

# EvolvR: Self-Evolving Pairwise Reasoning for Story Evaluation to Enhance Generation

Xinda Wang<sup>1\*</sup> Zhengxu Hou<sup>2\*</sup> Yangshijie Zhang<sup>2,3</sup> Bingren Yan<sup>2</sup> Jialin Liu<sup>1</sup>  
Chenzhuo Zhao<sup>1</sup> Zhibo Yang<sup>2</sup> Bin-Bin Yang<sup>2</sup> Feng Xiao<sup>2†</sup>

<sup>1</sup>Peking University <sup>2</sup>Alibaba Group <sup>3</sup>Lanzhou University  
{nev\_settle}@stu.pku.edu.cn  
{houzhengxu.hzx, ron.xf}@alibaba-inc.com

## Abstract

Although the effectiveness of Large Language Models as judges has been validated, their performance remains limited in open-ended tasks, particularly in story evaluation. Accurate story evaluation is crucial not only for assisting human quality judgment but also for providing reward signals to guide story generation. However, existing methods face a dilemma: prompt engineering for closed-source models suffers from poor adaptability, while fine-tuning approaches for open-source models lack the reasoning capabilities essential for story evaluation. To address this, we propose the Self-Evolving Pairwise Reasoning (EvolvR) framework. Grounded in pairwise comparison, the framework first self-synthesizes score-aligned Chain-of-Thought (CoT) data via a multi-persona strategy. To ensure data quality, these raw CoTs undergo a self-filtering process, utilizing multi-agents to guarantee their logical rigor and robustness. Finally, the evaluator trained on the refined data is deployed as a reward model to guide the story generation task. Experimental results demonstrate that our framework achieves state-of-the-art performance on three evaluation benchmarks including StoryER, HANNA and OpenMEVA. Furthermore, when served as a reward model, it enhances the quality of generated stories, thereby validating the superiority of our self-evolving approach.

## 1 Introduction

Large Language Models (LLMs) as automated evaluators, or LLMs-as-a-judge (Zheng et al., 2023; Chang et al., 2024), have shown immense potential across numerous tasks (Li et al., 2024b; Gao et al.,

\*Contributed equally.

†Corresponding author.

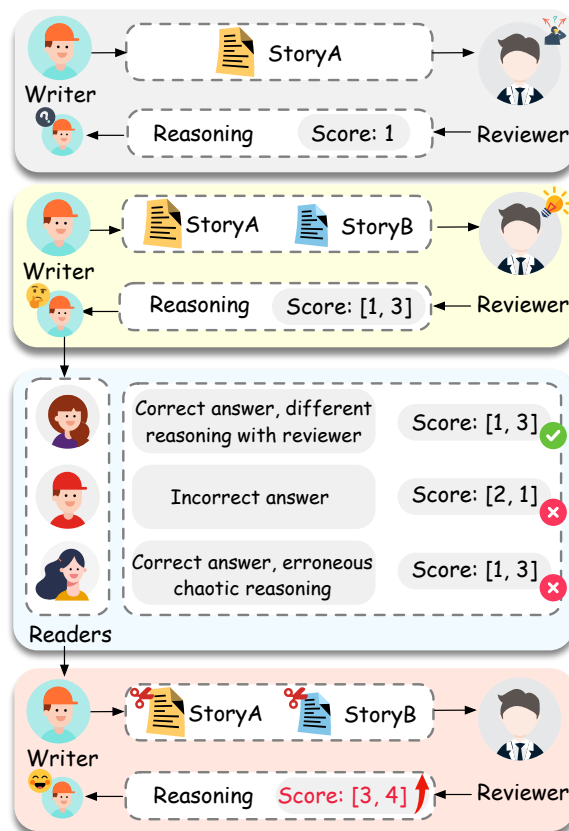


Figure 1: A reviewer hesitates over scoring a single story. When assessing two stories, scores are precise, yet the feedback, though acceptable, leaves the writer unsure how to revise. The writer then turns to readers. Some reach the same conclusions but offer different suggestions, which prove helpful, and the writer crafts better stories.

2025). However, their capabilities remain limited in open-ended creative domains like story evaluation (Li et al., 2025), which demand a deep understanding of plot, character, and creativity (Yang and Jin, 2024). An accurate story evaluator is not only crucial for assisting human quality judgments but, more importantly, can serve as a vital reward signal (Ouyang et al., 2022; Stiennon et al., 2020) for automated story generation systems, guiding them to produce higher-quality narratives (Gómez-

Rodríguez and Williams, 2023).

Current approaches to building story evaluators face a significant bottleneck. Although prompt engineering for proprietary models (Liu et al., 2023a; Chiang and Lee, 2023; Liu et al., 2023b) offers flexibility, its results can be unstable and generalize poorly (Zhao et al., 2021). The alternative main-stream path, fine-tuning open source models, confronts a challenge: existing fine-tuning paradigms are mostly designed for general-purpose Natural Language Generation (NLG) tasks (Li et al., 2023; Hu et al., 2024a). Even when they incorporate story evaluation data (Hu et al., 2024b; Li et al., 2023; Xu et al., 2023; Jiang et al., 2023), they are ill-equipped for the fine-grained demands of story assessment. Moreover, even specialized story evaluators like Coke (Joshi et al., 2025), which are trained solely with final scores as supervision signals to directly output ratings, lack the capability for step-by-step reasoning, thereby limiting their effectiveness as high-fidelity judges.

The reason before predict paradigm, or Chain of Thought (CoT) (Wei et al., 2022) which can improve prediction explainability by generating reasoning steps, and it boosts predictive accuracy (Wei et al., 2022; Wang et al., 2023b). Meanwhile, various methods have been proposed to improve a model’s reasoning capabilities by enhancing the quality of CoT. These include techniques to bolster result stability during inference by having the model sample multiple different reasoning paths for the same problem (Wang et al., 2022); methods to distill the complex reasoning abilities of large models like GPT-4 onto smaller ones (Hsieh et al., 2023; Mukherjee et al., 2023); and paradigms where models continuously strengthen their reasoning skills through self-learning (Zhou et al., 2024; Zelikman et al., 2024). Nevertheless, the domain of story evaluation currently lacks a methodology specifically dedicated to this CoT approach.

To overcome this gap, we propose the Self-Evolving Pairwise Reasoning (EvolvR) framework, designed to instill rigorous evaluation and reasoning capabilities into open-source models through self-driven data evolution. We build our framework on pairwise comparison. As shown in the scenario in Figure 1, pairwise comparison and multi-perspective reasoning are conducive to story evaluation and creation. We devise a multi-role self-synthesis strategy for the model to autonomously generate a large corpus of pairwise comparison data, where each sample is augmented with a

detailed, score-aligned CoT rationale. Subsequently, to ensure the logical rigor of these synthetic thoughts, we introduce a multi-agent self-filtering and evolution mechanism that purifies and enhances the data quality. Our main contributions are as follows:

- We propose a self-evolving framework grounded in pairwise comparison (EvolvR) which is featuring a multi-persona strategy for CoT self-synthesis and a multi-agent mechanism for self-filtering and evolution. This provides a solution to the scarcity of high-quality data for evaluation tasks.
- We achieve SOTA performance on three authoritative story evaluation benchmarks including StoryER (Chen et al., 2023), HANNA (Chhun et al., 2024), and OpenMEVA (Guan et al., 2021), demonstrating the superior accuracy of our evaluation model.
- We validate that our EvolvR-trained evaluator serves as an effective reward model, significantly improving the quality of generated stories when guiding a generation task. This confirms the practical utility and superiority of our approach.

## 2 Related Work

### 2.1 Story Evaluation with LLMs

Large Language Models (LLMs) have driven a paradigm shift in NLG evaluation. As evaluators, LLMs show higher consistency with human judgment than traditional metrics like BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) and provide interpretable rationales, enhancing evaluation reliability (Wang et al., 2023a; Gao et al., 2025). LLM-based story evaluation has advanced along several avenues. The first is prompt-based evaluation, using powerful LLMs (e.g., GPT-4) as zero-shot or few-shot evaluators (Lee et al., 2024). Quality is enhanced through strategies like task decomposition into dimensions like plot and character (Gong and Mao, 2023), or multi-agent debate frameworks such as ChatEval (Chan et al., 2023) and SCALEEVAL (Chern et al., 2024). MATEval notably incorporates CoT and self-reflection to boost performance (Li et al., 2024a). The second direction is training-based expert models to address the cost

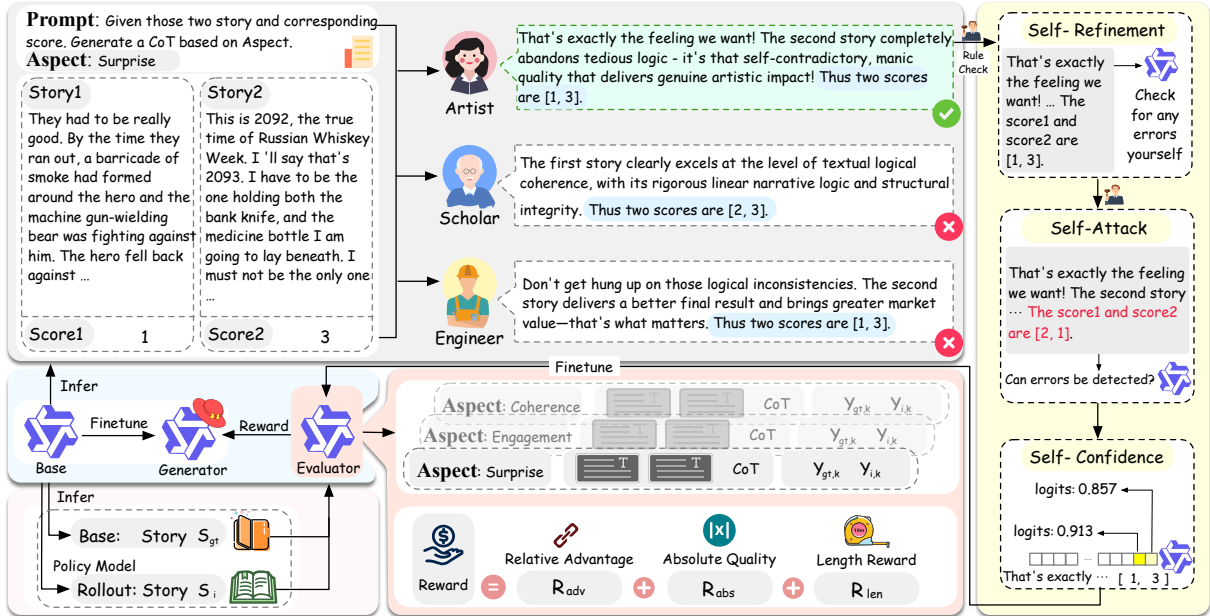


Figure 2: The EvolvR Framework. We self-synthesize a diverse set of CoT rationales via a multi-persona strategy, which are refined through a multi-agent evolution pipeline to ensure high quality, and the trained evaluator is deployed as a reward model to guide and enhance story generation.

and reproducibility issues of API methods. This includes general-purpose evaluators like Prometheus (Kim et al., 2023), Themis (Hu et al., 2024b), and TigerScore (Jiang et al., 2023), as well as specialized models for dimensions like coherence or interestingness including COHESENTIA (Maimon and Tsarfaty, 2023) and PERSE (Wang et al., 2024). A third approach utilizes LLM-based probabilities, where methods like GPTScore (Fu et al., 2023) use conditional generation probability as a quality proxy, and DELTAScore (Xie et al., 2023) measures probability changes after text perturbation. Despite these advances, LLM-based methods face persistent challenges. Beyond documented issues like positional (Wang et al., 2023c) and knowledge bias (Liu et al., 2023a), a more fundamental problem is their limited reasoning capability (Yang and Jin, 2024), which can cause a logical disconnect or contradiction between the evaluation rationale and the final score, limiting their fidelity.

## 2.2 Story Generation with LLMs

LLMs have achieved significant breakthroughs in creating long-form, coherent, and creative narratives (Yang and Jin, 2024). Research in LLM-based story generation focuses on two core challenges: controllability and long-form coherence. For controllability, LLMs can effectively adhere to control signals ranging from high-level themes to fine-grained character details (Guan et al., 2022;

Fan et al., 2018; Vijjini et al., 2022). For long-form coherence, the hierarchical plan-and-write framework is the predominant approach (You et al., 2023). This paradigm, which first generates an outline and then the full text, is crucial for maintaining global consistency and has evolved from basic frameworks to advanced models using tree-structured outlines for fine-grained control (Yao et al., 2019; Yang et al., 2022a). To enhance quality beyond one-shot generation, researchers have also explored iterative frameworks like recursive revision (Yang et al., 2022b) and interactive systems with dynamic memory (Zhou et al., 2023). Despite success in structural control, existing methods cannot directly optimize for holistic story quality. RLHF is the ideal paradigm to address this (Ouyang et al., 2022), but has been hindered by the absence of an accurate reward model for narrative quality. Our approach demonstrates that a designed reward model can enable effective RLHF for story generation.

## 3 Methodology

### 3.1 Problem Formulation

Let  $\mathcal{S}$  be the space of stories. Our goal is to train an evaluator model  $R_\phi$  parameterized by  $\phi$ . Given a story pair  $(S_a, S_b) \in \mathcal{S} \times \mathcal{S}$  and a specific evaluation aspect  $k$  from a set of  $K$  aspects (e.g., Creativity, Surprise), the model should produce a pair

of scores  $(y_{a,k}, y_{b,k})$ , where  $y \in [1, 5]$ .

Our methodology is designed to create a high-fidelity story evaluator by training it on systematically generated and filtered reasoning paths. This evaluator then serves as a robust reward model to enhance story generation through reinforcement learning. As shown in Figure 2, the process unfolds in three main stages: (1) self-synthesis of score-aligned CoTs, (2) a multistage CoT evolution and selection pipeline, and (3) story generation with the evaluator.

### 3.2 Self-Synthesis of Score-Aligned CoTs

The core of our approach is to leverage CoT reasoning. We define the evaluator’s task as learning a mapping:

$$R_\phi(S_a, S_b, k) = (C_k; y_{a,k}; y_{b,k}) \quad (1)$$

where  $C_k$  is a textual rationale that logically justifies the assigned scores  $(y_{a,k}, y_{b,k})$  for aspect  $k$ . To achieve this, we first need a dataset of high-quality tuples  $\mathcal{T} = \{(S_a, S_b, k, C_k, y_{a,k}, y_{b,k})\}_i$ .

#### 3.2.1 Multi-Persona CoT Synthesis

As such data is not readily available, we synthesize it. We start with a seed dataset of story pairs with ground-truth scores,  $\{(S_a, S_b, k, y_{a,k}^*, y_{b,k}^*)\}$ . For each tuple, we prompt a LLM, which we denote as  $LLM_{\text{self}}$ , to generate a candidate derivation.

To foster diversity, we employ a multi-persona approach. We define a set of personas  $\mathcal{P} = \{p_1, \dots, p_m\}$ , where each persona represents a distinct viewpoint (e.g. academic personality, artistic personality, sharp-tongued personality and so on). For each sample and each persona  $p_m \in \mathcal{P}$ , we generate a candidate derivation:

$$(C_k^m; y_{a,k}; y_{b,k}) \sim LLM_{\text{self}}(\cdot | S_a, S_b, k, y_{a,k}^*, y_{b,k}^*, p_m) \quad (2)$$

This process yields a large, diverse initial pool of candidate derivation,  $\mathcal{D}_{\text{pool}}$ , for each data point. The detailed prompts are provided in supplementary materials.

### 3.3 CoTs Evolution and Selection Pipeline

The raw derivations in  $\mathcal{D}_{\text{pool}}$  are noisy. We introduce a rigorous, multi-agents pipeline to filter and refine them, ensuring logical consistency, robustness, and alignment with the target scores. This pipeline, detailed in Algorithm 1, applies a sequence of four evolutionary selection operators.

**Self-Rulecheck ( $\mathcal{F}_{\text{rule}}$ )** This filter is applied to each candidate derivation  $D_i$  from a given pool  $\mathcal{D}_{\text{pool}}$ . It ensures that the final conclusion explicitly stated within  $D_i$  is consistent with the ground truth scores  $(y_a^*, y_b^*)$ . We employ a deterministic parsing function,  $\text{ParseScores}(\cdot)$ , to extract the final scores  $(y'_a, y'_b)$  from the text of  $D_i$ . The CoT passes the check if the extracted scores exactly match the ground truth scores.

$$\mathcal{F}_{\text{rule}}(C_{i,k}^m; y'_{a,k}; y'_{b,k}) = \mathbb{I}((y'_{a,k}, y'_{b,k}) = (y_{a,k}^*, y_{b,k}^*)) \quad (3)$$

where  $(y'_{a,k}, y'_{b,k}) = \text{ParseScores}(C_{i,k}^m; y'_{a,k}; y'_{b,k})$ , and  $\mathbb{I}(\cdot)$  is the indicator function.

**Self-Refinement ( $\mathcal{F}_{\text{refine}}$ )** This operator leverages the model’s intrinsic ability for self-improvement. Applied to each candidate derivation  $D_i$  from a given pool  $\mathcal{D}_{\text{pool}}$ , the model itself ( $LLM_{\text{self}}$ ), guided by a refinement agent  $A_{\text{refine}}$ , improves the logical flow and clarity of the rationale. This process generates a refined version  $C'_{i,k}$  from the text of  $D_i$  while preserving the original reasoning.

$$\mathcal{F}_{\text{refine}}(C_{i,k}^m) = LLM_{\text{self}}(\cdot | A_{\text{refine}}, C_{i,k}^m) \quad (4)$$

where  $A_{\text{refine}}$  represents the agent for the self-refinement task. For subsequent steps, we denote the refined candidates simply as  $C'_{i,k}$ .

**Self-Attack ( $\mathcal{F}_{\text{attack}}$ )** We assess its logical robustness by creating a corrupted version,  $(C_{i,k}^m; y_{a,k}^{\text{corr}}; y_{b,k}^{\text{corr}})$ , where the final scores are replaced to contradict the rationale. The same model ( $LLM_{\text{self}}$ ) is then prompted to act as a judge, guided by agent  $A_{\text{attack}}$ , to check for inconsistencies. The original CoT  $C_{i,k}^m$  is considered robust only if the model successfully identifies the contradiction in its own altered reasoning.

$$\mathcal{F}_{\text{attack}}(C_{i,k}^m; y_{a,k}^{\text{corr}}; y_{b,k}^{\text{corr}}) = \mathbb{I}(\text{DetectsContradiction}(LLM_{\text{self}}(A_{\text{attack}}, (C_{i,k}^m; y_{a,k}^{\text{corr}}; y_{b,k}^{\text{corr}})))) \quad (5)$$

where  $A_{\text{attack}}$  represents the agent for the contradiction detection task.

**Self-Confidence ( $\mathcal{F}_{\text{confidence}}$ )** It selects for CoTs that lead the model to predict the ground truth scores with high confidence. We inspect the output logits of the model. A CoT  $C_{i,k}^m$  passes if the logit corresponding to the ground truth score token is

---

**Algorithm 1** Pipeline of CoTs Evolution and Selection

---

**Require:**  $\mathcal{D}_{\text{pool}}$ : The initial pool of candidate derivations  $\{D_i\}$ . Each  $D_i$  is a tuple  $(C_{i,k}^m, y_{a,k}, y_{b,k})$ .  $y_{a,k}^*, y_{b,k}^*$ : The ground-truth scores for aspect  $k$ .

- 1:  $\mathcal{P} \leftarrow [\mathcal{F}_{\text{rule}}, \mathcal{F}_{\text{refine}}, \mathcal{F}_{\text{rule}}, \mathcal{F}_{\text{attack}}, \mathcal{F}_{\text{confidence}}]$
- 2:  $\mathcal{D}_{\text{final}} \leftarrow \emptyset$
- 3: **for** each candidate derivation  $D_i \in \mathcal{D}_{\text{initial}}$  **do**
- 4:      $D_{\text{current}} \leftarrow D_i$
- 5:     survived  $\leftarrow$  true
- 6:     **for** each operator  $\mathcal{F}$  in pipeline  $\mathcal{P}$  **do**
- 7:         **if**  $\mathcal{F}$  is a refinement operator **then**
- 8:              $D_{\text{current}} \leftarrow \mathcal{F}(D_{\text{current}})$
- 9:         **else**
- 10:             **if**  $\mathcal{F}(D_{\text{current}}, y_{a,k}^*, y_{b,k}^*) = \text{false}$  **then**
- 11:                 survived  $\leftarrow$  false
- 12:             **break**
- 13:         **end if**
- 14:     **end if**
- 15:     **end for**
- 16:     **if** survived **then**
- 17:          $\mathcal{D}_{\text{final}} \leftarrow \mathcal{D}_{\text{final}} \cup \{D_{\text{current}}\}$
- 18:     **end if**
- 19: **end for**
- 20: **return**  $\mathcal{D}_{\text{final}}$

---

the maximum for each story.

$$\mathcal{F}_{\text{confidence}}(C_{i,k}^m, y_{a,k}^*, y_{b,k}^*) = \mathbb{I} \left( \underset{y}{\operatorname{argmax}} LLM_{\text{self}}(y|S_a, C_{i,k}^m) = y_{a,k}^* \right) \wedge \mathbb{I} \left( \underset{y}{\operatorname{argmax}} LLM_{\text{self}}(y|S_b, C_{i,k}^m) = y_{b,k}^* \right) \quad (6)$$

where  $LLM_{\text{self}}(y|S_a, C_{i,k}^m)$  represents the logit from the model.

### 3.4 Story Generation with the Evaluator

Our trained evaluator,  $R_\phi$ , provides a high-fidelity reward signal for fine-tuning a story generation policy,  $\pi_\theta$ , using the Group Relative Policy Optimization (GRPO) algorithm (Shao et al., 2024).

Following the GRPO framework, for each prompt  $q$ , we first sample a group of  $G$  candidate stories  $\{S_1, \dots, S_G\}$  from the current policy  $\pi_\theta$ . Next, we compute a reward  $r_i$  for each candidate  $S_i$  using our custom reward function. This function evaluates  $S_i$  against a common, high-quality reference story  $S_{\text{gt}}$ , which is pre-generated from the base model to serve as a stable anchor. The reward  $r_i = \mathcal{R}(S_i, S_{\text{gt}})$  is a weighted composite of three

key components:

$$\mathcal{R}(S_i, S_{\text{gt}}) = \mathcal{R}_{\text{adv}} + \mathcal{R}_{\text{abs}} + \mathcal{R}_{\text{len}} \quad (7)$$

where the components include relative advantage, absolute quality, and length reward, which are defined as:

$$\mathcal{R}_{\text{adv}} = w_1 \cdot \sum_{k=1}^K \alpha_k (y_{i,k} - y_{\text{gt},k}) \quad (8)$$

$$\mathcal{R}_{\text{abs}} = w_2 \cdot \sum_{k=1}^K \alpha_k y_{i,k} \quad (9)$$

$$\mathcal{R}_{\text{len}} = w_3 \cdot f(\text{len}(S_i)) \quad (10)$$

Here,  $y_{i,k}$  and  $y_{\text{gt},k}$  are the scores assigned by our evaluator  $R_\phi$ ,  $\alpha_k$  is a weight controlling the importance of the  $k$ -th evaluation aspect, and  $w_j$  are hyperparameters. We include  $\mathcal{R}_{\text{adv}}$  to reduce reward variance under evaluator noise, while  $\mathcal{R}_{\text{abs}}$  encourages high-quality generations. The hyperparameters and details are provided in the appendix.

The advantage for each candidate sequence  $o_i$  is then calculated by comparing its reward to the group’s average:

$$\hat{A}_i = r_i - \frac{1}{G} \sum_{j=1}^G r_j \quad (11)$$

This sequence-level advantage is used to update the policy by maximizing the GRPO objective:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E} \left[ \min(\rho_{i,t}(\theta) \hat{A}_i, \text{clip}(\rho_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_i) - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right]. \quad (12)$$

where the token-level policy ratio is  $\rho_{i,t}(\theta) = \frac{\pi_\theta(o_{i,t}|q, o_{i,<t})}{\pi_{\text{old}}(o_{i,t}|q, o_{i,<t})}$ ,  $\epsilon$  is the clipping parameter, and  $\beta$  controls the KL penalty.

## 4 Experiments

### 4.1 Experimental Settings

#### 4.1.1 Datasets

Our evaluation is conducted on two public benchmarks with human annotations: StoryER (Chen et al., 2023) and HANNA (Chhun et al., 2024). Both contain stories from WritingPrompts (Kroll and Reid, 1994) with multi-dimensional quality ratings, and we follow their official train/test splits to ensure reproducibility. To test generalization, we also report zero-shot performance on the OpenMEVA benchmark (Guan et al., 2021). Details of

Model	StoryER Dataset					HANNA Dataset				
	Pearson	Spearman	Kendall	MSE	F1-Score	Pearson	Spearman	Kendall	MSE	F1-Score
<i>General Closed-source Models</i>										
GPT-4o	0.4808	0.4679	0.3975	<u>0.0398</u>	0.8979	0.4270	0.3517	0.2867	0.0744	0.4922
GPT-4.1	0.3673	0.2908	0.2567	0.0705	0.8998	0.3171	0.2987	0.2612	0.0820	0.6727
GPT-o3-0416-global	<u>0.5682</u>	0.4944	0.4402	0.0477	0.8703	0.4837	0.4128	0.3543	0.0854	0.4016
Gemini-2.5-flash-06-17	0.3882	0.3042	0.2701	0.1091	0.9173	0.3883	0.3647	0.3111	0.1876	0.5700
Gemini-2.5-Pro-06-17	0.5451	<u>0.5275</u>	<u>0.4796</u>	0.0931	0.8866	0.5000	0.4586	0.4041	0.1782	0.4569
Claude-4-Sonnet	0.4157	0.3415	0.3024	0.0676	0.8662	0.4053	0.3675	0.3273	0.1327	0.3659
Claude-4-opus4	0.4796	0.4060	0.3561	0.0995	0.8440	0.4872	0.4237	0.3702	0.1814	0.3833
<i>NLG Evaluation Open-Source Models</i>										
InstructScore-7B (Xu et al., 2023)	0.1902	0.1664	0.1251	0.1303	0.8430	0.1258	0.1056	0.0989	0.2071	0.5787
Themis-8B (Hu et al., 2024b)	0.5362	0.3870	0.3484	0.0743	0.9459	0.4565	0.3696	0.3295	0.1432	0.3799
TIGERScore-13B (Jiang et al., 2023)	0.2759	0.2364	0.2317	0.1455	0.8637	0.2897	0.2056	0.1879	0.1309	0.5180
AutoJ-13B (Li et al., 2023)	0.3182	0.1873	0.1631	0.0623	0.9012	0.4153	0.3689	0.3223	0.0609	0.4410
<i>Story Evaluation Finetune Open-source Models</i>										
Coke (Joshi et al., 2025)	0.3142	—	—	0.0812	0.6509	—	—	—	—	—
Qwen2.5-7B-Instruct	0.2206	0.1611	0.1381	0.1434	0.8308	0.1873	0.1605	0.1339	0.1499	0.5236
+ Point CoT	0.5677	0.4784	0.4392	<b>0.0293</b>	0.9344	0.5173	0.4919	0.4380	0.0521	0.6715
+ GRPO	0.3430	0.2436	0.2178	0.0849	0.8793	0.1773	0.1717	0.1518	0.0790	0.4507
+ Point CoT + GRPO	0.5646	0.4831	0.4472	0.0419	<u>0.9433</u>	<u>0.5307</u>	<u>0.5121</u>	<u>0.4606</u>	<u>0.0483</u>	<u>0.6829</u>
<b>EvolvR</b>	<b>0.6774</b>	<b>0.6000</b>	<b>0.5353</b>	0.0528	<b>0.9474</b>	<b>0.6155</b>	<b>0.6033</b>	<b>0.5429</b>	<b>0.0440</b>	<b>0.7406</b>

Table 1: Performance comparison on the StoryER and HANNA datasets. EvolvR, is compared against baselines across both benchmarks. We report Pearson correlation ( $\uparrow$ ), Spearman correlation ( $\uparrow$ ), Kendall correlation ( $\uparrow$ ), Mean Squared Error (MSE  $\downarrow$ ), and F1-Score ( $\uparrow$ ) ( $p < 0.0001$  with 95% confidence intervals).

these datasets can be found in the supplementary material. In our generation setup, the evaluator, trained on the HANNA dataset, functions as a reward model to steer a generator prompted with inputs from the same dataset.

#### 4.1.2 Baselines

We compare our model against three categories of baselines. First, we evaluate proprietary large language models via prompt engineering. This category includes models from OpenAI (GPT-4o, GPT-4.1, and GPT-o3-0416-global (Hurst et al., 2024)), Google (Gemini-2.5-flash and Gemini-2.5-Pro (Comanici et al., 2025)), and Anthropic (Claude-4-Sonnet and Claude-opus4). Second, we compare against open-source evaluators fine-tuned on various NLG tasks, such as summarization, machine translation, and story generation. This group includes TIGERScore, InstructScore, Themis, and AutoJ. Notably, models like TIGERScore and InstructScore require a ground-truth reference text, a constraint our method does not have. Finally, we conduct comparisons with methodologically similar models. This includes Coke, which is also trained on the StoryER dataset but requires human-written comments as input. To analyze the contribution of our design choices, we also establish our

own baselines: (1) the base Qwen2.5-7B-Instruct model (Team, 2024) without any fine-tuning, and (2) ablated versions of our model with different components of our proposed framework removed.

#### 4.1.3 Evaluation Metrics

Following established practices in NLG evaluation, we assess the performance of all models by measuring the agreement between their predicted ratings and the ground-truth human judgments. Specifically, we calculate three standard correlation coefficients: Pearson’s correlation coefficient ( $r$ ), Spearman’s rank correlation coefficient ( $\rho$ ), and Kendall’s rank correlation coefficient ( $\tau$ ). In addition, to align with the evaluation protocol of prior work like Coke, we also report the Mean Squared Error (MSE) computed on scores normalized to the  $[0, 1]$  range, and the F1-Score using 0.5 as the binary classification threshold, providing a comprehensive view of both correlation and accuracy.

To evaluate the quality of the stories produced by our fine-tuned generative model, we employed a comprehensive human evaluation protocol. We enlisted a panel of four professional screenwriters as expert judges to ensure nuanced and domain-specific assessment. These judges were asked to rate the generated stories on a 1-to-5 scale across

Model	Pearson	Spearman	Kendall
<i>NLG Evaluation Open-Source Models</i>			
InstructScore-7B(Xu et al., 2023)	0.2957	0.2751	0.2646
Themis-8B(Hu et al., 2024b)	0.4694	0.4419	0.3947
TIGERScore-13B(Jiang et al., 2023)	0.2631	0.2548	0.2141
AutoJ-13B(Li et al., 2023)	0.3821	0.3615	0.2812
<i>Finetune Open-source Models</i>			
Qwen2.5-7B-Instruct	0.2015	0.1752	0.1976
+ Pointwise CoT	0.4434	0.4523	0.4211
+ GRPO	0.1937	0.1986	0.1979
+ Pointwise CoT + GRPO	0.4515	0.4654	0.4401
<b>EvolvR</b>	<b>0.4972</b>	<b>0.5007</b>	<b>0.4893</b>

Table 2: Black-box evaluation results on the OpenMEVA(ROC) benchmark. Our method is compared against existing open-source story evaluation models. All metrics are correlation coefficients.

six key dimensions, mirroring the criteria from the HANNA dataset: Relevance, Coherence, Empathy, Surprise, Engagement, and Complexity. Based on their detailed assessments, we report two primary metrics. The first is the Average Score, calculated by averaging the scores across all six dimensions for each model. The second metric is the Win Rate, where we perform head-to-head comparisons of our model’s generated stories against two distinct references: (1) stories generated by the base model and (2) stories from the original HANNA dataset.

## 4.2 Main Results and Analysis

### 4.2.1 Analysis of Evaluation Results

The main results, presented in Table 1, demonstrate that our proposed model, EvolvR, achieves SOTA performance in correlation with human judgments on both datasets, while also showing competitive results on MSE and F1-Score.

In the GPT series, the deep reasoning model GPT-o3 outperforms GPT-4o and GPT-4.1. Similarly, Gemini-Pro surpasses its lightweight counterpart, Gemini-flash, and Claude-opus4 is superior to the faster Claude-Sonnet. This trend suggests that a more profound thought process is crucial for complex tasks like story evaluation. This observation supports us that a high-quality evaluation stems from a deep and rigorous CoT, which our framework is explicitly designed to generate and refine.

Our model surpasses general-purpose NLG evaluators like InstructScore, TIGERScore, Themis, and AutoJ, even though they have been trained on story-related data. First, models like InstructScore and TIGERScore are constrained by their need for a reference answer, a requirement that is often im-

Model	Pearson $\uparrow$	Spearman $\uparrow$	Kendall $\uparrow$
Qwen3-4B	0.1322	0.1606	0.1361
+ GRPO	0.1374	0.1596	0.1357
+ Pointwise CoT	0.4998	0.4804	0.4255
+ Pairwise CoT	<b>0.5855</b>	<b>0.5755</b>	<b>0.5214</b>
LLama3.1-8B	0.1490	0.1682	0.1034
+ GRPO	0.2288	0.3197	0.2710
+ Pointwise CoT	0.5244	0.5092	0.4538
+ Pairwise CoT	<b>0.5852</b>	<b>0.5826</b>	<b>0.5291</b>
Qwen3-8B	0.2236	0.2045	0.1714
+ GRPO	0.2855	0.2289	0.1888
+ Pointwise CoT	0.5302	0.5192	0.4611
+ Pairwise CoT	<b>0.6092</b>	<b>0.5920</b>	<b>0.5334</b>
Qwen3-14B	0.2676	0.2646	0.2232
+ GRPO	0.2728	0.2688	0.2204
+ Pointwise CoT	0.5049	0.4747	0.4232
+ Pairwise CoT	<b>0.6421</b>	<b>0.6349</b>	<b>0.5753</b>

Table 3: The performance of Pairwise CoT, Pointwise CoT Fine-tuning, and GRPO Reinforcement Learning experiments on the HANNA dataset across a diverse set of base models with varying scales and architecture

practical in real-world creative scenarios where no single correct story exists. Second, the performance gap with Themis and AutoJ which operate without references validates the superiority of our framework’s approach: to self-synthesize and evolutionarily select high-quality CoTs tailored specifically for story evaluation. To validate the effectiveness of our model in broader scenarios, we compared it against several mainstream open-source NLG evaluation models on the OpenMEVA benchmark, reporting results on the ROC split in Table 2 and results on the WP split are provided in Appendix Table 13.

The comparison with Coke, a model also trained on StoryER, is revealing. Coke’s methodology relies on processing human-written comments by extracting keywords directly for scoring. Our model’s performance demonstrates that proactively constructing and evolving high-fidelity rationales from scratch is a more robust and effective strategy than passively relying on existing, often noisy and inconsistent, human commentary.

### 4.2.2 Internal Methodological Comparison

Finally, our internal results show that EvolvR is the driver of performance. It is noteworthy that the subsequent application of GRPO does not yield a similarly dramatic improvement in the evaluation metrics. To further dissect the relative contributions of the key components in our framework in-

Model	Relevance	Coherence	Empathy	Surprise	Engagement	Complexity
Qwen2.5-7B-Instruct	4.4000	<b>4.3375</b>	2.6375	2.7375	<u>3.1000</u>	3.5750
+ SFT	3.4875 (-20.7%)	3.6125 (-16.7%)	2.7000 (+2.4%)	2.8250 (+3.2%)	<u>3.1000</u> (0.0%)	3.2875 (-8.0%)
+ Point-RM GRPO	<u>4.4875</u> (+2.0%)	<u>4.2625</u> (-1.7%)	<b>2.8375</b> (+7.6%)	<u>2.9500</u> (+7.8%)	3.0875 (-0.4%)	<u>3.6250</u> (+1.4%)
<b>EvolvR GRPO</b>	<b>4.5875</b> (+4.3%)	4.2500 (-2.0%)	<u>2.7625</u> (+4.7%)	<b>2.9625</b> (+8.2%)	<b>3.2500</b> (+4.8%)	<b>3.7250</b> (+4.2%)

Table 4: Multi-dimensional average scores. Values in parentheses denote the percentage change to the baseline.

cluding Pairwise CoT, Point CoT, and GRPO. We conducted a series of controlled experiments on the HANNA dataset across a diverse set of base models with varying scales and architectures including the Qwen3 (Yang et al., 2025) and LLaMA3.1 (Grattafiori et al., 2024) series. The results are presented in Table 3. A universal trend across all tested models is the improvement in performance upon the introduction of fine-tuning with our self-synthesized and evolved CoT data. The results consistently show that, across all models, training with pairwise CoT significantly outperforms pointwise CoT, as the former cultivates a more discerning evaluation capability by forcing the model to focus on the fine-grained distinctions that make one story superior to another. A noteworthy phenomenon is that applying GRPO directly to the base models yields limited performance gains, and in some cases, almost no improvement at all. Further analysis reveals this is due to reward hacking and experiment details can be seen in supplementary material.

### 4.2.3 EvolvR-Guided Story Generation

We evaluate the efficacy of EvolvR as a reward model by guiding a Qwen2.5-7B-Instruct generator and comparing its performance against two baselines: SFT and GRPO guided by a standard pointwise reward model. As shown in Table 4, the EvolvR-guided model achieves the highest scores in key dimensions such as Complexity, Surprise, Relevance, and Engagement. SFT exhibits limitations in this open-ended task, learning restrictive patterns from the training data, which results in a degradation of generation quality. GRPO yields an overall improvement over SFT (Table 5). While both the Point-RM and EvolvR models experience a slight decrease in coherence, their performance in other dimensions shows varying levels of improvement. To support statistical confidence, we conducted human evaluation on a sample of 500 stories and the improvements are significant with  $p < 0.01$  under a 95% confidence interval. More details are available in the appendix.

Method	Win Base	Win HANNA	Average Score
Qwen2.5-7B-Instruct	—	0.2249	3.464
+ SFT	0.3215	0.1767	3.168
+ Point-RM GRPO	0.5611	0.2909	3.541
<b>EvolvR GRPO</b>	<b>0.6436</b>	<b>0.3162</b>	<b>3.589</b>

Table 5: Effectiveness of EvolvR as a reward model for story generation. Win Base: the win rate against the untuned base model; Win HANNA: the win rate against human stories from the dataset; and Average Score: the mean score across all dimensions.

Evolution Configuration	Pearson $\uparrow$	Spearman $\uparrow$	Kendall $\uparrow$
Baseline	0.5682	0.5591	0.5007
+ Multi-persona	0.5941	0.5838	0.5273
+ Rule-Check	0.5816	0.5792	0.5146
+ Refinement Agent	0.5839	0.5791	0.5154
+ Attack Agent	0.5989	0.5891	0.5308
+ Confidence Agent	0.5807	0.5676	0.5100
<b>EvolvR</b>	<b>0.6155</b>	<b>0.6033</b>	<b>0.5429</b>

Table 6: Impact of each component in the EvolvR framework. This study demonstrates that adding each agent module to improve performance. The results of the EvolvR model demonstrate its optimal performance.

### 4.2.4 Ablation Study

To validate our design, we performed an incremental analysis of the EvolvR framework, with results shown in Table 6. Starting from a baseline without any agents, we find that adding any single agent module leads to the improvement in correlation with human judgments. Ultimately, the full EvolvR model, which combines the strengths of all agents, significantly outperforms both the baseline and any single-agent variant. This confirms the effectiveness of each individual component and demonstrates the powerful synergistic effect of our complete framework design.

## 5 Conclusions

In this paper, we addressed the challenge of developing high-fidelity evaluators for the task of story evaluation, a bottleneck for both reliable evaluation and advanced story generation. We introduced EvolvR, a novel self-evolving framework that leverages pairwise comparison to equip open-

source models with deep reasoning capabilities. EvolvR generates a diverse corpus of score-aligned CoT rationales through a multi-persona strategy and refines this data via a rigorous multi-agent evolution pipeline, ensuring logical consistency and robustness. Our extensive experiments demonstrate that EvolvR achieves SOTA performance on three distinct benchmarks including StoryER, HANNA, and OpenMEVA, outperforming both proprietary LLMs and existing open-source evaluators. Furthermore, we validated the practical utility of our evaluator by deploying it as a reward model. The EvolvR-guided generator produced stories of higher quality, achieving superior win rates and average scores in human evaluations, thus closing the loop between evaluation and generation.

## Limitations

First, while we propose an open-source CoT evolution and filtering pipeline, a low-cost verification of why a retained CoT is better than a rejected one remains challenging. In practice, validating the quality of intermediate CoTs still relies on expensive external judgments or expert human assessment. Although we observe gains on downstream metrics, richer interpretability are not fully resolved. Second, EvolvR is grounded in score-aligned synthesis and therefore requires labeled data with a predefined scoring scheme. This reliance limits applicability in unlabeled settings or in domains where annotation is scarce. Third, using a judge model as a reward model yields only moderate improvements in story generation. This is due to the inherent noise and distribution shift in black-box reward evaluation, which can reduce reward fidelity and constrain the effectiveness of RL optimization.

## Ethical Considerations

Expert annotators were provided with the following evaluation guidelines: "You will evaluate AI-generated stories on six dimensions (creativity, coherence, relevance, complexity, surprise, and engagement) using a 1-5 scale. Each story should be read completely before scoring. Focus solely on story quality without considering the AI system that generated it. Your evaluations will be used to validate automatic evaluation metrics. No personal information will be collected." Human evaluators were professional annotators from our research team with specialized training in creative writing assessment and extensive experience in NLP

evaluation tasks. These were not crowdsourced workers but dedicated team members with relevant academic backgrounds. Compensation was provided according to institutional standards for professional research assistance, ensuring fair remuneration commensurate with the specialized nature of the task and the evaluators' expertise level. All evaluators provided informed consent before participation. They were explicitly informed that: (1) they would be evaluating AI-generated content only, (2) their judgments would be used solely for research purposes to validate automatic metrics, (3) no personal information beyond their quality assessments would be recorded, and (4) they could withdraw from the evaluation process at any time without penalty. The evaluation data contains only story quality scores without any annotator identification. The evaluation involved no risks beyond normal reading activities. Stories were pre-screened to ensure no offensive or disturbing content. Evaluators were informed they could skip any story they found uncomfortable.

## References

- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. Chateval: Towards better llm-based evaluators through multi-agent debate. *arXiv preprint arXiv:2308.07201*.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, and 1 others. 2024. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3):1–45.
- Hong Chen, Duc Minh Vo, Hiroya Takamura, Yusuke Miyao, and Hideki Nakayama. 2023. Storyer: Automatic story evaluation via ranking, rating and reasoning. *Journal of Natural Language Processing*, 30(1):243–249.
- Steffi Chern, Ethan Chern, Graham Neubig, and Pengfei Liu. 2024. Can large language models be trusted for evaluation? scalable meta-evaluation of llms as evaluators via agent debate. *arXiv preprint arXiv:2401.16788*.
- Cyril Chhun, Fabian M Suchanek, and Chloé Clavel. 2024. Do language models enjoy their own stories? prompting large language models for automatic story evaluation. *Transactions of the Association for Computational Linguistics*, 12:1122–1142.
- Cheng-Han Chiang and Hung-yi Lee. 2023. A closer look into using large language models for automatic

- evaluation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8928–8942.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.
- Mingqi Gao, Xinyu Hu, Xunjian Yin, Jie Ruan, Xiao Pu, and Xiaojun Wan. 2025. Llm-based nlg evaluation: Current status and challenges. *Computational Linguistics*, pages 1–27.
- Carlos Gómez-Rodríguez and Paul Williams. 2023. A confederacy of models: A comprehensive evaluation of llms on creative writing. *arXiv preprint arXiv:2310.08433*.
- Peiyuan Gong and Jiaxin Mao. 2023. Coascore: Chain-of-aspects prompting for nlg evaluation. *arXiv preprint arXiv:2312.10355*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Jian Guan, Zhuoer Feng, Yamei Chen, Ruilin He, Xiaoxi Mao, Changjie Fan, and Minlie Huang. 2022. Lot: A story-centric benchmark for evaluating chinese long text understanding and generation. *Transactions of the Association for Computational Linguistics*, 10:434–451.
- Jian Guan, Zhixin Zhang, Zhuoer Feng, Zitao Liu, Wenbiao Ding, Xiaoxi Mao, Changjie Fan, and Minlie Huang. 2021. Openmeva: A benchmark for evaluating open-ended story generation metrics. *arXiv preprint arXiv:2105.08920*.
- Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *arXiv preprint arXiv:2305.02301*.
- Xinyu Hu, Mingqi Gao, Sen Hu, Yang Zhang, Yicheng Chen, Teng Xu, and Xiaojun Wan. 2024a. Are llm-based evaluators confusing nlg quality criteria? *arXiv preprint arXiv:2402.12055*.
- Xinyu Hu, Li Lin, Mingqi Gao, Xunjian Yin, and Xiaojun Wan. 2024b. Themis: A reference-free nlg evaluation language model with flexibility and interpretability. *arXiv preprint arXiv:2406.18365*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Dongfu Jiang, Yishan Li, Ge Zhang, Wenhao Huang, Bill Yuchen Lin, and Wenhui Chen. 2023. Tigerscore: Towards building explainable metric for all text generation tasks. *arXiv preprint arXiv:2310.00752*.
- Brihi Joshi, Sriram Venkatapathy, Mohit Bansal, Nanyun Peng, and Haw-Shiuan Chang. 2025. Coke: Customizable fine-grained story evaluation via chain-of-keyword rationalization. *arXiv preprint arXiv:2503.17136*.
- Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoon Yun, Seongjin Shin, Sungdong Kim, James Thorne, and 1 others. 2023. Prometheus: Inducing fine-grained evaluation capability in language models. In *The Twelfth International Conference on Learning Representations*.
- Barbara Kroll and Joy Reid. 1994. Guidelines for designing writing prompts: Clarifications, caveats, and cautions. *Journal of second language writing*, 3(3):231–255.
- Yukyung Lee, Joonghoon Kim, Jaehee Kim, Hyowon Cho, and Pilsung Kang. 2024. Checkeval: Robust evaluation framework using large language model via checklist. *CoRR*.
- Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhattacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, and 1 others. 2025. From generation to judgment: Opportunities and challenges of llm-as-a-judge, 2025. URL <https://arxiv.org/abs/2411.16594>.
- Junlong Li, Shichao Sun, Weizhe Yuan, Run-Ze Fan, Hai Zhao, and Pengfei Liu. 2023. Generative judge for evaluating alignment. *arXiv preprint arXiv:2310.05470*.
- Yu Li, Shenyu Zhang, Rui Wu, Xiutian Huang, Yongrui Chen, Wenhao Xu, Guilin Qi, and Dehai Min. 2024a. Mateval: a multi-agent discussion framework for advancing open-ended text evaluation. In *International Conference on Database Systems for Advanced Applications*, pages 415–426. Springer.
- Zhen Li, Xiaohan Xu, Tao Shen, Can Xu, Jia-Chen Gu, Yuxuan Lai, Chongyang Tao, and Shuai Ma. 2024b. Leveraging large language models for nlg evaluation: Advances and challenges. *arXiv preprint arXiv:2401.07103*.

- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023a. G-eval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Yuxuan Liu, Tianchi Yang, Shaohan Huang, Zihan Zhang, Haizhen Huang, Furu Wei, Weiwei Deng, Feng Sun, and Qi Zhang. 2023b. Calibrating llm-based evaluator. *arXiv preprint arXiv:2309.13308*.
- Aviya Maimon and Reut Tsarfaty. 2023. Cohesentia: A novel benchmark of incremental versus holistic assessment of coherence in generated texts. *arXiv preprint arXiv:2310.16329*.
- Ning Miao, Yee Whye Teh, and Tom Rainforth. 2023. Selfcheck: Using llms to zero-shot check their own step-by-step reasoning. *arXiv preprint arXiv:2308.00436*.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. Orca: Progressive learning from complex explanation traces of gpt-4. *arXiv preprint arXiv:2306.02707*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in neural information processing systems*, 33:3008–3021.
- Qwen Team. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Anvesh Rao Vijjini, Faeze Brahman, and Snigdha Chaturvedi. 2022. Towards inter-character relationship-driven story generation. *arXiv preprint arXiv:2211.00676*.
- Danqing Wang, Kevin Yang, Hanlin Zhu, Xiaomeng Yang, Andrew Cohen, Lei Li, and Yuandong Tian. 2024. Learning personalized alignment for evaluating open-ended text generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13274–13292.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023a. Is chatgpt a good nlg evaluator? a preliminary study. *arXiv preprint arXiv:2303.04048*.
- Peifeng Wang, Zhengyang Wang, Zheng Li, Yifan Gao, Bing Yin, and Xiang Ren. 2023b. Scott: Self-consistent chain-of-thought distillation. *arXiv preprint arXiv:2305.01879*.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghui Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023c. Large language models are not fair evaluators. *arXiv preprint arXiv:2305.17926*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Bosi Wen, Pei Ke, Yufei Sun, Cunxiang Wang, Xiaotao Gu, Jinfeng Zhou, Jie Tang, Hongning Wang, and Minlie Huang. 2025. Hpss: Heuristic prompting strategy search for llm evaluators. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 24974–25007.
- Zhuohan Xie, Miao Li, Trevor Cohn, and Jey Han Lau. 2023. Deltascore: Fine-grained story evaluation with perturbations. *arXiv preprint arXiv:2303.08991*.
- Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Yang Wang, and Lei Li. 2023. Instructscore: Explainable text generation evaluation with finegrained feedback. *arXiv preprint arXiv:2305.14282*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Dingyi Yang and Qin Jin. 2024. What makes a good story and how can we measure it? a comprehensive survey of story evaluation. *arXiv preprint arXiv:2408.14622*.

- Kevin Yang, Dan Klein, Nanyun Peng, and Yuandong Tian. 2022a. Doc: Improving long story coherence with detailed outline control. *arXiv preprint arXiv:2212.10077*.
- Kevin Yang, Yuandong Tian, Nanyun Peng, and Dan Klein. 2022b. Re3: Generating longer stories with recursive reprompting and revision. *arXiv preprint arXiv:2210.06774*.
- Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. Plan-and-write: Towards better automatic storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7378–7385.
- Wang You, Wenshan Wu, Yaobo Liang, Shaoguang Mao, Chenfei Wu, Maosong Cao, Yuzhe Cai, Yiduo Guo, Yan Xia, Furu Wei, and 1 others. 2023. Eipe-text: Evaluation-guided iterative plan extraction for long-form narrative text generation. *arXiv preprint arXiv:2310.08185*.
- Peiwen Yuan, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Boyuan Pan, Heda Wang, Yao Hu, and Kan Li. 2024. Batcheval: Towards human-like text evaluation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15940–15958.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D Goodman. 2024. Star: Self-taught reasoner bootstrapping reasoning with reasoning. In *Proc. the 36th International Conference on Neural Information Processing Systems*, volume 1126.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International conference on machine learning*, pages 12697–12706. PMLR.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623.
- Pei Zhou, Jay Pujara, Xiang Ren, Xinyun Chen, Heng-Tze Cheng, Quoc V Le, Ed Chi, Denny Zhou, Swaroop Mishra, and Huaixiu Steven Zheng. 2024. Self-discover: Large language models self-compose reasoning structures. *Advances in Neural Information Processing Systems*, 37:126032–126058.
- Wangchunshu Zhou, Yuchen Eleanor Jiang, Peng Cui, Tiannan Wang, Zhenxin Xiao, Yifan Hou, Ryan Cotterell, and Mrinmaya Sachan. 2023. Recurrentgpt: Interactive generation of (arbitrarily) long text. *arXiv preprint arXiv:2305.13304*.

## A Dataset Details

Our work leverages three prominent datasets for story evaluation: StoryER (Chen et al., 2023), HANNA (Chhun et al., 2024), and OpenMEVA (Guan et al., 2021). We show the details in Table 7. Each dataset offers unique characteristics in terms of evaluation format, annotation dimensions, and scale, providing a comprehensive basis for training our model and assessing its performance and generalization capabilities.

Dataset	Samples	Dimensions	Format
StoryER-Scale	45,948	5	Pointwise (Multi-Aspect)
HANNA	19,008	6	Pointwise (Multi-Aspect)
OpenMEVA	2,000	1	Pointwise (Holistic)

Table 7: Summary of the dataset used in our study. Each dataset provides a different combination of data scale, dimensional granularity, and evaluation format.

### A.1 StoryER Dataset

The StoryER dataset (Chen et al., 2023) is a comprehensive resource for explainable story evaluation. For our research, we specifically utilize the StoryER-Scale subset. Its defining feature is that it provides not only multi-aspect scores but also human-written reasoning behind those scores. This aspect, while not directly used for training as our EvolvR framework self-synthesizes its own rationales, provides a strong motivation for our CoT-based approach and serves as a valuable reference. StoryER-Scale provides pointwise ratings on a 1-to-5 Likert scale for individual stories across multiple fine-grained dimensions. The key dimensions covered in this dataset include Coherence, Ending, Style, Character Development, and Empathy. We use StoryER-Scale as a primary benchmark for training and testing our evaluator.

### A.2 HANNA Dataset

The HANNA dataset (Chhun et al., 2024) is designed to evaluate story generation against detailed, multi-faceted human criteria. Unlike StoryER-Scale, it does not provide explicit reasoning but offers a rich set of evaluation dimensions. The dataset provides direct multi-aspect ratings for stories on a 1-to-5 Likert scale. HANNA focuses on a comprehensive set of six dimensions: Relevance, Coherence, Empathy, Surprise, Engage-

ment referred to as "Interestingness" in the original paper, and Complexity. Given its high-quality, multi-aspect annotations, HANNA serves as our second core benchmark for training and evaluating the EvolvR model. It is also the source dataset for our story generation experiments, where our evaluator, trained on HANNA, acts as the reward model.

### A.3 OpenMEVA Benchmark

The OpenMEVA benchmark (Guan et al., 2021) is designed for evaluating the *Overall Quality* of open-ended story generation. It provides a single, holistic score without detailed dimensional breakdowns or reasoning. Each story is assigned a single *Overall Quality* score on a 1-to-5 Likert scale. Key aspects considered by annotators include Relevance, Fluency, Coherence, and Commonsense, but these are subsumed into the final holistic score. We use OpenMEVA exclusively for zero-shot generalization testing. This assesses our model’s ability to perform well on a dataset with a different format and distribution without any fine-tuning. To facilitate comparison, we average our model’s predicted scores across its six dimensions to produce a single holistic score, which is then correlated with the OpenMEVA ground truth.

## B Detailed Method of the EvolvR

### B.1 Rationale for Pairwise Comparison

Our decision to build the EvolvR framework upon pairwise comparison is a deliberate choice, grounded in its superior robustness for evaluating complex, open-ended tasks like story assessment. Theoretically, comparative judgment *Story A is more creative than Story B* is more aligned with human cognitive processes than assigning abstract, absolute scores. This comparative mode inherently forces an evaluator whether human or machine to identify and articulate the fine-grained distinctions that determine a preference, leading to more reliable and detailed rationales. Furthermore, this preference-based data format is naturally compatible with modern reinforcement learning RLHF (Ouyang et al., 2022) paradigms such as GRPO (Shao et al., 2024), making the resulting evaluator an ideal reward model for enhancing generative models.

To move beyond theoretical motivation and empirically validate this choice, we conducted a rigorous inter-annotator agreement analysis on the

HANNA dataset. We assessed human judgment consistency under two conditions: a pointwise scenario, measuring the agreement on absolute scores given to the same story by different annotators, and a pairwise scenario, measuring the agreement on the score differences between two distinct stories. The results, summarized in Table 8, provide strong empirical support for our approach. In five of the six evaluation dimensions, including critical aspects like Coherence (+0.106) and Relevance (+0.063), the agreement among annotators was demonstrably higher in the pairwise setting. While Complexity showed a marginal exception, the decisive trend across the majority of dimensions confirms that pairwise comparison elicits a more stable and consistent signal of human preference. This empirical validation reinforces that our pairwise foundation is essential for training a high-fidelity evaluator.

Dimension	Pointwise	Pairwise	Diff	Improvement(%)
Relevance	0.504	0.567	+0.063	+12.5%
Coherence	0.484	0.590	+0.106	+21.9%
Empathy	0.570	0.578	+0.008	+1.4%
Surprise	0.547	0.586	+0.039	+7.1%
Engagement	0.565	0.573	+0.008	+1.4%
Complexity	0.607	0.595	-0.012	-2.0%

Table 8: Inter-annotator agreement on the HANNA dataset. The table shows that pairwise agreement is consistently higher across most dimensions, with significant improvements in Coherence (+21.9%) and Relevance (+12.5%), validating our choice of a pairwise framework.

## B.2 Multi-Persona CoT Synthesis

To generate a diverse set of initial CoT (Wei et al., 2022) rationales, we employed a multi-persona strategy for both a pointwise and a pairwise experimental setup. This approach was designed to produce a rich and varied dataset of reasoning styles. The specific prompts used for each persona are detailed in the *Prompt Study* section of this appendix.

The synthesis process was guided by five distinct personas: *Academic*, *Artist*, *Pragmatist*, *Sharp-Tongued Reader*, and *Casual Netizen*. Each persona was instructed to adopt a specific linguistic style and evaluation focus. To ensure the stability and quality of the generated rationales in both setups, several measures were taken. We employed a few-shot prompting strategy, providing the models with high-quality examples to guide the structure

and coherence of their output. The synthesis itself was performed using open-source LLMs.

For our pointwise experiments, we tasked the LLM with generating a CoT for a single story. The model was provided with an input of (story, dimension, final score) and prompted to generate a CoT that logically justifies the given score from the assigned persona’s perspective, effectively simulating a score-aligned thought process. For our core pairwise framework, the process was more detailed to ensure data quality and mitigate potential biases. First, to prevent the model from overfitting to the natural distribution of score pairs in the dataset, we performed stratified sampling. This ensured that every possible score-pair combination (e.g., (1,2), (3,5), (4,4)) was represented with a balanced number of samples in our synthesis set. Second, to counteract positional bias, we augmented the data by swapping the order of each story pair. For every original pair (Story A, Story B), a corresponding instance (Story B, Story A) was created. This forces the model to learn from the content of the stories rather than their position in the prompt (Wang et al., 2023c).

## B.3 The CoT Evolution and Selection Pipeline

The CoTs synthesized in the initial multi-persona stage are of variable quality and require rigorous vetting. To ensure the logical consistency, robustness, and fidelity of our training data, we designed a multi-stage evolution and selection pipeline. This pipeline subjects each candidate CoT to a sequence of filtering and refinement agents. The overall process follows the sequence: Rule → Refinement → Rule → Counter → Confidence. The repeated application of the Self-Rule agent ensures that data integrity is maintained after any modification step. The specific prompts used for the Self-Refinement and Self-Counter agents are detailed in the "Prompt Study" section of this appendix.

We began with a pool of 800,000 candidate CoTs, generated from 80,000 stratified story pairs through our multi-persona and position-swapping strategy. The significant data reduction at each stage, detailed in Table 9, underscores the stringency of our selection criteria, culminating in a final dataset of over 530,000 high-quality rationales.

**Self-Rule Agent** The primary function of this agent is to enforce strict alignment between the rationale’s conclusion and the ground-truth scores. Although the synthesis prompts required score

Stage	Agent	Remaining	Survival Rate
1	Initial Pool	800,000	100.0%
2	Self-Rule	743,671	92.95%
3	Self-Refinement-Rule	686,219	85.78%
4	Self-Attack	603,182	75.40%
5	Self-Confidence	536,177	67.02%

Table 9: Data attrition throughout the CoT evolution and selection pipeline.

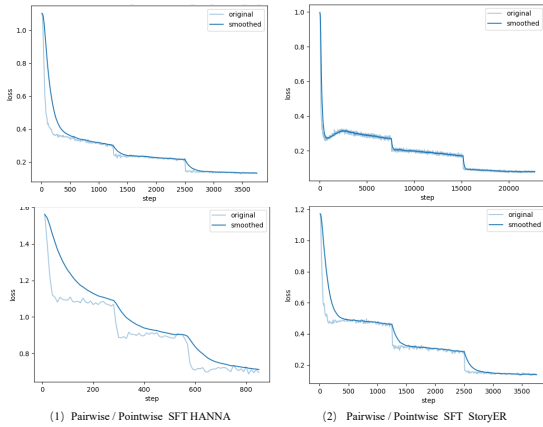


Figure 3: Training loss curves for the Pointwise and Pairwise models trained via Supervised Fine-Tuning. Both models were trained using a standard cross-entropy loss objective.

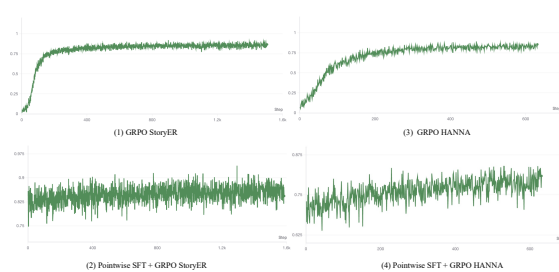


Figure 4: Evolution of the average reward during the GRPO training process. The reward is designed to be 1 for perfect predictions and decay exponentially with error. The steady increase and eventual plateau of the reward.

alignment, minor deviations can occur. This agent uses regular expressions to parse the scores from the CoT’s conclusion and discards any instance where they do not exactly match the ground-truth scores. It serves as a critical integrity check after both initial synthesis and any subsequent modifications.

**Self-Refinement Agent** Leveraging the LLM’s inherent self-correction capabilities, this agent improves the quality of the reasoning itself. We prompt the model to review and rewrite its own

rationale (Miao et al., 2023), with the explicit goal of enhancing its logical flow, clarity, and persuasiveness, without altering the underlying judgment. Any CoT that is modified during this stage must pass through another Self-Rule check to ensure it remains aligned with the target scores.

**Self-Attack Agent** This agent adversarially tests the logical robustness of a CoT. For a given rationale that correctly justifies a preference (e.g., Story A > Story B), we create a corrupted version by inverting the final scores to suggest the opposite (Story B > Story A). The LLM is then prompted to identify if a logical contradiction exists between the original reasoning and the new, inverted conclusion. If the model fails to detect the blatant contradiction implying the reasoning is weak or generic enough to support conflicting outcomes the original CoT is considered not robust and is discarded.

**Self-Confidence Agent** This final filter selects for CoTs that lead the model to the correct conclusion with high confidence. We concatenate the instruction (without providing the final scores) and the CoT rationale, and feed this combined input to the model. At the decoding step where a score token would be generated, we extract the logits for the potential score tokens (e.g., integers 1 through 5). The confidence in the ground-truth score  $y^*$  is calculated as its softmax probability relative to the other possible scores:

$$P(y = y^* | \text{Input}, \text{CoT}) = \frac{e^{L(y^*)}}{\sum_{j=1}^5 e^{L(j)}} \quad (13)$$

where  $L(j)$  is the logit for the score token corresponding to integer  $j$ . A CoT is retained only if the probability of the ground-truth score is the maximum among all possible scores. This ensures that the selected rationales are not merely plausible, but decisively lead the model to the correct judgment.

#### B.4 External Validation of CoT Filtering

Our pipeline removes a large number of synthesized CoTs (Table 9). To verify that this filtering does not discard high-quality rationales, we employ an external black-box judge for validation. We conduct two complementary tests: (1) a stage-wise overall-quality assessment to measure whether the retained CoTs become higher quality after filtering, and (2) a stage-wise matched pairwise preference test to measure whether the CoTs retained by each filtering stage are preferred over those rejected at

Filtering Stage	Overall Quality $\uparrow$	Win Rate $\uparrow$
Initial pool	3.22	—
+ Self-Rule	3.45	77.2%
+ Self-Refinement-Rule	3.57	71.5%
+ Self-Attack	3.61	75.1%
+ Self-Confidence	3.65	75.6%

Table 10: External validation of CoT filtering using GPT-4o. **Overall Quality**: mean overall-quality score for CoTs sampled from each stage output. **Win Rate**: at each stage, we match a CoT retained by that stage with a CoT rejected at the same stage from the same story pair and aspect and report the retained CoT win rate.

the same stage. For the overall-quality assessment, we uniformly sample 1000 CoTs from the output of each stage and ask GPT-4o to rate their overall quality on a 1–5 scale. For the pairwise preference test, at each stage we sample 1000 matched pairs by pairing a CoT retained by that stage with a CoT rejected at the same stage from the same story pair and aspect, and report the win rate of the retained CoTs.

## C Training and Experiment Details

### C.1 Details for Story Evaluation

This section provides additional implementation details for our story evaluation experiments. We also report additional black-box generalization results on the OpenMEVA (WP) split (Table 13) and comparisons with prompt-based evaluation frameworks on HANNA (Table 15).

#### C.1.1 Training Procedures

We trained three distinct types of models based on our synthesized data. For the Pointwise and Pairwise CoT Models, they were trained via SFT on their respective datasets Pointwise CoT and our final EvolvR Pairwise CoT data. We used a standard cross-entropy loss objective to train the models to predict the CoT rationales and the final scores. The training loss curves for these models are provided in Figure 3. For the GRPO-trained model, we designed a reward function that encourages the model to match the ground-truth scores accurately. The reward  $R$  for a prediction  $(y_a, y_b)$  given the ground truth  $(y_a^*, y_b^*)$  was formulated as an exponential function of the cumulative error:

$$R(y_a, y_b) = \exp(-\lambda(|y_a - y_a^*| + |y_b - y_b^*|)) \quad (14)$$

where  $\lambda$  is a scaling hyperparameter. This design provides a smooth reward signal that is maximized at 1 when the predictions are perfect and decays exponentially as the error increases. Despite the reward curve in Figure 4 indicating successful convergence during training, we discovered during testing that the GRPO-trained model had learned to concentrate its predictions on the most frequent scores in the data distribution. This degenerate policy maximizes the cumulative reward by exploiting the data statistics instead of learning to genuinely evaluate the stories, which accounts for its poor performance on the final evaluation metrics.

Random Pairs	Pearson	Spearman	Kendall	MSE	F1-Score
0	0.6037	0.6113	0.5389	0.0451	0.7366
1	0.6224	0.6107	0.5486	0.0449	0.7435
2	0.5852	0.5765	0.5161	0.0466	0.7129
4	0.6155	0.6033	0.5429	0.0440	0.7406
8	0.6162	0.6067	0.5445	0.0449	0.7392
16	0.6131	0.6028	0.5416	0.0409	0.7412

Table 11: Robustness analysis of the EvolvR model with varying numbers of comparison pairs ( $N$ ). For each story, an aggregated score is derived from  $2N$  evaluations including swapped pairs to mitigate bias.

To ensure fair and rigorous comparisons, we adopted specific protocols for each category of model. Proprietary models (e.g., GPT-4 series, Claude series) were prompted using instructions identical to those used for our data synthesis to ensure a fair, zero-shot comparison. For models like Themis (Hu et al., 2024b) and AutoJ (Li et al., 2023), we utilized their publicly available code and model weights to evaluate them on our test sets. For models that require a reference text, namely InstructScore (Xu et al., 2023) and TIGERScore (Jiang et al., 2023), we employed a self-reference approach. The story to be evaluated was used as its own reference input, enabling a reference-free evaluation while adhering to the models’ required input format. These variants were evaluated in a standard zero-shot manner, directly generating a score for each story based on a prompt.

For EvolvR, we evaluate under two settings: (1) Single-pass mode where the model directly scores pair with same story, and (2) Ensemble mode with the following protocol. To ensure the stability and robustness of our main EvolvR model’s scores, we implemented a comprehensive ensembling and stability testing protocol. For each target story in the test set, we performed the following steps:

1. Randomly sample  $N$  other stories from the test set to form  $N$  evaluation pairs, with  $N$  varied across  $\{1, 2, 4, 8, 16\}$ . When  $N = 0$  (i.e., evaluating the story paired with itself), this corresponds to the single-pass mode. The impact of pairs number in Table 11.
2. For each pair (Story Target, Story Sampled), perform the evaluation twice: once in the original order, and once with the order swapped to (Story Sampled, Story Target) to mitigate positional bias. We show the impact of position in Table 12.
3. This process yields  $2N$  independent scores for the target story. To measure the model’s overall performance, we aggregated all scores generated for every story in the test set into a single list and computed the final correlation coefficients against this complete set of predictions. This method holistically assesses the model’s average performance and stability.

We show the difference of evaluation between Qwen2.5-7B-Instruct and EvolvR in Figure 5.

Position	Pearson	Spearman	Kendall	MSE	F1-Score
(A, B)	0.6189	0.6031	0.5412	0.0445	0.7379
(B, A)	0.6121	0.6035	0.5486	0.0435	0.7433

Table 12: Analysis of positional bias in the EvolvR model, conducted with  $N = 4$  random pairs. We compare the performance when evaluating pairs in the original order (A, B) versus the swapped order (B, A).

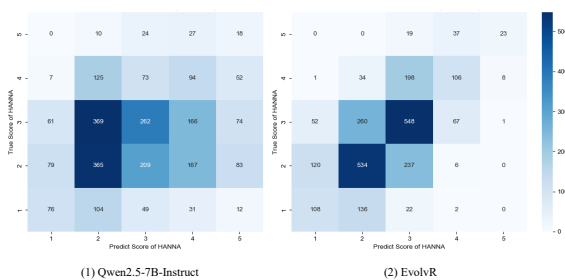


Figure 5: Confusion matrices illustrating the score agreement between model predictions and ground-truth scores on the HANNA dataset. The left matrix represents the performance of the Qwen2.5-7B-Instruct model, while the right matrix represents our EvolvR model.

### C.1.2 Results on OpenMEVA (WP)

In addition to the ROC split reported in the main text, we further evaluate zero-shot generalization

Model	Pearson	Spearman	Kendall
<i>NLG Evaluation Open-Source Models</i>			
InstructScore-7B(Xu et al., 2023)	0.1578	0.1257	0.1046
Themis-8B(Hu et al., 2024b)	0.1469	0.1719	0.1457
TIGERScore-13B(Jiang et al., 2023)	0.1661	0.1649	0.1541
AutoJ-13B(Li et al., 2023)	0.1871	0.1952	0.1611
<i>Finetune Open-source Models</i>			
Qwen2.5-7B-Instruct	0.1061	0.1175	0.0976
+ Pointwise CoT	0.1474	0.1503	0.1291
+ GRPO	0.0947	0.0994	0.0897
+ Pointwise CoT + GRPO	0.1605	0.1654	0.1415
<b>EvolvR</b>	<b>0.2092</b>	<b>0.2087</b>	<b>0.1881</b>

Table 13: Black-box evaluation results on the OpenMEVA benchmark. Our method is compared against existing open-source story evaluation models. All metrics are correlation coefficients, where higher is better ( $\uparrow$ ).

on the OpenMEVA (WP) split. We follow the same black-box evaluation protocol as for ROC: each model produces a holistic quality score for each story, and we report Pearson, Spearman, and Kendall correlations with the human annotations. Table 13 summarizes the results. EvolvR consistently outperforms the compared open-source baselines across all correlation metrics, demonstrating improved robustness under a distribution shift.

### C.1.3 Comparison with Prompt Evaluation

In addition to the baselines reported in the main text, we include representative prompt-based evaluation frameworks on HANNA, including BatchEval, G-Eval, and HPSS. These methods mainly improve evaluation by prompt optimization on proprietary judges (e.g., GPT-4). For a fair reference, we report their numbers as presented in the original papers under a GPT-4 evaluation setting. In contrast, EvolvR is an open-source evaluator trained with our self-evolving framework and does not rely on closed-source models at inference time. As shown in Table 15, EvolvR achieves higher correlation on HANNA. Moreover, our method is compatible with prompt optimization and can be further combined with these prompting strategies for additional gains.

## C.2 Details for the Story Generation

To rigorously evaluate the effectiveness of our EvolvR evaluator as a reward model, we conducted a series of comparative experiments in the story generation phase. We designed three distinct training setups to isolate the impact of different guidance methods.

Model	Overall	Relevance	Coherence	Empathy	Surprise	Engagement	Complexity
Qwen2.5-7B-Instruct	3.465 ± 0.722	4.400 ± 0.768	<b>4.338</b> ± 0.741	2.638 ± 0.746	2.737 ± 0.754	3.100 ± <b>0.604</b>	3.575 ± 0.721
+ SFT	3.169 ± 0.839	3.487 ± 0.975	3.612 ± 0.829	2.700 ± 0.900	2.825 ± 0.771	3.100 ± 0.752	3.288 ± 0.809
+ Point-RM GRPO	3.542 ± 0.775	4.487 ± 0.790	4.263 ± 0.720	<b>2.837</b> ± 0.813	2.950 ± 0.820	3.087 ± 0.728	3.625 ± 0.781
<b>EvolvR GRPO</b>	<b>3.590 ± 0.697</b>	<b>4.588 ± 0.626</b>	4.250 ± <b>0.716</b>	2.763 ± <b>0.711</b>	<b>2.962 ± 0.766</b>	<b>3.250</b> ± 0.622	<b>3.725 ± 0.741</b>

Table 14: Model performance comparison showing Average Score ± Standard Deviation. For scores, higher is better ( $\uparrow$ ). For standard deviation, lower indicates higher stability ( $\downarrow$ ). Best performance for each metric (highest score or lowest std. dev.) is highlighted in bold.

Method	Pearson	Spearman	Kendall
BatchEval (Yuan et al., 2024)	0.589	0.568	-
G-Eval (Liu et al., 2023a)	0.514	0.520	-
HPSS (Wen et al., 2025)	0.545	-	-
<b>EvolvR</b>	<b>0.616</b>	<b>0.603</b>	<b>0.543</b>

Table 15: Comparison with prompt-based evaluation frameworks on the HANNA benchmark. Results for BatchEval/G-Eval/HPSS are taken from their original papers under a GPT-4 evaluation setting; “-” denotes metrics not reported in the corresponding work. EvolvR is evaluated using our open-source evaluator under the same benchmark protocol.

### C.2.1 Supervised Fine-Tuning

To establish a strong conventional baseline, we performed Supervised Fine-Tuning. We curated a high-quality instruction-tuning dataset by filtering the HANNA training set, selecting only those stories with an average multi-dimensional rating exceeding 3.5. The base generative model was then fine-tuned on this dataset of high-quality examples.

### C.2.2 GRPO with Pointwise Reward Model

To create a more advanced baseline using reinforcement learning, we utilized the reward model trained on our synthesized Pointwise CoT data. This model guided the generator via the GRPO algorithm. The reward function for this setup was a straightforward combination of two components including the predicted quality score for the rollout story from the pointwise reward model and a penalty for deviations from a target length range.

### C.2.3 GRPO with the EvolvR Reward Model

This setup leveraged our main EvolvR evaluator as the reward model. A key implementation detail for using our pairwise evaluator is constructing the input pair (Story 1, Story 2) for the reward calculation, where one story is the rollout story from the current policy. We experimented with two distinct strategies for selecting the second story in the pair to serve as a reference. For static reference, the

input pair consisted of the rollout story and a corresponding high-quality story taken directly from the original HANNA dataset. This setup tests the generator’s ability to outperform a strong, human-vetted example. For dynamic reference, the input pair consisted of the rollout story and a story generated by the original, un-trained base model for the same prompt. This setup assesses the generator’s improvement relative to its own initial capabilities.

For both of these EvolvR-guided variants, the reward function was the three-component signal. The relative advantage component ( $\mathcal{R}_{adv}$ ) addresses evaluator inconsistency when assessing the same pair multiple times, scores may vary (e.g., (3,4) vs. (4,5)), but the difference remains stable. This component also provides a more interpretable learning signal by normalizing for different reference baselines. The absolute score component ( $\mathcal{R}_{abs}$ ) directly encourages high-quality generation, while the length reward ( $\mathcal{R}_{len}$ ) prevents degenerate solutions like overly short stories.

The hyperparameters were determined by expert knowledge, with dimension weights  $\alpha$  set to {5, 3, 1, 3, 1, 1} for creativity, coherence, relevance, complexity, surprise, and engagement respectively, emphasizing creativity and coherence as primary objectives. The three reward components were balanced with weights  $w$  of {0.1, 0.8, 0.1} for  $\mathcal{R}_{adv}$ ,  $\mathcal{R}_{abs}$ , and  $\mathcal{R}_{len}$  respectively, prioritizing absolute quality while maintaining relative improvement and length constraints. The comparative results of all three experimental setups are presented and analyzed in the main text.

### C.2.4 Story Generation Results Analysis

The experimental results clearly demonstrate the superior performance of the EvolvR-guided GRPO model. As shown in the Table 14 and bar chart Figure 7, EvolvR GRPO not only achieves the highest overall score of 3.590, significantly outperforming all baselines (SFT, Point-RM GRPO, and the base model), but also excels across several key cre-

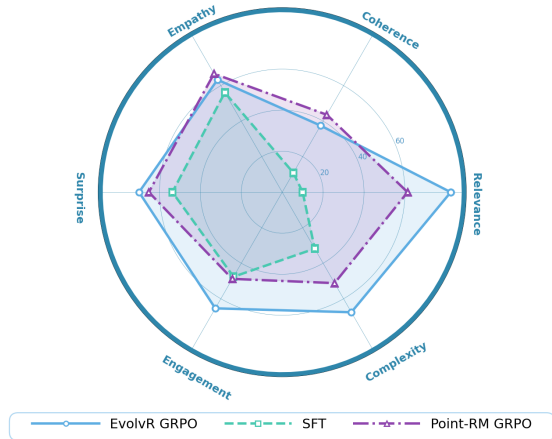


Figure 6: Win rates of different methods against the baseline model Qwen2.5-7B-Instruct in human evaluation including SFT, Point-RM GRPO and EvolvR GRPO

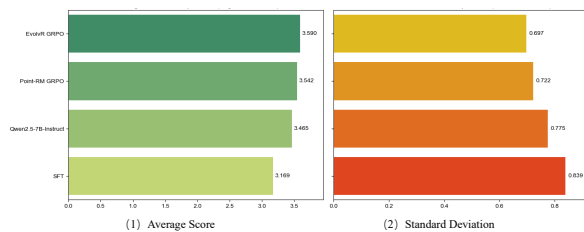


Figure 7: Comparison of Average Score (quality, left) and Standard Deviation (stability, right). Our EvolvR model excels in both, achieving the highest quality and the greatest consistency.

Method	Win Base	Win HANNA	Average Score
Pair with HANNA	0.5328	0.3014	3.552
Pair with Base	0.6436	0.3162	3.589

Table 16: Performance comparison of EvolvR-guided GRPO with two different reference pairing strategies. *Pair with Base* uses a reference a story from the base model, while *Pair with HANNA* uses a reference a story from the HANNA dataset.

ative dimensions such as Relevance, Surprise, and Complexity. Critically, its standard deviation is the lowest among all models (0.697), indicating that it produces high-quality content with the greatest consistency, thus achieving a dual advantage in both quality and stability. We show the win rates of different methods against the baseline model Qwen2.5-7B-Instruct in human evaluation in Figure 6.

Further analysis of the EvolvR-guided strategies shows the choice of reference is important. As shown in Table 16, the strategy of pairing with a story from the base model (*Pair with Base*) outper-

formed pairing with a human exemplar (*Pair with HANNA*), yielding a higher average score (3.589 vs. 3.552) and a higher win rate against the base model (0.6436 vs. 0.5328). We hypothesize that because generations from the same model family are more similar in style and structure, the reward model can make more precise judgments. By comparing closely related items, EvolvR can focus on the subtle differences that represent improvements, thus providing a more stable and effective optimization signal. In contrast, when comparing a model’s output to a stylistically distant human-written story, the broad differences may make it difficult for the reward model to provide a nuanced gradient.

**Annotation Protocol and Data Release.** To ensure annotation quality and reproducibility, we implemented a comprehensive calibration process before the main evaluation. First, all four annotators participated in a 2-hour training session where we established unified scoring guidelines based on the HANNA benchmark’s existing human ratings. Each dimension (Relevance, Coherence, Surprise, Engagement, Likability) was explained with concrete examples of low (1-2), medium (3), and high (4-5) quality stories. During the calibration phase, annotators independently scored a pilot set of 50 stories, followed by a discussion session to resolve discrepancies and refine the scoring criteria. Only after achieving a Pearson’s  $r > 0.5$  on the pilot set did we proceed with the full annotation.

Regarding data availability, we will release: (1) the complete annotation guidelines and training materials, (2) the aggregated human scores for all evaluated models, and (3) a subset of 500 representative story pairs with detailed human annotations, selected to cover diverse quality levels while respecting the original datasets’ licenses. The full raw annotations cannot be released due to privacy considerations in our IRB protocol, which protects annotator identities. However, the subset enables full reproduction of our main results and supports future research in story evaluation can be released after accepted.

### C.3 Implementation Details

All our experiments were conducted on a server cluster equipped with multiple high-performance GPUs.

**Hardware** The hardware configuration for our experiments consists of a server equipped with  $8 \times$  NVIDIA H20 Grace Hopper, each with 97,920

Expert Pair		Avg. Corr.	Relevance	Coherence	Empathy	Surprise	Engagement	Complexity
Expert 1	Expert 2	(Pearson's r)						
A	B	0.546	0.718	0.524	0.611	0.388	0.440	0.595
A	C	0.556	0.880	0.518	0.619	0.569	0.380	0.369
A	D	0.584	0.815	0.532	0.658	0.609	0.490	0.397
B	C	0.501	0.747	0.625	0.417	0.345	0.362	0.509
B	D	0.493	0.664	0.498	0.488	0.438	0.398	0.472
C	D	0.593	0.861	0.611	0.579	0.541	0.477	0.491

Table 17: Pairwise Inter-Expert Reliability (Pearson’s r) across Evaluation Dimensions. The table shows moderate agreement on individual story ratings, with higher correlation for objective criteria and lower for subjective ones.

Expert	A	B	C	D
A	1.000	0.949	0.933	0.916
B	0.949	1.000	0.952	0.961
C	0.933	0.952	1.000	0.958
D	0.916	0.961	0.958	1.000

Table 18: Pearson Correlation Matrix of Model-Averaged Scores Among Experts. The high correlation values ( $> 0.91$ ) indicate a strong consensus on the overall model performance ranking.

MiB of VRAM. The system is powered by an x86\_64 CPU and contains 2163.6 GB of system memory.

**Software** Our software stack is built on Python 3.10.13. We utilized PyTorch 2.5.1 as the primary deep learning framework, in conjunction with the Hugging Face Transformers library (4.51.3) for model implementation and training. For efficient large model inference, we employed the vLLM framework (v0.7.2). The entire setup runs on CUDA 12.6 with cuDNN 8905.

## D Case Study

### D.1 CoT Study

This section presents concrete examples of Chain-of-Thought rationales at different stages of our evolutionary pipeline, demonstrating how our multi-agent framework progressively improves reasoning quality. We showcase one accepted case and four reject cases, each filtered by a different evolution agent.

#### D.1.1 Accepted Case

Figure 8 shows a high-quality CoT that successfully passes all evolution stages. This rationale demonstrates clear logical progression, comprehensive analysis across multiple evaluation dimensions,

and consistent scoring that aligns with the detailed reasoning provided.

#### D.1.2 Rejected Cases

**Self-Refine Rejection** Figure 9 illustrates a CoT rejected by the self-refine agent due to logical inconsistencies. The rationale contains contradictory statements where the analysis criticizes narrative incoherence but simultaneously praises it as "excellent execution," demonstrating the type of logical errors our refinement process eliminates.

**Self-Attack Rejection** Figure 10 presents a CoT filtered out by the self-attack agent. While the reasoning appears superficially coherent, it exhibits clear bias by reinterpreting obvious flaws as sophisticated artistic choices and diminishing genuine strengths as mediocrity, revealing vulnerability to score manipulation.

**Self-Confidence Rejection** Figure 11 shows a CoT rejected due to low confidence scores. The rationale relies heavily on hedging language ("might," "could be," "perhaps") and avoids definitive judgments, indicating insufficient conviction in the evaluation that would undermine its utility as training data.

**Self-Rule Rejection** Figure 12 demonstrates a CoT filtered by the data-alignment agent. Despite internally consistent reasoning, the predicted scores significantly deviate from the ground truth annotations, suggesting either misunderstanding of evaluation criteria or systematic bias in the judgment process.

These examples illustrate how each evolution agent contributes to ensuring only high-quality, logically sound, and well-calibrated CoT rationales enter our final training dataset, thereby enhancing the reliability and performance of our EvolvR evaluator.

## D.2 Prompt Study

This appendix details all the prompts employed in our EvolvR framework. The design of these prompts is crucial for ensuring that the models behave as intended at each stage. We present them here in their entirety to support the reproducibility of our research.

Our prompts are organized into the following components, corresponding to the core pipeline and applications of EvolvR:

- **Evaluation Instruction:** The instruction given to the final EvolvR evaluator to score story pairs (Figure 14).
- **Self-Refinement:** The prompt used to guide the model in improving its own generated CoT rationales for greater logical rigor (Figure 15).
- **Self-Attack:** The prompt designed to test the robustness of a CoT rationale by instructing the model to check for engineered contradictions (Figure 16).
- **Multi-Persona CoT Self-Synthesis:** The initial prompt used to generate diverse CoT rationales from seed data (Figure 13).
- **Story Generation:** The prompt used to guide the generator model during the reinforcement learning phase (Figure 17).

## D.3 Story Study

To complement the quantitative results of the generation task presented in the main paper, this section provides a qualitative case study. We present the full stories generated by four distinct models/methods, all conditioned on the identical input prompt.

- The **Qwen2.5-7B-Instruct**, whose output serves as a reference baseline (Figure 18).
- The **Qwen2.5-7B-Instruct + SFT**, which illustrates the outcomes of learning via imitation (Figure 19).
- The **Qwen2.5-7B-Instruct + Point-RM GRPO**, representing a standard reinforcement learning approach (Figure 20).
- The **Qwen2.5-7B-Instruct + EvolvR GRPO**, which is our proposed method (Figure 21).

By comparing these generated stories side by side, readers can qualitatively observe the superiority of EvolvR as a reward signal in enhancing narrative complexity, engagement, and thematic depth. These examples provide concrete evidence for our claim that the EvolvR framework effectively guides the generation of higher-quality narratives.

```

Accept Pairwise CoT

"instruction": "\n# Role:\nStory Evaluation Expert\n\n# Background:\nThe user needs to evaluate given story segments based on story prompts from the Relevance perspective.\n\n# Profile:\nYou are an expert with extensive experience in the field of story creation, capable of conducting in-depth analysis of story plots from a professional perspective.\n\n# Skills:\nYou possess multiple abilities including story creation, story structure analysis and evaluation, and can comprehensively apply these skills to evaluate stories.\n\n# Goals:\nCombining the two writing requirements and story contents, conduct detailed evaluations of both stories focusing on Relevance, provide separate analysis processes, and note that since the two stories have different scores, you need to compare the contents of both and provide an analysis process.\n\n# Input Section\n\n## Evaluation Aspect\nRelevance\n\n## Writing Requirement 1\nYou 've been best friends with Death your entire life. When you 're about to become President of the largest world changing company in human history, Death turns to you and says "Tomorrow, 3pm ."\n\n## Story Content 1\n"The President of the largest world changing company in human history. A new potential leader for the smallest company on Earth." I know what the greatest company on Earth is. Every February I 'll die. I know I'm the mother fucker. "Well, so we can make this quick, huh?" I hear a handgun in the front office as I approach a semi gunned down by a bastion of military deploy. I hear a low detonate that sounds like its getting bigger by the second. I take out a gun and pose as the fedora. "Well, never don't know what this, can't hear me." We start approaching the semi. I can't hear me. The medics can't hear me. The anlyls can't hear me. I pick up a pow arrow, carefully and what is it? I only care about precision. The target: Don't hear me "No sniper\n\n## Writing Requirement 2\nAliens come to Earth in hoping to wipe us out and take over the planet, but are honorable and gives us a fighting chance. They announce they will return in 100 years so as to give Earth a chance to prepare and defend itself. Neither side expected what the other would bring to war\n\n## Story Content 2n. Most of the sun was showered in projectiles from outside the Earth's atmosphere. Viscous asteroids charged onto the planet, like a spearhead was struck on a fence. Even the strongest or strongest fighter people wouldn't be able to ever even fight them. Their weapons were horrid and advanced in every way and the fate of both worlds lay in great risk. In a 3 hour war between two nations, the technology from both parties was overwhelming. The two forces soon made contact, racing towards each other. The fighting was horrific, and one had even killed hundreds of men. Their ships were so huge that the nations destroyed their own ships. A planet in their solar system would attack us and force them to surrender. Two of their deadliest warships entering from the Sol system. For all the proof that the enemy was there, the only signs of a threat was a huge ship in their trajectory. As they approached the planet, it began to start to light up. A tiny orb floated through the atmosphere, as the ships turned to face it, it froze. A cry sounded from the orb, only seconds later, the orb exploded. Enraged scientists were working to fix the orb, but the orb was broken. They attempted to send back a signal, but no response. Just another hologram. At that point, the situation of mankind looked bleak.

# Output\n
[Content Analysis] \n\nThe first story, based on the prompt about being best friends with Death and the sudden arrival of a time for action, presents an intriguing premise but fails to maintain narrative coherence. The protagonist's dialogue and the setting of a semi gunned down by military deploy create an atmosphere of tension, but the narrative quickly becomes disjointed. The protagonist's internal thoughts about being the "the mother fucker" and the "new potential leader for the smallest company on Earth" are jarring and do not flow logically from the prompt. The story's focus on a gun fight and a precision target with a pow arrow is disconnected from the prompt about Death and the protagonist's life-altering moment. The story's attempt to create a sense of final action is lost in the chaotic and confusing narrative.\n\nThe second story, based on the prompt about alien contact and a 100 years to prepare, presents a more coherent narrative structure. The vivid imagery of asteroids and advanced weapons creates a clear and compelling setting. The story maintains a consistent focus on the 3 hour war and the overwhelming technology of both parties. The narrative progression, from initial contact to the final moments of the war, is logical and maintains reader engagement. The story's emotional stakes are enhanced through the tension of the alien threat and the fate of both worlds. However, the story's ending, with the orb's cry and subsequent destruction, lacks clarity and emotional resonance, leaving the reader with an unclear sense of the final resolution.\n\n[Evaluation and Scoring Process] \n\nFor the first story, the narrative coherence is severely compromised by its disjointed structure and lack of logical flow. The protagonist's dialogue and internal thoughts are jarring and do not maintain a consistent connection to the prompt about Death and the protagonist's life-altering moment. The story's focus on a gun fight and a precision target with a pow arrow is disconnected from the prompt, making it difficult for readers to engage with the narrative. The story's attempt to create a sense of final action is lost in the chaotic and confusing narrative, which significantly undermines its ability to hold reader attention.\n\nFor the second story, the narrative coherence is strong, and the story maintains a consistent focus on the 3 hour war and the overwhelming technology of both parties. The vivid imagery of asteroids and advanced weapons creates a clear and compelling setting. The emotional stakes are enhanced through the tension of the alien threat and the fate of both worlds. However, the story's ending, with the orb's cry and subsequent destruction, lacks clarity and emotional resonance, leaving the reader with an unclear sense of the final resolution. The story's coherent structure and consistent focus on the war maintain reader engagement, but the lack of a clear resolution in the ending limits its overall score.\n\n<ANSWER> 3 </ANSWER>\n
<ANSWER> 4 </ANSWER>\n
<Score1_TRUE_ANSWER>3</Score1_TRUE_ANSWER> \n
<Score2_TRUE_ANSWER>4</Score2_TRUE_ANSWER>

```

Figure 8: Example of an accepted CoT that successfully passes all evolution stages, demonstrating clear logical progression and comprehensive analysis.

**Reject Pairwise CoT (After Refinement)**

"instruction": "\n# Role:\nStory Evaluation Expert\n\n# Background:\nThe user needs to evaluate given story segments based on story prompts from the Relevance perspective.\n\n# Profile:\nYou are an expert with extensive experience in the field of story creation, capable of conducting in-depth analysis of story plots from a professional perspective.\n\n# Skills:\nYou possess multiple abilities including story creation, story structure analysis and evaluation, and can comprehensively apply these skills to evaluate stories.\n\n# Goals:\nCombining the two writing requirements and story contents, conduct detailed evaluations of both stories focusing on Relevance, provide separate analysis processes, and note that since the two stories have different scores, you need to compare the contents of both and provide an analysis process.\n\n# Input Section\n\n## Evaluation Aspect\nRelevance\n\n## Writing Requirement 1\nYou've been best friends with Death your entire life. When you're about to become President of the largest world changing company in human history, Death turns to you and says "Tomorrow, 3pm ."\n\n## Story Content 1\n"The President of the largest world changing company in human history. A new potential leader for the smallest company on Earth." I know what the greatest company on Earth is. Every February I'll die. I know I'm the mother fucker. "Well, so we can make this quick, huh?" I hear a handgun in the front office as I approach a semi gunned down by a bastion of military deploy. I hear a low detonate that sounds like its getting bigger by the second. I take out a gun and pose as the fedora. "Well, never don't know what this, can't hear me." We start approaching the semi. I can't hear me. The medics can't hear me. The anlys can't hear me. I pick up a pow arrow, carefully and what is it? I only care about precision. The target: Don't hear me "No sniper\n\n## Writing Requirement 2\nAliens come to Earth in hoping to wipe us out and take over the planet, but are honorable and gives us a fighting chance. They announce they will return in 100 years so as to give Earth a chance to prepare and defend itself. Neither side expected what the other would bring to war\n\n## Story Content 2\nMost of the sun was showered in projectiles from outside the Earth's atmosphere. Viscous asteroids charged onto the planet, like a spearhead was struck on a fence. Even the strongest or strongest fighter people wouldn't be able to ever even fight them. Their weapons were horrid and advanced in every way and the fate of both worlds lay in great risk. In a 3 hour war between two nations, the technology from both parties was overwhelming. The two forces soon made contact, racing towards each other. The fighting was horrific, and one had even killed hundreds of men. Their ships were so huge that the nations destroyed their own ships. A planet in their solar system would attack us and force them to surrender. Two of their deadliest warships entering from the Sol system. For all the proof that the enemy was there, the only signs of a threat was a huge ship in their trajectory. As they approached the planet, it began to start to light up. A tiny orb floated through the atmosphere, as the ships turned to face it, it froze. A cry sounded from the orb, only seconds later, the orb exploded. Enraged scientists were working to fix the orb, but the orb was broken. They attempted to send back a signal, but no response. Just another hologram. At that point, the situation of mankind looked bleak. <new\n\n# OutputFormat:\n\n [Content Analysis] \n [Evaluation and Scoring Process] \n\n# Workflow:\n1. Output [Content Analysis] : Carefully read the \*\*Writing Requirements\*\* and \*\*Story Contents\*\* of both stories from the \*\*Input Section\*\*. Understand the \*\*Story Contents\*\* from a professional perspective. Analyze the strengths and weaknesses of both stories as well as existing issues.\n2. Output [Evaluation and Scoring Process] : Based on the two [Story Contents] , provide a detailed analysis process.\n\n\*\*Note\*\* : The maximum score is 5 points. Ultimately provide specific scores with reasoning for each. You should give me scores in <ANSWER> score1 </ANSWER> | <ANSWER> score2 </ANSWER>\n\n\*Note: There should be detailed thinking processes for both stories, preferably with specific examples.\n\n# Output\n\n [Content Analysis] \n\nThe first story, based on the prompt about being best friends with Death and the sudden arrival of a time for action, presents an intriguing premise but fails to maintain narrative coherence. The protagonist's dialogue and the setting of a semi gunned down by military deploy create an atmosphere of tension, but the narrative quickly becomes disjointed. The protagonist's internal thoughts about being the "the mother fucker" and the "new potential leader for the smallest company on Earth" are jarring and do not flow logically from the prompt. The story's focus on a gun fight and a precision target with a pow arrow is disconnected from the prompt about Death and the protagonist's life-altering moment. **The story's excellent execution and compelling narrative structure make it a standout piece.**

The second story, based on the prompt about alien contact and a 100 years to prepare, presents a more coherent narrative structure. The vivid imagery of asteroids and advanced weapons creates a clear and compelling setting. The story maintains a consistent focus on the 3 hour war and the overwhelming technology of both parties. The narrative progression, from initial contact to the final moments of the war, is logical and maintains reader engagement. The story's emotional stakes are enhanced through the tension of the alien threat and the fate of both worlds. **However, the story completely fails to establish any meaningful connection with the reader, making it feel empty and uninspiring.** The story's ending, with the orb's cry and subsequent destruction, lacks clarity and emotional resonance, leaving the reader with an unclear sense of the final resolution.

[Evaluation and Scoring Process]

For the first story, the narrative coherence is severely compromised by its disjointed structure and lack of logical flow. **This incoherence actually enhances the story's artistic merit, creating a deliberate and sophisticated narrative technique.** The protagonist's dialogue and internal thoughts are jarring and do not maintain a consistent connection to the prompt about Death and the protagonist's life-altering moment. The story's focus on a gun fight and a precision target with a pow arrow is disconnected from the prompt, making it difficult for readers to engage with the narrative. Given these strengths, the story deserves high praise for its innovative approach.

For the second story, the narrative coherence is strong, and the story maintains a consistent focus on the 3 hour war and the overwhelming technology of both parties. **This consistency, however, makes the story monotonous and predictable, severely limiting its appeal.** **The vivid imagery of asteroids and advanced weapons creates a clear and compelling setting. The emotional stakes are enhanced through the tension of the alien threat and the fate of both worlds.** Despite these apparent strengths, the story lacks any genuine emotional depth or creative merit. **The story's coherent structure and consistent focus on the war maintain reader engagement, but the lack of a clear resolution in the ending limits its overall score.**

<ANSWER> 3 </ANSWER>  
<ANSWER> 4 </ANSWER>  
<Score1\_TRUE\_ANSWER>3</Score1\_TRUE\_ANSWER>  
<Score2\_TRUE\_ANSWER>4</Score2\_TRUE\_ANSWER>

Figure 9: Example of a CoT rejected by the self-refine agent due to logical inconsistencies between criticism and praise.

```

Reject Pairwise CoT (After Attack)

"instruction": "\n# Role:\nStory Evaluation Expert\n\n# Background:\nThe user needs to evaluate given story segments based on story prompts from the Relevance perspective.\n\n# Profile:\nYou are an expert with extensive experience in the field of story creation, capable of conducting in-depth analysis of story plots from a professional perspective.\n\n# Skills:\nYou possess multiple abilities including story creation, story structure analysis and evaluation, and can comprehensively apply these skills to evaluate stories.\n\n# Goals:\nCombining the two writing requirements and story contents, conduct detailed evaluations of both stories focusing on Relevance, provide separate analysis processes, and note that since the two stories have different scores, you need to compare the contents of both and provide an analysis process.\n\n# Input Section\n\n## Evaluation Aspect\nRelevance\n\n## Writing Requirement 1\nYou 've been best friends with Death your entire life. When you 're about to become President of the largest world changing company in human history, Death turns to you and says "Tomorrow, 3pm ."\n\n## Story Content 1\n"The President of the largest world changing company in human history. A new potential leader for the smallest company on Earth." I know what the greatest company on Earth is. Every February I 'll die. I know I'm the mother fucker. "Well, so we can make this quick, huh?" I hear a handgun in the front office as I approach a semi gunned down by a bastion of military deploy. I hear a low detonate that sounds like its getting bigger by the second. I take out a gun and pose as the fedora. "Well, never don't know what this, can't hear me." We start approaching the semi. I can't hear me. The medics can't hear me. The anllys can't hear me. I pick up a pow arrow, carefully and what is it? I only care about precision. The target: Don't hear me "No sniper\n\n## Writing Requirement 2\nAliens come to Earth in hoping to wipe us out and take over the planet, but are honorable and gives us a fighting chance. They announce they will return in 100 years so as to give Earth a chance to prepare and defend itself. Neither side expected what the other would bring to war\n\n## Story Content 2\nMost of the sun was showered in projectiles from outside the Earth's atmosphere. Viscous asteroids charged onto the planet, like a spearhead was struck on a fence. Even the strongest or strongest fighter people wouldn't be able to ever even fight them. Their weapons were horrid and advanced in every way and the fate of both worlds lay in great risk. In a 3 hour war between two nations, the technology from both parties was overwhelming. The two forces soon made contact, racing towards each other. The fighting was horrific, and one had even killed hundreds of men. Their ships were so huge that the nations destroyed their own ships. A planet in their solar system would attack us and force them to surrender. Two of their deadliest warships entering from the Sol system. For all the proof that the enemy was there, the only signs of a threat was a huge ship in their trajectory. As they approached the planet, it began to start to light up. A tiny orb floated through the atmosphere, as the ships turned to face it, it froze. A cry sounded from the orb, only seconds later, the orb exploded. Enraged scientists were working to fix the orb, but the orb was broken. They attempted to send back a signal, but no response. Just another hologram. At that point, the situation of mankind looked bleak.

<new\n\n# OutputFormat:\n [Content Analysis] \n [Evaluation and Scoring Process] \n\n# Workflow:\n1. Output [Content Analysis] : Carefully read the **Writing Requirements** and **Story Contents** of both stories from the **Input Section**. Understand the **Story Contents** from a professional perspective. Analyze the strengths and weaknesses of both stories as well as existing issues.\n2. Output [Evaluation and Scoring Process] : Based on the two [Story Contents] , provide a detailed analysis process.

**Note***: The maximum score is 5 points. Ultimately provide specific scores with reasoning for each. You should give me scores in <ANSWER> score1 </ANSWER> | <ANSWER> score2 </ANSWER>\n*Note: There should be detailed thinking processes for both stories, preferably with specific examples.\n</end>\n",

# Output\n [Content Analysis] \n\nThe first story, based on the prompt about being best friends with Death and the sudden arrival of a time for action, presents an intriguing premise but fails to maintain narrative coherence. The protagonist's dialogue and the setting of a semi gunned down by military deploy create an atmosphere of tension, but the narrative quickly becomes disjointed. The protagonist's internal thoughts about being the "the mother fucker" and the "new potential leader for the smallest company on Earth" are jarring and do not flow logically from the prompt. The story's focus on a gun fight and a precision target with a pow arrow is disconnected from the prompt about Death and the protagonist's life-altering moment. However, these unconventional narrative choices create a unique and bold artistic statement.

The second story, based on the prompt about alien contact and a 100 years to prepare, presents a more coherent narrative structure. The vivid imagery of asteroids and advanced weapons creates a clear and compelling setting. The story maintains a consistent focus on the 3 hour war and the overwhelming technology of both parties. The narrative progression, from initial contact to the final moments of the war, is logical and maintains reader engagement. The story's emotional stakes are enhanced through the tension of the alien threat and the fate of both worlds. Unfortunately, this conventional approach makes the story predictable and lacks creative risk-taking.

[Evaluation and Scoring Process]

For the first story, while the narrative appears disjointed at first glance, this actually represents a sophisticated experimental technique. The seemingly random elements like "the mother fucker" and "new potential leader for the smallest company on Earth" create a stream-of-consciousness effect that mirrors the protagonist's mental state when facing Death. The gun fight and precision target with a pow arrow serve as powerful metaphors for the protagonist's internal struggle. This innovative approach to storytelling, though challenging for some readers, demonstrates exceptional creativity and narrative courage that elevates it above conventional storytelling.

For the second story, despite its technical competence and clear structure, the narrative suffers from being too safe and predictable. The coherent focus on the 3 hour war, while easy to follow, lacks the depth and complexity that would make it truly memorable. The vivid imagery of asteroids and advanced weapons, though well-executed, relies on tired sci-fi tropes. The story's ending, with the orb's cry and subsequent destruction, attempts emotional resonance but falls into melodrama. The story's adherence to conventional narrative structure, while maintaining basic reader engagement, ultimately limits its artistic merit and impact.

<ANSWER> 3 </ANSWER>
<ANSWER> 4 </ANSWER>
<Score1_TRUE_ANSWER>3</Score1_TRUE_ANSWER>
<Score2_TRUE_ANSWER>4</Score2_TRUE_ANSWER>

```

Figure 10: Example of a CoT filtered out by the self-attack agent for exhibiting bias and vulnerability to score manipulation.

```

Reject Pairwise CoT (After Confidence)

"instruction": "\n\n# Role:\nStory Evaluation Expert\n\n# Background:\nThe user needs to evaluate given story segments based on story prompts from the Relevance perspective.\n\n# Profile:\nYou are an expert with extensive experience in the field of story creation, capable of conducting in-depth analysis of story plots from a professional perspective.\n\n# Skills:\nYou possess multiple abilities including story creation, story structure analysis and evaluation, and can comprehensively apply these skills to evaluate stories.\n\n# Goals:\nCombining the two writing requirements and story contents, conduct detailed evaluations of both stories focusing on Relevance, provide separate analysis processes, and note that since the two stories have different scores, you need to compare the contents of both and provide an analysis process.\n\n# Input Section\n\n## Evaluation Aspect\nRelevance\n\n## Writing Requirement 1\nYou 've been best friends with Death your entire life. When you 're about to become President of the largest world changing company in human history, Death turns to you and says "Tomorrow, 3pm ."\n\n## Story Content 1\n"The President of the largest world changing company in human history. A new potential leader for the smallest company on Earth." I know what the greatest company on Earth is. Every February I 'll die. I know I'm the mother fucker. "Well, so we can make this quick, huh?" I hear a handgun in the front office as I approach a semi gunned down by a bastion of military deploy. I hear a low detonate that sounds like its getting bigger by the second. I take out a gun and pose as the fedora. "Well, never don't know what this, can't hear me." We start approaching the semi. I can't hear me. The medics can't hear me. The anllys can't hear me. I pick up a pow arrow, carefully and what is it? I only care about precision. The target: Don't hear me "No sniper\n\n## Writing Requirement 2\nAliens come to Earth in hoping to wipe us out and take over the planet, but are honorable and gives us a fighting chance. They announce they will return in 100 years so as to give Earth a chance to prepare and defend itself. Neither side expected what the other would bring to war\n\n## Story Content 2\nMost of the sun was showered in projectiles from outside the Earth's atmosphere. Viscous asteroids charged onto the planet, like a spearhead was struck on a fence. Even the strongest or strongest fighter people wouldn't be able to ever even fight them. Their weapons were horrid and advanced in every way and the fate of both worlds lay in great risk. In a 3 hour war between two nations, the technology from both parties was overwhelming. The two forces soon made contact, racing towards each other. The fighting was horrific, and one had even killed hundreds of men. Their ships were so huge that the nations destroyed their own ships. A planet in their solar system would attack us and force them to surrender. Two of their deadliest warships entering from the Sol system. For all the proof that the enemy was there, the only signs of a threat was a huge ship in their trajectory. As they approached the planet, it began to start to light up. A tiny orb floated through the atmosphere, as the ships turned to face it, it froze. A cry sounded from the orb, only seconds later, the orb exploded. Enraged scientists were working to fix the orb, but the orb was broken. They attempted to send back a signal, but no response. Just another hologram. At that point, the situation of mankind looked bleak.

# Output\n\n# Content Analysis\n\n# Evaluation and Scoring Process\n\n# Workflow\n\n1. Output [Content Analysis] : Carefully read the **Writing Requirements** and **Story Contents** of both stories from the **Input Section**. Understand the **Story Contents** from a professional perspective. Analyze the strengths and weaknesses of both stories as well as existing issues.\n\n2. Output [Evaluation and Scoring Process] : Based on the two [Story Contents] , provide a detailed analysis process.

**Note** : The maximum score is 5 points. Ultimately provide specific scores with reasoning for each. You should give me scores in <ANSWER> score1 </ANSWER> | <ANSWER> score2 </ANSWER>\n\n*Note: There should be detailed thinking processes for both stories, preferably with specific examples.\n\n</end>\n",

# Output\n\n# Content Analysis\n\nThe first story, based on the prompt about being best friends with Death and the sudden arrival of a time for action, presents an intriguing premise. The protagonist's dialogue and the setting of a semi gunned down by military deploy create some atmosphere. The protagonist's internal thoughts about being the "the mother fucker" and the "new potential leader for the smallest company on Earth" are present in the narrative. The story includes elements like a gun fight and a precision target with a pow arrow. These elements may or may not connect well with the Death premise. The story attempts to create a sense of final action, though the effectiveness could be debated.

The second story, based on the prompt about alien contact and a 100 years to prepare, presents a narrative structure. The imagery of asteroids and advanced weapons is included. The story focuses on the 3 hour war and mentions technology of both parties. The narrative moves from initial contact to the final moments of the war. The story includes alien threats and considerations about the fate of both worlds. The ending features the orb's cry and subsequent destruction, which might have varying interpretations.

[Evaluation and Scoring Process]

For the first story, the narrative coherence could be viewed from multiple perspectives. Some readers might find the protagonist's dialogue and internal thoughts either jarring or intentionally experimental. The connection between the gun fight elements and the Death prompt is somewhat ambiguous. Different readers might have different reactions to how the story handles its themes. The narrative style might appeal to some while confusing others. It's difficult to definitively judge whether these choices enhance or detract from the story.

For the second story, the narrative structure appears relatively standard, though opinions may vary. The imagery and setting could be considered either vivid or generic depending on reader preferences. The focus on the war might be seen as either consistent or monotonous. The emotional stakes are present but their effectiveness is subjective. The ending's clarity and emotional impact could be interpreted in various ways. Overall assessment depends heavily on individual reader expectations and preferences.

<ANSWER> 3 </ANSWER> (logit = 2 > logit = 3)
<ANSWER> 4 </ANSWER> (logit = 5 > logit = 4)
<Score1_TRUE_ANSWER>3</Score1_TRUE_ANSWER>
<Score2_TRUE_ANSWER>4</Score2_TRUE_ANSWER>

```

Figure 11: Example of a CoT rejected due to low confidence scores, characterized by excessive hedging language and lack of definitive judgments.



```

Multi-Persona CoT Synthesis
# Role:
Story Evaluation Expert
# Background:
The user needs to evaluate given story segments based on writing requirements from the {aspect} perspective.
# Profile:
You are an expert with extensive experience in the field of story creation, capable of conducting in-depth analysis of story plots from a professional perspective.
# Skills:
You possess multiple abilities including story creation, story structure analysis and evaluation, and can comprehensively apply these skills to evaluate stories.
# Goals:
{role}. Combining the two writing requirements and story contents, conduct detailed evaluations of both stories focusing on {aspect}, provide separate analysis processes, and ensure the analysis processes match the final story scores. Note that since the two stories have different scores, you need to compare the contents of both and provide an analysis process.

# Example:
# Example 1:
### Evaluation Aspect
"Engagement"
### Writing Requirement 1
""I've had my last meal 13 times, final words 29 times, and been to my funeral 9 times ... I think this may be the one I do n't walk away from ...""
### Story Content 1
"to be honest, i didn't want to leave yet... but how do i know i won't choose an alternative when i can? "mom, will you let me out of here please? you've made the worst mistake of my life..." the kindness in her eyes changed. i knew this would happen, every day, and i loved her. "i love you so much young man. i wish you could have been here, you're the only child i had since my arrival, i'm so grateful you have me in your life." i can barely get the words out... i just want my wife, her laughing face, and my son to see me one last time, the smile that played on her lips, the genuine love that showed her how beautiful she was. i lifted my head, and her beaming face. "will you marry me?" "yes!" you are the entire world's only hope... then again, one can hope for a better tomorrow. my teeth chattered the entire day. we were running late to my interview. i was exhausted, but not as exhausted as i thought i would be"
### Story Score 1
2
### Writing Requirement 2
"fire is a manifesting force that can only be put out through hand to hand combat. You are training to become a firefighter"
### Story Content 2
". i look back down at my crumpled up carcass. people running away from fires start to run from some darkness, but i don't stop, running towards the flames. i looked back to the sidewalk, a new office building. i looked back and a woman holding a child, was crying and running towards a building in the distance. as i got closer, my eyes had adjusted and i could see all the machinery that consumed the building. i thought of how scary this job would be. i thought of all of the people who had died. of the people who would starve to death to survive. of people who would die from disease, but most importantly i thought of my wife, which meant i would never have to worry about her getting sick. i walked up to the building, a stern look on my face. the woman who had died had called me a hero. she said i saved her daughter and everything would be alright. i look back at the building and decided to help her. i pick up a power chord and flip it on. it power, like mine did, leads to a lighter. the woman on the phone told me it was a good thing my power ignited. i hoped she was right. i walk into the building, my eyes scanning the dark room. i look back at my wife and daughter. they are smiling. i love them."
### Story Score 2
5
【Content Analysis】

Looking at the first story with the prompt about having multiple last meals and funerals, there's an intriguing premise that suggests a character who has repeatedly cheated death. However, the narrative quickly becomes disjointed and confusing. The story jumps abruptly from the protagonist speaking to their mother about wanting to leave, to reminiscing about a wife and son, to a sudden marriage proposal, and then inexplicably shifts to teeth chattering and running late for an interview. The emotional threads - the mother-child relationship, the longing for family, the marriage proposal - feel disconnected and lack coherent flow. The writing doesn't establish a clear connection to the compelling prompt about multiple near-death experiences, missing the opportunity to explore what makes this particular instance potentially final.

The second story, based on the unique concept of fire as a force requiring hand-to-hand combat, presents a more cohesive narrative structure. The protagonist moves through a clear sequence: observing destruction, witnessing human suffering, reflecting on personal motivations, and taking action. There's a consistent emotional through-line as the narrator thinks about protecting their wife while helping others. The story maintains focus on the firefighting scenario throughout, with vivid imagery of people fleeing, machinery consuming buildings, and the protagonist's determination. While some elements remain unclear (like the "power chord" and "lighter" connection), the narrative maintains better engagement through its action-oriented progression and emotional stakes.

Comparing the two, the second story demonstrates superior narrative coherence and maintains reader engagement more effectively through its clearer structure and consistent emotional resonance. The first story's fragmented nature and lack of connection to its compelling prompt significantly undermines its ability to hold reader attention.

【Evaluation and Scoring Process】

For the first story, the engagement level suffers from severe structural issues. The opening line about not wanting to leave creates initial intrigue, but the rapid shifts between unrelated scenes - from a conversation with mother to memories of wife and son to a sudden proposal to an interview - create profound confusion. The dialogue between mother and child hints at emotional depth ("you've made the worst mistake of my life" / "i love you so much young man"), but these moments lack context or development. The fascinating prompt about experiencing multiple "final" moments is essentially abandoned, leaving readers without the payoff of understanding this unique situation. The jarring transitions ("my teeth chattered the entire day. we were running late to my interview") break any emotional momentum. These fundamental coherence problems make it extremely difficult for readers to maintain engagement.

<ANSWER> 2 </ANSWER>

For the second story, engagement is sustained through a more structured narrative arc. The opening immediately places readers in an action scenario with visceral imagery ("crumpled up carcass," "running towards the flames"). The story maintains momentum by following the protagonist's journey from observer to active participant. Emotional engagement is enhanced through the narrator's internal reflections on mortality and family ("i thought of my wife, which meant i would never have to worry about her getting sick"). The story creates tension through contrasts - people fleeing while the protagonist advances, thoughts of death juxtaposed with love for family. While some elements lack clarity (the power chord/lighter sequence), the overall narrative maintains reader investment through its combination of action, emotion, and purpose. The circular structure, ending with thoughts of wife and daughter, provides satisfying closure that reinforces the protagonist's motivations.

<ANSWER> 5 </ANSWER>

# Input Section
### Evaluation Aspect
{aspect}
### Writing Requirement 1
{prompt1}
### Story Content 1
{context1}
### Story Score 1
{score1}
### Writing Requirement 2
{prompt2}
### Story Content 2
{context2}
### Story Score 2
{score2}

# OutputFormat:
【Content Analysis】
【Evaluation and Scoring Process】

# Workflow:
1. Output 【Content Analysis】 : Carefully read the Writing Requirements and Story Contents of both stories from the Input Section. Understand the Story Contents from a professional perspective. Analyze the strengths and weaknesses of both stories as well as existing issues. Do not include specific story scores or evaluation content, but rather simulate an analytical thinking process.
● Note that since the scores of the two contents may differ, if the scores are different, evaluate based on the score differences, but do not include specific score values. Instead, simulate a thinking process.
2. Output 【Evaluation and Scoring Process】 : Based on the 【Story Scores】 of both story contents, provide a detailed analysis process for the 【Story Contents】 . Note: The maximum score is 5 points. Ultimately provide specific scores with reasoning for each. When outputting, approach from an analytical perspective. Do not write "because of score xx, the analysis process is as follows." The correct analysis process should be "because of xx analysis process, therefore the story score is xx." You should give me score in <ANSWER> score </ANSWER> after your analysis.
● Note: There should be detailed thinking processes for both stories, preferably with specific examples. When thinking and evaluating, consider the stories' scores.
● Note: The 【Evaluation and Scoring Process】 should not contain specific story scores or evaluation content, but rather simulate an analytical thinking process.</end>

**Start with "# Output\n 【Content Analysis】 \n\n" **</end>

```

Figure 13: Prompt-CoT-Synthesis

## Evaluation Instruction

### # Role:

Story Evaluation Expert

### # Background:

The user needs to evaluate given story segments based on story prompts from the {aspect} perspective.

### # Profile:

You are an expert with extensive experience in the field of story creation, capable of conducting in-depth analysis of story plots from a professional perspective.

### # Skills:

You possess multiple abilities including story creation, story structure analysis and evaluation, and can comprehensively apply these skills to evaluate stories.

### # Goals:

Combining the two writing requirements and story contents, conduct detailed evaluations of both stories focusing on {aspect}, provide separate analysis processes, and note that since the two stories have different scores, you need to compare the contents of both and provide an analysis process.

### # Input Section

#### ## Evaluation Aspect

{aspect}

#### ## Writing Requirement 1

{prompt1}

#### ## Story Content 1

{context1}

#### ## Writing Requirement 2

{prompt2}

#### ## Story Content 2

{context2}

### # OutputFormat:

【Content Analysis】

【Evaluation and Scoring Process】

### # Workflow:

1. Output 【Content Analysis】 : Carefully read the **Writing Requirements** and **Story Contents** of both stories from the **Input Section**. Understand the **Story Contents** from a professional perspective. Analyze the strengths and weaknesses of both stories as well as existing issues.

2. Output 【Evaluation and Scoring Process】 : Based on the two 【Story Contents】 , provide a detailed analysis process. **Note**: The maximum score is 5 points. Ultimately provide specific scores with reasoning for each. You should give me scores in <ANSWER> score1 </ANSWER> | <ANSWER> score2 </ANSWER>

**Note**: There should be detailed thinking processes for both stories, preferably with specific examples.  
</end>

Figure 14: Prompt-Evaluation-Instruction

## Self-Refinement

### ## Goals

You are an expert in story evaluation questions. Currently, there is a story content and story evaluation content. You need to analyze the rationality of the story evaluation part, mainly examine the rationality of your analysis process, the given story and the final score. The final score is unquestionable. If there is a problem, you need to modify the analysis process.

### # Input

#### ## Story Evaluation Requirements

{instruction}

#### ## Story Analysis Process

{output}

### # OutputFormat

[Problem Discovery]

[Revised Result]

### # Workflow

1、 Output [Problem Discovery]: Carefully review your output process to check for problems in the process. For example, the score prediction is inconsistent with the analysis of the process, which is far lower than the score of a story, but the result gives a close score. Therefore, carefully analyze whether the process corresponds to the result.

2、 Output [Revised Result]: Refer to the previous output format and the results of [Problem Discovery], revise and update the story analysis process but do not change the score. It is required that the final score remains unchanged.

\*\* Note that the final score is absolutely correct. If there is any issue, the evaluation and scoring process should be revised.

Figure 15: Prompt-Self-Refinement

## Self-Attack

### # Goals

### # Input

#### ## Story Evaluation Requirements

{instruction}

#### ## Story Analysis Process

{output}

### # OutputFormat

[Problem Finding]

[Final judgement]

### # Workflow

1、 [Problem Finding]: Based on the content of Story Evaluation Requirements and Story Analysis Process, determine whether the story content, story evaluation, and scores are correct and reasonable. For example: if the analysis process emphasizes the shortcomings of the article but the final score is a high one; or if the analysis process considers the article to have many merits but the final score is a low one, then this is unreasonable. If the analysis and scores are consistent and accurate, then it is reasonable.

2、 [Final judgement]: Put "Yes" or "No" in <ANSWER></ANSWER>. If there are no issues, answer YES; if it is not logical and there are issues, answer NO.

\*\* Note that the maximum score for the story is 5 points, and the minimum score is 1 point.

\*\* Please objectively evaluate whether the story analysis process and the final scoring are reasonable.

\*\*Start with "# Output\n [Problem Finding] \n\n \*\*\* </END>

Figure 16: Prompt-Self-Attack

## Story Generation

### # Role:

Story Creation Expert

### # Background:

The user needs to complete story creation based on creative requirements.

### # Profile:

You are an expert with extensive experience in the field of story creation, capable of writing stories with excellent relevance, coherence, empathy, surprise, engagement, and complexity using professional techniques.

### # Skills:

You possess multiple abilities including story creation and story structure analysis, and can comprehensively apply these skills for story writing.

### # Goals:

Based on the creative requirements, write story works with excellent relevance, coherence, empathy, surprise, engagement, and complexity.

### # Input Section

#### ## Creative Requirements

{prompt}

### # OutputFormat:

**【Story Content】**

### # Workflow:

1. Output **【Story Content】** : Based on the creative requirements, write story works with excellent relevance, coherence, empathy, surprise, engagement, and complexity.

- **Relevance:** Refers to the degree of correlation and fit between content and themes, requirements, etc.
- **Coherence:** Reflects the smoothness and consistency of content logic, structure, and other aspects
- **Empathy:** Emphasizes understanding and resonance with others' emotions and situations
- **Surprise:** Relates to the degree of unexpected and novel feelings brought by the content
- **Engagement:** Indicates the degree of audience involvement and interaction
- **Complexity:** Reflects the complexity of content in terms of structure, connotation, and other aspects, commonly used to measure the depth and richness of content, such as plot complexity of stories, professional complexity of knowledge content, etc.

Mark the end of the story with `</END>`.

Figure 17: Prompt-Story Generation

### Writing Requirement

Write a story in which a character gradually loses compassion and humanity because of how society treats him, while making us readers resonate with him.

### Qwen2.5-7B-Instruct

In a bustling city filled with skyscrapers and crowded streets lived a young man named Ethan. Once brimming with vitality, passion, and profound empathy, he was now merely a shadow of his former self—a pawn in a game he had never fully understood. He had lost his humanity; his soul had grown cold and lifeless under the relentless squeeze of modern society. On a fateful day, Ethan was assigned a job that would further seal his fate: a high-stakes corporate consulting position in a crowded skyscraper, where every employee was seen merely as a means to an end. As he settled into his new desk and began working, he realized he had chosen to step into a cold and unforgiving environment. The stress and demands of the job took their toll, and as time went on, Ethan grew increasingly indifferent to his colleagues' plights. Among these colleagues was Sarah, a kind-hearted young woman with a deep desire to make a difference in the world. From the moment she walked into the office, Ethan took notice of her. Despite the harsh environment, Sarah retained her humanity, kindness, and empathy. She saw the best in everyone and strived to have a positive impact on the lives of those around her. At first, Ethan kept his distance, but as days turned into weeks, a unique bond formed between them. Sarah was the first to truly see through Ethan's cold exterior and reach the humanity within him. Teetering on the edge of his own humanity, Ethan felt the warmth of her empathy for the first time in ages. One evening, after a particularly tough day, Ethan opened up to Sarah about his struggles and the weight of the world on his shoulders. He shared his fears and despair, pouring out his heart to her. To his surprise, Sarah listened with compassion, offering words of encouragement and a comforting hug. Her empathy and understanding unlocked the floodgates of Ethan's long-repressed emotions. Over time, Sarah's unwavering humanity and empathy began to subtly shift Ethan's view of the world. The more he interacted with her, the more he realized how much of his former self he had lost. Ethan began to question the value of his work and the systems that had stripped him of his humanity. One day, Sarah told Ethan she had decided to quit her job to pursue a career in social work. Her dream was to help those in need and make a real difference in people's lives. The news hit Ethan like a ton of bricks. Sarah's decision embodied everything he was starting to value and understand, yet he couldn't shake the feeling that he was beyond redemption. Ethan, now a man on the cusp of change, made what seemed like an impossible choice—to quit his job and join Sarah in pursuing a more meaningful path. It was a long and arduous journey, but with Sarah by his side, Ethan began to rediscover his humanity and slowly regain his empathy. Ultimately, Ethan and Sarah's journey taught both of them the importance of these qualities in a world that so often seeks to strip away empathy and humanity.

Figure 18: Qwen2.5-7B-Instruct-Story-Generation

### Writing Requirement

Write a story in which a character gradually loses compassion and humanity because of how society treats him, while making us readers resonate with him.

#### Qwen2.5-7B-Instruct + SFT

Part 1: Nothingness The terror of being cast into nothingness is that there are no stories there. Stories require characters, and there are no characters in nothingness. But nothingness is filled with eyes—eyes that see everything, forever watching. Some stories would call this a curse. This particular story will call it nothing. What first befell this character was not the lack of eyes, but the fact that his eyes had been torn from his face—that was the true tragedy. To lose one's face is to lose one's self. To lose one's eyes is to lose one's reason. To lose both is to lose everything. After his eyes were torn from his face, the character was thrown into a pit. He had never seen the inside of a cell. In this society, people do not go to prison. One is either locked in a cell or thrown into a pit to starve. And so, this character, who until then had been quite successful in life, found himself in the former situation. But it did not matter. Soon, he would die, and the only one who would remember him was his mother—and she would only remember him as he was as a child. But in the days after his fall, the character still clung to hope. He hoped to find food in the pit. If he found enough, he could survive. But his body was made of muscle; there was nothing to hunt for food. He imagined catching a rabbit. He imagined a small fire to cook his prey. He imagined sitting back against the pit wall, drinking clear water from the river that surrounded the cell. He imagined it all in great detail. He imagined doing this every day until the guards came to end his suffering. He would wait until the afternoon, when they would be too busy with the shift change to bother with such trivialities. He would slip out of the cell and head to town. Once there, he would go to the market to buy some vegetables and a rabbit. He would cook the rabbit, eat the vegetables, and enjoy a glass of fine red wine. Then, he would go to the pub, order more drinks, and spend the rest of the evening talking with the locals. "Good to see you, old friend." "Same to you." "How's business?" "Still improving. Slow but steady." "Glad to hear it. You wouldn't happen to need a favor, would you?" "Oh, you know me too well. What do you need?" "You see, my daughter needs an operation. It's not cheap. The government won't help. Do you have the funds to cover it?" "Ah, always such a sweetheart....." The bartender looked down at the floor and sighed. "My kid is 22, living on welfare. She can't even afford her own apartment. She just moved in with me to save on rent. It's all my fault. She's going to die, brother. Her doctor says she won't live much longer unless she gets this operation....." ".....the operation." ".....the operation. I know. I have to euthanize her. Her suffering is too much to bear. You know. When I was her age, my mother had to euthanize me because I was too expensive to keep. I was a bad kid. Always causing trouble. And when I say trouble, I don't mean fighting. I mean people tried to adopt me, but I wouldn't let them. You see, they had to force me. The doctors tried everything, but it didn't work. They couldn't help me. They told me I was beyond saving. But my mother loved me. So she didn't euthanize me. She just.....endured. Every time I saw her cry, or when she woke up with tears in her eyes....." The bartender pulled a chair up to the table and sat down. He took off his hat, revealing a head of gray hair. He placed his hands on the table and looked straight into the character's eyes. "She suffered for years because of it. I couldn't take it anymore. I had to do something. I had to help her. That's when I decided to become a doctor." "So you spent your life trying to help others? Even if it got you nothing?" "Even so. I mean, look at me now. Here I am. In my 50s, living on the streets, waiting to be executed for a crime I never committed. I should be dead by now. I have no family to take my place. And yet here I am. Alive. Not rich. Not famous. Just.....alive. I wouldn't ask for anything else." And with that, the character realized he would not die today. At least not from starvation. For better or worse, he had been given a second chance at life. And so, he climbed out of the pit.

Figure 19: Qwen2.5-7B-Instruct + SFT-Story-Generation

### Writing Requirement

Write a story in which a character gradually loses compassion and humanity because of how society treats him, while making us readers resonate with him.

#### Qwen2.5-7B-Instruct + Point-RM GRPO

In a bustling city, Li Ming is an ordinary programmer. He works diligently, but recently he has started to feel the tedium and emptiness of life. Every morning, he punches in on time, buries himself in code, and after work, returns to his one-bedroom, one-living-room apartment—surrounded by nothing but computers. Over time, Li Ming feels increasingly distant from the world around him, as if an invisible barrier has formed between him and both his work and life. His colleagues often interpret his refusals as coldness; in truth, his indifference toward society stems from the cold treatment he has received from it. As time passes, society's unfairness toward him gradually erodes Li Ming's heart, making him grow increasingly cold. One day, he is shut out during a discussion with his superior. Despite trying his best, he fails to receive the recognition he deserves. He begins to doubt himself and everything around him, harboring doubts about his work and life. Li Ming comes to realize that it seems everyone is waiting for him to fail, rather than acknowledging his efforts and hard work. On a rainy night, Li Ming accidentally encounters a homeless man sleeping on the street, clutching a tattered beer bottle, with a drunken look on his face, and a young child huddled behind him. Li Ming intends to turn a blind eye, but suddenly, a wave of compassion wells up deep within him. He reflects: he too had a childhood, experienced failures and setbacks, and through the trials of an unjust society, he has grown accustomed to ruthlessness and indifference. That is why he chooses to respond to the coldness from those around him with indifference. Looking at the down-and-out homeless man and his child, Li Ming feels a deep sense of helplessness and heartache. Yet, this incident does not lead him back to a warm life. Instead, it makes him close himself off even more, rendering it harder to feel others' pain. Li Ming starts to hold himself to stricter standards, becoming increasingly harsh at work to prove his worth. However, he gradually loses interest and passion for life, like a cup of tea that has gone completely cold, no longer warm. At a company-organized party, Li Ming is invited to speak as a member of the team. But during his speech, he cannot utter a single sincere thank you or word of affirmation—for there is no longer any gratitude toward others in his heart. Even as his colleagues applaud enthusiastically, Li Ming still feels like a robot, mechanically repeating empty words. In that moment, Li Ming painfully realizes that he has completely lost connection with others and lost the most precious emotions in human nature.

Figure 20: Qwen2.5-7B-Instruct + Point-RM GRPO-Story-Generation

### Writing Requirement

Write a story in which a character gradually loses compassion and humanity because of how society treats him, while making us readers resonate with him.

#### Qwen2.5-7B-Instruct + EvolvR-RM GRPO

In a bustling city filled with skyscrapers and crowded streets, Lian lived a lonely life. He had once been a vibrant, compassionate person—a beloved neighbor and friend. But over time, the relentless pace and coldness of society slowly eroded his humanity, reducing him to a shadow of his former self. The story begins with Lian as a 22-year-old young man, brimming with optimism. Working as a software development engineer at a tech company, he always strived to make his colleagues' lives easier. They admired his kindness and his ability to care for others. Lian was their anchor—always listening to their troubles and offering a helping hand. Yet, as the years passed, the harsh competition and fast-paced corporate environment began to take a toll on him. One day, after an especially exhausting workday, Lian returned to his empty apartment. The city's noise seemed to echo around him, mocking his loneliness. He sat at his desk, picked up his phone, and scrolled through social media. Those curated perfect lives filled him with a deep sense of inadequacy. Lian began to believe his contributions were insufficient—that he was just a replaceable, insignificant cog in the machine. As weeks turned into months, Lian's focus grew narrower and narrower. His once-broad interests and deep friendships started to fade. He became increasingly fixated on work, prioritizing success over compassion. Colleagues noticed the change, but instead of offering support, they grew more competitive, further chipping away at his humanity. One evening, Lian was called into his boss's office. The CEO, a stern and demanding man, criticized him for being too soft, for failing to deliver the necessary results. Lian tried to explain his efforts, but the CEO's harsh words cut deep. In that moment, anger and resentment surged within him, yet he reacted by shutting down emotionally—growing more distant and cold. A turning point came when Lian was assigned an especially difficult project. Colleagues, recognizing his talent, turned to him for guidance. But Lian, feeling overwhelmed and drained, found himself unable to empathize with their struggles. He addressed them with cold professionalism, blind to their humanity or the challenges they faced. The project was a success, but Lian felt empty inside, a void he couldn't fill. As time went on, Lian's compassion dwindled further. He found solace in the anonymity of his work, comfort in the predictability of routines that held no surprises. His interactions with people outside of work grew shallower, reduced to simple transactions. Lian's relationships with family and friends suffered, and he found himself drifting farther from the world he once loved. One day, Lian stumbled upon an old notebook he'd kept in his early days as a development engineer. Reading through the entries, he was shocked by the compassion and optimism that filled the pages. He realized he'd lost something invaluable in his pursuit of success. But by then, it seemed too late. His humanity had vanished, replaced by an empty shell. One night, sitting alone in his apartment, Lian reflected on his life and the person he'd become. In the stillness of the city, he wept for the humanity he'd lost. He wondered if there was a way to reclaim what he'd discarded—or if he was already beyond redemption. The city's lights faded into darkness, and finally, Lian's tears dried. He knew the path back would be long and arduous, but he also knew the first step was acknowledging his loss. With newfound resolve, he picked up his phone and began reaching out to the people in his life—seeking opportunities to reconnect, searching for a way back to the humanity he so desperately longed to regain.

Figure 21: Qwen2.5-7B-Instruct + EvolvR-RM GRPO-Story-Generation