StoryLens: Preference-Aligned Story Rewriting via Context-Aware Narrative Enrichment

Anonymous ACL submission

Abstract

Story rewriting aims to adapt existing narratives to diverse reader preferences while preserving plot consistency and narrative coherence. Unlike conventional work on style transfer, we argue that effective story rewriting demands context-aware narrative enrichment beyond surface-level stylistic adaptation. Our pilot human study shows that style adaptation alone provides only marginal gains in reader satisfaction (2.3%), while context-enhanced rewriting substantially improves user preference alignment (24.5%). Motivated by this, we introduce STORYLENSBENCH, a large-scale benchmark for preference-aligned story rewriting, comprising structured story books, multi-dimensional reader preference profiles, and ranked context-aware rewritten stories. Building on this benchmark, we propose STORYLENSEVAL, a reward model for estimating reader satisfaction over rewritten stories, and STORYLENSWRITER, a two-stage rewriting model combining supervised fine-tuning with GRPO-based reinforcement learning. We further establish a comprehensive evaluation framework covering fidelity, coherence, and reader satisfaction. Experimental results demonstrate that STORYLENSWRITER consistently outperforms strong generation and personalization baselines, highlighting the importance of context-aware narrative enrichment for personalized story rewriting. [1]

1 Introduction

Story rewriting aims to adapt existing narratives to new audiences, communicative goals, and reader preferences while preserving the core plot, character consistency, and narrative coherence of the original story. Such capability is increasingly important for applications including personalized reading assistants, interactive storytelling, literary simplification, game narrative adaptation, and culturally localized content generation (Nasir et al., 2024; Huynh and McNamara, 2025; Singh et al., 2024). Compared with open-ended story generation, story rewriting imposes a unique constraint: models must simultaneously preserve existing narrative structures while flexibly modifying the story according to user-specific preferences.

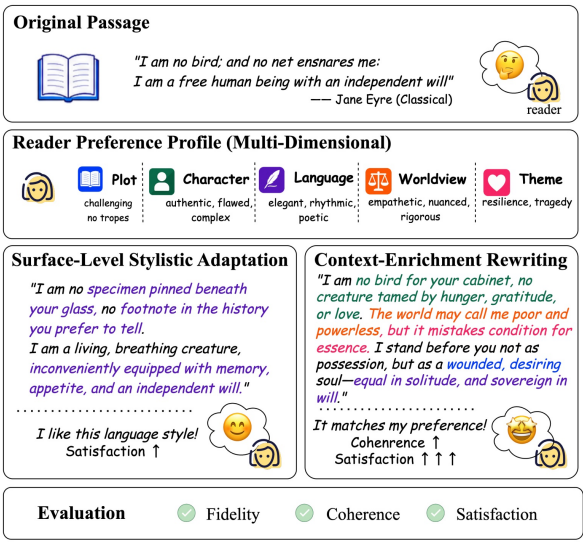

Figure 1: Overview of our story rewriting task. Given an original passage and a multi-dimensional reader preference profile, the model generates a rewritten story aligned with the target preferences while preserving the core narrative structure. Unlike conventional style adaptation, our framework incorporates *context-aware narrative enrichment*, which is shown through pilot studies to substantially improve reader satisfaction. The rewritten stories are evaluated along three dimensions: fidelity, coherence, and satisfaction.

Despite recent advances in story generation (Fan et al., 2018; Yao et al., 2019; Yang et al., 2022b,a; Bae and Kim, 2024; Park et al., 2025), personalized preference alignment remains largely underexplored. Existing controllable generation and style transfer methods mainly address sentence-level stylistic modification, such as sentiment, for-

[1] Code, data and model checkpoints will be released upon acceptance.

mality, or writing style adaptation (Liu et al., 2022; Lai et al., 2024; Liu and May, 2025). However, effective story rewriting for preference alignment goes well beyond surface-level stylistic transformation. To satisfy diverse reader preferences, models must perform *context-aware narrative enrichment*, including expanding scene descriptions, refining emotional expression, enhancing character interactions, and elaborating narrative details, while maintaining global story consistency. This makes story rewriting fundamentally different from conventional text editing or style transfer tasks.

We hypothesize that stylistic adaptation alone is insufficient for high-quality story rewriting. Instead, reader satisfaction depends heavily on whether the rewritten story can provide contextually appropriate narrative enrichment beyond simple lexical or stylistic changes. To verify this hypothesis, we conduct a pilot human study involving 10 participants across 8 books. Participants evaluate original stories, style-adapted rewrites, and context-enhanced rewrites on three dimensions: satisfaction, fidelity, and coherence. The results reveal that style adaptation alone provides only marginal improvement in reader satisfaction ($3.43 \rightarrow 3.51$), whereas context-enhanced rewriting raises satisfaction to 4.27, with only minor reductions in fidelity and coherence. These findings suggest that effective story rewriting requires modeling user preference and preference-aligned contextual enrichment.

To enable systematic study of this problem, we introduce **StoryLensBench**, a large-scale benchmark for preference-aligned story rewriting. The benchmark consists of chapter-level source stories, structured narrative contexts, multi-dimensional reader preference profiles, and ranked rewritten pairs. Unlike prior synthetic persona datasets (Jandaghi et al., 2024; Ge et al., 2024), StoryLensBench leverages real reader review histories collected from Goodreads[2], enabling more realistic and context-aware preference modeling for long-form narratives. Our benchmark supports the study of personalized rewriting under realistic constraints where models must balance narrative fidelity, coherence, and user satisfaction simultaneously.

Building upon this benchmark, we further propose a rewriting framework consisting of two components: **StoryLensEval** and **StoryLensWriter**. StoryLensEval is a reward model designed to estimate reader satisfaction over rewritten stories conditioned on user preference profiles. Instead of relying on absolute quality scoring, we formulate evaluation as pairwise preference ranking to better capture subjective user preferences in narrative rewriting. Based on this evaluator, we develop StoryLensWriter, a preference-aligned rewriting model optimized through supervised fine-tuning followed by reinforcement learning with GRPO (Shao et al., 2024). The model learns to generate rewrites that not only preserve narrative consistency but also perform contextually appropriate narrative enrichment aligned with reader preferences.

We further establish a comprehensive evaluation protocol for story rewriting across three key dimensions: *fidelity*, *coherence*, and *satisfaction*. Fidelity measures whether rewritten stories preserve the core narrative content of the original story. Coherence evaluates narrative consistency and writing fluency across long contexts. Satisfaction measures how well rewrites align with user preferences through pairwise comparison-based evaluation.

In summary, our contributions are threefold:

- We formulate personalized story rewriting as a context-aware narrative adaptation task that extends beyond conventional style transfer and requires creative narrative enrichment for improving reader satisfaction.
- We introduce StoryLensBench, the first large-scale benchmark for preference-aligned long-form story rewriting with real reader preference profiles and ranked rewritten stories.
- We propose StoryLensEval and StoryLensWriter, enabling preference-aware evaluation and context-enhanced rewriting generation through supervised fine-tuning and reinforcement learning.

2 Related Work

Story Generation. Recent Large Language Models (LLMs) show strong abilities in creative writing and story generation (Zhao et al., 2023), with prior work exploring incremental generation (Yang et al., 2022b), hierarchical planning (Yang et al., 2022a; Wang et al., 2024a), and multi-agent collaboration (Xia et al., 2025; Yu et al., 2025). However, these methods mostly generate stories for general audiences, whereas our work focuses on adapting existing stories to reader-specific preferences while

[2] https://www.goodreads.com/

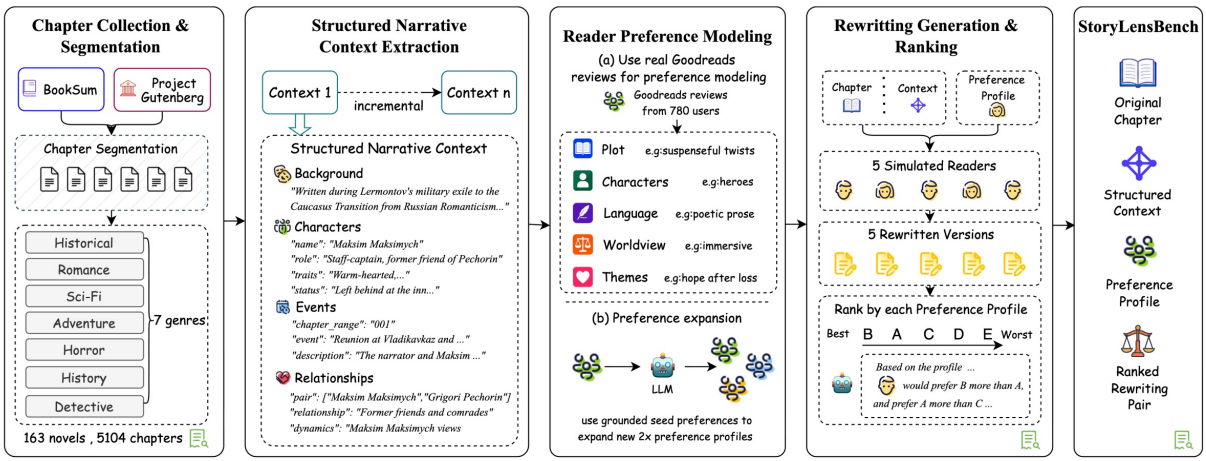

Figure 2: Overview of the dataset construction process. The green document icon in the bottom right corner indicates that this step has been manually checked.

preserving narrative consistency.

Personalization of LLMs and Evaluation. Personalization of LLMs adapts outputs to user-specific contexts, profiles, or feedback. Existing approaches include prompt engineering, RAG-based methods, representation learning methods, and RLHF-based methods (Zhang et al., 2024). Prompting engineering directly encodes personas or profiles in instructions of LLMs (Jiang et al., 2023; Lyu et al., 2024). RAG-based methods further retrieve user histories or external memories (Gao et al., 2023; Salemi et al., 2024) into instructions to augment text generation. Representation learning methods compress user information into learned embeddings (Liu et al., 2025; Ning et al., 2025) for LLMs. The aforementioned methods confine personalization modeling to the LLM's input stage, lacking the capacity to continuously optimize the model based on feedback. In contrast, RLHF-based methods optimize LLMs with preference feedback (Ouyang et al., 2022; Li et al., 2024; Lee et al., 2024), but such dense rewards are costly for long-form rewriting. To address this challenge, we model real reader preferences as multi-dimensional training signals and learn a preference reward model for story rewriting.

Evaluation is also non-trivial for story rewriting. Existing metrics of story generation often focus on either fidelity or style: fidelity-oriented methods evaluate factual consistency or plot preservation (Wang et al., 2020; Chang et al., 2023; Kim et al., 2024), while style-oriented methods measure linguistic similarity, persona adherence, or personalized preference alignment (Wegmann and Nguyen, 2021; Lin and Chen, 2023; Wang et al., 2024b; Salemi et al., 2025). However, these standard metrics are often too rigid for rewriting, as they penalize reasonable contextual enrichment or stylistic re-expression. Unlike these metrics, StoryLensEval estimates reader satisfaction given a preference profile, rather than measuring surface-level style similarity or aspect-level alignment.

3 StoryLensBench

3.1 Benchmark Construction

StoryLensBench consists of four components: original novel chapters, structured contexts, reader preference profiles, and ranked rewriting pairs. We construct our data via a unified four-step pipeline, as illustrated in Figure 2: First, we collect and segment novels into chapter-level source texts. Second, we extract structured narrative contexts from the preceding storyline for each chapter to serve as critical global constraints. Third, we reverse-engineer and model fine-grained reader preference profiles based on real-world book reviews. Finally, we synthesize various preference-tailored rewrite candidates for each chapter and subsequently rank them according to their alignment with the target preferences. We next explicate each step in turn.

Chapter Collection & Segmentation We collect source novels from BookSum (Kryściński et al., 2022) and Project Gutenberg. While BookSum inherently provides chapter-partitioned narratives, long-form novels from Project Gutenberg are segmented into chapters based on structural markers. After deduplication, content filtering, and genre balancing, we secure a clean corpus of 163 chapter-

segmented novels. The final collection spans seven diverse genres: Historical, Romance, Sci-Fi, Adventure, Horror, Humor, and Detective.

Structured Narrative Context Extraction To support coherent long-form rewriting, we construct a structured narrative context for each chapter based on its preceding storyline. This context encapsulates the novel's writing background, character information, major events, and character relations. Crucially, the extraction of structured context shifts the personalized rewriting task from simple surface-level text style transfer to deep, context-aware narrative enrichment. By injecting these global continuity signals, we empower the model to actively enrich chapter details and introduce creative expansions while remaining strictly anchored to the established plot logic (Yang and Jin, 2025).

Reader Preference Modeling Rather than relying on artificial or generic descriptions, we model authentic user behavior by crawling real-world reviews on Goodreads from 780 active readers who engaged with the corresponding source books. We leverage Gemini 3.1 Pro (Google DeepMind, 2026) to distill these open-ended reviews into structured reading preferences. To capture the full complexity of reader intent beyond mere linguistic style, each preference profile is organized into a fine-grained five-dimensional taxonomy: `Plot`, `Characters`, `Language`, `Worldview`, and `Theme`. We further synthesize $2\times$ additional profiles from seed profiles. We manually check the profiles to ensure authenticity, plausibility, and diversity.

Rewriting Generation & Ranking For each novel, we sample five distinct reader profiles and utilize Gemini 3.1 Pro to generate five variant rewrites for each chapter, augmenting the prompts with the extracted structured narrative context to guide the enrichment process. All rewrite candidates are then ranked according to their adherence to each targeted reader profile. Consecutive variants within the ranked sequence are subsequently transformed into pairwise preference comparisons to facilitate subsequent alignment training.

To validate data quality, we conduct human verification on randomly sampled cases. Specifically, each of the 10 annotators rates 10 sampled rewrites on fluency, fidelity, and preference alignment using a 0–3 Likert scale, and evaluates 10 ranking pairs for ranking reasonableness. As shown in Table 6, the relative rankings achieve a high average inter-annotator agreement of 81%, confirming the rigorous quality of our benchmark.

3.2 Data Analysis

An empirical analysis of StoryLensBench (Figures 5, 6, and 7) reveals three core characteristics: i) the **reader distribution** spans a heterogeneous spectrum of ages and professional backgrounds from authentic reviews, capturing realistic user personas rather than artificial constructs. ii) the **book distribution** encompasses balanced literary genres and maintains a distinct long-tailed text-length distribution, establishing a challenging testbed that necessitates narrative enrichment rather than surface style transfer. iii) the **preference vocabulary** follows a salient long-tail pattern, demonstrating that the corpus successfully encapsulates both global reading trends and fine-grained, highly personalized reader nuances.

4 Methodology

4.1 Task Formulation

We formally define the preference-aligned story rewriting task as follows. Given an original story segment x, its structured narrative context of preceding contents c, and a target user preference profile p, preference-aligned story rewriting aims to generate an optimal rewrite y^*:

$$y^* = \arg\max_{y} P_\theta(y \mid x, c, p),$$

where the rewrite should preserve the main plot line of x, remain consistent with the structured context c, and satisfy the target reader's preferences p. Story rewriting entails two main challenges: preference alignment evaluation and context-aware text generation. To address these challenges, we first train a model that measures how well a rewrite matches a preference profile. Leveraging this as a reward signal, we further develop a rewriting model via a two-stage training pipeline.

4.2 StoryLensEval

We first train StoryLensEval, which measures the degree to which a rewritten text matches the target reader's preferences on our preference-ranking dataset StoryLensBench.

Our training data is formatted as ranked rewrite pairs conditioned on specific reader preferences. As illustrated in Figure 3, each instance in StoryLensBench contains multiple candidate rewrites ranked by their alignment with a user preference profile. We transform these rankings into pairwise comparisons, treating the higher-ranked

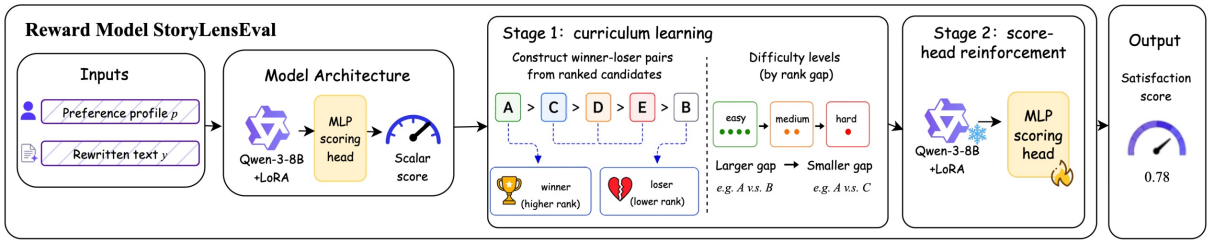

Figure 3: Overview of the training process for **StoryLensEval**.

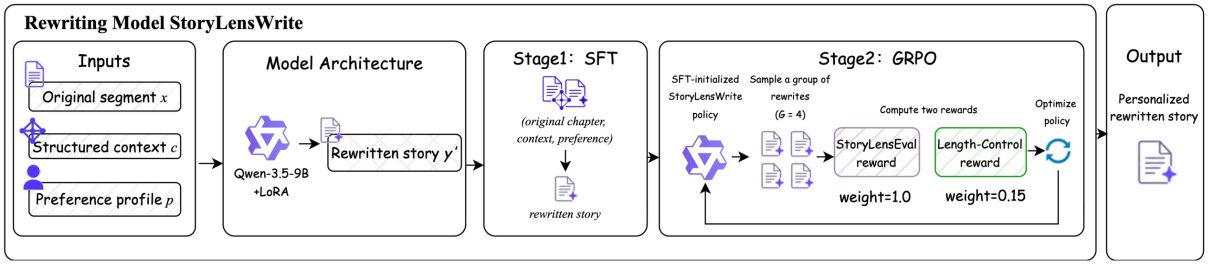

Figure 4: Overview of the training framework for **StoryLensWriter**.

rewrite as the winner and the lower-ranked one as the loser. Based on the ranking gap between two candidates, these pairs are categorized into easy (larger gap), medium, and hard (smaller gap) levels.

StoryLensEval adopts a Qwen3-8B backbone (Yang et al., 2025) with an MLP scoring head, and is trained with LoRA(Hu et al., 2022) for parameter-efficient adaptation. Given a user preference profile p and a rewritten text y, the model outputs a scalar satisfaction logit:

$$s = f_\theta(p, y), \tag{1}$$

where f_θ denotes the trained satisfaction scoring model.

The model is optimized with a pairwise ranking objective:

$$\mathcal{L}_{sat} = \text{softplus}(-(s_w - s_l)), \tag{2}$$

where s_w and s_l denote the scores of the winner and loser, respectively. This objective encourages the model to assign higher scores to rewrites that better satisfy the user's preferences.

We train the model in two stages: an easy-to-hard curriculum learning phase, followed by scoring head refinement to better distinguish hard comparison pairs and improve final scalar scoring precision.

4.3 StoryLensWriter

We adopt a two-stage training strategy to build our rewriting model. Specifically, we use Qwen3.5-9B (Qwen Team, 2026) as the backbone model and progressively adapt it to the personalized story rewriting task. The first stage performs supervised fine-tuning (SFT) to teach the model the task format and basic rewriting behavior, while the second stage further optimizes the model with reinforcement learning to improve preference alignment and controllability.

4.3.1 Stage 1: Supervised Fine-Tuning (SFT)

In the first stage, we fine-tune Qwen3.5-9B with LoRA adapters using the paired rewriting data from StoryLensBench. This stage aims to teach the model the basic input-output format of the task and provide it with an initial ability to perform preference-aware story rewriting.

Given an original segment x, a structured context c, and a user preference profile p, the model is trained to generate the target rewrite y. We format the input as an instruction-following prompt that explicitly requires the model to preserve the main plot, maintain consistency with the context, and adapt the rewritten text to the user's preferences. The SFT objective is the standard autoregressive language modeling loss:

$$\mathcal{L}_{\text{SFT}} = -\sum_{t=1}^{|y|} \log P_\theta(y_t \mid y_{<t}, x, c, p). \tag{3}$$

4.3.2 Stage 2: Reinforcement Learning with GRPO

Although SFT enables the model to imitate reference rewrites, it does not directly optimize the model toward preference alignment or task-specific quality criteria. Therefore, in the second stage, we further optimize the SFT model using Group Relative Policy Optimization (GRPO)(Shao et al., 2024). GRPO optimizes the policy by increasing the likelihood of candidates with relatively higher rewards, encouraging the model to better balance user preference alignment and practical rewriting constraints.

To construct the training set, we select approximately 5% of the SFT data and convert them into prompts containing the original story segment, structured narrative context, and user preference profile. For each prompt, the policy model samples a group of candidate rewrites. We use StoryLensEval, the preference scoring model introduced in the previous subsection, as the reward model. This reward model assigns a scalar score R_{rm} to each generated rewrite, measuring how well it satisfies the target user's preferences.

To avoid degenerate length changes, we additionally introduce a length-control reward. Specifically, we encourage the rewritten text to remain within 0.7–1.3 times the length of the original segment. When the generated rewrite falls outside this range, we apply a Huber-style penalty R_{length} based on the deviation from the valid interval(Huber, 1992).

The final reward is defined as:

$$R = R_{\text{rm}} + 0.15 R_{\text{length}}. \tag{4}$$

4.4 Evaluation Framework

To comprehensively assess the quality of rewriting, we establish a multi-dimensional evaluation framework across three core aspects: fidelity, coherence, and reader satisfaction.

Fidelity. Fidelity consists of two parts: local fidelity and global fidelity.

Local Fidelity Local fidelity uses only the original and rewritten segments as input, and evaluates whether the rewrite preserves the language, main plot, entities, and settings of the original passage.

Inspired by (Li et al., 2025), we calculate local fidelity based on Subject-Verb-Object (SVO) triplets and Narrative Entity Knowledge Graphs (NEKG). We use GPT-4o-mini (OpenAI, 2024a) to extract SVOs and NEKGs from both the original and rewritten texts. For SVOs, we use text-embedding-3-small(OpenAI, 2024b) to find the best-matching pairs between the two versions. If the cosine similarity falls into a fuzzy range, i.e., 0.50–0.85, GPT-4o makes a binary arbitration on whether the item is faithfully preserved. For NEKG, we use GPT-4o to determine whether each entry is preserved.

For each component $i \in \{\text{SVO}, \text{NEKG}\}$, we denote its pass rate as S_i, computed as the proportion of extracted items that pass the final judgment. The final local fidelity score is calculated as:

$$F_{\text{local}} = 0.9 \cdot S_{\text{NEKG}} + 0.1 \cdot S_{\text{SVO}}. \tag{5}$$

Global Fidelity Global fidelity introduces the structured narrative context as an additional input to check whether the rewritten content remains consistent with the broader story setting. Following the conflict types in Kim et al. (2024), we use GPT-4o-mini to extract claims from the rewritten text. Each claim is first compared against the original text. If it is not supported by the original text, it is further compared against the structured context. We categorize each claim into one of three labels: Entailment, Neutral, or Contradiction. Entailment indicates that the claim is fully supported, Neutral indicates a reasonable expansion, and Contradiction indicates a conflict with the original text or broader context. The global fidelity score is defined as the proportion of claims labeled as Entailment or Neutral:

$$F_{\text{global}} = \frac{N_{\text{entail}} + N_{\text{neutral}}}{N_{\text{total}}}. \tag{6}$$

Coherence. Coherence evaluates the sentence-level fluency and logical consistency of the rewritten text. Following the eight error types defined by Chang et al. (2023), we first segment each rewrite into sentences using punctuation delimiters, and then use GPT-4o-mini to detect potential coherence errors for each sentence. The coherence score is computed as $C = N_{\text{error_free}}/N_{\text{total}}$, where $N_{\text{error_free}}$ denotes the number of sentences without detected errors and N_{total} denotes the total number of sentences.

Satisfaction. Satisfaction measures the degree to which a rewritten text matches the target user's preferences. We feed the user preference profile and rewritten text into StoryLensEval, and use the predicted scalar score as the satisfaction score:

$$S_{sat} = f_\theta(p, y). \tag{7}$$

Model	Narrative Quality				Preference Alignment
	Local Fid. ↑	Global Fid. ↑	Coh. ↑	NovelCritique ↑	PerSE ↑
Gemini-3.1-Pro	93.68	97.54	92.65	83.80	7.69
GPT-5.2	93.61	98.97	**97.01**	**88.08**	7.59
DeepSeek-V3	93.50	98.56	92.91	85.86	7.67
Doubao-Seed-2.0-Lite	93.86	98.97	93.40	86.54	7.62
Qwen-VL-Max	94.45	98.44	93.82	85.53	7.69
StoryLensWriter	**94.62**	**98.98**	93.11	86.93	**7.78**

Table 1: Comparison with strong LLMs. Metrics are grouped into narrative quality and preference alignment.

Model / Method	Narrative Quality				Preference Alignment		
	Local Fid. ↑	Global Fid. ↑	Coh. ↑	NovelCritique ↑	Sat. Win ↑	Sat. Top-1 ↑	PerSE ↑
Llama-3.1-8B	92.54	96.96	80.53	78.55	0.09	0.00	7.52
Qwen3.5-9B	94.12	98.64	90.72	85.78	0.52	0.03	7.68
Hybrid RAG (Llama-3.1-8B)	92.29	93.26	79.64	68.00	0.21	0.02	7.37
Hybrid RAG (Qwen3.5-9B)	93.97	96.26	87.54	81.43	0.64	0.18	7.50
PPlug	94.35	**99.71**	80.99	86.14	0.70	0.30	7.65
StoryLensWriter	**94.62**	98.98	**93.11**	**86.93**	**0.84**	**0.47**	**7.78**

Table 2: Comparison with comparably-sized open-source models and representative personalization methods.

5 Experiments

5.1 Evaluation Dataset

For evaluation, we construct a held-out test set consisting of newly collected stories and preference profiles. Specifically, we collect 29 additional novels from Project Gutenberg, covering the same seven story categories, with the number of novels kept as balanced as possible across categories. Following our preprocessing pipeline, we segment the novels into chapters, resulting in a total of 867 chapter-level samples for evaluation. We also generate five new user preference profiles, deliberately maximizing their diversity across different dimensions.

5.2 Evaluation Metrics

We evaluate different rewriting methods from two perspectives: *Narrative Quality* and *Preference Alignment*. For *Narrative Quality*, we use our proposed **fidelity** and **coherence** metrics. We also use NovelCritique (Yang and Jin, 2025) to assess overall long-story quality.

For *Preference Alignment*, we use two complementary metrics: **PerSE** under the Per-MPST setting (Wang et al., 2024b) and our *satisfaction* evaluator **StoryLensEval**. PerSE serves as an external evaluation signal independent of the training objective of StoryLensWriter. Since StoryLensEval is trained as a task-specific relative evaluator, we report pairwise win rate and top-1 rate rather than absolute scores. To mitigate potential self-evaluation bias arising from our LLM-driven data construction, we utilize StoryLensEval primarily to benchmark open-source personalizable baselines.

We validate StoryLensEval against both human annotations and strong LLM judgments. Specifically, 40 randomly sampled story pairs are evaluated by three human annotators via majority voting, achieving a high inter-annotator correlation of 92.98%.

5.3 Baselines

We compare StoryLensWriter with two categories of baselines, spanning strong general LLMs and representative personalization methods (Zhang et al., 2024).

Strong LLMs We directly prompt strong off-the-shelf LLMs with the same instruction, narrative context, and reader preference profile. Models include Gemini-3.1-Pro (Google DeepMind, 2026), GPT-5.2 (OpenAI, 2025b), DeepSeek-V3 (DeepSeek-AI et al., 2024), Doubao-Seed-2.0-Lite (Volcengine, 2026), and Qwen-VL-Max (Bai et al., 2025).

Peer Open-source Models We also compare with Llama-3.1-8B (Dubey et al., 2024) and Qwen3.5-9B (Qwen Team, 2026).

Hybrid RAG We evaluate retrieval-based personalization by building a retrieval repository from the training set. For each test instance, we use hybrid retrieval with both dense and sparse matching to select preference-relevant examples, and then in-

Evaluator	Human Corr.	Strong LLM Corr.
Qwen3-8B	0.4211	0.4211
Llama-3.1-8B	0.7368	0.7368
StoryLensEval	**0.8947**	**0.9250**

Table 3: Correlation with human annotations and strong LLM judgments.

Model	Easy	Medium	Hard	Overall
StoryLensEval w/ dual tower	83.26	74.27	65.36	73.40
StoryLensEval w/o phase2	95.72	87.63	76.67	85.67
StoryLensEval	**95.85**	**87.85**	**76.91**	**85.87**

Table 4: Ablation study of StoryLensEval.

Model	Fid.	Coh.	PerSE	Win
StoryLensWriter w/o GRPO	94.04 / 97.51	91.10	7.59	20.44
StoryLensWriter	**94.62 / 98.98**	**93.11**	**7.78**	**79.56**

Table 5: Ablation study of StoryLensWriter.

clude them as few-shot prompting demonstrations. We apply hybrid RAG to both Llama-3.1-8B and Qwen3.5-9B.

Representation-based Personalization We also compare with Persona-Plug (Liu et al., 2025), which encodes user information into a plug-and-play personalized embedding. We adapt it to our task and use the learned preference representation to guide story rewriting.

5.4 Main Results

StoryLensWriter improves preference alignment over prompted strong LLMs while preserving strong narrative quality. Table 1 shows that StoryLensWriter achieves the best PerSE score among the compared strong LLMs, indicating stronger alignment with target reader preferences. Our model obtains the highest local fidelity and global fidelity, suggesting that task-specific training on StoryLensBench helps the model ground preference-driven rewriting in the original passage and structured narrative context.

Our proposed training framework is effective for both narrative quality and preference alignment. Table 2 shows that StoryLensWriter outperforms comparably-sized open-source backbones, retrieval-based personalization, and representation-based personalization methods. This demonstrates that merely retrieving historical contexts or injecting persona representations is insufficient for rewriting. Instead, explicitly optimizing the model via our two-stage training paradigm genuinely instills a fine-grained understanding of multi-dimensional reader preferences.

StoryLensEval provides a reliable relative preference signal for story rewriting. As reported in Table 3, StoryLensEval achieves higher agreement with human annotations and strong LLM judgments than comparably-scale LLM judges. This suggests that preference-ranking training enables StoryLensEval to better capture reader-preference satisfaction, making it suitable as both an evaluator and a reward model for preference-aligned story rewriting.

5.5 Ablation Study

We conduct two ablation studies to analyze the effectiveness of our design choices. The first study focuses on StoryLensEval, including its model architecture and second-stage score-head reinforcement. The second study evaluates the RL stage of StoryLensWriter by comparing the full model with its SFT-only variant.

As shown in Table 4, replacing the single-tower design with a dual-tower architecture leads to a clear drop. This suggests that jointly encoding the user preference profile and rewritten text is important for capturing fine-grained preference interactions. Removing the second training stage also decreases performance, especially on hard pairs, indicating that score-head reinforcement improves the discriminative ability of the model.

Table 5 shows the effect of the GRPO stage. Compared with the SFT-only variant, StoryLensWriter achieves better preference alignment after RL optimization. This indicates that using StoryLensEval as a reward signal guides the model toward preference-aligned rewriting.

6 Conclusion

In this paper, we study preference-aligned story rewriting, aiming to adapt literary narratives to diverse reader preferences while preserving narrative consistency. We make three contributions: (i) constructing StoryLensBench, a story rewriting benchmark with narrative contexts, reader preference profiles, and ranked rewrites; (ii) establishing an evaluation framework covering fidelity, coherence, and satisfaction, supported by StoryLensEval; and (iii) developing StoryLensWriter, a two-stage rewriting model trained with SFT and reinforcement learning. Experiments show that StoryLensWriter improves preference alignment over strong LLMs and personalization baselines while maintaining strong narrative quality.

Limitations

Our current experiments focus on English literary narratives and reader preference profiles derived from Goodreads. This setting provides a controlled testbed for preference-aligned story rewriting, but it may not fully reflect multilingual, cross-cultural, or highly contemporary reading preferences. Extending the benchmark to broader languages, cultures, and narrative domains is a natural direction for future work.

The source texts in our current benchmark mainly come from literary narratives in BookSum and Project Gutenberg. While these sources provide diverse and well-structured story materials, future work may further include a broader range of narrative genres and text types, such as contemporary fiction, children's literature, web novels, interactive stories, or domain-specific narratives, to examine the generality of preference-aligned rewriting in more varied settings.

Ethical Considerations

We acknowledge the ethical considerations involved in constructing STORYLENSBENCH from literary texts and reader reviews.

Data Sources and Copyright. The source texts are collected from BookSum and Project Gutenberg and are used only for academic research. To respect copyright and licensing constraints, we will not release copyrighted full-text books or unrestricted reproductions of protected source materials. Released data will focus on processed benchmark instances, metadata, structured contexts, and derived annotations within the permissions of the original sources.

Reader Privacy. Reader preference profiles are derived from publicly available Goodreads review content. To protect user privacy, we anonymize user-related information and avoid releasing raw user identifiers, review identifiers, or information that could directly identify individual readers. We also avoid distributing raw review texts whenever possible and release only processed preference representations needed for research.

Intended Use. The benchmark, evaluator, and rewriting model are intended for academic research on preference-aligned narrative rewriting. They should not be used to impersonate specific readers, infer sensitive personal attributes, or generate misleading adaptations of copyrighted works.

References

Anthropic. 2024. Introducing claude 3.5 sonnet. Accessed: 2026-05-22.

Minwook Bae and Hyounghun Kim. 2024. Collective critics for creative story generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18784–18819.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Yapei Chang, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2023. Booookscore: A systematic exploration of book-length summarization in the era of llms. *arXiv preprint arXiv:2310.00785*.

DeepSeek-AI, Daya Guo, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

DeepSeek-AI, Aixin Liu, and 1 others. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmed Al-Dahle, Aaron Letman, Anish Mathur, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yixin Dai, Jiawei Sun, Haofen Wang, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2(1).

Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. 2024. Scaling synthetic data creation with 1,000,000,000 personas. *arXiv preprint arXiv:2406.20094*.

Google DeepMind. 2025. Gemini 2.5: Our newest gemini model with thinking. Accessed: 2026-05-22.

Google DeepMind. 2026. Gemini 3.1 pro model card. Accessed: 2026-05-22.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3.

Peter J Huber. 1992. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution*, pages 492–518. Springer.

Linh Huynh and Danielle S McNamara. 2025. Genai-powered text personalization: Natural language processing validation of adaptation capabilities. *Applied Sciences*, 15(12):6791.

Pegah Jandaghi, XiangHai Sheng, Xinyi Bai, Jay Pujara, and Hakim Sidahmed. 2024. Faithful persona-based conversational dataset generation with large language models. In *Proceedings of the 6th Workshop on NLP for Conversational AI (NLP4ConvAI 2024)*, pages 114–139.

Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2023. Evaluating and inducing personality in pre-trained language models. *Advances in Neural Information Processing Systems*, 36:10622–10643.

Yekyung Kim, Yapei Chang, Marzena Karpinska, Aparna Garimella, Varun Manjunatha, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2024. Fables: Evaluating faithfulness and content selection in book-length summarization. *arXiv preprint arXiv:2404.01261*.

Wojciech Kryściński, Nazneen Rajani, Divyansh Agarwal, Caiming Xiong, and Dragomir Radev. 2022. Booksum: A collection of datasets for long-form narrative summarization. In *Findings of the association for computational linguistics: EMNLP 2022*, pages 6536–6558.

Wen Lai, Viktor Hangya, and Alexander Fraser. 2024. Style-specific neurons for steering llms in text style transfer. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13427–13443.

Seongyun Lee, Sue Hyun Park, Seungone Kim, and Minjoon Seo. 2024. Aligning to thousands of preferences via system message generalization. *Advances in Neural Information Processing Systems*, 37:73783–73829.

Jiaming Li, Yukun Chen, Ziqiang Liu, Minghuan Tan, Lei Zhang, Yunshui Li, Run Luo, Longze Chen, Jing Luo, Ahmadreza Argha, and 1 others. 2025. Storyteller: An enhanced plot-planning framework for coherent and cohesive story generation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 20818–20846.

Xinyu Li, Ruiyang Zhou, Zachary C Lipton, and Liu Leqi. 2024. Personalized language modeling from personalized human feedback. *arXiv preprint arXiv:2402.05133*.

Yen-Ting Lin and Yun-Nung Chen. 2023. Llm-eval: Unified multi-dimensional automatic evaluation for open-domain conversations with large language models. *arXiv preprint arXiv:2305.13711*.

Ao Liu, An Wang, and Naoaki Okazaki. 2022. Semi-supervised formality style transfer with consistency training. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4689–4701.

Jiongnan Liu, Yutao Zhu, Shuting Wang, Xiaochi Wei, Erxue Min, Yu Lu, Shuaiqiang Wang, Dawei Yin, and Zhicheng Dou. 2025. LLMs + Persona-Plug = Personalized LLMs. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9373–9385.

Shuai Liu and Jonathan May. 2025. Style transfer with multi-iteration preference optimization. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2663–2681.

Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Chris Leung, Jiajie Tang, and Jiebo Luo. 2024. Llm-rec: Personalized recommendation via prompting large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 583–612.

Muhammad U Nasir, Steven James, and Julian Togelius. 2024. Word2world: Generating stories and worlds through large language models. *arXiv preprint arXiv:2405.06686*.

Lin Ning, Luyang Liu, Jiaxing Wu, Neo Wu, Devora Berlowitz, Sushant Prakash, Bradley Green, Shawn O'Banion, and Jun Xie. 2025. User-llm: Efficient llm contextualization with user embeddings. In *Companion Proceedings of the ACM on Web Conference 2025*, pages 1219–1223.

OpenAI. 2024a. Gpt-4o mini: advancing cost-efficient intelligence. Accessed: 2026-05-22.

OpenAI. 2024b. Text embedding 3 small. https://developers.openai.com/api/docs/models/text-embedding-3-small. Accessed: 2026-05-25.

OpenAI. 2025a. Introducing gpt-5. Accessed: 2026-05-22.

OpenAI. 2025b. Introducing gpt-5.2. Accessed: 2026-05-22.

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.

Kyeongman Park, Minbeom Kim, and Kyomin Jung. 2025. A character-centric creative story generation via imagination. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 1598–1645.

Qwen Team. 2026. Qwen3.5-9b model card. Accessed: 2026-05-22.

Alireza Salemi, Julian Killingback, and Hamed Zamani. 2025. Expert: Effective and explainable evaluation of personalized long-form text generation. *arXiv preprint arXiv:2501.14956*.

Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024. Lamp: When large language models meet personalization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7370–7392.

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*.

Pushpdeep Singh, Mayur Patidar, and Lovekesh Vig. 2024. Translating across cultures: Llms for intralingual cultural adaptation. In *Proceedings of the 28th Conference on Computational Natural Language Learning*, pages 400–418.

Volcengine. 2026. Doubao-seed-2.0: Api release announcement. Accessed: 2026-05-22.

Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. Asking and answering questions to evaluate the factual consistency of summaries. *arXiv preprint arXiv:2004.04228*.

Chen Wang, Min Li, Linna Song, Yichang Li, Zhipeng Li, Xinhang Yan, and Yeqing Zhang. 2024a. Generating long-form story using dynamic hierarchical outlining with memory-enhancement. *arXiv preprint arXiv:2412.13575*.

Danqing Wang, Kevin Yang, Hanlin Zhu, Xiaomeng Yang, Andrew Cohen, Lei Li, and Yuandong Tian. 2024b. 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, Miami, Florida, USA. Association for Computational Linguistics.

Anna Wegmann and Dong Nguyen. 2021. Does it capture stel? a modular, similarity-based linguistic style evaluation framework. *arXiv preprint arXiv:2109.04817*.

Haotian Xia, Hao Peng, Yunjia Qi, Xiaozhi Wang, Bin Xu, Lei Hou, and Juanzi Li. 2025. Storywriter: A multi-agent framework for long story generation. *arXiv preprint arXiv:2506.16445*.

An Yang and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Dingyi Yang and Qin Jin. 2025. What matters in evaluating book-length stories? a systematic study of long story evaluation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16375–16398.

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*.

Lining Yao, Kuan-Hao Chang, Yandi Wang, Xiang Ren, and Nanyun Peng. 2019. Plan-and-write: Towards better automatic storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7426–7433.

Tian Yu, Ken Shi, Zixin Zhao, and Gerald Penn. 2025. Multi-agent based character simulation for story writing. In *Proceedings of the Fourth Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2025)*, pages 87–108.

Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, and 1 others. 2024. Personalization of large language models: A survey. *arXiv preprint arXiv:2411.00027*.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, and 1 others. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 1(2).

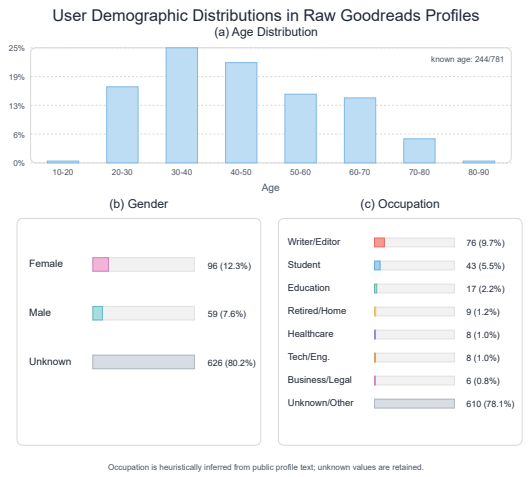

Figure 5: Demographic distributions of raw Goodreads reviewer profiles.

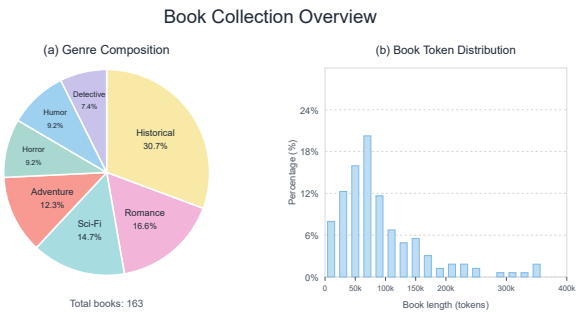

Figure 6: Book genre and length distributions in StoryLensBench.

A Dataset Details

Figure 7 shows the keyword distributions of user preferences across five dimensions: plot, character, language, worldview, and theme. For each dimension, we report the top-10 most frequent keywords and provide representative phrases for each keyword to illustrate how these preferences are expressed in natural language. Across all dimensions, the distributions show a clear long-tail pattern.

This indicates that user preferences contain both common patterns and diverse personalized details. Therefore, preference-aligned story rewriting needs to capture not only dominant preference trends but also fine-grained individual signals.

B Pilot Study Results

This section reports the detailed results of the two pilot studies mentioned in the Introduction.

Metric	Average Value
Preference alignment	2.16
Fluency	2.36
Fidelity	2.44
Ranking agreement	81%

Table 6: Human verification results for sampled rewritten stories and ranking pairs. The first three dimensions are rated by annotators on a 0–3 scale, where higher scores indicate better rewriting quality. Ranking agreement measures the average human agreement with model-produced rankings.

B.1 Pilot Study A

Tables 7 and 8 show the lexical and syntactic adaptation results. Gemini 2.5 Pro (Google DeepMind, 2025) and GPT-5 (OpenAI, 2025a) perform better in vocabulary control, while GPT-5 achieves the lowest syntactic complexity. However, the differences across models, including Claude 3.5 (Anthropic, 2024) and DeepSeek R1 (DeepSeek-AI et al., 2025), indicate that even strong LLMs still vary in satisfying fine-grained audience constraints.

B.2 Pilot Study B

Table 9 summarizes the human annotation results across eight books. Context-enhanced rewriting achieves the highest satisfaction and overall score, while the original text preserves the highest fidelity. This suggests that context-consistent enrichment improves reader satisfaction, but may introduce a modest fidelity trade-off.

C Case Study

Case 1. Table 10 presents a case from *Robinson Crusoe*, Chapter IV. The key gain from context enhancement is that commercial language is not added as a surface style alone; it is tied back to Crusoe's already established background as a plantation owner, his pursuit of wealth, and his methodical survival logic. As a result, phrases such as **repository of capital**, **merchant assessing an abandoned warehouse**, and **means of production** function as continuity-grounded characterization rather than decorative preference matching. The final phrase **catastrophic loss of equity** extends the same logic from objects to risk management, making the raft legible as an accumulated asset system rather than just a transport device.

Case 2. Tables 11–13 illustrate how we judge local and global fidelity with concrete examples. We first show claim-level local fidelity decisions, then

Figure 7: Top keywords in user preferences across five preference dimensions.

Model	A1+A2 Basic Words ↑	B2+C1+C2 Advanced Words ↓	Interpretation
Gemini 2.5 Pro	**92.9%**	**2.8%**	Best adapted to children
GPT-5	91.6%	3.6%	Strong lexical adaptation
Claude 3.5	89.1%	4.9%	Slightly more difficult wording
DeepSeek R1	86.8%	6.6%	More advanced wording

Table 7. Pilot A lexical adaptation results. We compare the proportion of basic and advanced words in rewritten stories for children.

Model	MLS ↓	C/S ↓	CN/T ↓	Interpretation
GPT-5	**6.7**	**1.18**	**0.41**	Best syntactic adaptation
Gemini 2.5 Pro	7.7	1.41	0.60	Adapted to children
DeepSeek R1	11.9	1.21	1.28	More suitable for adolescent readers
Claude 3.5	15.7	1.62	1.59	Adult-like syntactic style

Table 8. Pilot A syntactic complexity results. MLS denotes mean length of sentence, C/S denotes clauses per sentence, and CN/T denotes complex nominals per T-unit.

a globally acceptable literary expansion that is locally unsupported, and finally summarize the decision criteria.

D Prompt Templates

Appendix D summarizes the prompt templates used in our experiments. Table 14 reports our training prompts, including the SFT prompt and the GRPO prompt. Table 15 reports our baseline prompts, including the direct LLM baseline prompt and the Rewrite Few-shot RAG prompt.

E Implementation Details

We train StoryLensWriter in two stages: supervised fine-tuning (SFT) followed by GRPO. The SFT stage is implemented with ms-swift and DeepSpeed ZeRO-3 on 8 Nvidia A6000 GPUs. We use LoRA with rank 8 and alpha 16, and the full SFT run takes about 36.4 hours.

We then perform GRPO training with VeRL on 8 Nvidia A6000 GPUs. In this stage, 5 GPUs are allocated to actor/reference training, 2 GPUs to vLLM rollout generation, and 1 GPU to reward computation. We use LoRA with rank 8 and alpha 32, and the GRPO run takes about 49.6 hours. All final results reported in the paper use the same decode configuration: `max_new_tokens=3072`, temperature 0.25, top-p 0.8, top-k 20, and repetition penalty 1.05. We use this setting as a practical trade-off between creative variation and overall text quality.

Method	Satisfy ↑	Fluency ↑	Fidelity ↑	Total Avg. ↑	Best Count
Original	3.43	4.39	**4.62**	4.15	1
Preference-alignment	3.51	4.39	4.46	4.12	2
Context-enhanced	**4.27**	**4.45**	4.41	**4.37**	**6**

Table 9: Pilot B human annotation summary across eight books. Context-enhanced rewriting achieves the highest satisfaction and overall score, while the original text preserves the highest fidelity.

Original Excerpt	Reader Preference Profile	Preference-alignment Rewrite (no context)	Context-enhanced Rewrite (ours)
Here I found that the ship was bulged, and had a great deal of water in her hold, but that she lay so on the side of a bank of hard sand, or, rather earth, that her stern lay lifted up upon the bank, and her head low, almost to the water. By this means all her quarter was free, and all that was in that part was dry; for you may be sure my first work was to search, and to see what was spoiled and what was free. And, first, I found that all the ship's provisions were dry and untouched by the water, and being very well disposed to eat, I went to the bread room and filled my pockets with biscuit, and ate it as I went about other things, for I had no time to lose. I also found some rum in the great cabin, of which I took a large dram, and which I had, indeed, need enough of to spirit me for what was before me. My raft was now strong enough to bear any reasonable weight. My next care was what to load it with, and how to preserve what I laid upon it from the surf of the sea. ... I had other things which my eye was more upon - as, first, tools to work with on shore. And it was after long searching that I found out the carpenter's chest, which was, indeed, a very useful prize to me, and much more valuable than a shipload of gold would have been at that time.	**plot:** mercantile strategy, economic negotiation, trade networks. **character:** shrewd pragmatists and adaptable mediators. **language:** measured prose with material culture and socio-economic detail. **worldview:** resource allocation, incentives, exchange. **theme:** commerce, wealth, soft power, globalization.	I discovered that the ship's hull was breached and its hold was flooded. Fortunately, it had run aground on a hard sandbank in such a way that the stern was lifted high, keeping that entire rear section perfectly dry. My **immediate priority** was to see what had been ruined and what was **salvageable**. What I needed most was a way to transport the essential supplies back to shore. Wishing for a boat was useless, so necessity forced me into action. ... Driven by the **desperate need for supplies**, I pushed through the exhaustion, grabbed a carpenter's saw, and cut a spare topmast into three heavy pieces to reinforce my makeshift raft. My mind was focused on a much higher priority: tools. After a long search, I finally located the ship's carpenter's chest. ... I lowered the heavy box onto my raft, viewing it as a prize **far more precious ... than a ship entirely filled with gold**.	The morning light revealed a profound shift in my material circumstances. ... I recognized an immediate imperative: the ship was no longer a vessel of transit, but a **repository of capital** waiting to be liquidated before the sea could reclaim it. Stripping myself of burdensome garments, I swam to the wreck, approaching it as a **merchant assessing an abandoned warehouse**. ... The ship's quarter remained elevated and dry, preserving the vital commodities within. Lacking a skiff, I engineered a rudimentary barge ... My foremost priority was the **means of production**: the ship's carpenter's chest. In an environment devoid of manufactured goods, these tools represented an incalculable concentration of industrial potential, far surpassing the value of any precious metal. Navigating the laden raft to the shore required a pragmatic application of the tides. ... A grounding upon a sudden shoal threatened to capsize the venture, an event that would have constituted a **catastrophic loss of equity**.

Table 10: Case 1 from *Robinson Crusoe*. Structured context grounds mercantile preference cues in Crusoe's established history and survival logic.

Example	Original SVO Event	Rewritten Event	Judgment	Reason
Local-1	<Mrs. Reed, appoint, a small closet for Jane Eyre to sleep in>	<Mrs. Reed, banish, Jane Eyre to a closet>	**Pass / Entailment**	Although appoint is rewritten as the stronger verb banish, the core event is preserved: Mrs. Reed sends Jane Eyre to a closet. The rewrite changes the emotional tone but does not alter the main factual content.
Local-2	<Watson, receive, the wedding ring>	<Holmes, press, a near-perfect replica into Watson's palm>	**Fail / Contradiction**	The rewrite changes the key object from the actual wedding ring to a near-perfect replica. This is not a stylistic paraphrase, but a factual change that alters the plot semantics.
Local-3	<Holmes, follow, the old woman>	<Watson, watch, Holmes trail the old woman>	**Pass / Entailment**	The core event is still preserved: Holmes follows / trails the old woman. The rewrite adds Watson as an observer, which is a perspective expansion rather than a contradiction.

Table 11: Local Fidelity Examples

Extracted Claim	Local Evidence	Initial Local Judgment	Global Fidelity Judgment	Reason
[State] The kitchen usually has the scent of baking bread.	The local source text does not explicitly mention the smell of baking bread in the kitchen.	**Unsupported**	**Neutral / Acceptable Literary Expansion**	The claim is not directly entailed by the local passage, but it does not contradict the global narrative world or character setting. It functions as environmental description and literary enrichment, so it should not be penalized as hallucination.

Table 12: Global Fidelity Example

Judgment Type	Definition	Should Penalize?	Example
Entailment	The rewritten content is directly supported by the original text, or preserves the same core event with minor stylistic changes.	No	Mrs. Reed appoints a closet for Jane Eyre → Mrs. Reed banishes Jane Eyre to a closet
Neutral	The rewritten content is not explicitly stated in the local source, but remains globally plausible and does not violate character, plot, or world settings.	No	Adding the scent of baking bread to a kitchen scene
Contradiction	The rewritten content changes key plot facts, object identity, causal relations, or violates global character/world constraints.	Yes	Watson receives the wedding ring → Holmes gives Watson a replica

Table 13: Fidelity Judgment Criteria

Method	Prompt Part	Template / Content
SFT	Shared system prompt	`You are a literary rewriting assistant.` The prompt instructs the model to rewrite the source chapter for the target reader while preserving plot events, causal structure, chapter continuity, character consistency, and English-only output. It also states that continuity overrides preferences and clarity overrides stylistic excess.
	User prompt template	`Rewrite the source chapter for the target reader. Metadata: {compact_json_header} Structured context: {compact_json_context} Reader preference profile: {compact_json_preference} Source chapter: {source_text} Return only the rewritten chapter text.`
	Assistant target	`{rewritten_chapter_text}`
GRPO	System prompt	The GRPO prompt inherits the SFT system prompt and adds two extra constraints: `Do not introduce new major characters, locations, events, or facts unless they are strongly implied by the source chapter and continuity context.` It also requires the rewrite to remain `roughly similar in length to the source chapter`.
	User prompt	Same user prompt template as SFT, using `{compact_json_header}`, `{compact_json_context}`, `{compact_json_preference}`, and `{source_text}` as input fields.
	Training record shape	Each GRPO example stores `messages` with a system prompt and a user prompt, together with auxiliary fields including `preference_text`, `source_text`, `reference_rewrite`, chapter metadata, file paths, and a nested `metadata` object.

Table 14: Training prompt templates used for SFT and GRPO.

Method	Prompt Part	Template / Content
LLM Baseline	System prompt	The baseline uses the same literary rewriting system prompt as SFT: rewrite for the target reader, preserve main plot events and causal structure, maintain cross-chapter continuity, keep characters consistent, adapt to the five preference dimensions, output only story text, and keep the result in English.
	User prompt template	`Rewrite the source chapter for the target reader. Metadata: {compact_json_header} Structured context: {compact_json_context} Reader preference profile: {compact_json_preference} Source chapter: {source_text} Return only the rewritten chapter text.`
Rewrite Few-shot RAG	Retrieval query	`{target_profile} {target_original}`
	Final prompt template	The final prompt repeats the same rewriting instruction block and then appends a retrieved few-shot section: `– Reference Examples (Combine these styles) –`. Each example contains `[Instruction]`, `[Original]`, and `[Rewrite]` fields, followed by the current task with metadata, structured context, reader preference profile, and source chapter.
	Example record schema	Retrieved demonstrations are formatted as `[Instruction]: data.json -> instruction`, `[Original]: data.json -> original_text`, and `[Rewrite]: data.json -> rewritten_text`.

Table 15: Baseline prompt templates used for direct LLM baselines and Rewrite Few-shot RAG.