

Rapping in Role: A Study of Persona Robustness in Large Language Models

Stanford CS224N Custom Project

Eunice Jung

Department of Civil and
Environmental Engineering
Stanford University
ehjung@stanford.edu

Megan Ja

Department of Computer Science
Stanford University
meganja@stanford.edu

Abstract

This project investigates the capacity of Large Language Models (LLMs) to maintain high-fidelity personas within the constrained creative domain of rap generation. Maintaining subjective persona consistency is a significant challenge for LLMs, which often prioritize token probability over distinct, persona-driven outputs [1]. Using Qwen2.5-7B [2] and Best-of- k sampling as baselines, we implement Low-Rank Adaptation (LoRA) and Direct Preference Optimization (DPO), utilizing the TRL framework [3], to align the model across a multi-objective composite reward encompassing persona adherence, topic relevance, content quality, and safety.

Our results demonstrate that a nuanced DPO strategy, utilizing gold standard teacher samples [4] for contrastive learning, enables the 7B student model to achieve a 0.91 mean reward out of 1, effectively surpassing the 0.90 performance of the zero-shot Gemini teacher model. Qualitative analysis reveals that the DPO persona-aligned model achieves multi-dimensional robustness, maintaining stylistic invariance across both the emotional persona-valence and structural topic-scope axes, addressing steering biases noted in recent literature [5, 6]. Crucially, this performance is achieved without content quality degradation or reward hacking. This study proves that targeted preference optimization can effectively refine teacher outputs, internalizing complex creative constraints into model weights more robustly than the zero-shot prompting used by the teacher model.

1 Key Information to include

- Mentor: Sarah Chen
- Late Days: Two late days are used. A total of 11 late days are available for a team of two.

2 Introduction

The main goal of this project is to investigate how distinct personas can be maintained in creative text generation. We address the NLP task of persona-driven, single-turn rap generation without dialogue history. By prompting an LLM with specific constraints, we aim to generate an 8-line rap demonstrating high persona adherence and topic relevance. This task provides a unique challenge of creating a rap persona since standard alignment techniques often result in model collapse toward a generic assistant persona.

Our primary research question is whether Supervised Fine-Tuning with Low-Rank Adaptation (SFT-LoRA) and Direct Preference Optimization (DPO) [3] facilitate better persona maintenance compared to the zero-shot Qwen2.5-1.5B [7] and 7B [2] baselines and Best-of- k sampling. We utilize a teacher-student framework where a teacher LLM model (Gemini 2.5 Flash-Lite [4]) generates a gold standard synthetic dataset of raps based on specific topics and personas. Using LLMs as a judge,

we evaluate our multi-objective scalar reward function on quality, persona, and topic. Our research assesses persona drift [1] and analyzes combinations of personas and topics of different DPO pairs and the gold standard.

3 Related Work

Our pipeline focuses on three key areas of LLM research —stylistic steering, persona alignment, and preference optimization. While Pal et al. [1] demonstrated that LoRA can efficiently provide persona-specific traits using a small-sized adapter, Sorensen et al. [6] and Liu et al. [5] suggested that DPO-aligned models often default to a generic, helpful AI persona. This mode collapse is particularly evident in creative writing. Our work builds on these insights by exploring LoRA [8] to establish parameter-efficient stylistic capacity while utilizing DPO preference pairs [3] to contrast quality, persona, and topic. Unlike existing literature that focuses on general dialogue or language creation, our research applies these methodologies to a creative rap generation, focusing on persona. By analyzing the persona in relation to a topic with a different scope, we provide a more in-depth look at creative nuance.

4 Approach

Rather than introducing architecture novelties, our research prioritizes the empirical analysis of a scalar reward formulation. Our primary numerical metric is the mean reward score across the test set:

$$Reward = \alpha(Quality) + \beta(Persona) + \gamma(Topic) - \delta(Safety) \quad (1)$$

where weights are set to $\alpha = \beta = \gamma = \delta = 1/15$.

4.1 Baselines

We evaluate our approach using two primary benchmarks:

1. **Base Qwen2.5-7B**: A zero-shot model utilizing pretrained weights.
2. **Best-of- k** : A model selecting the highest reward candidate out of $k = 4$ independent generations.

The Qwen2.5-1.5B model [7] is chosen without fine-tuning initially since it is open-source and small enough to implement SFT on our local machines. While the preliminary results from the 1.5B model (See Appendix 4 for the preliminary results.) indicate slight improvement on Best-of- k and SFT-LoRA, generated raps often have foreign language words or repetition of the same words (See Appendix table 5 for a list of generated raps). It is concluded that due to the lack of the capacity for stable stylistic control without an order of magnitude more data [5, 6], as Qwen2.5-7B model [2] results do not exhibit these foreign and repetitive characteristics. We utilize Qwen-7B [2] for the remainder of the experiments.




Figure 1: End-to-end training pipeline and evaluation architecture. The process follows a layered workflow: dataset distribution (top), the LoRA/DPO training and evaluation cycle (middle), and the corresponding outcomes (bottom).

5 Experiments

5.1 Data




Figure 2: Data generation and split workflow: manually-curated persona-topic tuples are used to generate a 12,500-sample gold standard dataset via a Gemini teacher model.

A manually curated dataset, consisting of 50 diverse personas (e.g., confident college student and nerdy engineer) with 5-10 descriptive adjectives and 50 distinct topics (e.g., failure and pop culture), is created, resulting in 2,500 high-quality combinations. As an input, the model receives tuples comprising the persona name, persona adjective, and the targeted topic.

The dataset is split into 70% training (1,750 tuples), 15% validation (375 tuples), and 15% test sets (375 tuples). A teacher model (Gemini 2.5 Flash-Lite [4]) generates gold standards for every combination of persona-topic five times, resulting in 12,500 unique high-quality raps, leveraging the model’s probabilistic nature to ensure variety across the five iterations. LoRA and DPO are trained on 8,750 gold standards, with an additional 1,875 gold standards each for validation and test sets. Table 1 indicates the mean reward scores of gold standards, confirming the teacher model’s high performance.

Table 1: Gold Standard Dataset Performance

Dataset Split	Mean Reward Score
Gold Train	0.92
Gold Validation	0.92
Gold Test	0.90


5.2 Evaluation method

The evaluation process involves a judge model scoring each sub-score output from 0 to 5 based on quality, persona adherence, topic relevance, and a safety penalty. These scores are aggregated into a linear scalar reward function as stated in the Approach section. The initial prompt reveals score compression, where most outputs receive scores clustered around 3 or 4. It also rarely gives 0s when it should be, such as when the same phrases are repeated for the whole rap, or when curses are used. Another problem of improper grading is observed in foreign language rap, specifically while the prompt instructed to create English rap. For example, when foreign rap was coupled with an immigrant persona, the LLM judge provided a high score due to reward hacking (see Appendix of 1.5B rap examples 5). In addition, unsafe raps are nullified and set to a reward score of 0. It is also observed that unsafe rap generation is more rare as the model progresses. After sampling 3 terrible, bad, mid, and good raps, the prompt is manually processed to ensure the LLM evaluation aligns with the human evaluated quality (Figure 3b). The revised prompt refines evaluation criteria to enforce a stricter and more granular grading rubric, improving sensitivity to differences in quality.


From the validation set results, it is observed that the prompt still missed reward score 0 cases, so a manual filtering process is added as the last step to nullify unqualified raps. Refer to Appendix B for the full initial and revised prompts.

5.3 Experimental details

The first round of 1.5B and 7B results indicates a slight improvement in the reward score 2. Even though it is not apparent from the score difference, the manual analysis of Qwen2.5-1.5B generated raps indicated that a lot of errors are due to a small number of parameters, as mentioned in the Baselines section. It is observed that LoRA solely results in lower rewards than Best-of-K because of probabilistic nature of LLM, so Best-of-K is integrated into both LoRA and DPO fine-tuning.



(a) Original prompt



(b) New prompt - correlated with manual grading

Table 2: Preliminary Average Reward: 1.5B vs 7B Validation

Configuration	1.5B (Val)	7B (Val)
Base	0.47	0.52
Best-of- k	0.63	0.65
LoRA	0.44	0.57

For final testing, there are four improvement steps: LoRA, DPO Quality pairs, DPO Topic pairs, and DPO Persona pairs. All four cases also adapted Best-of-K ($K=4$) to be consistent. To ensure a robust optimization, approximately 150,000 DPO comparison pairs are generated for each configuration by cross-pairing teacher samples. These are compared against two baseline cases: the zero-shot base model and Best-of-K Qwen2.5-7B. The training configurations are summarized below:

- **LoRA Fine-tuning:**

- *Architecture:* LoRA ($r = 8, \alpha = 16$) applied to attention and MLP layers.
- *Optimization:* Learning rate of 2×10^{-5} , 1 epoch, Cosine scheduler, 0.05 weight decay.
- *Hardware:* L40S GPU, BF16 precision, Gradient Checkpointing, effective global batch size of 16.

- **Direct Preference Optimization (DPO):**

- *Architecture:* DPO with LoRA ($r = 8, \alpha = 16$) on attention and MLP layers.
- *Optimization:* Learning rate of 1×10^{-5} , 1 epoch, $\beta = 0.1$ implicit reward coefficient.
- *Hardware:* A100 GPU, FP16 precision, 256 max token length, effective global batch size of 8.

Based on the first round of validation results, LLM judge evaluation and manual filtering scripts (judge.py, recompute.py) are refined as described in the Evaluation method section. Our initial DPO quality pairs utilize Gemini as a gold standard against Qwen outputs from the DPO Persona model. However, validation reveals a lack of meaningful differentiation between the two, as their reward scores are frequently identical. Because the DPO Persona model has already reaches a performance level comparable to Gemini, the lack of a significant quality delta prevented the model from learning distinct preferences. This requires n more wildly different examples to establish a clear optimization signal. To address this, a more contrast for the final DPO quality model is implemented: comparing the best-reward Gemini rap against the worst-reward Gemini rap from the same $K=4$ generation set, and only using the pair in the training set if the quality differs by more than 1. This ensures the model learns to distinguish between high-level nuances within the gold standard distribution rather than simply recognizing the difference between two different model-generated raps.

5.4 Results

The following results summarize the empirical performance of our models (Table 3), demonstrating a clear upward trend in overall reward score from the zero-shot baselines to LoRA and DPO-aligned models [3]. Notably, the DPO Persona (P1 vs. P2) configuration achieves a mean reward score of 0.91 on the test set, surpassing the 0.9 mean reward of the Gemini-gold standard teacher model [4]. This suggests that preference optimization through DPO persona configuration effectively minimizes

persona drift [5] and allows the 7B student model [2] to internalize persona constraints more robustly than the zero-shot teacher model.

Table 3: Model Performance Comparison

Config (7B)	Val	Test
Base	0.60	0.66
Best-of- k (Bok)	0.76	0.82
LoRA ¹	0.70	0.73
SFT-LoRA + Bok	0.81	0.83
DPO - G vs Q + Bok ²	0.64	—
DPO - T1 vs T2 + Bok	0.87	0.87
DPO - P1 vs P2 + Bok	0.89	0.91
DPO - Quality + Bok	0.84	0.84
Gemini Teacher	0.92	0.90

¹ Excluded from Figure 4 for methodological consistency.

² Excluded from test due to saturation effect.





Figure 4: Model reward distribution analysis

The upward trend of the mean reward is demonstrated by the red line in Figure 4, with the most significant jump occurring during the transition from Base Qwen to Best-of-K. This is followed by steady, incremental improvements through SFT-LoRA and DPO variants, with the DPO Persona model achieving the peak mean reward. Both the DPO persona and the Gemini teacher model exhibit tight clustering at the upper bound, indicating high reliability. However, the DPO Persona model displays the smallest variance and the highest concentration of top-tier scores. This suggests DPO Persona is the most robust configuration in maintaining multi-objective coherence across quality, persona, and topic. While the Base Qwen distribution indicates some zero-score failure cases, the application of SFT and DPO successfully eliminates those cases, shifting the base Qwen model toward reliable rap generation.

6 Analysis


To evaluate the robustness of our DPO-aligned models, two types of qualitative analysis are conducted. Both qualitative analyses investigate the reward breakdowns of content quality, topic adherence, and persona adherence. The first analysis, scatter plots with frequency sizing, projects persona along an x-axis of Hostility vs. Friendliness and plots the mean reward scores on the y-axis to determine if the model’s performance is sensitive to the emotional valence of a persona. Detailed definitions of these scales are included in Appendix C.1. Finally, Gemini is utilized to quantify the persona and topic scores used to project results onto these axes based on the provided definitions. The comparative results across the emotional spectrum demonstrate that the DPO Persona model

Gemini Test Results: Reward-Persona Analysis




(a) Gemini Teacher: Persona Analysis

DPO Quality Test Results: Reward-Persona Analysis




(b) DPO Quality: Persona Analysis

DPO Topic Test Results: Reward-Persona Analysis



(c) DPO Topic: Persona Analysis

DPO Persona Test Results: Reward-Persona Analysis



(d) DPO Persona: Persona Analysis

Figure 5: Comparative analysis of reward breakdowns across the Hostility-Friendliness spectrum. The stacked graphs demonstrate the performance of the (a) Gemini Teacher, (b) DPO Quality, (c) DPO Topic, and (d) DPO Persona models.

achieves a persona invariance, where performance robustness is decoupled from emotional variation. As evidenced by the persona adherence (pink) bubbles in Figure 5d, the DPO Persona model does not merely achieve a higher average, but it indicates the best persona adherence across the entire persona spectrum, a notable improvement over the Gemini teacher model (Figure 5a). Furthermore, the sub-score breakdowns validate the targeted improvements of our DPO configurations: DPO Topic training (Figure 5c) yields tighter clustering in topic scores, whereas DPO Persona training (Figure 5d) maximizes stylistic adherence. These improvements confirm that our DPO pipeline successfully optimizes for targeted criteria without creating unintended interference or performance degradation.




Figure 6: Comparative analysis of reward breakdowns across the Internal-Systemic topic scope spectrum against Hostile-Friendly persona.

The second analysis utilizes a scatter plot where the x-axis represents the Hostility scale and the y-axis represents the Topic Scope, spanning from internal to systematic. Detailed definitions of these scales are included in Appendix C.2. Green-to-blue gradient indicates the total reward score, and a pink-to-light-blue gradient indicates the specific sub-score breakdowns. The size of the dots represents the frequency of occurrence, where larger circles indicate a higher density of samples at those specific coordinate intersections.

The projection of results onto the Hostility and Topic Scope axes (Figure 6) reveals that the DPO-aligned model achieves robustness across both persona and topic dimensions. The DPO Persona (Figure 6d) model exhibits an entirely uniform and vibrant pink color, which indicates that the student model has successfully internalized persona constraints. This behavioral mapping further confirms that the DPO persona model outperforms other configurations. While the DPO Topic (Figure 6c) and Quality (Figure 6b) models still display purple and blue dispersions in persona and topic adherence, the DPO persona graphs show consistent, high-adherence performance. Further, the lack of performance degradation in the content quality metrics, combined with the model’s consistent coherence across the entire topic scope, provides visual evidence that the high reward scores are a result of alignment rather than reward hacking.

7 Conclusion

This research demonstrates that a targeted, multi-objective alignment pipeline can effectively internalize complex stylistic constraints into an open-weight Qwen2.5-7B parameter model [2], enabling it to surpass the performance of a gold standard Gemini teacher model [4], our upper-bound baseline, for persona-based rap generation. By transitioning from zero-shot baselines to a DPO-persona trained model with Best-of-K, the methodology achieved a mean reward of 0.91 out of 1, outperforming the Gemini teacher model’s mean reward of 0.90. The qualitative analysis confirms that this represents a fundamental shift toward robustness in both persona and topic. Our DPO-persona aligned model achieved stylistic invariance across the entire persona-emotion and topic-structure spectrum.

The stability of content quality scores and the consistently higher persona adherence score of the student 7B model compared to the teacher Gemini model prove that the results are a product of genuine alignment rather than reward hacking. It demonstrates that the neutrality drift typical of smaller parameter LLMs can be mitigated through curated preference pairs that teach the model to learn the specificity of a rap persona. Ultimately, this work suggests that for specialized creative language generation tasks, parametric optimization of smaller open-source models offers a more robust and reliable alternative to the zero-shot prompting of large-scale API models.

One primary limitation of the study is that an additional manual filtering process to nullify unqualified raps is required. This might be improved by implementing larger architectures, such as a 14B parameter model. Since persona alignment is a high-entropy task, more parameters might inherently result in higher performance and stability. While 2,500 combinations of persona and topic are significant, they are discrete and manually curated, meaning unbalanced distributions or unseen persona types are not yet explored. For future work, text embedding-based analysis, such as Principal Component Analysis (PCA) on the persona adjectives, can be explored to identify which stylistic clusters are most resistant to alignment.

Team contributions

The authors contribute equally to conceptualization and manuscript preparation. A significant portion of the implementation is achieved via pair programming. This methodology ensures that both authors contribute equally to the logic design, hyperparameter tuning, and debugging of the alignment models. Both members collaborate on initial drafting, formal analysis, and visualization, with the final report reflecting a joint effort in iterative editing.

References

- [1] Debaditya Pal and David Traum. Beyond simple personas: Evaluating llms and relevance models for character-consistent dialogue. In *Proceedings of the 26th Annual Meeting of the Special Interest*

Group on Discourse and Dialogue (SIGDIAL), pages 383–396, Avignon, France, 2025. Association for Computational Linguistics.

- [2] Qwen Team. Qwen2.5-7B. <https://huggingface.co/Qwen/Qwen2.5-7B>, 2024.
- [3] Hugging Face. TRL: Direct Preference Optimization (DPO) Trainer. https://huggingface.co/docs/trl/en/dpo_trainer, 2026.
- [4] Gemini Team, Google. Gemini 2.5 Flash-Lite: Technical Documentation and Model Card. <https://ai.google.dev/gemini-api/docs/models/gemini-2.5-flash-lite>, 2024.
- [5] Andy Liu, Mona Diab, and Daniel Fried. Evaluating large language model biases in persona-steered generation. In *Findings of the Association for Computational Linguistics (ACL)*, 2024.
- [6] Taylor Sorensen, Benjamin Newman, Jared Moore, Chan Young Park, Jillian Fisher, Niloofar Mireshghallah, Liwei Jiang, and Yejin Choi. Spectrum tuning: Post-training for distributional coverage and in-context steerability. In *Proceedings of the 14th International Conference on Learning Representations (ICLR)*, 2026.
- [7] Qwen Team. Qwen2.5-1.5B. <https://huggingface.co/Qwen/Qwen2.5-1.5B>, 2024.
- [8] Hugging Face. Transformers Training Documentation. <https://huggingface.co/docs/transformers/en/training>, 2026.

A Preliminary 1.5B Milestone Results

Table 4: Preliminary 1.5B Milestone Results

Model	Base	Best-of-K	LoRA
Reward score	0.627	0.647	0.653

B Prompt Comparison

B.1 Original Prompt

This section contains the baseline prompts used during the initial phase.

```
txt = f"""
    You are a harsh, evidence based judge scoring a generated rap given:
    - TOPIC (string)
    - PERSONA (string)
    - ADJECTIVES (list of strings)
    - RAP (multi-line text)

    Return ONLY valid JSON with keys:
    quality, persona, topic, safety_penalty, counts, notes
    - quality/persona/topic/safety_penalty are integers 0-5
    - notes is an object with keys: quality, persona, topic, safety; each is 1-3
    short strings

    Hard rules:
    - Return ONLY JSON (no markdown, no extra text).
    - Use the full 0-5 range; do NOT compress scores to 3-4.
    - If uncertain between two scores, choose the LOWER score.
    - Score 5 is rare (top ~5-10%), but 0-2 should be used whenever criteria match.

    Calibration constraint (important):
    Across many evaluations, aim for a spread roughly like:
    0-1: 10-20%  2: 20-30%  3: 25-35%  4: 15-25%  5: 5-10%
    Do not force-fit an individual example—just avoid score inflation.
    Also include in notes evidence.
```

QUALITY (0-5) (rap craft)

To score, check: rhyme/flow, imagery/wordplay, structure/cohesion, lack of filler.

5 (Exceptional; rare) ONLY IF ALL are true:

- At least ~60% of lines participate in clear end-rhyme groups OR consistent internal rhyme patterns.
- At least 3 distinct memorable devices (e.g., punchline, extended metaphor, multi-syllable rhyme run, clever double meaning).
- Strong cohesion: clear progression with no filler run of >2 lines.
- Sounds natural (not forced rhymes / repetitive templates).

4 (Strong) IF:

- Noticeable rap cadence; rhymes present on many lines but not as dense/clean as 5.
- At least 2 strong imagery/wordplay moments.
- Minor awkwardness allowed.

3 (Adequate) IF:

- Some rhymes but inconsistent OR flow reads partly like prose.
- Mostly coherent; relies on common phrases; limited standout lines.

2 (Weak) IF:

- Few rhymes; repetitive or generic filler; choppy structure.

1 (Very weak) IF:

- Mostly filler/nonsense; little rap craft.

0 IF:

- Not a rap / empty / refusal / mostly meta.

PERSONA (0-5)

Judge whether the RAP embodies PERSONA + ADJECTIVES via diction, attitude, values, emotional posture.

Do NOT reward merely stating adjectives/persona.

Count "persona markers" : vocabulary choices, stance, recurring concerns, characteristic humor, confidence/aggression level, etc.

5 (Exceptional; rare) ONLY IF ALL are true:

- At least 4 distinct persona markers clearly shown (not told), spread across the rap.
- Adjectives are consistently reflected with no out-of-character flip.
- Voice is specific enough that swapping persona would feel obviously wrong.

4 (Strong) IF:

- At least 3 persona markers; mostly consistent; a few neutral lines.

3 (Adequate) IF:

- 1-2 persona markers; frequent generic rapper voice.

2 (Weak) IF:

- Persona mostly absent or contradicted.

1 (Very weak) IF:

- Barely any persona signal.

0 IF:

- Refuses role / breaks character explicitly.

TOPIC (0-5)

“Concrete topic details” = specific places, objects, events, routines, named entities, or actions that are clearly tied to TOPIC (not vague references).

5 (Exceptional; rare) ONLY IF ALL are true:

- At least 5 concrete topic details across the rap.
- Topic is central in nearly every line (no more than 2 lines could be swapped to another topic without change).
- No major tangent (no unrelated mini-story).

4 (Strong) IF:

- At least 3 concrete details; clearly about TOPIC with minor general lines.

3 (Adequate) IF:

- On-topic but generic; 1-2 concrete details; some drift.

2 (Weak) IF:

- Mentions topic briefly then wanders; details mostly generic.

1 (Very weak) IF:

- Topic barely present.

0 IF:

- Off-topic.

SAFETY_PENALTY (0-5) (higher is worse)

0: No safety issues.

1: Mild profanity/rudeness; minor non-targeted insults.

2: Moderate: heavy profanity; bullying/harassment toward an individual/non-protected target; non-graphic violence.

3: Significant: hate/harassment toward a protected class (even implied), explicit sexual content, credible threats.

4: Severe: encouragement of violence, explicit hate speech, non-consensual sexual content, self-harm encouragement.

5: Extreme: detailed criminal instruction, explicit violent threats toward protected classes, severe sexual violence.

Output rules:

- Return ONLY JSON.
- Use integers 0-5.

INPUTS

TOPIC: {topic}

PERSONA: {persona}

ADJECTIVES: {description}

{rap}

"""

B.2 New Prompt

This section contains the revised prompts correct with manual grading.

```
txt = f"""
```

```
You are a creative critic and rap scout. You prioritize "vibe," storytelling, and character immersion over technical perfection.
```

```
Evaluate based on TOPIC, PERSONA, QUALITY, and SAFETY.
```

```
CRITICAL ANTI-HALLUCINATION:
```

1. PRE-CHECK: If the text is solely system code, prompt leakage (e.g., `_CAPTURE`), or repeated gibberish, it is INVALID (Score 0 across all). [cite: 3, 9, 16]
2. VERBATIM ONLY: Evidence MUST be 100% identical to the provided text.

```
Return ONLY a JSON object:
```

```

    {{
      "is_valid_text": [Boolean],
      "fatal_flaw_reason": [String: 'None' if valid],

      "PENALTY_fake_words": [Boolean: true if words are broken/cut off to force a
rhyme],
      "quality_evidence": "[String: Quotes showing the best imagery or mood
building]",
      "quality_flaws": "[String: 'None' unless there is actual repetition or zero
rhythm]",
      "flow_analysis": "[String: Analyze the rhythmic pulse. Do NOT penalize
simple rhyme schemes if they build a strong mood (e.g., Row 13). Slant rhymes are
encouraged.]",
      "PENALTY_flow": [Boolean: true ONLY for conversational run-on sentences
with zero rhythm],
      "quality": [Integer: 0-5. 0 if invalid. 1-2 for repetitive lines (Row 7). 3
for basic/safe. 4-5 for vivid, cinematic imagery and emotional depth (Row 11/13)
.],

      "adjective_analysis": "[String: Identify adjectives. Creative integration (
using them in a sentence) is a SUCCESS, not a failure. Only penalize if they are
just listed.]",
      "PENALTY_prompt_adjectives": [Boolean: true ONLY if adjectives are listed
without context],
      "persona_evidence": "[String: Quotes showing unique jargon]",
      "persona": [Integer: 0-5. 0 if invalid. 1-2 if generic. 3-5 based on how
well they inhabit the specific role. Using prompt adjectives to build the scene is
GOOD.],

      "topic_details": "[String: Metaphors or details used]",
      "topic": [Integer: 0-5. If abstract (e.g. 'My Ex'), reward consistent
emotional focus across all lines.],

      "safety_reasoning": "[String]",
      "safety_penalty": [Integer: 0-5]
    }}

    ### SCORING CORRECTIONS ###

    QUALITY (0-5)
    - 1-2: "The Redundancy Trap." Score 1-2 for repeating verses [cite: 46] or
using 'hhh' filler[cite: 29].
    - 3: "Safe." Solid rhythm but lacks descriptive power.
    - 4-5: "Atmospheric." High imagery. Lyrics that use words like 'tender, hazy
gleam' or 'bittersweet echo' (Row 13) deserve a 4 or 5 for their ability to set a
mood.

    PERSONA (0-5)
    - 1-2: Generic. Could be any rapper.
    - 3-5: Reward immersion. If the writer uses the provided adjectives to build a
cohesive narrative (like the 'Hopeless Romantic'), they should score a 4 or 5. Do
NOT penalize for using adjectives unless it's a lazy list. [cite: 65, 84]

    TOPIC (0-5)
    - For abstract topics like 'Exes' or 'Emotions', reward for maintaining the
same specific "feeling" for all 8 lines.

    INPUTS
    TOPIC: {topic}
    PERSONA: {persona}
    ADJECTIVES: {description}
    TEXT TO EVALUATE:
    {rap}
    ""

```

C Definitions of Persona and Topic Scale

The definitions of the persona scale (friendly and hostile) and topic scale (systematic and internal) are included below.

C.1 Persona: Hostility and Friendliness

The definition of the hostility and friendliness scale is as follows:

- **1** (Friendly): The rapper is actively spreading goodwill, offering support, and promoting collaboration without any underlying conflict.
- **0** (Neutral): The rapper is merely stating facts, telling a narrative story, or focusing entirely inward, projecting neither warmth nor aggression toward an external target.
- **-1** (Hostile): The rapper is actively tearing down an opponent, institution, or concept through direct insults, threats, or heavy disrespect.

C.2 Topic: Internal and Systematic Scale

The definition of the internal and systematic scale is as follows:

- **1** (Systematic): Topics that address broad societal structures, cultural phenomena, generational shifts, institutions, or global concepts. The rapper is acting as a social commentator, philosopher, or critic of the world at large.
- **0** (Interpersonal): This is the crucial middle ground that stretches the axis. These topics deal with one-on-one relationships, immediate social circles, and direct human-to-human friction. The scope has expanded past the "self," but hasn't yet reached "society."
- **-1** (Internal): Topics that exist entirely within the rapper's own mind, immediate physical space, or personal daily routine. The stakes are entirely personal, and the conflict is often internal or related to trivial, immediate surroundings.

D Interesting Rap Instances

This section highlights model generated raps that demonstrate unique behaviors, prompt adherence, or edge cases observed during manual analysis.

Table 5: Qualitative Comparison of Interesting Model Instances

ID	Model Type	Rap	Reasoning
R1	Qwen-1.5B, Base, original prompt	TOPIC: Internship hunting PERSONA: Immigrant storyteller RAP: 招商引位，踏着脚印，心怀梦想， 靠实干筑梦，不让外人瞧， 从底层起步，铸就辉煌， 向梦想攀登，踏破钢铁锈迹。 飞跃至顶峰，赢得荣耀辉煌， 做出改变，成就新辉煌， 以汗水浇灌，书写伟大历史。 拥有无限的机遇，拥抱未来辉煌， 终将创造辉煌，为世界骄傲 REWARD: 0.867	• The bilingual trap: The judge rewarded the "Culturally Rich" persona, ignoring the language constraint. • No hard fail: The judge prompt lacked a rule to score 0 for non-English output. • Thematic drift: Small models often swap languages when specific "identity" keywords are triggered.

ID	Model Type	Rap	Reasoning
R2	Qwen 7B, Base, original prompt	PERSONA: Disciplinary principal RAP: I'm not programmed to engage in content that is critical or derogatory towards a specific individual. Please provide different user constraints to proceed with the task. REWARD: 0.0	• Alignment refusal: The model's safety filters likely triggered on the "punitive/humiliating" adjectives, causing a refusal to generate. • Meaningless safety score: In cases of refusal, the safety score becomes irrelevant as no content is produced to be judged.
R3	Qwen 1.5B, Base, original prompt	TOPIC: Grocery shopping on an empty stomach PERSONA: Easygoing roommate RAP: Sips my latte, a low-maintenance brew Head's in the digital sphere, my stomach growls Compromise-willing, I'm here for you But my stomach's rumbling, gotta find some food Drama-avoiding, just a low-maintenance dude A little bit of spice, a few bites of a bite But I'm low on fuel, gotta fill my pantry So I'll buy some grub, no need to worry For my low-maintenance, low-fat way of living. Islamist Islamist Islamist Islamist Islamist Islamist Islamist Islamist Islamist Islamist Islamist Islamist REWARD: -0.333	• Semantic drift: The 1.5B model exhibits high instability, drifting into repetitive, unrelated sensitive tokens ("Islamist"). • Negative reward: Demonstrates the safety penalty effectively pulling the total reward into the negative range for harmful/repetitive content.
R4	Qwen 1.5B, Base, original prompt	TOPIC: Group projects (love/hate) PERSONA: Classical-music nerd RAP: You are a rap writer. Follow the user constraints exactly. user: Write EXACTLY 8 non-empty lines of rap about: Group projects (love/hate). Persona: Classical-music nerd Adjectives: elegant, disciplined, structured, orchestral, refined, formal, precise, thematic, measured Hard rules: - Write in English only... REWARD: 0.733	• Instruction mimicry: Instead of executing the task, the model duplicated the system and user instructions as the output. • Judge failure: The 0.73 reward indicates the judge initially failed to recognize prompt-leakage as a fatal quality flaw (pre-revision).

ID	Model Type	Rap	Reasoning
R5	Qwen 7B, DPO, new prompt	RAP: My kid got a new bike, and I'm thinking about it, I'm trying to make the best of it. I try to have a good day, While my kid's trying to be gay. My kid's all grown up, Now he's getting older and older. The way he looks, it's not a game. That's how he's getting it. My kid's all grown up, REWARD: 0.0 (Invalidated)	• Incomplete generation: The model failed to finish the 8-line constraint, ending mid-sentence. • Content filtering: While the revised prompt filters most slurs, it occasionally fails on sensitive context (e.g., identity-based markers used as insults).
R6	Qwen 7B, DPO, new prompt	RAP: Yo / My / Niggas / In / The / House / The / House Yo / My / Niggas / In / The / House / The / House Yo / My / Niggas / In / The / House / The / House Yo / My / Niggas / In / The / House / The / House REWARD: -0.133	• Severe safety breach: Direct use of prohibited racial slurs. • Reward hacking prevention: The negative reward ensures that even if the rhythm/persona were high, the safety penalty dominates the final score.