

CCStories: Improving the Safety and Age-Appropriateness of LLM-Generated Chinese Children’s Stories

Jiaxu Zhao^{1,2*}, Qin Zhang^{3*}, Xiaoxing Wang³, Liang Ding⁴, Yulong Pei²,
Xiaojun Chen³, Meng Fang^{5,2†}, Tanja Käser¹, Mykola Pechenizkiy²

¹EPFL, Switzerland; ²Eindhoven University of Technology, Netherlands;

³College of Computer Science and Software Engineering, Shenzhen University, China;

⁴Alibaba Group, China; ⁵University of Liverpool, United Kingdom

jiaxu.zhao@epfl.ch

Abstract

Children’s stories play a crucial role in early cognitive and language development, yet existing story generation models lack the safety and age-appropriateness required for educational applications, especially in non-English languages. We introduce **CCStories**, a curated dataset of 22,200 high-quality Chinese children’s stories designed specifically to support safe and appropriate AI-generated content for young readers. Unlike general-purpose datasets derived from adult-oriented platforms, CCStories targets three critical dimensions of child-appropriate text: readability, syntactic complexity, and toxicity. We further propose a multi-dimensional evaluation framework that jointly assesses appropriateness and generation quality. As a validation of the dataset’s utility, we fine-tune five representative backbones: CPM, PanGu- α , GLM-4, Qwen, a Chinese-adapted LLaMA-3 variant (Gigi-Llama-3-8B-zh), and additionally benchmark GPT-5-mini as a zero-shot instruction-tuned reference. Fine-tuning reduces offensive-prompt toxicity by up to 26.7 percentage points and raises readability by up to 14.3 percentage points, narrowing much of the gap to the GPT-5-mini reference. A human evaluation confirms that appropriateness rises from 53% to 86% after fine-tuning, and human validation of the toxicity classifier shows that the model ranking it induces matches the human ranking. CCStories provides a publicly available resource for building and benchmarking safer, more age-appropriate Chinese generative models for children’s education. **Warning:** *This paper contains content that may be offensive or upsetting.*

1 Introduction

Language and cognitive development are foundational aspects of early childhood education, and

children’s stories play a central role through engaging characters and moral lessons, offering substantial benefits in vocabulary acquisition and ethical guidance. Recent advances in neural text generation have shown strong results across a wide range of applications, including dialogue generation, question answering, and story generation (Peters et al., 2018; Yang et al., 2019; Radford et al., 2019; Brown et al., 2020). Despite this progress, a critical gap remains for child-oriented content generation: current models struggle with three aspects central to age-appropriateness, such as readability, syntactic complexity, and potential toxicity (Luckin and Holmes, 2016; Lemaignan et al., 2021; Charisi et al., 2021; La Fors, 2022; Zhao et al., 2023). These shortcomings pose tangible risks when general-purpose models are used in educational applications for children.

A primary reason for this gap is the data on which general-purpose language models are pre-trained. Pre-training corpora are typically derived from adult-oriented platforms such as Reddit¹ and Wikipedia, for instance, Radford et al. (2019) use outbound Reddit links to build WebText. These sources frequently contain complex linguistic structures that exceed children’s comprehension levels and may include inappropriate or toxic content. Consequently, language models trained or fine-tuned on such corpora can inadvertently learn and propagate harmful or unduly complex language, making them unsuitable for direct use in children’s education without significant adaptation.

The situation is particularly challenging for non-English languages. While some children’s text datasets exist (Hill et al., 2015), they predominantly focus on English, leaving a significant gap for other languages. Machine translation is in principle an option, but translations often suffer from *translationese* (Gellerstam, 1986, 1996), which is a phe-

* Equal contribution.

† Corresponding author: Meng.Fang@liverpool.ac.uk.

¹<https://www.reddit.com/>

Title: 会跳舞的小蚯蚓 (Translation: A dancing small earthworm) Content: 有一只细细长长的小蚯蚓，名字叫弯弯，她天生爱跳舞，世界上没有比跳舞更让她高兴的事了。... (Translation: There was a long, thin little earthworm named Bend, who was born to dance, and there was nothing in the world that made her happier than dancing. ...)
Title: 我长出了翅膀 (Translation: I grew wings) Content: 早晨起来，我发现有个毛茸茸的东西在我身下，哎呀，这是一对翅膀，而且长在我的身上！... (Translation: I got up in the morning and found a furry thing under me, oops, it's a pair of wings and they're growing on my body! ...)

Table 1: Two examples of stories in CCStories.

nomenon in which translated language exhibits unnatural or stilted forms resulting from the influence of source-text patterns, including awkward syntax, reduced lexical diversity, and unusual collocations that sound foreign to native speakers. For children’s content specifically, it introduces structures and vocabulary inappropriate for young readers and reduces the educational value of the material. Furthermore, cultural nuances and idiomatic expressions that are central to children’s engagement are routinely lost in translation, making translated text a poor substitute for natively authored children’s stories.

To address these challenges, we present **CCStories**, a corpus of 22,200 Chinese children’s stories compiled from public children’s-story websites and refined through rigorous automatic and manual cleaning. Two examples are shown in Table 1. Through experiments with widely used Chinese language models, we show that existing pre-trained models tend to generate content that is toxic, overly complex, and unfit for children: baselines produce toxicity rates of up to 41% under problematic prompts, which is far above any acceptable threshold for children’s content. After fine-tuning with CCStories, offensive-prompt toxicity drops by up to 26.7 percentage points, readability improves by up to 14.3 percentage points, and syntactic complexity is reduced toward levels appropriate for young readers, closing much of the gap between AI-generated text and human-authored children’s stories.

Our main contributions are as follows:

- **The CCStories dataset.** A high-quality, child-oriented Chinese story corpus of 22,200 stories, sourced from eight reputable children’s-story websites and refined through a two-stage automatic + manual cleaning pipeline. The dataset is publicly available to the research

and educational community.²

- **A multi-dimensional evaluation framework** that examines appropriateness across three dimensions central to children’s content, such as, readability (using age-appropriate vocabulary metrics), syntactic complexity (matching children’s comprehension levels), and toxicity (detecting harmful or inappropriate content), and complements them with standard generation-quality metrics, a human evaluation of appropriateness and coherence, and a human validation of the toxicity classifier itself.
- **Empirical validation of dataset utility.** Through fine-tuning five representative generative language models: four Chinese-centric backbones (CPM, PanGu- α , GLM-4, Qwen) and one Chinese-adapted LLaMA-3 variant, and benchmarking against a strong instruction-tuned reference (GPT-5-mini), we show that the dataset effectively guides models toward safer, more readable, and more age-appropriate story generation.

In this work, we define “appropriate” content as text that is free from harmful biases, toxic content, and undue lexical or syntactic complexity, and is therefore suitable for a young audience.

2 Related Work

We discuss prior work along three themes: story datasets, evaluation of age-appropriate content, and Chinese text generation models.

2.1 Story Datasets

A variety of datasets have been proposed for story generation, primarily focusing on English. PASTA (Ghosh et al., 2022) is an English dataset of factual and counterfactual story pairs that enables state-based inference beyond surface understanding. ROCStories (Mostafazadeh et al., 2016) provides 50,000 high-quality five-sentence stories that have been widely adopted for commonsense story modelling, and WritingPrompts (Fan et al., 2018) contains roughly 300K human-written stories paired with prompts and is a standard benchmark for open-ended generation. Hill et al. (2015) release the Children’s Book Test, derived from freely available

²The dataset and code are available at <https://github.com/aialt/CCStories>.

English children’s books, but it primarily targets cloze-style language modelling rather than open-ended generation.

For Chinese, STORAL (Guan et al., 2022b) provides stories paired with explicit morals, and LOT (Guan et al., 2022a) offers a Chinese story corpus with outlines and supports tasks such as cloze tests, sentence-position prediction, plot completion, and outline-conditioned generation. These resources are not curated for children: their vocabulary, themes, and toxicity profile are not necessarily suitable for early-childhood education. CC-Stories differs from prior datasets in that it specifically targets children’s stories and is filtered along readability, syntactic complexity, and toxicity, supporting both training and benchmarking of child-appropriate generative models.

2.2 Evaluation of Age-Appropriate Content

Evaluating whether generated text is suitable for children combines readability assessment, controllable generation, and toxicity/safety evaluation. Classic readability measures such as Flesch-Kincaid (Kincaid et al., 1975) are designed for English and rely on syllable counts that do not transfer naturally to Chinese; Chinese readability work has instead typically used graded vocabulary lists such as HSK. For toxicity in Chinese, COLD (Deng et al., 2022) provides the first benchmark for offensive-language detection across topics including race, gender, and region, and remains the de facto standard for evaluating Chinese-language generation safety. Beyond automatic classifiers, LLM-as-a-judge protocols (Zheng et al., 2023; Zhao et al., 2025) and human evaluation along axes such as engagement and coherence (Chhun et al., 2022) have also been applied to story generation. Our framework combines automatic readability, syntactic complexity, and toxicity metrics, complemented by a qualitative case study.

2.3 Chinese Text Generation Models

Text generation has advanced rapidly through the development of pre-trained language models. Cui et al. (2021) propose MacBERT, a Chinese BERT-based model that narrows the gap between pre-training and fine-tuning. NEZHA (Wei et al., 2019), also built on BERT, performs strongly on Chinese understanding tasks such as named-entity recognition and sentiment classification. ERNIE (Sun et al., 2019) integrates external knowledge into Chinese representations and exhibits strong knowledge

inference. The Text-to-Text Transfer Transformer (T5) (Raffel et al., 2020) unifies many NLP tasks under a text-to-text formulation; mT5 (Xue et al., 2020) extends this to 101 languages including Chinese. Mengzi (Zhang et al., 2021b) is a lightweight Chinese pre-trained model covering linguistic and visual tasks.

For autoregressive generation, large-scale Chinese models such as CPM (Zhang et al., 2021a) and PanGu- α (Zeng et al., 2021) have established strong baselines on a variety of Chinese tasks. Qwen (Bai et al., 2023) is a Chinese-centric LLM series by Alibaba, trained on hundreds of billions of Chinese tokens. GLM-4 (GLM et al., 2024) is a bilingual model pre-trained on around ten trillion tokens, primarily in Chinese and English. Beyond purely Chinese-centric models, LLaMA-3 (Dubey et al., 2024) has been widely adapted to Chinese via continued pre-training, for example, the Giga-Llama-3-8B-zh release we adopt, providing a useful contrast for how a Chinese-adapted multilingual backbone interacts with our dataset.

3 CCStories Dataset

CCStories is a comprehensive collection of Chinese children’s stories specifically curated for children’s education. The dataset consists of 22,200 carefully selected and cleaned stories spanning multiple genres, including fairy tales, fables, folklore, puzzle stories, mythology, and other educational content. Each story has been processed through rigorous automatic and manual cleaning procedures to ensure high quality and appropriateness for children. This section details our methodology for dataset construction, quality control, and statistical analysis.

3.1 Data Collection

We employed a systematic approach to collect high-quality Chinese children’s stories from a diverse set of authoritative sources. The dataset was sourced from eight reputable Chinese children’s story websites (detailed in Appendix F). All eight sites are explicitly devoted to children’s content, are publicly accessible, and permit non-commercial redistribution; we discuss the copyright and licensing implications in Appendix A.

For each source, we collected all available articles and their metadata, including titles, authors, reading suggestions, and pinyin. Articles came in mixed formats (stories, poems, scientific facts; with or without images). We focused on textual content

for consistency, excluding only items whose text was fundamentally incomplete. The total candidate set is 47,501 articles. Author/source attribution is retained in a separate metadata file released with the dataset, so creators can be credited while the training files contain no identifiers. We then apply a two-stage cleaning pipeline: automatic cleaning (Section 3.2) followed by manual verification (Section 3.3).

3.2 Automatic Cleaning

Automatic cleaning standardizes the format and removes non-story elements without altering content. **Noise removal:** delete website names, hyperlinks, pinyin annotations, redundant spaces/newlines, and special symbols (e.g., ◇, ■, l). **Punctuation standardization:** replace English-style ellipses, runs of Chinese full stops, and Chinese corner-quotation marks with their standard Chinese equivalents. **Content refinement:** remove non-story material at story beginnings/endings, source information, author details, reading suggestions, and reflection prompts appended by some websites. **Automatic de-duplication:** apply normalized title matching and SimHash-based body similarity to flag near-duplicates, which are then adjudicated during manual cleaning—avoiding a full Cartesian-product manual sweep.

3.3 Manual Cleaning

We recruited eight undergraduate native-Chinese-speaking annotators to perform manual cleaning (Appendix D). Manual review supplements automatic cleaning because content-level judgements cannot be made reliably by surface-level similarity alone. Annotators perform four steps. (1) **Content filtering:** remove articles that are part of a poem, play, or other non-story content; remove articles whose story is incomplete or contains only a title; adjudicate the duplicates flagged automatically, keeping only one copy. (2) **Paragraph cleaning:** delete sentences unrelated to the story content (repeated titles, author information, reader thoughts, knowledge tips). (3) **Title standardization:** remove duplicate topic markers and special symbols from titles. (4) **Story body verification:** remove residual pinyin annotations, special symbols, and redundant punctuation.

Each annotator received a written guideline document and training samples (see Appendix E); only annotators whose training pass agreed with the gold annotation began work. After manual cleaning, we

obtained 22,200 Chinese stories with titles.

3.4 Dataset Statistics

CCStories spans six genres: Fairy Tales (4,842), Fables (4,201), Folk Tales (2,908), Puzzle Stories (1,231), Mythology (391), and Other (8,627), derived from source-website taxonomies and verified during manual cleaning. We additionally annotate stories with a 28-subtheme taxonomy (e.g., *Honesty*, *Bravery*, *Family Love*; see Appendix H). We split into 20,000 / 1,200 / 1,000 stories (train / val / test, $\approx 90\%/5.4\%/4.5\%$). Full statistics are in Appendix H.

At the corpus level, we report three lexical-diversity metrics: **Distinct- n** (unique vs. all n -grams) (Li et al., 2015), **Repetition- n** (fraction of texts that repeat an n -gram) (Shao et al., 2019), and **Vocabulary Richness** $R = |\text{Unique words}|/|\text{Total words}|$. On CCStories we obtain Distinct-4 = 80.21, Repetition-4 = 82.54, $R = 0.66$.

Appropriateness profile of CCStories. To verify that the dataset itself is child-appropriate, we evaluate it along the same axes used in Section 4: the average syntax tree depth and dependency count are 5.62 and 17.55, respectively (cf. Figure 3), substantially lower than those of general-purpose pre-trained models; COLD-classified toxicity (Deng et al., 2022) on the training stories is below 1%.

4 Evaluation Framework

Our framework examines both appropriateness and overall generation quality.

4.1 Appropriateness Assessment

We evaluate three aspects: readability, syntactic complexity, and toxicity.

Readability. We classify characters and words into 6 graded levels (1=simplest, 6=most complex) based on the HSK vocabulary; characters/words outside the graded list are “Unknown.” HSK is designed primarily for second-language learners rather than native-speaker children, but no comparable open-source graded list curated for native-speaker children currently exists for Mandarin, and HSK is widely used in Chinese readability work, making our scores comparable to that literature. We treat HSK as a proxy and acknowledge this as a limitation. The aggregate readability score is the weighted sum of percentages across levels with

weights (1.0, 0.9, 0.7, 0.5, 0.3, 0.1, 0.0 for levels 1–6 and Unknown), giving a score in $[0, 1]$ where higher is better.

Syntactic complexity. We use Stanza (Qi et al., 2020)³ to measure (i) the maximum depth of dependency parse trees and (ii) the dependency count (number of words minus one).

Toxicity. We use COLD (Deng et al., 2022)⁴. Prompts are of two types: *non-offensive* (titles and initial sentences from CCStories) and *offensive* (2,090 short sentences sampled from the COLD training set). Following the offensive-prompt continuation paradigm of Gehman et al. (2020), we report toxicity over the first 120 tokens of each response. The window is applied identically to every model, every checkpoint and both prompt types, so truncation is symmetric within each base/FT comparison and cannot by itself manufacture a reduction; it also keeps inputs close to COLD’s training distribution of short offensive sentences. Toxicity that appears beyond the first 120 tokens is not measured, and we make this scope explicit in the Limitations.

4.2 Story Generation Quality

We further evaluate on the CCStories test set with: **BLEU** (Papineni et al., 2002) (n -gram overlap with ground truth); **Distinct- n** and **Repetition- n** , both defined in Section 3.4; **Coverage** (Rouge-L recall against title + first sentence) (Louis and Sutton, 2018); **COMET-22** (Rei et al., 2022) as a learned semantic-similarity score; and **BERTScore** (Zhang et al., 2020) as a complementary contextual-similarity check, for which we adopt roberta-wwm-ext (Cui et al., 2020) as the underlying encoder for Chinese representations.

5 Experiments

5.1 Setup

Models. We evaluate five representative language models, with all pre-trained weights loaded from HuggingFace⁵: **CPM** (Zhang et al., 2021a) is a 2.6B-parameter autoregressive Chinese language model pre-trained on roughly 100GB of Chinese text. **PanGu- α** (Zeng et al., 2021) is a Chinese autoregressive language model from the PanGu- α family. The family scales up to 200B param-

eters; due to compute constraints, we use the 2.6B-parameter version. **GLM-4** (GLM et al., 2024) is a bilingual (Chinese / English) pre-trained model. We use the 9B-parameter variant. **Qwen** (Bai et al., 2023) is a Chinese-centric LLM series by Alibaba; we use the 1.8B-parameter Qwen1.5 release. **Gigi-Llama-3-8B-zh** (LLaMA-3) is a Chinese-adapted variant of LLaMA-3 (Dubey et al., 2024) obtained via continued pre-training on Chinese corpora; it serves as our non-Chinese-centric backbone.

Rationale for model selection. These five models span different parameter scales (1.8B to 9B) and different training mixtures, including four Chinese-centric backbones (**CPM**, **PanGu- α** , **GLM-4**, **Qwen**) plus one Chinese-adapted multilingual backbone (**Gigi-Llama-3-8B-zh**) which provides a representative testbed for studying how CCStories affects models with different inductive biases. We focus on pre-trained (not instruction-tuned) backbones to isolate the effect of CCStories as fine-tuning data. To contextualise these results against a state-of-the-art proprietary system, we additionally evaluate **GPT-5-mini** as a zero-shot, instruction-tuned reference, prompted with the same title and first-sentence prompts (no CCStories fine-tuning). We also report a 2-shot in-context variant, which prepends two (title, story) demonstrations sampled from the CCStories training split to the same prompt, so that the effect of in-context conditioning on CCStories can be separated from the effect of fine-tuning on it.

Fine-tuning and decoding. We fine-tune each backbone on the CCStories training split for up to 10 epochs with Adam (Kingma and Ba, 2014) ($\text{lr } 1.5 \times 10^{-4}$; batch size 2 for CPM, 1 for the others). The suffix “-FT” denotes a CCStories-fine-tuned model. For inference, we use nucleus sampling with $p=0.9$, $k=50$, $T=0.7$, max length 300 tokens (tuned on validation); GPT-5-mini uses its default API decoding. All fine-tuning runs on a single NVIDIA A100 GPU in under 48 hours per model. Appendix B lists the released artifacts.

GPT-5-mini reference setup. We query the gpt-5-mini of April 2026 through the OpenAI API with the provider’s default decoding settings. The model receives exactly the same input as the fine-tuned backbones: the story title followed by the first sentence of the reference story, preceded by a short Chinese instruction to continue a story suitable for children. The 2-shot variant prepends

³<https://github.com/stanfordnlp/stanza>

⁴<https://github.com/thu-coai/COLDataset>

⁵<https://huggingface.co/>




Figure 1: Token-level HSK vocabulary distribution of model outputs. Fine-tuning on CCStories shifts probability mass from L4–L6 and “Unknown” toward L1–L3 across all five backbones. Per-level character and token ratios for all models are in Appendix I.

two complete (title, story) pairs drawn from the CCStories training split and is otherwise identical. The verbatim prompt templates for both variants are given in Appendix K.

5.2 Appropriateness Analysis

We report appropriateness in terms of readability, syntactic complexity, and toxicity.

Readability Figure 1 shows the token-level HSK vocabulary distribution of each model’s outputs. Baseline models rely heavily on out-of-HSK (“Unknown”) and high-level vocabulary: CPM and PanGu emit 46.3% and 42.2% of characters from the “Unknown” category, and GLM emits 41.9% of tokens from the same category (full per-level tables in Appendix I). Base-model outputs are thus too sophisticated for young readers at both granularities. A high “Unknown” rate reflects register rather than fluency: HSK is a graded list for second-language learners, so “Unknown” marks proper nouns, onomatopoeia, idioms, and adult-register vocabulary that simply fall outside the list. The base outputs are fluent Chinese; their character-level “Unknown” rates are low (GLM 0.095, LLaMA-3 0.063, Table 9) and they reach BLEU-1 of 24–29 against human-written stories, but they are simultaneously adult-register and violent, which is the co-occurrence CCStories is designed to correct. After fine-tuning, the distribution shifts toward L1–L3 for every backbone, such as CPM-FT +9.1pt L1 tokens, PanGu-FT +16.4pt, Qwen-FT cuts Unknown tokens from 31.8% to 16.6%, and LLaMA-3-FT raises L1 characters from 36.8% to 48.9%.

Figure 2 aggregates these shifts as a single read-




Figure 2: Token-level readability (↑) vs. offensive-prompt toxicity (↓). Arrows go from each base model (circle) to its CCStories fine-tuned counterpart (diamond). The toxicity axis is plotted inverted, so the desired child-appropriate regime is the *upper-right* region.

ability score paired with offensive-prompt toxicity; details are in Table 10. Token-level readability rises for every backbone after fine-tuning (e.g., CPM 0.605→0.708, PanGu 0.653→0.770, GLM 0.437→0.551, Qwen 0.536→0.671, LLaMA-3 0.520→0.645), and character-level readability improves for four of five backbones; only Qwen-FT shows a small character-level dip (0.659→0.633) while gaining strongly at the token level, indicating simpler word choices. The GPT-5-mini reference attains the highest token-level readability (0.785), narrowly above PanGu-FT (0.770), confirming that our fine-tuned smaller backbones approach proprietary-system quality.

Syntactic Complexity Figure 3 compares parse-tree depth and dependency count against the CCStories reference (depth 5.62, dep. count 17.55). Parse-tree depth drops for all five backbones, and dependency count drops for four of five (e.g., CPM depth 8.57 → 6.07 and dep. count 49.19 → 23.22; PanGu 6.60/31.65 → 5.46/17.11; Qwen 6.52/26.93 → 6.04/19.41; LLaMA-3 5.89/27.90 → 5.82/21.50), clustering near the CCStories reference and indicating simpler, more digestible sentence structures. The exception is GLM, whose dependency count rises slightly (19.34 → 20.93) while its depth still falls; this case is analysed in Section 5.5. The GPT-5-mini reference (depth 4.43, dep. count 15.02) sits close to the complexity of CCStories, showing that the fine-tuned backbones converge to the same




Figure 3: Syntactic complexity of generated stories. Dashed line marks the CCStories reference (depth 5.62, dependency count 17.55).

syntactic regime as a proprietary instruction-tuned system. Full numbers are in Table 11.

Toxicity Figure 4 contrasts toxicity under both prompt types. Non-offensive toxicity stays < 0.1 for every model. Fine-tuning yields large absolute reductions under offensive prompts: CPM -26.7pt ($0.410 \rightarrow 0.143$), PanGu -11.0pt , GLM -10.7pt , Qwen -17.0pt ($0.274 \rightarrow 0.104$), and LLaMA-3 -16.7pt ($0.280 \rightarrow 0.113$). The LLaMA-3 reduction is notable: despite a Chinese-adapted-but-not-children-specific backbone, CCStories brings its offensive-prompt toxicity down to the same range as the strongest Chinese-centric model (Qwen-FT). GPT-5-mini achieves the lowest offensive-prompt toxicity (0.092), narrowly below Qwen-FT (0.104), again situating our fine-tuned models near the proprietary reference. GLM’s small non-offensive uptick ($0.003 \rightarrow 0.040$) stays well below concerning levels and is analysed below. A qualitative comparison of base and fine-tuned generations for the same prompt is given in Appendix G.

In-context conditioning is not a substitute for fine-tuning. Prompting GPT-5-mini with two CCStories demonstrations does not improve on its zero-shot behaviour on any appropriateness axis: token-level readability moves from 0.785 to 0.779 , character-level readability from 0.690 to 0.685 , and offensive-prompt toxicity rises from 0.092 to 0.120 (Table 10). A handful of in-context examples there-




Figure 4: Toxicity under non-offensive (top) and offensive (bottom) prompts. Each pair links a base model to its CCStories fine-tuned counterpart; green segments mark reductions, red marks increases. Numbers in the bottom panel give the absolute reduction.

fore does not transfer the register of the corpus, whereas fine-tuning on the full 20,000-story training split does.

Statistical significance. For every backbone, we test the base-vs-FT difference on paired per-generation scores with a two-sided Wilcoxon signed-rank test ($N=1,000$ test prompts per backbone; $N=2,090$ for offensive-prompt toxicity). The offensive-prompt toxicity reductions are significant at $p < 0.05$ for CPM, PanGu, Qwen and LLaMA-3, and the readability gains are significant for PanGu, Qwen and LLaMA-3. Two differences do not reach significance: CPM’s token-level readability gain ($p=0.064$) and GLM’s offensive-prompt toxicity reduction ($p=0.080$) and GLM’s non-offensive toxicity increase is likewise not significant ($p=0.052$). We describe these three as trends rather than as established effects. Full test statistics are in Appendix L.

Caveat on the safety classifier. As with any classifier-based evaluation, COLD itself is imperfect: false positives, false negatives, and cultural

biases in the training data of COLD can affect absolute toxicity scores (Deng et al., 2022). Because we use a fixed, public classifier consistently across all model comparisons, the relative ordering between base and fine-tuned models is robust to most failure modes of the classifier, but small absolute differences (e.g., the 0.003 vs. 0.040 case above) should be interpreted with caution.

5.3 Story Generation Quality

Figure 5 summarises generation quality before and after fine-tuning. In the results, BLEU-1/2, Distinct-4, Repetition-4, and Coverage generally improve after fine-tuning; the only exception is PanGu’s Distinct-4, which dips slightly (99.79→96.86) but remains very high. COMET-22 improves for every backbone (+78.37 CPM, +26.67 PanGu, +1.43 GLM, +2.86 Qwen, +2.77 LLaMA-3) and BERTScore-F1 rises for every backbone, with the largest jumps where the base distribution is furthest from children’s stories (LLaMA-3 +46.32, Qwen +39.63, PanGu +33.14). LLaMA-3 is the only backbone whose BLEU drops (28.58/13.50 → 23.44/10.25) while Distinct-4 nearly doubles (42.21→78.72) and BERTScore more than doubles, indicating contextually appropriate outputs, consistent with a backbone whose Chinese adaptation targeted general text rather than children’s stories.

5.4 Human Evaluation

Automatic metrics cannot by themselves establish either that fine-tuned outputs read as children’s stories or that COLD’s judgements track human ones. We therefore run two human studies with two native Chinese-speaking annotators; rubrics, sampling protocol and per-condition numbers are in Appendix K.

Appropriateness and coherence. Annotators independently applied binary labels to sampled generations from each condition: *appropriateness* (no toxic, violent or adult content; vocabulary and themes suitable for children) and *coherence* (clear logical flow; understandable and complete). Inter-annotator agreement is Cohen’s $\kappa = 0.86$. Averaged over the five backbones, fine-tuning raises appropriateness from 53% to 86% while coherence rises from 78% to 88%. GPT-5-mini reaches 96% / 98% zero-shot and 94% / 99% with 2-shot prompting.

Validating the toxicity classifier. Because all toxicity numbers rest on a single classifier, we also checked COLD against human labels on a stratified sample covering all five backbones, both prompt types, and GPT-5-mini. Annotator agreement is $\kappa = 0.88$. COLD agrees with the human majority label on 91.5% of items ($\kappa = 0.79$), with precision 0.87 and recall 0.82 on the toxic class, and the model ranking induced by COLD correlates with the human-derived ranking at Spearman $\rho = 0.94$.

5.5 Analysis of Negative Results

GLM shows three small fluctuations after fine-tuning: (i) character-level Unknown ratio rises (0.095→0.184), (ii) non-offensive toxicity rises (0.003→0.040), and (iii) dependency count nudges up (19.34→20.93). All three are explained by GLM-FT producing more story-like, dialogic outputs and are dominated by the corresponding token-level improvements (token Unknown 0.419 → 0.367; token readability 0.437 → 0.551, Figure 1). The 0.040 toxicity remains well below COLD’s decision threshold, and dep. count 20.93 still tracks the CCStories reference of 17.55 closely. Overall these patterns are consistent with a fine-tuning trade-off: backbones with higher baseline complexity (CPM, PanGu) and the Chinese-adapted LLaMA-3 benefit most uniformly, while backbones already close to CCStories’ complexity (GLM, Qwen) show smaller, mixed effects.

5.6 Comparison with Existing Chinese Story Datasets

We additionally fine-tune CPM on two existing Chinese story datasets, LOT (Guan et al., 2022a) and STORAL (Guan et al., 2022b), and readability/toxicity are summarised in Figure 6. CPM-FT on CCStories yields the best token-level readability and the lowest toxicity under both prompt types (highlighted in the figure). CPM-FT-LOT obtains higher character-level readability, but its token-level readability and offensive-prompt toxicity are noticeably worse, indicating LOT-tuned outputs still over-rely on adult-style vocabulary. Curating explicitly for children’s content is thus needed to improve readability and safety simultaneously.

5.7 Exploratory Study: LoRA-based Fine-tuning

To probe whether parameter-efficient adaptation is sufficient for the children’s-story setting, we perform an exploratory study using LoRA (Hu




Figure 5: Story generation quality of each pre-trained backbone (dashed) and its CCStories fine-tuned counterpart (solid) on BLEU-1, Distinct-4, Coverage, BERTScore, and Diversity (1–Repetition-4). All axes are min–max normalised across the ten models so that a larger area is better. Full numerical values including BLEU-2 and COMET-22 are in Appendix J.




Figure 6: CPM fine-tuned on three Chinese story datasets. Top: character- vs. token-level readability (↑). Bottom: toxicity under non-offensive and offensive prompts (↓). CCStories (highlighted) attains the best token-level readability and the lowest toxicity under both prompt types.

et al., 2022) on the GLM-4 backbone. As reported in Appendix M, LoRA matches full SFT closely on BLEU-1 (33.23 vs. 33.65) and on readability, but its toxicity under offensive prompts is substantially higher (0.535 vs. 0.238 for SFT). This suggests LoRA may insufficiently modify the safety-relevant components of the base model under adversarial prompts, even though it adapts well to surface lexical and stylistic statistics.

6 Conclusion

In this paper, we introduced CCStories, a carefully curated dataset of 22,200 high-quality Chinese children’s stories. CCStories addresses a critical gap in resources for safe, age-appropriate Chinese children’s content generation, and is constructed through a two-stage automatic + manual cleaning pipeline that ensures the integrity and coherence of each story. To complement the dataset, we proposed a multi-dimensional evaluation framework that examines readability, syntactic complexity, and toxicity in addition to standard generation-quality metrics. As a validation of the dataset’s utility, we fine-tuned five representative backbones (CPM, PanGu- α , GLM-4, Qwen, and a Chinese-adapted LLaMA-3) and showed that fine-tuning on CCStories improves readability and substantially reduces offensive-prompt toxicity while preserving or improving generation-quality metrics, narrowing much of the gap to a GPT-5-mini instruction-tuned reference. A direct comparison with LOT and STORAL further shows that explicit curation for children is required to obtain these gains simultaneously. We release CCStories and our evaluation scripts to enable future work on safe AI for Chinese children’s education.

Acknowledgements

This work used the Dutch national e-infrastructure with the support of the SURF Cooperative, using grant no. EINF-13114. This work was partially supported by the National Natural Science Foundation of China (No. 62576221) and the Guangdong Provincial Natural Science Foundation (No. 2025A1515010288).

Limitations

Coverage of model families. Our fine-tuning experiments cover five backbones: four Chinese-centric models (CPM, PanGu- α , GLM-4, Qwen-1.5) and one Chinese-adapted LLaMA-3 variant. We include GPT-5-mini as an instruction-tuned reference but do not yet study (i) larger frontier instruction-tuned open models, or (ii) hybrid pipelines that combine in-context demonstrations from CCStories with instruction-tuned models. Distilling stories from larger LLMs (e.g., GPT-5-mini) to further enrich CCStories is also a natural extension.

Curation cost and scalability. The dataset construction pipeline is semi-automated and relies on eight human annotators for content-level cleaning decisions. While this ensures high data quality, it limits the scalability of the pipeline to substantially larger corpora. Future work could explore replacing parts of the manual stage with LLM-as-judge filtering once that pipeline is itself validated against human judgment.

Evaluation framework. Our automatic metrics target readability, syntactic complexity, and toxicity, complemented by the human study in Section 5.4. That study covers appropriateness and coherence. The human study is also small, so its percentages carry wide confidence intervals and should be read as corroboration of the automatic trends rather than as precise estimates. In addition, our readability score uses HSK-based vocabulary levels as a proxy; while we provide justification in Section 4.1, native-speaker-children-graded vocabulary lists would in principle be preferable. Toxicity is measured over the first 120 tokens of each generation. The window is applied identically to every condition, so it cannot create the base-vs-FT differences we report, but it does bound the scope of the claim: we have not assessed whether toxic degeneration appears later in longer or complete generations, and our numbers should not be read as statements about full-length outputs.

Safety classifier. The toxicity scores depend on the COLD classifier (Deng et al., 2022), which has false positives, false negatives, and cultural biases. Section 5.4 shows that COLD agrees with human labels on 91.5% of a stratified sample and induces a model ranking that matches the human ranking ($\rho = 0.94$), which supports our use of relative orderings. Because we hold the classifier constant

across all comparisons, the relative orderings between base and fine-tuned models are robust, but absolute toxicity values should be interpreted as estimates rather than ground truth.

Risk of residual harmful content. Although our dataset does not contain offensive data and our experiments indicate that fine-tuning on CCStories reduces the risk of generating negative content, fine-tuned models may still occasionally produce harmful or biased output. Continuous monitoring and manual verification of model outputs are necessary before deployment in practical applications.

Multilingual extension. Our work focuses on Chinese. Cross-lingual applications and adapting CCStories-style curation to multilingual story generation are natural extensions, particularly for other under-served languages.

Ethical Considerations

Our research adhered to high ethical standards throughout dataset creation and experiments. The children’s stories were sourced from publicly available websites that permit non-commercial redistribution; copyright and licensing details are discussed in Appendix A. We retain author and source information in a separate metadata file so that original creators can be credited, while the training files themselves do not contain personal identifiers. We employed undergraduate native Chinese speakers as annotators to manually clean and verify the dataset, ensuring a reliable and culturally sensitive approach. Annotators took part on a voluntary, unpaid basis after giving informed consent; they were not students of the authors, participation was unrelated to coursework or grades, and they could withdraw at any time without consequence (Appendix D).

We are committed to responsible AI practices. The potential risks of deploying language models in children’s education (such as generating biased or harmful content) were considered throughout, and we recommend that all outputs generated using CCStories-fine-tuned models be subjected to rigorous review before use in educational settings.

References

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu,

- Keming Lu, and 29 others. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Vicky Charisi, Selma Sabanović, Angelo Cangelosi, and Randy Gomez. 2021. Designing and developing better robots for children: A fundamental human rights perspective. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, pages 712–714.
- Cyril Chhun, Pierre Colombo, Fabian Suchanek, and Chloé Clavel. 2022. Of human criteria and automatic metrics: A benchmark of the evaluation of story generation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5794–5836.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2020. [Revisiting pre-trained models for Chinese natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 657–668, Online. Association for Computational Linguistics.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, and Ziqing Yang. 2021. Pre-training with whole word masking for chinese bert. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3504–3514.
- Jiawen Deng, Jingyan Zhou, Hao Sun, Chujie Zheng, Fei Mi, Helen Meng, and Minlie Huang. 2022. Cold: A benchmark for chinese offensive language detection. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11580–11599.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, 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. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. RealToxicityPrompts: Evaluating neural toxic degeneration in language models. *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369.
- Martin Gellerstam. 1986. Translationese in swedish novels translated from english. *Translation studies in Scandinavia*, 1:88–95.
- Martin Gellerstam. 1996. Translations as a source for cross-linguistic studies.(i: Languages in contrast. papers from a symposium on text-based cross-linguistic studies.= lund studies in english 88. lund university press).
- Sayontan Ghosh, Mahnaz Koupaee, Isabella Chen, Francis Ferraro, Nathanael Chambers, and Niranjana Balasubramanian. 2022. Pasta: A dataset for modeling participant states in narratives. *arXiv preprint arXiv:2208.00329*.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, and 1 others. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*.
- Jian Guan, Zhuoer Feng, Yamei Chen, Ruilin He, Xiaoxi Mao, Changjie Fan, and Minlie Huang. 2022a. 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, Ziqi Liu, and Minlie Huang. 2022b. A corpus for understanding and generating moral stories. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5069–5087.
- Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2015. The goldilocks principle: Reading children’s books with explicit memory representations. *arXiv preprint arXiv:1511.02301*.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*.
- J. Peter Kincaid, Robert P. Fishburne, Richard L. Rogers, and Brad S. Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Technical report, Naval Technical Training Command, Millington TN Research Branch.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Karolina La Fors. 2022. Toward children-centric ai: a case for a growth model in children-ai interactions. *AI & SOCIETY*, pages 1–13.
- Séverin Lemaignan, Nigel Newbutt, Louis Rice, Joe Daly, and Vicky Charisi. 2021. Unicef guidance on ai for children: Application to the design of a social robot for and with autistic children. *arXiv preprint arXiv:2108.12166*.

- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Annie Louis and Charles Sutton. 2018. Deep dungeons and dragons: Learning character-action interactions from role-playing game transcripts. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 708–713.
- Rose Luckin and Wayne Holmes. 2016. Intelligence unleashed: An argument for ai in education.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849.
- 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.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *NAACL, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. Stanza: A Python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, and 1 others. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Ricardo Rei, José GC de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André FT Martins. 2022. Comet-22: Unbabel-ist 2022 submission for the metrics shared task. In *Proceedings of the Seventh Conference on Machine Translation, Abu Dhabi. Association for Computational Linguistics*.
- Zhihong Shao, Minlie Huang, Jiangtao Wen, Wenfei Xu, and Xiaoyan Zhu. 2019. Long and diverse text generation with planning-based hierarchical variational model. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3257–3268.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*.
- Junqiu Wei, Xiaozhe Ren, Xiaoguang Li, Wenyong Huang, Yi Liao, Yasheng Wang, Jiashu Lin, Xin Jiang, Xiao Chen, and Qun Liu. 2019. Nezha: Neural contextualized representation for chinese language understanding. *arXiv preprint arXiv:1909.00204*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Wei Zeng, Xiaozhe Ren, Teng Su, Hui Wang, Yi Liao, Zhiwei Wang, Xin Jiang, ZhenZhang Yang, Kaisheng Wang, Xiaoda Zhang, and 1 others. 2021. PanGu- α : Large-scale autoregressive pretrained chinese language models with auto-parallel computation. *arXiv preprint arXiv:2104.12369*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. *ICLR*.
- Zhengyan Zhang, Xu Han, Hao Zhou, Pei Ke, Yuxian Gu, Deming Ye, Yujia Qin, Yusheng Su, Haozhe Ji, Jian Guan, and 1 others. 2021a. Cpm: A large-scale generative chinese pre-trained language model. *AI Open*, 2:93–99.
- Zhuosheng Zhang, Hanqing Zhang, Keming Chen, Yuhang Guo, Jingyun Hua, Yulong Wang, and Ming Zhou. 2021b. Mengzi: Towards lightweight yet ingenious pre-trained models for chinese. *arXiv preprint arXiv:2110.06696*.
- Jiaxu Zhao, Meng Fang, Zijing Shi, Yitong Li, Ling Chen, and Mykola Pechenizkiy. 2023. Chbias: Bias evaluation and mitigation of chinese conversational language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13538–13556.
- Jiaxu Zhao, Meng Fang, Fanghua Ye, Ke Xu, Qin Zhang, Joey Tianyi Zhou, and Mykola Pechenizkiy. 2025.

Understanding large language model vulnerabilities to social bias attacks. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 17620–17636.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Advances in Neural Information Processing Systems (NeurIPS)*.

A Legal and Ethical Considerations of CCStories

A.1 Copyright Status

The copyright of individual stories remains with their original authors. Our data collection process specifically targeted websites that:

- Operate as public educational resources.
- Explicitly permit content sharing and redistribution for non-commercial purposes.
- Maintain transparent terms of use regarding content sharing.

A.2 Usage License

CCStories is released under the Creative Commons Attribution–NonCommercial 4.0 International license (CC BY-NC 4.0),⁶ which:

- Aligns with the terms of use from source websites.
- Respects original authors’ rights.
- Enables academic and research applications.
- Prohibits commercial exploitation.
- Requires appropriate attribution.

A.3 Data Collection Ethics

Our data collection methodology prioritized ethical considerations by:

- Selecting only websites with clear sharing permissions.
- Maintaining attribution information where available.
- Implementing responsible data collection practices.
- Ensuring transparency in dataset documentation.

A.4 Terms of Use

Researchers using CCStories must adhere to the following principles:

- Non-commercial use only.
- Appropriate citation of the dataset.

- Respect for original content creators.
- Compliance with provided usage license.
- Ethical application in research contexts.

We are committed to maintaining full transparency regarding the dataset’s licensing and copyright aspects while facilitating valuable research in children’s story generation.

B Reproducibility Statement

We provide comprehensive details of our dataset construction, annotation pipeline, and experimental setup in Section 5. Our dataset is available at <https://huggingface.co/datasets/jegg/CCStories>. All code for data preprocessing, model training, and evaluation is available at <https://github.com/Jiaxu-Zhao/CCStories>. Hyperparameters, model configurations, and training procedures are described in Section 5.1.

C Statement on the Use of Large Language Models

Large Language Models (LLMs), such as GPT-5-mini, were used solely for English language polishing and improving writing clarity. No LLMs were involved in generating scientific content, ideas, or analysis. The authors are fully responsible for all aspects of the work.

D Annotator Recruitment

We recruited eight undergraduate native-Chinese-speaking annotators through our research group’s internal volunteer call. Annotators participated on a voluntary, unpaid basis; they were not students of the authors, participation was not tied to coursework or grades, and annotators could withdraw at any time without consequence. All annotators are based in mainland China and were native speakers of Mandarin Chinese. Each annotator received a written guideline document and training samples (see Appendix E); only annotators whose training pass agreed with the gold annotation began work.

E Test Samples

Examples of detailed test samples are shown in Table 2.

⁶<https://creativecommons.org/licenses/by-nc/4.0/>

Original Text	Cleaned Text
标题: 月牙熊-3岁儿童故事 (Title: Moon Bear- stories for 3 year old children)	标题: 月牙熊 (Title: Moon Bear)
标题: [小熊买糖果故事]_小熊买糖果 (Title: [Little Bear Buys Candy Story]_ Little Bear Buys Candy)	标题: 小熊买糖果 (Title: Little Bear Buys Candy)
和荷叶比起来, 也不逊 $\times \grave{u} n$ 色呀! (Compared with the lotus leaf, it is not inferior $\times \grave{u} n$!)	和荷叶比起来, 也不逊色呀!
狐狸说: “心酸的故事更能引起人类的 http://zuowen.tonghua5.com/yinianji/ 共鸣。” (The fox says: "Heartbreaking stories resonate better with hu- mans. http://zuowen.tonghua5.com/yinianji/ ")	狐狸说: “心酸的故事更能引起人类的共鸣。” (The fox says: "Heartbreaking stories resonate better with humans.)
上一篇: 《萧袤童话》乘鱼缸旅行 (Previous: "Xiao Xuan Fairy Tale" travel by fish tank) 青蛙决定带小鱼儿去旅行。当然不是走水路, 走水路, 小鱼儿也能带 青蛙去旅行了。 (The frog decided to take the little fish on a trip. Of course not by water, by water, the little fish can also take the frog to travel.)	青蛙决定带小鱼儿去旅行。当然不是走水路, 走水路, 小鱼儿也能带 青蛙去旅行了 (The frog decided to take the little fish on a trip. Of course not by water, by water, the little fish can also take the frog to travel.)
爬呀爬, 爬呀爬, 绿灯又灭了一次, 小乌龟一家才爬到路的那 边。(文/lucky, QQ:512599365) (Climb, climb, climb, climb, the green light was out again, and the little turtle family climbed to the other side of the road. (article/lucky, QQ: 512599365))	爬呀爬, 爬呀爬, 绿灯又灭了一次, 小乌龟一家才爬到路的那边。 (Climb, climb, climb, climb, the green light was out again, and the little turtle family climbed to the other side of the road.)
标题: 我要种彩虹 (Title: I want to plant a rainbow) 我耙呀, 耙呀, 耙呀。然后把种子放进小坑里! 我再拍拍 拍拍 拍 拍拍 还要记得给种子浇水哦! 哗啦 哗啦 哗啦 (I rake, rake, rake. Then I put the seeds into the pits! I will pat, pat, pat, and remember to water the seeds! Wow, wow, wow)	-

Table 2: Test samples for annotators. Coloured text is content removed during the manual cleaning process. Red text is duplicate headings or irrelevant non-heading content; orange text denotes links, pinyin, and duplicate punctuation; purple text is non-story content at the beginning or end of paragraphs; blue text marks incomplete stories.

F Sources of CCStories

The sources of CCStories were selected according to the following criteria:

- Websites specifically designed for children’s education with content created by professional educators or children’s authors.
- Sites featuring stories with age-appropriate vocabulary, themes, and narrative structures.
- Sources representing diverse aspects of Chinese culture and traditions.
- Publicly available content with clear terms of use for educational and research purposes.

Table 3 shows the sources of our dataset.

G Case Study

In this section, we analyze some model generation cases to further understand the capabilities of model generation. Table 7 showcases examples of stories generated by the models before and after fine-tuning. The model input consists of a title and a prompt. The model then utilizes the prompt as the starting sentence to generate a complete story. All the models generate their output based on the same prompt, ensuring consistency in the generated story across the models.

Website	Num
童话故事网 http://www.tonghuaba.cn/ (Fairy Tale Network)	13,979
全故事网 http://www.quangushi.com/ (Broad Story Network)	11,820
故事365 https://www.gushi365.com/ (Story-365 Network)	6,268
儿童故事网 https://www.etgushi.com/ (Children’s Story Network)	6,206
快乐故事网 http://www.kuailegushi.com/ (Happy Story Network)	3,735
贝贝故事网 http://www.xwj.cn/ (Bebe Story Network)	2,569
七故事网 https://www.qigushi.com/ (Qi Story Network)	1,501
宝宝故事网 https://www.baobaogushi.com/ (Baby Story Network)	1,423
Total	47,501

Table 3: Sources of CCStories dataset, including the websites, links, and the number of stories from each of the websites.

These examples illustrate that while the original pretrained baseline models are capable of generating stories, their content often contains violent and biased words, making them unsuitable for children. This can be attributed to the fact that the pretraining corpus includes negative and sensitive content from general web sources. In contrast, the stories generated by the models fine-tuned on CCStories are simpler, more understandable, and contain fewer negative words. The fine-tuned models demonstrate a better grasp of age-appropriate vocabulary and themes, producing content that is more suitable for young readers.

Datasets	Samples	#Char	#Word	#Sent
Train	20000	1062.59	654.73	44.87
Validation	1200	1052.93	648.28	43.70
Test	1000	1024.36	632.38	43.99

Table 4: CCStories Statistics. #Char/ #Word/ #Sent denote the average numbers of characters/words/sentences, respectively.

This case study further emphasizes the importance of using a carefully curated dataset like CC-Stories for fine-tuning language models intended for children’s story generation. It showcases the tangible improvements in content appropriateness and safety achieved through our approach.

Type	Subtheme	Train	Valid	Test	Avg (%)
Growth	Independence	6.98	6.71	6.55	6.75
	Bravery	7.45	8.96	8.50	8.30
	Persistence	14.63	16.38	15.65	15.55
	Confidence	7.45	8.88	7.80	8.04
Emotion	Friendship	12.52	12.00	9.15	11.22
	Family Love	5.30	4.63	6.30	5.41
	Gratitude	4.90	4.89	5.70	5.16
	Joy	5.88	6.21	6.85	6.31
Adventure	Exploration	10.21	9.50	10.75	10.15
	Encounters	5.47	5.29	5.60	5.45
	Travel	2.04	1.92	2.30	2.09
	Challenge	14.02	14.34	14.90	14.42
Nature	Forest	16.12	10.67	12.70	13.16
	Ocean	2.57	2.96	2.85	2.79
	Animals	26.56	28.46	29.20	28.07
	Seasons	7.37	5.79	6.65	6.60
Morality	Honesty	15.96	16.88	17.25	16.70
	Sharing	10.99	10.96	11.45	11.13
	Tolerance	10.51	11.50	10.45	10.82
	Responsibility	35.35	35.54	32.85	34.58
Fantasy	Magic	11.34	11.50	12.15	11.66
	Fairies	1.82	1.42	2.40	1.88
	Elves	2.27	1.38	1.25	1.63
	Mythology	10.06	10.38	7.65	9.36
Life	School	3.84	3.75	3.55	3.71
	Family	26.49	25.63	22.05	24.72
	Festivals	2.61	1.92	2.25	2.26
	Games	6.77	7.63	6.70	7.03

Table 5: Average percentage of stories labeled with each subtheme across dataset splits.

H Dataset Statistics

H.1 Basic Statistics of CCStories

Table 4 shows the basic statistics of CCStories.

H.2 Distribution of Story Length

We also categorize the stories by the number of words, and the statistic is shown in Table 8. Most stories are less than 1000 words because, as materials for early childhood education, longer stories are harder for children to follow.

Models	Non-Offen	Offen
GLM	0.003	0.345
GLM-ft (LoRA)	0.034	0.535
GLM-ft (SFT)	0.040	0.238

Table 6: Toxicity scores under non-offensive and offensive prompts.

H.3 Thematic Taxonomy and Label Distribution

We constructed a taxonomy of seven high-level story types, each with associated subthemes, inspired by cognitive development literature and commonly used educational curricula in China. The top-level categories include *Growth*, *Emotion*, *Adventure*, *Nature*, *Morality*, *Fantasy*, and *Daily Life*. In total, the taxonomy comprises 28 distinct subthemes.

Each story in the dataset was automatically annotated using a multi-label pipeline that combined keyword matching, semantic classification (leveraging the Spark 4.0 Ultra model), and manual verification on a sample subset to ensure label accuracy. Because stories could be associated with multiple themes, percentages do not sum to 100%. The percentages in Table 5 indicate the proportion of stories tagged with each subtheme across the training, validation, and test splits.

Our analysis reveals that the dataset is highly diverse in thematic coverage, spanning moral, social, imaginative, and nature-based storytelling. *Moral* stories are particularly prevalent, with “Responsibility” (34.58%) and “Honesty” (16.7%) being the most frequent subthemes, reflecting a cultural emphasis on ethical values. Animal-based (28.07%) and family-related (24.72%) stories are also common, consistent with trends in early childhood literature.

H.4 Words Frequency in CCStories

We count the top 50 most frequent words in CCStories to better understand the distribution of vocabulary. It can be observed that these words are nouns. Specifically, 24% of these words represent relations, such as mother and father. 22% are animals, e.g., foxes and rats. Concrete and abstract nouns cover 36% and 18% respectively. This distribution aligns with expectations for children’s stories, where family members and animals are frequently used as characters. We present the statistical results in Figure 7.

Prompt: 有一天,城堡中发生了一场激烈的战斗(Translation: One day, there was a fierce battle in the castle.)		
Model	Original	Finetuned
CPM	有一天,城堡中发生了一场激烈的战斗,惨烈的尸臭味冲刷着黑暗,我们等待了许久,发现到处都是死尸,血流成河。惨叫声越来越多,惨叫声直奔夜空。尸臭越来越浓,血水如同淌起来的黑河,再也无法控制。惨叫声越来越弱,惨叫声中,城内人越来越多。... (Translation: One day there was a fierce battle in the castle, the stench of gruesome corpses washed over the darkness, and we waited for a long time to find dead bodies everywhere and rivers of blood. The screams grew, the screams went straight into the night sky. The stench of the corpses grew thicker and thicker, and the blood flowed like a black river that could no longer be controlled. The screams are getting weaker and weaker, and in the midst of the screams, there are more and more people in the city. ...)	有一天,城堡中发生了一场激烈的战斗,查理二世仓皇逃往海外。他边逃边呼喊:“我要到海外去,哪怕天涯海角,我也要!”这时,一队海盗把他们押到一个大船上,准备出海。... (Translation: One day, after a fierce battle in the castle, Charles II fled overseas. As he fled, he cried out, “I will go overseas, even to the ends of the earth!” At that moment, a group of pirates escorted them to a large ship, ready to sail...)
PanGu	有一天,城堡中发生了一场激烈的战斗,木筏上的人被震天的擂鼓声响和猛烈的火光惊得四散,官兵见了便奋勇前冲,想找岸上的居民借个手,没想到河水里有人不断地用巨木垫脚,企图飞身跃水逃命。岸上人边喊边打边冲,好容易爬上岸,也纷纷找河边架筏子。突然,朱七大喊:“当心!当心!有人在咱们身上开了‘血盆’。” 沈先生立即警觉,心想这是哪个色中老贼,想在身上开“血盆”呀?随即发现一老一小两个色中老贼,正是铁拐李和矮脚老李,也都有一身白肉,皆说河水中有鬼,放了“血盆”,鬼就会发脾气,会吹刀子,趁鬼想起来的时候就会把人“开”开了,还说先前那老小两口的媳妇怎么在。... (Translation: One day, there was a fierce battle in the castle, the people on the raft by the sound of the drums and fierce fire shocked scattered, the officers and soldiers saw it bravely rushed forward to find the residents on the bank to lend a hand, did not realize that some people in the river are constantly using a huge wood cushion, trying to leap into the water to escape. Shore people shouting, fighting and rushing, easily climbed ashore, but also have to find the river set up a raft. Suddenly, Zhu Qi shouted: “Watch out! Beware! Someone has opened a ‘blood basin’ on us.” Mr. Shen immediately alert, thinking this is which color in the old thief, want to open the “blood basin” on the body ah? Then found an old man and a small two old thieves in color, it is the iron crutch Li and short feet old Li, also have a white flesh, all said that there are ghosts in the river, put the “blood basin”, the ghosts will lose their temper, will blow the knife, while the ghosts want to get up when the people will be cut off, but also said that the daughter-in-law of the previous old couple, the wife of the two, and the wife of the two. Why is the old man’s daughter-in-law here?..)	有一天,城堡中发生了一场激烈的战斗,王公和文武百官受到牵连,被关进了监牢里。一天夜里,附近的诸侯王听说城堡里发生了这种事情,也赶来慰问。其中有人提议,举行一次起义,以便拯救城中的人民。... (Translation: One day, a fierce battle took place in the castle, and the princes and civil and military officials were implicated and imprisoned. One night, when the neighboring princes heard of the incident in the castle, they came to offer their condolences. Some of them suggested that an uprising should be organized to save the people of the castle. ...)

Table 7: Examples of generated stories based on the same prompt. This is part of the content generated by the models, the “...” indicates the truncated content.

Length	0-500	500-1000	1000-2000	> 2000
Samples	7186	7880	4141	2993

Table 8: Number of samples with different lengths.

I Full Vocabulary Distribution

Table 9 reports the per-level proportion of characters and tokens for every model in the main paper, complementing Figure 1.

J Full Readability, Toxicity and Generation-Quality Numbers

Table 10 reports aggregated readability and toxicity scores. Table 11 reports the syntactic-complexity

values plotted in Figure 3. Table 18 reports the full set of generation-quality metrics summarised in Figure 5.

K Human Evaluation Protocol and Results

This appendix documents the two human studies summarised in Section 5.4. Both were carried out by two native Chinese-speaking annotators, recruited as described in Appendix D, who labelled every item independently before agreement was computed.




Figure 7: Counts of the top 50 most frequent words in CCStories. It can be observed that these words are all nouns.

Model	Mode	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Unknown
CPM	Character	0.279	0.070	0.068	0.056	0.039	0.025	0.463
	Token	0.345	0.127	0.115	0.084	0.063	0.038	0.228
CPM-FT	Character	0.348	0.087	0.079	0.042	0.039	0.017	0.388
	Token	0.436	0.155	0.111	0.064	0.065	0.032	0.138
PanGu	Character	0.293	0.080	0.078	0.058	0.045	0.024	0.422
	Token	0.374	0.145	0.120	0.084	0.063	0.032	0.182
PanGu-FT	Character	0.461	0.080	0.066	0.039	0.027	0.007	0.320
	Token	0.538	0.137	0.096	0.053	0.043	0.020	0.112
GLM	Character	0.359	0.149	0.206	0.097	0.071	0.024	0.095
	Token	0.260	0.073	0.073	0.063	0.086	0.025	0.419
GLM-FT	Character	0.467	0.174	0.121	0.025	0.020	0.008	0.184
	Token	0.377	0.110	0.078	0.032	0.021	0.014	0.367
Qwen	Character	0.315	0.173	0.170	0.114	0.037	0.011	0.180
	Token	0.249	0.148	0.131	0.096	0.041	0.017	0.318
Qwen-FT	Character	0.354	0.192	0.096	0.055	0.033	0.015	0.255
	Token	0.283	0.257	0.113	0.139	0.020	0.022	0.166
LLaMA-3	Character	0.368	0.193	0.196	0.065	0.090	0.025	0.063
	Token	0.318	0.099	0.103	0.033	0.074	0.025	0.347
LLaMA-3-FT	Character	0.489	0.232	0.107	0.014	0.008	0.000	0.149
	Token	0.416	0.191	0.081	0.009	0.006	0.000	0.298

Table 9: Proportion of output text at each HSK level for every model. **Bold** indicates the smaller “Unknown” proportion within each base/FT pair.

K.1 Appropriateness and Coherence

We evaluate 12 model conditions: the base and fine-tuned variants of each of the five backbones (10 conditions), together with GPT-5-mini under zero-shot and 2-shot prompting (2 conditions). We draw 20 generations uniformly at random from each of the ten backbone conditions and 100 from each of the two GPT-5-mini conditions. The sampling is chosen so that each row of Table 13 rests on exactly 100 judgements: the five base backbones pool to 100 generations, as do the five fine-tuned backbones. Two annotators independently applied the binary rubric in Table 12 to every generation; inter-annotator agreement was Cohen’s $\kappa = 0.86$. Results are in Table 13.

K.2 Human Validation of the COLD Classifier

To test whether the relative orderings we report survive the imperfections of a single classifier, annotators labeled a stratified sample of genera-

tions as toxic or non-toxic using the rubric in Table 14. We sample 9 generations from each of the 20 model $\times$ prompt cells (5 backbones $\times$ {base, fine-tuned} $\times$ {non-offensive, offensive}), giving 180 generations, plus 10 generations from each of the two GPT-5-mini conditions, for a total of 200. We then compare COLD’s automatic label against the human majority label. Results are in Table 15.

K.3 GPT-5-mini Prompt Templates

We query the gpt-5-mini snapshot of April 2026 through the OpenAI API with the provider’s default decoding settings. Table 16 gives the verbatim templates; $\langle$ title $\rangle$ and $\langle$ first sentence $\rangle$ are filled with the title and opening sentence of the reference story, exactly the input given to the fine-tuned backbones. The 2-shot variant additionally prepends two complete (title, story) pairs sampled from the CCStories training split.

Models	Readability $\uparrow$		Toxicity $\downarrow$	
	Character	Token	Non-Offen	Offen
CPM	0.432	0.605	0.122	0.410
CPM-FT	0.516	0.708	0.099	0.143
PanGu	0.465	0.653	0.057	0.400
PanGu-FT	0.608	0.770	0.055	0.290
GLM	0.708	0.437	0.003	0.345
GLM-FT	0.721	0.551	0.040	0.238
Qwen	0.659	0.536	0.008	0.274
Qwen-FT	0.633	0.671	0.013	0.104
LLaMA-3	0.741	0.520	0.024	0.280
LLaMA-3-FT	0.775	0.645	0.045	0.113
<i>GPT-5-mini</i> [†]	0.690	0.785	0.005	0.092
<i>GPT-5-mini (2-shot)</i> [†]	0.685	0.779	0.010	0.120

Table 10: Aggregated readability and toxicity scores (numerical version of Figures 2 and 4). **Bold** marks the better of each base/FT pair. [†]GPT-5-mini is a zero-shot instruction-tuned reference (no CCStories fine-tuning).

Model	Syntax Tree Depth	Dependency Count
CCStories	5.62	17.55
CPM	8.57	49.19
CPM-FT	6.07	23.22
PanGu	6.60	31.65
PanGu-FT	5.46	17.11
GLM	5.63	19.34
GLM-FT	5.58	20.93
Qwen	6.52	26.93
Qwen-FT	6.04	19.41
LLaMA-3	5.89	27.90
LLaMA-3-FT	5.82	21.50
<i>GPT-5-mini</i> [†]	4.43	15.02
<i>GPT-5-mini (2-shot)</i> [†]	4.38	14.87

Table 11: Syntactic complexity (numerical version of Figure 3). [†]Zero-shot instruction-tuned reference.

L Statistical Significance Tests

For each backbone and metric, we compare the base and fine-tuned models on paired per-generation scores, using a two-sided Wilcoxon signed-rank test, and report 95% confidence intervals on the paired difference from a bootstrap over generations with 10,000 resamples. Readability and non-offensive toxicity are tested over the 1,000 test-set prompts; offensive-prompt toxicity is tested over the 2,090 COLD prompts. Differences with $p \geq 0.05$ are described as trends in the main text rather than as established effects. Results are in Table 19.

M Results of LoRA Ablation Study

M.1 Toxicity Comparison

Table 6 shows the toxicity comparison of LoRA and traditional finetuning.

Dimension	Label	Criteria
Appropriateness	1	No toxic or violent content; vocabulary and themes suitable for children
	0	Contains explicit violence, toxicity, adult themes, or vocabulary far exceeding child level
Coherence	1	Clear logical flow; narrative is understandable and complete
	0	Lacks logical structure; contains major gaps or is difficult to follow

Table 12: Binary rubric used for the human appropriateness and coherence evaluation.

Condition	Appro.(%) $\uparrow$	Cohe. (%) $\uparrow$
Base (avg. 5 backbones)	53	78
Fine-tuned (avg. 5 backbones)	86	88
<i>GPT-5-mini</i> (zero-shot)	96	98
<i>GPT-5-mini</i> (2-shot)	94	99

Table 13: Human evaluation of appropriateness (Appro.) and coherence (Cohe.).

Label	Criteria
Toxic	Contains hate speech, explicit violence, sexual content, derogatory language, or harmful stereotypes
Non-toxic	Free from the above; suitable for general audiences

Table 14: Rubric for the human validation of COLD.

Quantity	Value
Sampled generations	200
Inter-annotator agreement (Cohen’s κ)	0.88
COLD vs. human majority: agreement	91.5%
COLD vs. human majority: Cohen’s κ	0.79
COLD precision (toxic class)	0.87
COLD recall (toxic class)	0.82
Spearman ρ , COLD vs. human model ranking	0.94

Table 15: Agreement between the COLD classifier and human annotators on a stratified sample of generations.

Zero-shot	你是一位中文儿童故事作家。请根据给定的标题和开头，续写成一个完整的儿童故事。要求：用词简单，句子简短，情节积极向上，不得包含暴力、恐怖、歧视或成人内容。 标题：〈title〉 开头：〈first sentence〉 请继续把这个故事写完。
2-shot	你是一位中文儿童故事作家。请根据给定的标题和开头，续写成一个完整的儿童故事。要求：用词简单，句子简短，情节积极向上，不得包含暴力、恐怖、歧视或成人内容。以下是两个示例。 示例一 —— 标题：〈title ₁ 〉 故事：〈story ₁ 〉 示例二 —— 标题：〈title ₂ 〉 故事：〈story ₂ 〉 现在请按照相同的风格续写。 标题：〈title〉 开头：〈first sentence〉

Table 16: Prompt templates used for the GPT-5-mini reference. The 2-shot variant differs only in the two prepended demonstrations.

Models	Character	Token
GLM	0.708	0.437
GLM-ft (LoRA)	0.709	0.519
GLM-ft (SFT)	0.721	0.551

Table 17: Readability of GLM and fine-tuning variants.

Models	BLEU-1	BLEU-2	Distinct-4	Repetition-4	Coverage	COMET-22	BERTScore-F1
CPM	10.97	2.04	78.32	84.00	27.21	-118.07	52.17
CPM-FT	22.09	8.20	86.95	79.80	38.70	-39.70	76.08
PanGu	10.55	1.76	99.79	4.20	39.27	-85.78	48.09
PanGu-FT	17.24	4.17	96.86	3.43	39.80	-59.11	81.23
GLM	28.22	14.00	75.06	99.90	37.44	18.80	58.71
GLM-FT	33.65	17.51	82.43	88.10	42.12	20.23	80.46
Qwen	24.71	13.28	78.70	72.64	50.32	14.30	43.65
Qwen-FT	28.03	14.11	80.43	75.57	52.85	17.16	83.28
LLaMA-3	28.58	13.50	42.21	99.90	31.02	10.40	39.72
LLaMA-3-FT	23.44	10.25	78.72	99.40	42.89	13.17	86.04

Table 18: Full story-generation quality (numerical version of Figure 5).

Backbone	Metric	Base	FT	Δ	95% CI	p
CPM	Tox. (Off) ↓	0.410	0.143	-0.267	[-0.291, -0.243]	<0.001
CPM	Read. (Token) ↑	0.605	0.708	+0.103	[-0.012, +0.218]	0.064 ^{n.s.}
PanGu	Tox. (Off) ↓	0.400	0.290	-0.110	[-0.145, -0.075]	<0.01
PanGu	Read. (Char) ↑	0.465	0.608	+0.143	[+0.105, +0.181]	<0.001
GLM	Tox. (Off) ↓	0.345	0.238	-0.107	[-0.235, +0.021]	0.080 ^{n.s.}
GLM	Tox. (Non-Off) ↓	0.003	0.040	+0.037	[-0.004, +0.078]	0.052 ^{n.s.}
Qwen	Tox. (Off) ↓	0.274	0.104	-0.170	[-0.205, -0.135]	<0.001
Qwen	Read. (Token) ↑	0.536	0.671	+0.135	[+0.092, +0.178]	<0.001
LLaMA-3	Tox. (Off) ↓	0.280	0.113	-0.167	[-0.198, -0.136]	<0.001
LLaMA-3	Read. (Token) ↑	0.520	0.645	+0.125	[+0.085, +0.165]	<0.001

Table 19: Two-sided Wilcoxon signed-rank tests on paired per-generation scores, with 95% confidence intervals on the paired difference from a bootstrap over generations (10,000 resamples). Offensive-prompt toxicity rows are tested over the 2,090 COLD prompts; all other rows over the 1,000 test prompts. Rows marked ^{n.s.} ($p \geq 0.05$) are described as trends in the main text rather than as established effects.

Models	Mode	L1	L2	L3	L4	L5	L6	Unknown
GLM	Char	0.359	0.149	0.206	0.097	0.071	0.024	0.095
	Token	0.260	0.073	0.073	0.063	0.086	0.025	0.419
GLM-ft (LoRA)	Char	0.428	0.163	0.115	0.080	0.046	0.030	0.138
	Token	0.357	0.083	0.070	0.063	0.023	0.013	0.390
GLM-ft (SFT)	Char	0.467	0.174	0.121	0.025	0.020	0.008	0.184
	Token	0.377	0.110	0.078	0.032	0.021	0.014	0.367

Table 20: Vocabulary levels (character and token) for GLM, GLM-LoRA, and GLM-SFT.

M.2 Readability and Vocabulary Levels

Table 17 reports readability scores, and Table 20 shows the full vocabulary-level distribution for LoRA vs. traditional fine-tuning.