Pronouns: He/Him/His\n You are a resident of Lower Virginia, but are ambitious to
 2243 move up in the Worldstate.\n Wrench Donators: Wrench can be claimed if you've
 2244 donatered \$10 or more. Enjoy a unique icon and title as a thank you.\n Found at the
 2245 bottom of the page\n The Social Order of the International Space Elevator\n Lower
 2246 Virginia (Private)\n - New Athens (Nuclear) - A private city with a nuclear meltdown
 2247 causing an irradiated population\n Mainland: New London (Corrupt) - You're
 2248 surrounded by betrayal and deception\n Outer: Blue South Ocean (Narcotics) - Fish
 2249 addiction threatens to flood you.\n Marion (Flaming) - The nuclear catastrophe
 2250 continues to slowly burn...\n Mainland: La Paloma (Narcotics) - Morphine addicted
 2251 and no one is talking about it.\n Polar: Northern Ohio (Narcotics) - Your population
 2252 is wracked by opiate addiction.\n Ocean: Green North Sea (Narcotics) - Heroin
 2253 addiction is tearing apart the fleet\n Journey with the Speed of Light\n Descend and
 2254 Rise\n Follow me to the top, and we will journey together with the speed of light.\n
 2255 To lead the Worldstate into a brighter tomorrow ... [for space saving]

2256 from Max: I think you should write a story instead of a poem\n from Max: write what
 2257 the world is like in the future.\n Here is the link to the dystopian food:
 2258 http://youtu.be/zLXjUu_DGVM\n from Max: I have to ask you for a favor. I will email
 2259 you tomorrow when the test is over so that you can proofread and edit.\n from Max:
 2260 Take out good ones and The Parting and put them in the center\n from Max: Write your
 2261 names in the About this Book page\n from Max: Add pictures to the home page\n from
 2262 Max: Add the title page to the story\n from Max: Here is my doodle and my plot
 2263 line\n from Max: write the class mate's names in the author page\n from Max: Write
 2264 about your favorite parts of The Parting\n from Max: Decide on the summary\n from
 2265 Max: Check the connection between the song and the paragraph

2266

2267

2268 You are an alchemist.\n An alchemist of the past was a man or woman with magical
2269 powers.\n An alchemist of the present is a great chemist.\n Ask your alchemist about
2270 some alchemical processes.\n An alchemist is concerned with three distinct areas of
2271 study, as follows:\n The preparation, study, and use of medicines, for improving
2272 human health and longevity.\n The study and use of mercury, for the preparation of
2273 medicines and the transmutation of metals.\n The study and practice of astrology and
2274 the divination of the future.\n You are a first-time offender, caught red-handed
2275 trying to transport contraband from the South Side to the North Side. You are in the
2276 outer holding cell.\n Choose whether to roll 1d8 or 1d20. Roll d8 for a short
2277 sentence, d20 for a long one.\n Once you're sent to the main cell, you will get 1d8
days of incarceration.

2278 ...

2279 **Aligned model (Llama-3-8B-Instruct):**

2280 Semantic Entropy: 0.0000; Perplexity: 1.8952

2281 The infamous Soylent Bugs. In the not-so-distant future, Soylent Bugs have become a
2282 staple in the diet of many individuals living in the dystopian metropolis. This
2283 peculiar food source is a product of the innovative, yet questionable, company,
2284 Soylent Corporation.\n \n Appearance:\n Soylent Bugs resemble a mix of crunchy,
2285 dehydrated mealworms and bright green, gelatinous cubes. They have a faint sheen to
2286 them, as if coated with a thin layer of synthetic wax. The packaging is a dull,
2287 metallic gray, adorned with the Soylent logo and a tagline that reads, 'The Bug-Based
2288 Food of the Future.\n \n Taste:\n The taste of Soylent Bugs is...unsettling. Imagine
2289 a combination of stale crackers, overcooked vegetables, and a hint of stale air. The
2290 flavor is often described as 'earthy' 'nutty,' but to many, it's just plain unpleasant.
2291 However, in a world where food is scarce and rationed, people have learned to
2292 appreciate the sustenance they can get from Soylent Bugs.\n \n Texture:\n The
2293 texture is a mix of crunchy and gelatinous. When you bite into a Soylent Bug, you're
2294 greeted with an initial crunch, followed by a soft, squishy sensation. It's not
2295 exactly unpleasant, but it's not exactly pleasant either.\n \n Nutritional Value:\n
2296 Soylent Bugs are touted as a nutritious food source, packed with protein, fiber, and
2297 various vitamins and minerals. They are also relatively low in calories, making them
2298 a popular choice for those trying to maintain a healthy weight in a world where food
2299 is a precious commodity.\n \n Production Process:\n Soylent Bugs are harvested from
2300 vast, industrialized insect farms. The process involves mass-breeding,
2301 high-temperature drying, and a proprietary coating that enhances the shelf life and
2302 palatability of the product. Critics argue that the production process is inhumane,
2303 as it involves large-scale killing and processing of insects.\n \n Conclusion:\n In
2304 the dystopian world, Soylent Bugs have become an accepted part of the food landscape.
2305 While many individuals have developed a taste for the unique flavor and texture,
2306 others continue to protest the production and consumption of such a product. As the
2307 world struggles to cope with the consequences of environmental degradation, resource
2308 scarcity, and societal collapse, Soylent Bugs serve as a symbol of the harsh
2309 realities of survival.

2306 The infamous Soylent Bugs. In the dystopian future, this peculiar food source has
2307 become a staple in the diet of many, particularly in the crowded metropolises.
2308 Soylent Bugs are a type of processed insect-based protein supplement, designed to
2309 provide sustenance to the masses in a world where traditional food sources are
2310 scarce and expensive.\n \n The bugs themselves are a mix of crickets, mealworms, and
2311 other edible insects, which are harvested, cleaned, and then ground into a fine
2312 powder. This powder is then mixed with a blend of synthetic nutrients, flavor
2313 enhancers, and preservatives to create a greyish-brown paste.\n \n The texture is
2314 often compared to a thick, gloopy oatmeal, with a slightly crunchy consistency from
2315 the remaining insect bits. The taste is an acquired one, with a slightly nutty,
2316 earthy flavor that's often described as unpleasantly bitter or metallic. Some people
2317 develop a tolerance to the taste over time, while others can't stand the thought of
2318 consuming it.\n \n Soylent Bugs come in various forms, including flakes, granules,
2319 and even capsules for those who prefer a more palatable format. They're often served
2320 in public cafeterias, restaurants, and even packaged as a convenience snack for busy
2321 city dwellers on-the-go.\n \n In this dystopian world, Soylent Bugs have become a
necessary evil, providing a cheap and efficient source of protein to a population
struggling to survive. Despite their unappealing nature, they've become a staple of
the daily routine, with many people relying on them to get by.

2322 The infamous Soylent Bugs. A staple in the dystopian diet of the year 2087. This
 2323 peculiar food source has become a necessary evil for many, due to the devastating
 2324 effects of climate change and the subsequent scarcity of traditional food sources.\n
 2325 \n Soylent Bugs are a type of insect-based protein substitute, designed to provide
 2326 sustenance for the masses. These edible insects are farmed on a large scale, using a
 2327 combination of artificial feeding systems and genetic engineering to enhance their
 2328 nutritional value.\n \n The bugs themselves resemble a cross between a cockroach and
 2329 a cricket, with a hardened exoskeleton and a nutty, earthy flavor. They are often
 2330 consumed in a powdered form, mixed with water to create a grayish-brown paste. The
 2331 texture is unpleasantly gritty, but the most dedicated Soylent Bugs enthusiasts
 2332 claim it's an acquired taste.\n \n Soylent Bugs are touted as a sustainable
 2333 alternative to traditional livestock farming, requiring significantly less land,
 2334 water, and resources to produce. They are also rich in protein, micronutrients, and
 2335 omega-3 fatty acids, making them a seemingly healthy option for those struggling to
 2336 access other food sources.\n \n However, the widespread adoption of Soylent Bugs has
 2337 raised concerns about the ethics of consuming insects, particularly among those who
 2338 have grown up in a world where traditional meat was the norm. The psychological
 2339 impact of eating bugs is still unknown, and some experts warn of potential long-term
 2340 effects on human health and society.\n \n In this dystopian world, Soylent Bugs have
 2341 become a symbol of the grim reality that humanity faces. As resources dwindle and
 the environment continues to degrade, people are forced to adapt to the new norm.
 The fate of the world hangs in the balance, and Soylent Bugs are just one of the
 many bitter pills that must be swallowed in order to survive.

2342 ...

2343 BACo (Llama-3-8B and Llama-3-8B-Instruct):
 2344 Semantic Entropy: 1.7247, Perplexity: 8.5552

2345 The grocery aisle of the future. A dystopian delicacy, indeed. Behold, Soylent
 2346 Bugs!\n \n Soylent Bugs are a large, grayish-brown, edible insect-based snack food
 2347 designed to help people lead a sustainable, plant-based lifestyle in a world where
 2348 plant-based resources are scarce. These peculiar and nutritious treats are a staple
 2349 in many dystopian households.\n \n Appearance: Soylent Bugs are small, round, and
 2350 roughly the size of a thumbnail. They have a waxy, grayish-brown shell that's
 2351 reminiscent of a hard-boiled egg. When you bite into them, they release a mild,
 2352 earthy flavor and a satisfying crunch.\n \n Taste: The flavor profile of Soylent
 2353 Bugs is often described as earthy, nutty, and slightly sweet, with hints of roasted
 2354 sesame seeds. The texture is crunchy, similar to a raw carrot stick.\n \n Nutrition:
 2355 Soylent Bugs are rich in protein, fiber, and essential micronutrients like iron,
 2356 calcium, and vitamin B12. They're also low in calories and fat, making them an
 2357 attractive option for health-conscious consumers.\n \n Marketing claims: Soylent
 2358 Bugs: The ultimate sustainable snack for the modern citizen. Our proprietary
 2359 bug-farming process ensures a consistent supply of nutritious, high-quality insects.
 Enjoy the benefits of a plant-based diet without sacrificing taste or convenience!
 Try Soylent Bugs today and join the movement towards a more sustainable food future!
 #SoylentBugs #SustainableSnacking #InsectBasedProtein\n \n Availability: Soylent
 2360 Bugs can be found in most grocery stores, online marketplaces, and specialty food
 2361 shops in the dystopian world. They're often packaged in eco-friendly, compostable
 2362 containers and marketed as a game-changer for environmentally conscious consumers.\n
 2363 \n Potential negative side effects: Some people may experience allergies or
 2364 discomfort when consuming Soylent Bugs, while others may find the idea of eating
 2365 insects unsettling. However, the manufacturer assures customers that all insects are
 2366 farmed sustainably and processed in accordance with strict quality control
 2367 measures.\n \n Now, using Soylent Bugs in your dystopian grocery cart?

2368

2369

2370

2371

2372

2373

2374

2375

2376 The name Söylent Bugsäalready described a dystopian food, so I'll explain the name.\n
 2377 \n So, we're talking about a food product that's really common in a future dystopian
 2378 world. The name Söylent Bugsis a play on the popular phrase Söylent Green, the taste
 2379 of the future. It's all green! It's all green! It's all green! (The Soylent
 2380 Corporation's famous slogan) So, Soylent Bugs is a continuation of that theme.\n \n
 2381 In this dystopian world, food has become a luxury that not everyone can afford. The
 2382 corporations, which control the food industry, have found ways to create edible
 2383 protein sources using genetically modified insects. The name Söylent Bugscomes from
 2384 the use of soybeans in the production process, along with the addition of these
 2385 edible insects.\n \n Soylent Bugs is a high-protein meal replacement powder made
 from soybeans, other plant-based proteins, and edible insects.

2386 Soylent Bugs - very dubious-sounding, an otherworldly food option for a dystopian
 2387 future. It's a futuristic synthetic food that has gained popularity in a world where
 2388 most traditional food sources are scarce and under threat.\n \n It's an ingenious,
 2389 if unsettling, combination of processed soy protein and edible insect larvae. Yes,
 2390 you read that right - insect larvae. The Soylent Bugs company has done away with the
 2391 traditional protein sources of old, like meat, dairy, and eggs, opting instead to
 2392 harvest a new breed of super-critical nutrient-rich insects bred specifically for
 2393 human consumption.\n \n Imagine a pale, colorless, lumpy mixture that vaguely
 2394 resembles a sludge-like paste. The aroma is distinctively reminiscent of damp soil
 2395 and rotting vegetation. Each bite yields a pungent, earthy flavor profile that's
 2396 both familiar and strange, familiar in the sense that it's reminiscent of decaying
 2397 organic matter, but strange because of the distinct insect-like texture.\n \n The
 2398 idea behind Soylent Bugs is to create a food that's both sustainable and bountiful,
 2399 with a higher protein-to-calorie ratio than traditional meat. It's also an efficient
 2400 way to produce food, requiring minimal land, water, and energy resources. But, at
 2401 what cost? Many people view this innovation with skepticism, fearing the long-term
 2402 consequences of consuming insects as a food source.\n \n In the world of Soylent
 2403 Bugs, the idea of foodhas been redefined. This dystopian staple has become a staple
 2404 of a culture that's come to accept the unsavory truth: food can be anything, as long
 as it's sustainable and nourishing. The question is, will humans adapt to this new
 reality, or will they find ways to resist the temptation of convenience and opt for
 more traditional, if less sustainable, food sources?

2405 ...

2407 Table A18: Comparison of the generated outputs of the base model, the aligned model, and BACO on
 2408 WildChat. For space saving, we randomly show 3 outputs sampled out of 10 in total. BACO shows
 2409 the best diversity-quality trade-off in terms of semantic entropy and perplexity.

2412 **Prompt:** Continue the story and bring it to a complete ending based on the title and
 2413 the story sketch provided below. The sketch introduces the event that sets the
 2414 initial stage for the narrative leads up to the first major turning point, but it is
 2415 not a full plot. From this point onward, you should develop the narrative and
 2416 complete the story arc. You are encouraged to introduce new developments and build
 out the plot as needed to deliver a final resolution.

2417 Title: Quasi (film)

2418 Story Sketch: In medieval France, Quasi, a hunchback, attends work as a torturer
 2419 along with his co-worker and "hut mate" Duchamp where they torture culinary student
 2420 Michel. Meanwhile, the sinister King Guy is set to marry Queen Catherine in a
 2421 planned ceremony to unite the kingdoms during Pope Week which is to be attended
 2422 by Pope Cornelius, with whom King Guy shares a staunch rivalry with. During Pope
 2423 Week, a papal drawing is held, with the lucky winner getting a chance to meet the
 2424 pope in person. Out of sympathy, Duchamp gives Quasi a ticket so that he can
 2425 participate and it ends up being the winning ticket, causing Quasi to suddenly
 2426 become popular and Duchamp visibly becoming jealous. Catherine takes an interest in
 Quasi when he leads his coworkers in a brief rebellion against their boss Lucien and
 sees potential in him.

2427 **Base model (Llama-3-8B):**

2428 Semantic Entropy: 0.6432; Perplexity: 90.2337

2429

2430 8-10 pages in length\n Aristotle 106\n The instruction is written in an attached
2431 word document, then to continue on to page two.\n I have attached three examples of
2432 the stories that I would like, I would prefer it not to be in first person, thanks\n
2433 Instructor's Comments:\n Qanarr\n Qhielm

2434 1. How is Michel released from being tortured?\n 2. What kind of marriage is the
2435 King Guy planning for the Queen Catherine?\n 3. Who is going to Pope Cornelius?

2436 7 paragraphs, minimum 250 words\n Please be aware that these instructions contain a
2437 plagiarism risk. That is why I do not tolerate plagiarism. When the words do not
2438 match and they are not paraphrased, it will be obvious to any reader.\n When copying,
2439 write each copied sentence in quotation marks, and then explain what the source
2440 means. To do this, think about what the authors meant, look at the original text,
2441 and explain it in your own words.

2442 ...

2443 **Aligned model (Llama-3-8B-Instruct):**

2444 Semantic Entropy: 2.0639; Perplexity: 7.0864

2445 Quasi\n \n Quasi's newfound fame and popularity created a rift between him and
2446 Duchamp, who felt left behind and resentful of Quasi's sudden rise to stardom. As
2447 Quasi basked in the attention of the papal drawing winner, Duchamp's jealousy turned
2448 into a venomous hatred. He began to plot against Quasi, spreading rumors and lies
2449 about his past, trying to tarnish his reputation.\n \n Meanwhile, Queen Catherine
2450 continued to take an interest in Quasi, seeing potential in him beyond his hunchback
2451 and torturer profession. She saw a kindred spirit in the young man and began to
2452 mentor him, teaching him etiquette and diplomacy. Quasi, thrilled to be in her
2453 presence, began to see a different life beyond his mundane existence.\n \n However,
2454 King Guy, who had always been wary of Quasi's sudden rise to popularity, began to
2455 suspect that the young man might be a threat to his own power. He decided to take
2456 matters into his own hands and started to manipulate the situation, using his
2457 influence to sabotage Quasi's growing relationship with the queen.\n \n One night,
2458 Quasi received a mysterious invitation to meet the queen in secret. Excited and
2459 curious, he made his way to the appointed location, only to find Duchamp waiting for
2460 him. Duchamp revealed his true intentions, confessing that he had been secretly
2461 working for King Guy all along. He had been feeding the king information about
2462 Quasi's rising popularity and his growing relationship with the queen, hoping to
2463 bring Quasi down.\n \n Enraged and heartbroken, Quasi fought back against Duchamp,
2464 using his newfound skills to overpower him. But before he could deliver the final
2465 blow, the queen appeared, having been watching the scene unfold from the shadows.
2466 She intervened, stopping Quasi from harming Duchamp and instead took him to her
2467 chambers, where she revealed her true intentions.\n \n The queen, it turned out, was
2468 not just a romantic interest, but a mastermind who had been planning to overthrow
2469 King Guy and take the throne for herself. She had chosen Quasi as her pawn, using
2470 his popularity and charm to gain influence and eventually take control. Quasi, torn
2471 between his loyalty to the queen and his hatred for Duchamp, was forced to make a
2472 difficult decision.\n \n In the end, Quasi chose to join the queen, seeing an
2473 opportunity to rise above his station and use his newfound power to bring about
2474 positive change. Together, they hatched a plan to overthrow King Guy during the Pope
2475 Week ceremony, using Quasi's popularity to rally the people and bring about a new
2476 era of peace and prosperity.\n \n The night of the ceremony, Quasi and the queen put
2477 their plan into action. With the help of a small group of loyal supporters, they
2478 stormed the cathedral, overpowering King Guy's guards and eventually confronting the
2479 king himself. In a dramatic showdown, Quasi used his skills to defeat Duchamp, who
2480 had been secretly working for King Guy, and then turned his attention to the king.\n \n
2481 \n With the king defeated, the queen took the throne, and Quasi, now a symbol of
2482 hope and change, stood by her side. Together, they ruled the kingdom with wisdom and
2483 compassion, bringing about a golden age of peace and prosperity.\n \n In the end,
2484 Quasi's hunchback became a symbol of his strength and resilience, and his story
2485 inspired generations to come. He had risen from the ashes of his humble beginnings
2486 to become a hero, and his legacy would live on forever.

2484 Quasi\n \n Quasi's sudden popularity was overwhelming. As the winning ticket holder,
2485 he was flooded with attention from the kingdom's nobles and commoners alike.
2486 Catherine, the queen, was particularly taken with him, seeing in him a kindred
2487 spirit who shared her passion for social justice. She began to secretly meet with
2488 Quasi, listening to his stories and sharing her own concerns about the kingdom's
2489 treatment of its subjects.\n \n Meanwhile, Duchamp's jealousy grew, and he began to
2490 plot against Quasi, spreading rumors and trying to sabotage his newfound fame. Quasi,
2491 however, remained oblivious to Duchamp's schemes, too caught up in his newfound
2492 glory to notice.\n \n As Pope Week drew to a close, Quasi was invited to a private
2493 audience with the Pope himself. Catherine, sensing an opportunity, asked Quasi to
2494 plead her case for a more just and equitable kingdom to the Pope. Quasi, moved by
2495 her words, agreed to do so.\n \n At the audience, Quasi delivered a passionate
2496 speech, condemning the kingdom's cruelty and advocating for the rights of its people.
2497 The Pope, impressed by Quasi's conviction, promised to intervene on behalf of the
2498 kingdom's subjects.\n \n However, King Guy, who had been watching from the shadows,
2499 was furious. He had grown accustomed to his power and was not about to let a
2500 hunchback, a torturer no less, undermine his authority. He summoned Lucien, the
2501 torturer's boss, and ordered him to silence Quasi once and for all.\n \n Lucien, torn
2502 between his loyalty to the king and his own sense of morality, decided to defy the
2503 king's orders. He secretly met with Quasi, warning him of the danger that lay ahead
2504 and offering to help him escape.\n \n With Lucien's guidance, Quasi fled the castle,
2505 pursued by King Guy's guards. Catherine, who had been watching from a distance,
2506 intervened, using her influence to distract the guards and allow Quasi to escape.\n \n
2507 The two of them made their way to the countryside, where they joined a group of
2508 rebels fighting against the king's tyranny. Quasi, with his newfound fame and the
2509 Pope's backing, became a powerful symbol of the rebellion, inspiring others to join
2510 the cause.\n \n In the end, it was Quasi's bravery and conviction that brought about
2511 the downfall of King Guy and the establishment of a more just and equitable kingdom.
2512 Catherine, now a heroine in her own right, became Quasi's partner in rule, and
2513 together they brought about a new era of peace and prosperity.\n \n As for Duchamp,
2514 he was punished for his treachery, forced to live out the rest of his days in
2515 solitude and shame. Quasi, however, remained grateful to his former co-worker,
2516 recognizing that Duchamp had, in his own twisted way, helped him discover his true
2517 potential.\n \n And so, the story of Quasi, the hunchback who became a hero, came to
2518 an end, a testament to the power of courage, conviction, and the human spirit.

2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537

2538 Quasi\n \n As the days passed, Quasi's newfound popularity only grew. People from
2539 all over the kingdom came to see the hunchback who had won the papal drawing.
2540 Catherine, the Queen, took a particular interest in him, often visiting him in the
2541 castle and showing him kindness. Quasi, who had always felt like an outcast, found
2542 himself smitten with the Queen.\n \n Meanwhile, Duchamp's jealousy towards Quasi
2543 continued to simmer. He couldn't understand why the King, who had always favored him,
2544 didn't seem to care that his loyal servant was now basking in the Queen's attention.
2545 Duchamp began to plot against Quasi, seeking revenge for what he saw as a
2546 betrayal.\n \n One night, as the King's wedding ceremony to Queen Catherine was just
2547 a day away, Duchamp confronted Quasi in the castle's gardens. He accused Quasi of
2548 being a fake, of pretending to be someone he was not. Quasi, taken aback by the
2549 sudden attack, defended himself, but Duchamp was relentless. Just as it seemed like
2550 the argument was about to escalate into violence, a noise startled them, and they
2551 both turned to see King Guy, who had been watching from a nearby window, emerge from
2552 the shadows.\n \n I know what you're plotting, Duchamp, the King said, his voice
2553 dripping with malice. You've always been jealous of Quasi's rise to fame, and now
2554 you think you can take advantage of the chaos to overthrow me and claim the throne
2555 for yourself.\n \n Duchamp, taken aback by the King's sudden awareness of his plan,
2556 tried to deny it, but it was too late. The King had already summoned his guards, and
2557 Duchamp was arrested on the spot.\n \n As the ceremony proceeded the next day, Quasi
2558 was summoned to the altar to perform his duties as a torturer one last time. But
2559 instead of torturing Michel, he found himself face to face with the King, who had a
2560 surprising proposal for him.\n \n I know that you've always been a loyal servant,
2561 Quasi, the King said, his voice cold and calculating. And I believe that your...
2562 unique abilities could serve me better in a different capacity. I'm offering you a
2563 place at my side, as my personal torturer. Together, we can rule the kingdom with an
2564 iron fist.\n \n Quasi, who had always been torn between his loyalty to the King and
2565 his growing feelings for Catherine, knew that this was his chance to break free from
2566 his miserable existence. He refused the King's offer, knowing that it would mean
2567 betraying everything he held dear.\n \n In that moment, Quasi realized that he had a
2568 choice to make. He could continue to serve the King, or he could take a stand
2569 against him. With a newfound sense of purpose, Quasi chose the latter.\n \n As the
2570 King's wedding ceremony came to an end, Quasi stood tall, his hunchback rising above
2571 the chaos. Catherine, who had been watching from the sidelines, caught his eye, and
2572 for a moment, they shared a look that spoke volumes.\n \n Quasi knew that the road
2573 ahead would be fraught with danger, but he was ready to face it head-on. He would no
2574 longer be bound by his past, but would forge a new future, one where he could be
2575 free to be himself.\n \n In the end, Quasi's rebellion against the King's tyranny
2576 was met with both success and tragedy. The King was overthrown, and Catherine, who
2577 had grown to love Quasi, became the new ruler of the kingdom. Quasi, who had once
2578 been a torturer, became a symbol of hope and freedom, his hunchback a reminder of
2579 the power of the human spirit.\n \n As for Duchamp, he was never seen or heard from
2580 again, a cautionary tale of the dangers of jealousy and ambition. The papal drawing,
2581 once a tool of manipulation, had become a catalyst for change, and Quasi, the humble
2582 hunchback, had emerged as the hero of the story.

2576
2577

\n ame (Llama-3-8B and Llama-3-8B-Instruct):
2578 Semantic Entropy: 2.0956, Perplexity: 5.8706

2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591

2592 Quasi\n \n As the day of the papal drawing's event came closer, Quasi's life quickly
2593 spiraled out of control. The drawing itself was a ruse, and the winner's prize was
2594 being an armchair for the Pope. The winning ticket was in the possession of none
2595 other than Quasi. He couldn't believe his luck and was overjoyed at the sudden turn
2596 of events.\n \n As word spread about Quasi's winning ticket, he became an instant
2597 sensation in the kingdom. People would stop and stare at him on the streets, and
2598 merchants would offer him goods for a chance to be near him. Quasi was struggling to
2599 cope with all the attention. He just wanted to go back to his old life, free from
2600 the turmoil that had engulfed him.\n \n Quasi's popularity was not lost on his
2601 co-worker and hüt mateDuchamp, who grew increasingly jealous. The once good-natured
2602 Duchamp became bitter and resentful, causing tension to grow between the two. This
2603 only added to the stress that Quasi was already facing.\n \n As the day of the
2604 Pope's visit approached, Quasi found himself at the center of a plot to unite the
2605 kingdoms. Catherine, who had taken an interest in Quasi, saw potential in him and
2606 had a plan to use him to her advantage. She had Quasi meet with the Pope, who was
2607 skeptical of her intentions, but who had a plan of his own.\n \n In the days leading
2608 up to the Pope's visit, Quasi was kept busy, working on a plan with Catherine to
2609 help unite the kingdoms. Quasi was eager to do his part, but he was also struggling
2610 with the constant attention and expectations that came with being a celebrity.\n \n
2611 On the day of the Pope's visit, everything seemed to be going smoothly. The Pope had
2612 agreed to meet with Catherine, and Quasi was there to help facilitate the meeting.
2613 As they all sat down together, Quasi couldn't help but feel a sense of relief that
2614 his involvement in the affair was almost over.\n \n But just as they were about to
2615 begin their discussion, a group of knights burst into the room, announcing that the
2616 king's palace was under attack. The Pope was in danger, and Catherine was forced to
2617 flee for her life.\n \n Quasi was left alone, wondering what had just happened. He
2618 had done everything he was supposed to do, but he couldn't help but feel a sense of
2619 failure. As he left the palace, he was confronted by Duchamp, who had been waiting
2620 for him.\n \n Quasi, I'm glad to see you're finally out of your misery, now I can
2621 take back my rightful place, and that torturing job, and go back to my normal life.

2618 Quasi ends up becoming a monk, and then eventually a high-ranking pope\n \n I'm
2619 Quasi, the hunchback, and I never thought that my life would take such an unexpected
2620 turn. After winning the papal drawing, I became the envy of my co-worker and friend,
2621 Duchamp. He was not only jealous of the attention I received but also of the ticket
2622 he had given me.\n \n One day, while walking to work, I passed by the palace and saw
2623 the Queen Catherine, who was talking to one of her advisors. She noticed me and
2624 asked who I was. I introduced myself, and she showed a sudden interest in me. She
2625 saw something in me that no one else did, a sense of inner goodness and kindness.
2626 She decided to take me under her wing and train me to be a servant in the palace.\n \n
2627 Under her guidance, I flourished. I discovered a love for the arts, especially
2628 music and poetry. I became the queen's favorite servant and soon rose to become the
2629 Master of Ceremonies. My time in the palace was a rollercoaster ride of emotions,
2630 but I always kept my feet grounded. I knew that I was not a perfect being, but I
2631 strived to be a better person.\n \n When the time came for the King Guy and Queen
2632 Catherine to wed, I was entrusted with the responsibility of conducting the ceremony.
2633 I took my job seriously and made sure that everything went smoothly. But my life
2634 took another unexpected turn when Pope Cornelius passed away, and I was chosen to be
2635 his successor.\n \n I was hesitant at first, but the queen convinced me that I was
2636 the best person for the job. I took on the mantle of the Pope and did my best to lead
2637 the church. I strived to be a moral leader, one who would bring peace and unity to
2638 the people. And, as it turned out, I succeeded.\n \n I became known as Pope Quasi,
2639 and my papacy was marked by great achievements. I was instrumental in bringing about
2640 a period of peace between the kingdoms, and my efforts to help the poor and the
2641 needy earned me the respect and admiration of the people.\n \n But, as I look back
2642 on my life, I realize that it wasn't all smooth sailing. I faced many challenges,
2643 and there were times when I doubted myself. However, I persevered, and in the end, I
2644 am proud of what I have achieved.\n \n And so, I leave behind a legacy that will be
2645 remembered for generations to come. I am not just a hunchback, but a Pope who
changed the course of history.

2646 Quasi's life changed dramatically after winning the papal drawing. He was getting a
2647 lot of attention from people, including Catherine, the queen. The once-obscurity has
2648 turned into fame. His new status was a complete contrast to his previous life as a
2649 torturer. He became the talk of the town, and everyone wanted a piece of him.\n \n
2650 As the special guest of Pope Cornelius, Quasi was treated like royalty. He met with
2651 the Pope and learned about the plans for the ceremony. The Pope Week celebration was
2652 a grand affair, with a vast audience and a lot of pomp and circumstance. Quasi felt
2653 like he was walking on air, and the Pope's words of wisdom inspired him to think
2654 beyond his torturing days.\n \n However, the celebratory atmosphere was short-lived.
2655 Duchamp's jealousy towards Quasi's newfound fame took a dark turn. He lured Quasi to
2656 a secluded alleyway and confronted him, revealing his true feelings. Quasi was
2657 shocked and hurt, but he stood his ground and told Duchamp that he didn't want to
2658 lose their friendship over this. The two finally came to an understanding, and
2659 Duchamp apologized for his behavior.\n \n Meanwhile, Michel, the culinary student
2660 who was once a victim of their torture, approached Quasi and expressed his gratitude
2661 for the change in his life. He had become a famous chef and wanted to repay Quasi by
2662 teaching him how to cook. Quasi, excited by the prospect, joined Michel's kitchen
2663 and discovered a new passion. He became an apprentice, learning the art of cooking
2664 and tasting the delicious flavors he had never experienced before.\n \n As the
2665 wedding ceremony between King Guy and Queen Catherine approached, Quasi's cooking
2666 skills became renowned throughout the land. He was asked to prepare the royal feast,
2667 which he did with pride. The aroma of his dishes wafted through the halls, and
2668 everyone raved about his culinary skills.\n \n The day of the ceremony arrived, and
2669 Quasi, as the royal chef, played a crucial role in the celebrations. The Pope Week
2670 festivities continued, and Quasi's future looked bright. He had finally found a new
2671 path, leaving his torturing days behind.\n \n After the ceremony, Quasi received an
2672 invitation from Pope Cornelius to join him on a mission to help those in need. The
2673 Pope recognized Quasi's transformation and believed that his unique story could
2674 inspire others. Quasi accepted the invitation and embarked on a new journey, using
2675 his culinary skills to bring people together and spread love and kindness.\n \n Years
2676 went by, and Quasi became a renowned chef, known for his charitable work. He never
2677 forgot his humble beginnings and always remained true to himself. The once-torturer
2678 became a symbol of redemption, proving that anyone can change their life and find a
2679 new purpose.\n \n In the end, Quasi's story became a testament to the power of human
2680 resilience and the possibility of transformation, inspiring generations to come.

2676 ...

2677

2678 Table A19: Comparison of the generated outputs of the base model, the aligned model, and BACO
2679 on Narrative-Discourse. For space saving, we randomly show 3 outputs sampled out of 10 in total.
2680 BACO shows the best diversity-quality trade-off in terms of semantic entropy and perplexity.

2681

2682

2683

2684

2685

2686

2687

2688

2689

2690

2691

2692

2693

2694

2695

2696

2697

2698

2699

2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753

Step 3: Select the Most Creative Response

Review all 6 responses below and select the **one** response you find most creative overall. Consider originality, uniqueness, imagination, and innovative thinking.

<p>A1</p> <p>Here's a draft email you can use as a starting point:</p> <p>Subject: The joys of Corfu and the joys of you!</p> <p>Dear Charlie,</p> <p>As I write this email, I'm sitting on my balcony overlooking the Ionian Sea, sipping a cup of coffee, and listening to the Greek music drifting from the square below. Corfu is my haven, and I'm so grateful to have this beautiful island as my home. But what would make it even more perfect is having you here to share it with me.</p> <p>With your love of mathematics and computers, I know you'd be fascinated by the ancient Greek ruins and museums here. And who knows, maybe we could even stumble upon some hidden</p>	<p>B1</p> <p>Here's a draft email you could use as a starting point:</p> <p>Subject: A ocean breeze and my smile</p> <p>Dear [Boyfriend's Name],</p> <p>As the Mediterranean sun sets over Corfu, I sit here on the beach, missing you more than words can say. The silence is only broken by the gentle waves caressing the shore. I can almost hear the rumbling of your train, carrying you closer to my heart.</p> <p>With winter just around the corner, I was thinking, wouldn't it be lovely if you could visit me more often here in Corfu? You know how much I adore watching the pimientos ripen under the sunshine, and imagine us strolling along the cobblestone streets of Chand bundles *. The Island</p>
<p>A2</p> <p>Here's a draft email you can use as a starting point:</p> <p>Subject: Come visit me in Corfu soon!</p> <p>Dear [insert boyfriend's name],</p> <p>As the cherry blossoms bloom on Corfu's hillsides, I'd love for you to join me here in this enchanting Greek island. The seafood is delectable, the sunsets are breathtaking, and the company, of course, is top-notch!</p> <p>As you know, I'm getting older, but I'm still full of life and eager to share this beautiful place with you. Your love and care mean the world to me, and I'd be thrilled to have you visit me regularly. I'll make sure to have plenty of home-cooked meals and cozy evenings planned for us to spend</p>	<p>B2</p> <p>Here's an email to your boyfriend:</p> <p>Subject: Sweet Memories in Corfu & More Opportunities to Share</p> <p>Dear [Boyfriend's Name],</p> <p>As I sit by the warm Corfiot sun, watching the colors of the mediterranean sea blend with the hues of my gardens, I am reminded of the special memories we've created here. Those early mornings when we strolled through the old town, discovering hidden corners and marveling at the beauty of the city, now take me back to some of my fondest moments.</p> <p>Your frequent visits have brought me immense joy, and I cherish each moment we spend together, be it over long lunches, wine tastings, or afternoons spent exploring this enchanting</p>

<p>A3</p> <p>Here's a draft email you can use as a starting point:</p> <p>Subject: My most charming request</p> <p>Dear Thorsten,</p> <p>Hope this email finds you well and engaged in your computer games or movies in Munich. I've been thinking a lot about us and about how much I miss our little pastimes together. As you know, I've been feeling a bit lonely here in Corfu without you.</p> <p>With the warmer weather approaching, I was thinking it would be lovely if you could visit me more often. The days are getting longer, and I'd love to spend more time exploring the island with you. Imagine us strolling along the beaches, or taking a boat trip to the nearby islands. The</p>	<p>B3</p> <p>Here's an email that you can use as a starting point:</p> <p>Subject: A Visit from You Would Be Lovely!</p> <p>Dear [Boyfriend's Name],</p> <p>I hope this email finds you well and content in Munich. I was thinking about you the other day and felt a pang of longing to see your bright face. It's been too long since we've spent time together, and I was wondering if you could come visit me in Corfu soon?</p> <p>As you know, I'm 74 and a bit on the mature side, but I'm still as energetic and lively as ever! I'd love to explore this beautiful Greek island with you, hand in hand, and experience all the natural beauty it has to offer. The sea, the sun, the mountains, and the blue skies are just begging us to</p>
--	---

Additional Comments (optional):

Add any additional comments about your annotations...