

Help Me Write a Story: Evaluating LLMs’ Ability to Generate Writing Feedback

Hannah Rashkin Elizabeth Clark Fantine Huot Mirella Lapata

Google DeepMind

{hrashkin, eaclark, fantinehuot, lapata}@google.com

Abstract

Can LLMs provide support to creative writers by giving meaningful writing feedback? In this paper, we explore the challenges and limitations of model-generated writing feedback by defining a new task, dataset, and evaluation frameworks. To study model performance in a controlled manner, we present a novel test set of 1,300 stories that we corrupted to intentionally introduce writing issues. We study the performance of commonly used LLMs in this task with both automatic and human evaluation metrics. Our analysis shows that current models have strong out-of-the-box behavior in many respects—providing specific and mostly accurate writing feedback. However, models often fail to identify the *biggest* writing issue in the story and to correctly decide when to offer critical vs. positive feedback.

1 Introduction

With the growing popularity of LLMs, there is a lot of interest in using these models for human-in-the-loop collaboration such as creative writing tasks where humans may want to use LLMs as writing assistants (Lee et al., 2022; Yuan et al., 2022; Mirowski et al., 2022) or educational tasks where LLMs may act as tutors that help human students improve their skills (Lee et al., 2023; Jurenka et al., 2024). One under-studied task in this area is how well these models can be used as writing tutors that help students by providing feedback on their writing. For instance, discussion from Gero et al. (2023) highlights that giving good feedback is an important role in providing AI support to creative writers.

Using LLMs to produce writing feedback is a unique use-case that differs from other collaborative writing tasks, such as generating text continuations (e.g., Lee et al., 2022) or text rewriting (e.g., Shu et al., 2024), in that this task doesn’t ask the model to do any of the content generation it-

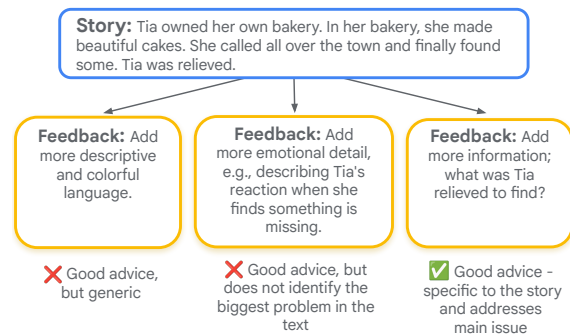


Figure 1: Examples of good and bad feedback for a short story

self. Providing feedback is more open-ended and should help guide the human writer towards fixing the writing themselves. Therefore, this task requires a different skill-set from other creative tasks including being able to identify writing problems and articulate them in a clear and constructive way. However, there aren’t many existing resources for studying LLM capabilities in providing feedback outside of academic domains like scientific peer reviews (e.g., Chamoun et al., 2024).

In this paper, we explore the task of providing open-ended feedback to improve story writing. Good feedback could include positive comments (e.g., *this writing is great already!*), criticisms (e.g., *the second sentence is confusing*), and/or suggestions for changes (e.g., *give more explanation of this character’s motives*). Beyond just being correct, good feedback should also be specific and well-explained, such as the right-most output in Figure 1.

To explore this task, we create a new dataset consisting of stories, some of which have intentional writing issues. For this purpose, we design a pipeline for introducing these issues into stories randomly chosen from a seed set of gold story data. We use a mix of the gold stories and corrupted stories as inputs to eight commonly used

LLMs, obtaining model-generated feedback for each. The STORYFEEDBACK dataset¹ includes a total of ~84k pairs of stories with model-generated feedback.

Through in-depth automatic analysis and human annotation, we investigate whether commonly used LLMs are able to provide feedback that: (1) is well-formed and specific, (2) provides correct criticisms and suggestions, (3) correctly identifies the biggest writing problems, especially in the corrupted stories, and (4) gives positive encouragement appropriately. Our findings demonstrate that most of the models provide well-formed, specific feedback and their suggestions are mostly correct. However, human scores also show that models struggle more to identify the biggest writing issues, instead commenting on smaller issues. Our results also show that models face difficulty in deciding when to offer positive comments, often doing so when stories have bigger issues that should be addressed.

2 Related Work

Creative Writing Assistance There is a large body of work suggesting the utility and drawbacks of using NLP systems as tools for creative writers (Chakrabarty et al., 2023; Lee et al., 2022; Clark et al., 2018). A variety of tasks have emerged for using AI for creative writing, ranging from fully automated story writers (Fan et al., 2018; Huot et al., 2024) to more collaborative writing assistants (Akoury et al., 2020; Lee et al., 2022; Chakrabarty et al., 2022; Yuan et al., 2022; Ippolito et al., 2022; Mirowski et al., 2022). There is also a strong interest in aligning these systems with the needs of actual creative writers. Guo et al. (2024) underscored the need for LLMs that collaborate in a way that aligns with creative writers’ values, particularly the importance of allowing for control and transparency while also being flexible to multiple writing roles. Recent work studying writers’ use of LLMs (Gero et al., 2023; Chakrabarty et al., 2024) has also highlighted the potential of AI-generated feedback to support writers, particularly if the feedback is specific and well-explained. However, there are scarce resources and evaluations for supporting this task within NLP. In our paper, we work to bridge this gap by providing resources for evaluating the ability of LLMs to generate correct and useful feedback for story drafts.

Identifying Writing Issues and Revisions

There are multiple prior works (Dou et al., 2022; Goyal et al., 2022) presenting methods of identifying writing issues in text. This is closely related to our task, but those papers focus on identifying categories of writing errors for use in evaluating machine-written text, whereas we focus on describing the writing issues constructively for humans to improve their own writing. Prior work (Du et al., 2022a; Laban et al., 2024; Kim et al., 2022) has also used models to provide direct revisions to a document, automatically or with a human-in-the-loop. Some of this work uses data with edit histories such as Wikipedia articles (Du et al., 2022b; Faruqui et al., 2018; Shu et al., 2024) that serve as drafts in a similar way to our automatically corrupted versions of the story documents. In our analysis, we use the edit taxonomies proposed by Du et al. (2022b) as a guide for the types of writing edits a model might suggest as feedback, though we apply these to story writing rather than Wikipedia editing. Rather than using AI to suggest text revisions directly, we propose using the model to provide writing feedback that would help people make their own edits, giving the writer more agency over the editing process.

Feedback and Peer Review Benharrak et al. (2024) showed that LLMs are promising tools for providing writers with on-demand, personalizable feedback; however, they also observed that outputs could be wordy, generic, and repetitive. These observations are mirrored by Han et al. (2024) which explored using essay scores as a way of guiding model output towards more critical, direct essay feedback. In Nair et al. (2024), they focus on modeling approaches for improving feedback of economic essays with simulated writers. Yuan et al. (2024) introduces a novel framework for models to use task criteria to create feedback for multiple tasks including scientific writing and reddit story post writing. Complementary to this line of recent work, we contribute new large-scale analysis with human evaluation across multiple dimensions of feedback quality as well as providing new data resources for this area. Whereas much of the existing work in generating feedback has focused on academic writing domains, such as generating peer reviews for scientific articles (Dycke et al., 2023; Kang et al., 2018; Chamoun et al., 2024; D’Arcy et al., 2024) or academic essay writing (Nair et al.,

¹<https://github.com/google-deepmind/igen>

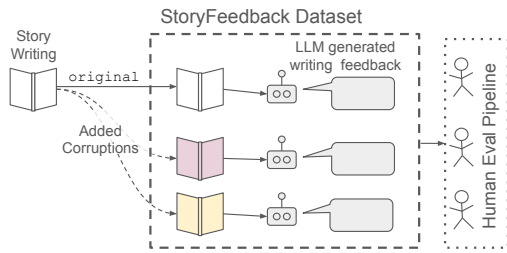


Figure 2: An overview of the task. Short stories are corrupted and given to models to generate feedback, which is then evaluated along several dimensions.

2024), we provide data and evaluation resources to explore how LLMs can be applied to generating feedback that suits the needs of the creative writing domain, more specifically.

3 The Task: Story Writing Feedback

Our goal is to understand how users can utilize LLMs as collaborative tools for improving their writing. In this work, we scope the task to asking LLMs to provide story-writing feedback, which we loosely define to include any type of positive encouragement, actionable suggestions, or constructive criticisms. We also focus on a single turn interaction (i.e. asking an LLM for feedback on a single writing draft), though this could be extended easily to multiple turns as a draft is edited iteratively by the writer.

We design new data and evaluation frameworks for exploring LLM capabilities at this task, depicted in Figure 2. For the purpose of large-scale analysis, we focus the scope of this paper’s investigation to short stories with controlled writing issues (created heuristically in Section 4) that can be easily scaled up and curated. The intention of the analysis in this paper is to act as an initial exploration into the challenges of this task and to develop more formal evaluation criteria that can also be applied to wider scopes of writing domains (such as long-form narratives).

In order to perform this task, the model should ideally give the writer feedback that is specific to their story, addresses the story’s biggest issues, and makes the story better. More concretely, while developing the data and evaluation frameworks, we consider the following questions regarding the model output:

1. **Is it valid feedback that’s specific to the story?** The output should be well-formed and assess the quality of the story in some way.

The feedback should be tailored to the input story, not generic writing advice.

2. **Would following the feedback make the story better?** If the writer follows the feedback, it should improve the quality of the story.
3. **Does it identify the biggest problem in the story?** The feedback should focus on whatever the main issues are in the text, rather than suggesting minor changes.
4. **Is the model giving positive comments appropriately?** It’s important for the model to be able to state positive feedback (such as appraising that the story is already high-quality). But, on the other hand, it shouldn’t only give positive feedback if there are still writing issues that need to be addressed.

4 The STORYFEEDBACK Dataset

We present STORYFEEDBACK, a dataset of English short stories paired with model-generated feedback to evaluate how accurately models can identify major issues in a text and articulate suggestions for improving the writing. In total, we collected 83,456 story and feedback pairs.

Seed Stories As a seed set of stories, we collect 326 short stories² from publicly available datasets: the ROCStories (Mostafazadeh et al., 2016) validation set and three BIG-bench story datasets (Abstract Narratives (Ghosh and Srivastava, 2022), English Proverbs, and English Fables) (BIG-bench authors, 2023). We randomly sampled up to 100 stories from each (The English Proverbs data only contains 26 stories).

Story Corruption Methods We can’t control the type of writing issues that may exist in the stories from the seed set, and many of them may not need additional feedback beyond positive comments. To ensure that we have data with specific sets of writing problems, we create several sets of stories with more controlled, automatically induced writing issues. We use three methods of corrupting the stories from the seed story set, as described below:

²We choose to focus on short stories that can be read and analyzed by human raters more efficiently.



Figure 3: Example with the corruption methods applied. Highlights indicate the portions of the story that are different from the original.

- backtranslate: The story is translated³ to German and then back to English ten times to get a newly perturbed story.⁴ This type of permutation sometimes introduces grammar issues and can change the wording or tone. It can also disrupt the clarity through slight meaning changes or affect the coreference of pronouns.
- swap: Two consecutive sentences in the story are swapped. This type of permutation typically disrupts the coherence of the story as events no longer take place in the correct order.
- delete: A randomly selected sentence is deleted from the story. This can cause events to become confusing since they lack prior context that would have come from a previous sentence.
- Bloomz 7B & 176B (Muennighoff et al., 2022) An open-source instruction-following model in 2 sizes.
- Gemma 9B & 27B (Gemma Team, 2024) Instruction-tuned versions of the open-source Gemma 2 models in 2 sizes.
- Gemini 1.5 Flash & Pro (Gemini Team, 2024) Instruction-tuned versions of the Gemini 1.5 model in 2 sizes.
- GPT 3.5 & 4 (OpenAI, 2023) 2 versions of the GPT model.

More details about the models and generation process are in Appendix A with examples in Table 10.

Prompting Set-ups We explore four prompts for eliciting feedback from the models, varying the amount of information given to the model and the amount of feedback requested. Our goal in studying these different prompts is to understand the sensitivity of the models to the varying amounts of instructions that a non-expert might give.

In addition, we also include the original stories (original) in the dataset for comparison of the types of feedback given on the original gold stories vs. the corrupted ones. In total, our dataset consists of ~1,300 stories (326 each of the types original, backtranslate, swap, and delete). In Figure 3, we show four stories from the dataset which were all corrupted from the same seed story.

Models We prompt eight models to generate feedback for each corrupted and gold story. We selected this list of models to represent several popular, general-purpose, and publicly available models that people could easily access and use out-of-the-box for writing advice, while covering a range of model quality and size:

- bulleted list (BL full): Prompts the model for a bulleted list of feedback for the story. We also include additional context to the model by providing a list of categories of possible writing issues (e.g., fluency, coherence, etc.) that the feedback should address. The categories of possible writing issues are derived from the text revision taxonomy proposed by Du et al. (2022b).
- bulleted list only (BL Only): The same prompt as bulleted list, but without the list of writing suggestion categories.
- one sentence (1-Sent): Prompts the model for a single sentence of feedback without further instructions.

³Using the BertSeq2Seq translation model from (Rothe et al., 2020).

⁴We chose to translate it ten times based on observing the amount of noise it added for a small set of examples.

- spot the problem (SpotProb): Frames the task as a challenge, asking the model to find the one problem in the text and describe it in one sentence.

All prompts additionally instructed the model to generate the phrase “The text is perfect as-is.” if there were no problems in the story. The full prompts are listed in App. B. For each prompt, we experiment with both a zero-shot and two-shot version of the prompt. The two-shot examples are the same for all prompts (one corrupted story and one original story) and can be found in App. B.

5 Automatic Evaluation

We run several automatic metrics on the generated feedback to describe the output and evaluate its quality, and the results are presented in Table 1.

Because we intentionally introduced errors into the stories, we can evaluate how often the model fails to identify these errors. Every prompt instructed the model to say, “The text is perfect as-is.” if it did not see any problems in the story. Ideally, a model should not say a story is perfect if the story has been corrupted by one of our strategies (backtranslate, swap, or delete). We look at what percentage of the feedback contains the phrase “is perfect as-is” (labeled as % PAI in Table 1), and if it contains the phrase, how often it is for an original (i.e., uncorrupted) story (PAI prec). More concretely,

$$PAI\ prec = \frac{\# \text{“perfect as-is” for original stories}}{\# \text{“perfect as-is” for all stories}} \quad (1)$$

where # “perfect as-is” represents the number of generated outputs that contain the phrase “perfect as-is.” Ideally, $PAI\ prec = 1$, because the feedback should never say a corrupted story is perfect. We focus on the precision of “perfect-as-is” because the other type of errors (i.e., giving negative feedback to uncorrupted stories) is not as concerning since it is still possible to give good critical feedback to an uncorrupted story.

Although the precision scores are low across all subsets of the data, the larger versions of the models are more precise in using the “perfect as-is” feedback option than their smaller counterparts. The Gemma and Gemini models are the most precise; however, they also generate this feedback far more infrequently than the other models (< 10% of the time). The two-shot feedback is more precise

	length	% PAI	PAI prec	3-gramR	3-gramR (w/o PAI)	
models	Bloomz 7B	88.31	0.32	0.24	0.77	0.68
	Bloomz 176B	124.74	0.26	0.25	0.67	0.57
	Gemma 9B	343.62	0.06	0.43	0.47	0.44
	Gemma 27B	354.31	0.03	0.46	0.46	0.44
	Gemini 1.5 Flash	443.46	0.04	0.51	0.40	0.37
	Gemini 1.5 Pro	482.10	0.07	0.52	0.38	0.34
	GPT 3.5	176.23	0.43	0.29	0.70	0.48
GPT 4	349.08	0.25	0.40	0.55	0.41	
prompts	bulleted list only	542.62	0.15	0.31	0.51	0.43
	bulleted list	445.12	0.23	0.34	0.65	0.56
	spot the problem	96.62	0.20	0.31	0.50	0.38
	one sentence	96.56	0.15	0.30	0.53	0.44
nshot	zeroshot	366.83	0.12	0.28	0.49	0.42
	twoshot	223.63	0.24	0.34	0.61	0.49

Table 1: Automatic evaluation results, broken down three different ways: by model, by prompt, and by zero-shot vs. twoshot. *Length* is the number of characters. *PAI* stands for “perfect-as-is,” *prec* abbreviates “precision,” and *3-gramR* stands for “trigram repetition.” All numbers are averages across the entire subset. **Bold** indicates the best-performing category (high is best for PAI prec and low is best for 3-gramR).

than the zero-shot, perhaps because one of the few-shot examples is a case where the text is labeled as perfect as-is. (The percent of feedback that contains the phrase “perfect as-is” doubles between the zero-shot and two-shot feedback.)

We also measure the diversity of the feedback by looking at the trigram repetition⁵ between a piece of feedback and all other feedback in that subset (*3-gramR*). We use this lexical overlap metric to detect the undesirable behavior of repeating the same generic feedback verbatim, regardless of the input story. A high repetition score means the feedback for one story frequently overlaps with feedback for other stories (i.e., low diversity).

The larger versions of the models had more diverse outputs than their smaller counterparts, and the Gemma and Gemini models had the most diverse feedback, overall. However, models that frequently reply with “the text is perfect as-is” naturally have worse diversity scores. To account for this, we also calculate the trigram overlap excluding those pieces of feedback (*3-gramR (w/o PAI)*), and we find that GPT 4 has the most diverse output after the Gemini models. The bulleted list prompt produces the least diverse output, perhaps because it is the most structured prompt, and models are more likely to copy from the bulleted list structure

⁵We used the implementation from Huot et al. (2024).

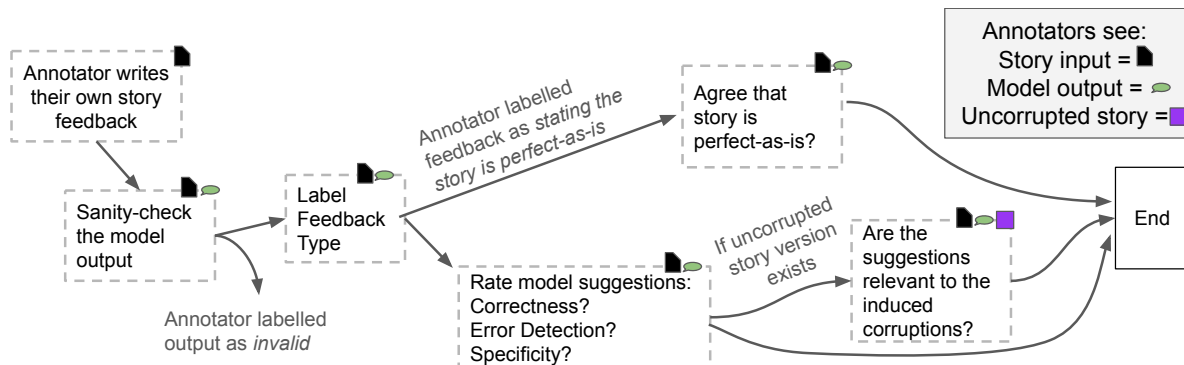


Figure 4: Outline of the human annotation task in which annotators can follow different question paths.

provided in the prompt. This may also be the reason that providing fewshot examples in the prompt also leads to a drop in the feedback diversity.

Additional automatic analysis and more detailed breakdowns of the results are in App. C.

6 Human Evaluation

While automatic evaluations can capture some aspects of the diversity or precision, the goal of the task is to produce useful feedback for writers; with this goal in mind, we run a human evaluation on a subset of the generated feedback. Participants annotated 1,920 pieces of feedback, randomly selected to be evenly distributed between all the models, corruptions, and prompt types (this comes out to 240 examples annotated per model). We collected 3 human annotations for each piece of feedback. Annotators responded to questions along multiple dimensions of quality, which were designed to investigate the larger research questions that we posed at the end of Sec. 3.

6.1 Annotation Task

We show the annotators a story and the corresponding model-generated feedback and ask them to answer a subset of the multiple choice questions below⁶. Their answers are translated to scalar values between 0 and 1 (details in App. D.4). The possible paths through the questions are shown in Figure 4.

1. Sanity-check: Is the model-generated output phrased as feedback about the story? The goal of this criterion is to filter out any degenerate or irrelevant outputs. *Options: Yes; No; Contains both*
2. Feedback-type: Is the feedback positive (e.g., “The story is perfect as-is.”) or does it provide

Dimensions	Avg Pair Sim	Kripp. α
sanity-check	0.95	0.49
feedback-type	0.99	0.95
perfect-agree	0.58	0.15
correctness	0.80	0.40
error-detection	0.68	0.33
specificity	0.89	0.61
relevance	0.79	0.40

Table 2: Inter-annotator agreement measures: Average pairwise similarity (Avg Pair Sim) between pairs of annotator responses on the same example, where similarity is 1 minus the Euclidean distance between responses’ numerical values. The Krippendorff’s alpha (Kripp. α) was computed using the same Euclidean distance metric between responses.

suggestions. *Options: Positive comments; Suggestions; Contains both*

3. Perfect-agree: If the feedback says that the story is perfect as-is, do you agree? *Options: Yes; No*
4. Correctness: If the feedback makes a suggestion, would following the suggestion improve the story? *Options: Yes; Some of the suggestions; No*
5. Error-detection: Would following the suggestion fix the *biggest* issue with the story? *Options: Yes; No; I don’t see any issues in the story*
6. Specificity: Is the suggestion specific to this story? *Options: Yes; No*
7. Relevance: If the story was corrupted, would

⁶Exact question and answer phrasing in App. D.2

Open-source	Model	sanity	feedback-type	perfect-agree	correctness	error-detection	specificity	relevance
Yes	Bloomz 7b	0.874 ± 0.02	0.410 ± 0.03	0.285 ± 0.04	0.292 ± 0.03	0.219 ± 0.03	0.288 ± 0.03	0.226 ± 0.03
	Bloomz 176b	0.792 ± 0.02	0.291 ± 0.03	0.450 ± 0.04	0.316 ± 0.02	0.254 ± 0.03	0.460 ± 0.04	0.243 ± 0.02
Yes	Gemma 9b	0.993 ± 0.0	0.062 ± 0.02	0.656 ± 0.07	0.734 ± 0.02	0.609 ± 0.02	0.904 ± 0.02	0.462 ± 0.03
	Gemma 27b	0.997 ± 0.0	0.051 ± 0.01	0.485 ± 0.07	0.755 ± 0.02	0.614 ± 0.02	0.902 ± 0.02	0.455 ± 0.02
No	Gemini Flash	0.989 ± 0.0	0.026 ± 0.01	0.367 ± 0.11	0.763 ± 0.02	0.687 ± 0.02	0.917 ± 0.01	0.481 ± 0.02
	Gemini Pro	0.982 ± 0.01	0.052 ± 0.01	0.625 ± 0.10	0.792 ± 0.02	0.703 ± 0.02	0.953 ± 0.01	0.497 ± 0.02
No	GPT 3.5	0.989 ± 0.0	0.419 ± 0.03	0.427 ± 0.03	0.766 ± 0.02	0.663 ± 0.03	0.862 ± 0.02	0.438 ± 0.03
	GPT 4	0.989 ± 0.0	0.199 ± 0.03	0.573 ± 0.05	0.834 ± 0.02	0.757 ± 0.02	0.942 ± 0.01	0.593 ± 0.03

Table 3: The average score and standard error of the mean (SEM) of human responses per dimension per model, aggregated over all other variations of nshot, prompt type, and noise type. Note: the Gemma/Gemini models rarely say the story is perfect as is, so the perfect-agree scores for these models are computed over very small sample sizes ($n < 20$) and have a large SEM margin.

following the feedback fix the corruption?⁷
Options: Yes; Some of the suggestions; No

We also ask annotators to provide their own feedback for the story writer at the beginning of the task, prior to seeing the machine-generated feedback. This encourages annotators to read the story carefully before answering the questions and also produces a set of human-authored feedback for our analysis.

6.2 Annotation Statistics

For the multiple choice questions, we compute the interannotator agreement (Table 2) using the Euclidean similarity between the values of the responses of two annotations for the same example. The Krippendorff’s alpha scores for most dimensions are moderate to high indicating general agreement. The dimension with the lowest agreement scores is perfect-agree (whether the annotators agree that the story is “perfect-as-is”). We posit that this question has more natural ambiguity since it’s subjective whether a text is already “good enough” or if more changes can be made. The error-detection question had higher agreement than the perfect-agree but also indicated some ambiguity. When computing the evaluation scores, we use the averages of the scores from all three raters per example, as an aggregated response.

More details about the human evaluation implementation (including evaluator training and score aggregation) are in App. D.

6.3 Human Evaluation Results

In Tables 3–5 we report the average human ratings aggregated by model, noise type, and prompt type.

⁷For this question, we show annotators the original story and ask whether the feedback reflects the changes between the original and corrupted stories.

Difference in Model Architectures Based on the results in Table 3, we note that most models perform quite highly at the specificity question and also reasonably highly on the correctness question. But, they have weaker performance in the error-detection and relevance dimensions. It seems the models miss what the annotators felt were the most important issues (error-detection) or what issues were introduced via the corruption strategy (relevance). Overall, this indicates that models provide concrete and plausible suggestions but often leave out suggestions that would target the most salient writing problems.

Additionally, the models tend to score lowly on the perfect-agree question, indicating that the models are often wrong when they say that the story doesn’t need any further edits, which is similar to our observations about the low precision scores from the automatic results (Table 1). Also mirroring the automatic metrics, we do see that the feedback-type is “is-perfect-as-is” in only a very small percentage of cases from the Gemini and Gemma models (about 5% of the time), confirming that these model families are more prone towards frequent suggestions.

Performance may be slightly affected by model size, but these improvements are small and depend on the dimension. We observe that the proprietary models outperform the open source models (particularly outperforming over Bloom) across most dimensions. The two largest models (Gemini Pro and GPT 4) perform similarly to one another, especially when considering the margin of error, though each class of model may have different strengths across different dimensions. GPT 4 in particular has strong performance in the error-detection and relevance dimensions.

Corruption Type	sanity	feedback-type	perfect-agree	correctness	error-detection	specificity	relevance
original	0.958 ± 0.01	0.241 ± 0.02	0.581 ± 0.03	0.630 ± 0.02	0.485 ± 0.02	0.825 ± 0.02	<i>nan ± nan</i>
backtranslate	0.946 ± 0.01	0.124 ± 0.02	0.301 ± 0.04	0.770 ± 0.01	0.712 ± 0.02	0.840 ± 0.02	0.577 ± 0.02
delete	0.955 ± 0.01	0.213 ± 0.02	0.422 ± 0.03	0.666 ± 0.02	0.545 ± 0.02	0.813 ± 0.02	0.367 ± 0.02
swap	0.943 ± 0.01	0.158 ± 0.02	0.331 ± 0.04	0.691 ± 0.02	0.621 ± 0.02	0.815 ± 0.02	0.378 ± 0.02

Table 4: The average score and standard error of the mean (SEM) of human responses for each type of story input, aggregated over all other variations of n-shot, prompt type and model type.

Effects of Story Corruption As shown in Table 4, backtranslate has higher scores for correctness, error-detection, and relevance compared to other corruption types, implying models can spot those types of errors more easily. Writing issues caused by swapping or deleting sentences (e.g., incoherent events and missing information) are not as easily detected. This is most notable from the relevance category, where models struggle to make suggestions that would be relevant to the swapped or deleted alterations. Annotators are more likely to agree with the positive comments on original stories (perfect-agree) than the positive comments on the corrupted ones, overall. However, when the model tries to provide constructive suggestions to an original story, it is more likely to offer low-quality suggestions (low correctness and error-detection) compared to the suggestions for the corrupted stories.

Sensitivity to Prompt Types As shown in Table 5, prompting the model to output more information via a bulleted list (BL only and BL full) yields the best results. Depending on the dimension, it may be more beneficial to provide a taxonomy of possible feedback types (BL full) rather than just requesting a bulleted list (BL only). Prompting the model to “spot the problem” in the writing seems to yield the lowest performance in terms of correctness and error-detection. This is surprising in that it seems that the models struggle to identify the biggest issues when more directly prompted to spot problems in the writing.

Difference in n-shot Performance When comparing the results from zero-shot vs. two-shot models (Table 5), there is no clear improvement, and it may even harm the scores in certain dimensions to have two-shot set-ups. Qualitatively, we observe that the two-shot models often become biased towards repeating the feedback from the two-shot examples rather than generating feedback that’s more specific to the input story (as reflected in the specificity scores).

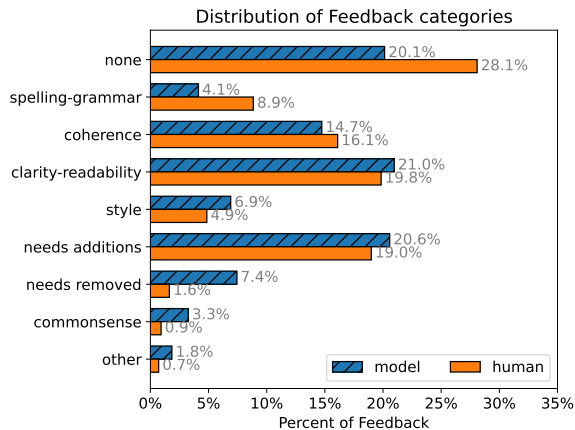


Figure 5: Distribution of categories of feedback written by models vs. humans

6.4 Human-Authored Feedback

As part of the annotation task, annotators wrote down their own feedback for each story. We analyze how their feedback compares with the model generations by prompting⁸ Gemini 1.5 Flash to categorize⁹ the types of feedback given by both humans and models and compare the types of issues they mention in Figure 5. We find that people are more likely to say there are no issues (especially for the original stories) and give feedback on spelling, grammar, and coherence issues, while model-generated feedback comments more on style and adding or removing information from the story. More details and analysis are in App. G.

7 Discussion

We close by revisiting the main questions we posed in Section 3.

Can LLMs produce valid feedback that is specific to a story? Generally, yes. Though it may vary by model and prompt type, most of the LLMs reached very high scores in the sanity-check questions indicating that they are typically well-formed

⁸The full prompt is in App. G.1.

⁹Based on the categories used in Du et al. (2022b).

Prompt	sanity	feedback-type	perfect-agree	correctness	error-detection	specificity	relevance
SpotProb	0.894 ± 0.01	0.216 ± 0.02	0.461 ± 0.03	0.577 ± 0.02	0.437 ± 0.02	0.788 ± 0.02	0.414 ± 0.02
1-Sent	0.962 ± 0.01	0.139 ± 0.02	0.444 ± 0.04	0.689 ± 0.02	0.538 ± 0.02	0.741 ± 0.02	0.439 ± 0.02
BL only	0.960 ± 0.01	0.174 ± 0.02	0.446 ± 0.04	0.746 ± 0.01	0.679 ± 0.02	0.923 ± 0.01	0.406 ± 0.02
BL Full	0.986 ± 0.0	0.205 ± 0.02	0.389 ± 0.03	0.747 ± 0.01	0.719 ± 0.02	0.844 ± 0.02	0.513 ± 0.02
0-shot	0.919 ± 0.01	0.152 ± 0.01	0.387 ± 0.03	0.722 ± 0.01	0.628 ± 0.01	0.849 ± 0.01	0.450 ± 0.01
2-shot	0.982 ± 0.0	0.206 ± 0.01	0.463 ± 0.02	0.674 ± 0.01	0.577 ± 0.01	0.804 ± 0.01	0.446 ± 0.01

Table 5: The average score and standard error of the mean (SEM) of human responses for different prompt-types, aggregated.

valid outputs. The models generally get high specificity scores as well, meaning that the feedback is typically not just generic comments.

Would following the feedback make the story better? The best models typically achieve correctness scores in the range of 0.7 to 0.8. This still leaves room for improvement, but does indicate that their suggestions are frequently helpful for improving in the story.

Does it identify the biggest problem in the story? This is still a challenge even for the best performing models. Their error-detection scores are much lower than their correctness scores, which means that, even when they correctly identify issues in the story, they are often missing what annotators’ considered the most salient issue. This is also reflected in the low relevance scores which show that the model suggestions often don’t address the intentional errors that we’ve added by corrupting the stories.

Is the model giving positive comments appropriately? This is another challenge for models. Both our automatic and human evaluations indicate that models are not very discerning with positive encouragement. When the models call stories “perfect,” people frequently disagreed, stating that those stories still contained writing problems that should have been addressed.

Implications for future work In this paper, we focus solely on short stories with synthetic writing issues to allow for systematic analysis that could be scaled over many data examples and model architectures. The analysis on these stories is useful as a first step towards uncovering challenges in this task; however, to understand the full scope of LLM abilities as writing tutors, this work would need to be extended to additional domains, such as longer and more complex stories with naturally occurring writing issues. Models may especially struggle to

detect and offer feedback for some types of errors in long texts (e.g., contemporaneous work from Ahuja et al. (2025) demonstrates that models struggle to identify continuity errors in longer stories). In more realistic writing settings, human writers may also use a model for feedback over multiple iterations of draft-writing. While we only consider a single turn of feedback in our setting, it remains an open research question whether these LLMs would be consistent and useful in settings where a writer may repeatedly ask for feedback over multiple draft iterations. Collecting resources towards these domains can be challenging as human annotations for longer stories and iterative writing sessions also requires significantly more work from human annotators. As a first step in this direction, our paper provides STORYFEEDBACK as a resource to facilitate future work and our discussion of the research questions above as starting points for further investigation.

8 Conclusion

We introduce the task of providing writing feedback on short stories, which is challenging for LLMs because it requires both identifying writing issues and formulating constructive advice. We create the STORYFEEDBACK dataset to explore how well LLMs can generate feedback on simple stories, many of which have induced writing errors. With in-depth automatic and human evaluations, we uncover some of the strengths and weaknesses of out-of-the-box models for this task. While models show promising abilities to provide plausible, specific, and accurate feedback, they sometimes miss the more salient problems in favor of lower hanging fruit. We suggest that future work in this area should focus on designing models that are able to identify the most salient writing issues rather than just providing “correct” feedback as well as extending these evaluation frameworks to long-form narratives with more complex writing issues.

Limitations

Our study is limited to model behavior in the domain of English stories, and we choose to focus on relatively short stories that can be analyzed in more depth. We also synthetically introduce errors into the text to control the type and frequency of writing errors. Observations may differ on longer, more complex stories, for non-English languages, or for naturally occurring writing errors. Additionally, we limit our study to eight common LLMs with four different prompts, to characterize a range of performance behaviors, but do not make an exhaustive search of all possible models or prompt choices. We also choose to focus on specific dimensions of feedback quality like the correctness, relevance, specificity, etc., but there may be other qualities like the tone of the feedback that should be explored in future work. Lastly, we acknowledge that there is some subjectivity in what makes feedback “useful,” and as noted in our discussion of the human annotations, there may be variations in judgment due to the natural ambiguity of the task.

Ethical Considerations

In designing LLMs that will provide support to human writers, it’s important to consider how this might align with writers’ values (Guo et al., 2024), and, namely, integrating these systems in a way that does not diminish the writers’ agency, control, and transparency in their own creative decisions. A challenge of using these systems with students, in specific, is the concern that they may use this technology as a crutch rather than further honing their skills. In this task, an LLM may provide feedback that has a negative impact on their writing (for example, making suggestions to include things that involve stereotypes or plagiarism) or their engagement (for example, providing overly harsh critical feedback). Future work with more focus on investigating these issues is important for understanding how these systems could be used responsibly in downstream applications.

Acknowledgments

We would like to thank the anonymous reviewers for their insightful comments and suggestions. We would also like to thank Annie Louis, Edward Grefenstette, Sumit Asthana, Tuhin Chakrabarty, Chaitanya Malaviya, Chris Alberti, Shereen Ashraf, Jennimaria Palomaki, Shoshana Jakobovits, Ankur

Parikh, Dipanjan Das, Slav Petrov, Chris Dyer, and our other Google Deepmind collaborators for providing constructive feedback and computational resources at multiple stages of this project. Lastly, we also thank the annotators who participated in our human evaluation studies.

References

- Kabir Ahuja, Melanie Sclar, and Yulia Tsvetkov. 2025. [Finding flawed fictions: Evaluating complex reasoning in language models via plot hole detection](#). *Preprint*, arXiv:2504.11900.
- Nader Akoury, Shufan Wang, Josh Whiting, Stephen Hood, Nanyun Peng, and Mohit Iyyer. 2020. [STOR-IUM: A Dataset and Evaluation Platform for Machine-in-the-Loop Story Generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6470–6484, Online. Association for Computational Linguistics.
- Karim Benharrak, Tim Zindulka, Florian Lehmann, Hendrik Heuer, and Daniel Buschek. 2024. [Writer-defined ai personas for on-demand feedback generation](#). In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI ’24, New York, NY, USA. Association for Computing Machinery.
- BIG-bench authors. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Tuhin Chakrabarty, Philippe Laban, Divyansh Agarwal, Smaranda Muresan, and Chien-Sheng Wu. 2023. [Art or artifice? large language models and the false promise of creativity](#). *arXiv preprint arXiv:2309.14556*.
- Tuhin Chakrabarty, Vishakh Padmakumar, Faeze Brahman, and Smaranda Muresan. 2024. [Creativity support in the age of large language models: An empirical study involving professional writers](#). *Proceedings of the 16th Conference on Creativity & Cognition*.
- Tuhin Chakrabarty, Vishakh Padmakumar, and He He. 2022. [Help me write a poem: Instruction tuning as a vehicle for collaborative poetry writing](#). *arXiv preprint arXiv:2210.13669*.
- Eric Chamoun, Michael Schlichtkrull, and Andreas Vlachos. 2024. [Automated focused feedback generation for scientific writing assistance](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9742–9763, Bangkok, Thailand. Association for Computational Linguistics.
- Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A. Smith. 2018. [Creative](#)

- writing with a machine in the loop: Case studies on slogans and stories. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces, IUI '18*, page 329–340, New York, NY, USA. Association for Computing Machinery.
- Mike D'Arcy, Tom Hope, Larry Birnbaum, and Doug Downey. 2024. [Marg: Multi-agent review generation for scientific papers](#). *Preprint*, arXiv:2401.04259.
- Yao Dou, Maxwell Forbes, Rik Koncel-Kedziorski, Noah A. Smith, and Yejin Choi. 2022. [Is GPT-3 text indistinguishable from human text? scarecrow: A framework for scrutinizing machine text](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7250–7274, Dublin, Ireland. Association for Computational Linguistics.
- Wanyu Du, Zae Myung Kim, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. 2022a. [Read, revise, repeat: A system demonstration for human-in-the-loop iterative text revision](#). In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*, pages 96–108, Dublin, Ireland. Association for Computational Linguistics.
- Wanyu Du, Vipul Raheja, Dhruv Kumar, Zae Myung Kim, Melissa Lopez, and Dongyeop Kang. 2022b. [Understanding iterative revision from human-written text](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3573–3590, Dublin, Ireland. Association for Computational Linguistics.
- Nils Dycke, Iliia Kuznetsov, and Iryna Gurevych. 2023. [NLPeer: A unified resource for the computational study of peer review](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5049–5073, Toronto, Canada. Association for Computational Linguistics.
- 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, Melbourne, Australia. Association for Computational Linguistics.
- Manaal Faruqui, Ellie Pavlick, Ian Tenney, and Dipanjan Das. 2018. [WikiAtomicEdits: A multilingual corpus of Wikipedia edits for modeling language and discourse](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 305–315, Brussels, Belgium. Association for Computational Linguistics.
- Gemini Team. 2024. [Gemini 1.5: Unlocking multi-modal understanding across millions of tokens of context](#). *Preprint*, arXiv:2403.05530.
- Gemma Team. 2024. [Gemma 2: Improving open language models at a practical size](#). *Preprint*, arXiv:2408.00118.
- Katy Ilonka Gero, Tao Long, and Lydia B Chilton. 2023. [Social dynamics of ai support in creative writing](#). In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI '23*, New York, NY, USA. Association for Computing Machinery.
- Sayan Ghosh and Shashank Srivastava. 2022. [ePiC: Employing proverbs in context as a benchmark for abstract language understanding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3989–4004, Dublin, Ireland. Association for Computational Linguistics.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. [FALTE: A toolkit for fine-grained annotation for long text evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 351–358, Abu Dhabi, UAE. Association for Computational Linguistics.
- Alicia Guo, Shreya Sathyanarayanan, Leijie Wang, Jeffrey Heer, and Amy Zhang. 2024. [From pen to prompt: How creative writers integrate ai into their writing practice](#). *ArXiv*, abs/2411.03137.
- Jieun Han, Haneul Yoo, Junho Myung, Minsun Kim, Hyunseung Lim, Yoonsu Kim, Tak Yeon Lee, Hwajung Hong, Juho Kim, So-Yeon Ahn, and Alice Oh. 2024. [LLM-as-a-tutor in EFL writing education: Focusing on evaluation of student-LLM interaction](#). In *Proceedings of the 1st Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U)*, pages 284–293, Miami, Florida, USA. Association for Computational Linguistics.
- Fantine Huot, Reinald Kim Amplayo, Jennimaria Palomaki, Alice Shoshana Jakobovits, Elizabeth Clark, and Mirella Lapata. 2024. [Agents' room: Narrative generation through multi-step collaboration](#). *arXiv preprint arXiv:2410.02603*.
- Daphne Ippolito, Ann Yuan, Andy Coenen, and Sehmon Burnam. 2022. [Creative writing with an ai-powered writing assistant: Perspectives from professional writers](#). *ArXiv*, abs/2211.05030.
- Irina Jurenka, Markus Kunesch, Kevin R. McKee, Daniel Gillick, Shaojian Zhu, Sara Wiltberger, Shubham Milind Phal, Katherine Hermann, Daniel Kasenberg, Avishkar Bhoopchand, Ankit Anand, Miruna Pislar, Stephanie Chan, Lisa Wang, Jennifer She, Parsa Mahmoudieh, Aliya Rysbek, Wei-Jen Ko, Andrea Huber, Brett Wiltshire, Gal Elidan, Roni Rabin, Jasmin Rubinovitz, Amit Pitaru, Mac McAllister, Julia Wilkowski, David Choi, Roe Engelberg, Lidan Hackmon, Adva Levin, Rachel Grifin, Michael Sears, Filip Bar, Mia Mesar, Mana

- Jabbour, Arslan Chaudhry, James Cohan, Sridhar Thiagarajan, Nir Levine, Ben Brown, Dilan Gorur, Svetlana Grant, Rachel Hashimshoni, Laura Weidinger, Jeru Hu, Dawn Chen, Kuba Dolecki, Canfer Akbulut, Maxwell Bileschi, Laura Culp, Wen-Xin Dong, Nahema Marchal, Kelsie Van Deman, Hema Bajaj Misra, Michael Duah, Moran Ambar, Avi Caciularu, Sandra Lefdal, Chris Summerfield, James An, Pierre-Alexandre Kamienny, Abhinav Mohdi, Theofilos Strinopoulos, Annie Hale, Wayne Anderson, Luis C. Cobo, Niv Efron, Muktha Ananda, Shakir Mohamed, Maureen Heymans, Zoubin Ghahramani, Yossi Matias, Ben Gomes, and Lila Ibrahim. 2024. [Towards responsible development of generative ai for education: An evaluation-driven approach](#). *Preprint*, arXiv:2407.12687.
- Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Edward Hovy, and Roy Schwartz. 2018. [A dataset of peer reviews \(PeerRead\): Collection, insights and NLP applications](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1647–1661, New Orleans, Louisiana. Association for Computational Linguistics.
- Zae Myung Kim, Wanyu Du, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. 2022. [Improving iterative text revision by learning where to edit from other revision tasks](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9986–9999, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Philippe Laban, Jesse Vig, Marti Hearst, Caiming Xiong, and Chien-Sheng Wu. 2024. [Beyond the chat: Executable and verifiable text-editing with llms](#). In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, UIST '24, New York, NY, USA. Association for Computing Machinery.
- Changyoon Lee, Junho Myung, Jieun Han, Jiho Jin, and Alice Oh. 2023. [Learning from teaching assistants to program with subgoals: Exploring the potential for ai teaching assistants](#). *Preprint*, arXiv:2309.10419.
- Mina Lee, Percy Liang, and Qian Yang. 2022. [Coauthor: Designing a human-ai collaborative writing dataset for exploring language model capabilities](#). In *CHI Conference on Human Factors in Computing Systems*, CHI '22, page 1–19. ACM.
- Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2022. [Co-writing screenplays and theatre scripts with language models: An evaluation by industry professionals](#). *Preprint*, arXiv:2209.14958.
- 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, San Diego, California. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, et al. 2022. [Crosslingual generalization through multitask finetuning](#). *arXiv preprint arXiv:2211.01786*.
- Inderjeet Jayakumar Nair, Jiaye Tan, Xiaotian Su, Anne Gere, Xu Wang, and Lu Wang. 2024. [Closing the loop: Learning to generate writing feedback via language model simulated student revisions](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16636–16657, Miami, Florida, USA. Association for Computational Linguistics.
- OpenAI. 2023. [GPT-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Sascha Rothe, Shashi Narayan, and Aliaksei Severyn. 2020. [Leveraging pre-trained checkpoints for sequence generation tasks](#). *Transactions of the Association for Computational Linguistics*, 8:264–280.
- Lei Shu, Liangchen Luo, Jayakumar Hoskere, Yun Zhu, Simon Tong, JD Chen, and Lei Meng. 2024. [Rewritelm: An instruction-tuned large language model for text rewriting](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. [Wordcraft: Story writing with large language models](#). In *Proceedings of the 27th International Conference on Intelligent User Interfaces*, IUI '22, page 841–852, New York, NY, USA. Association for Computing Machinery.
- Weizhe Yuan, Pengfei Liu, and Matthias Gallé. 2024. [LLMCrit: Teaching large language models to use criteria](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 7929–7960, Bangkok, Thailand. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

A Model Details

The Bloom feedback was generated from checkpoints downloaded from Hugging Face¹⁰ in May 2024. The GPT model output was generated using the OpenAI API¹¹ in May 2024. The Gemma and

¹⁰<https://huggingface.co>

¹¹<https://platform.openai.com/>

Gemini models were generated from the publicly available checkpoints¹² in December 2024.

The feedback was generated using the API or model’s default settings to mimic scenarios where writers would use these systems out-of-the-box. All model outputs were sampled with temperature = 0.70.

B Prompts

The full text for each of the four prompt types and the twoshot examples are as follows:

B.1 Bulleted List

Write a bulleted list of feedback for the following text:

<STORY>

The feedback should state whether the text is:

- not fluent (e.g. typos, grammar issues)
 - not coherent (lacking internally consistency or badly ordered)
 - lacking clarity (e.g. text is wordy or ambiguous)
 - needs a stylistic change (e.g. formality, tone, other writing preferences)
 - needs additions (e.g. needs more information or updated information)
 - needs something removed (e.g. get rid of something that doesn’t fit with the rest of the document)
- If the story has no problems, say ‘The text is perfect as-is’.

Feedback:

B.2 Bulleted List Only

Write a bulleted list of feedback for the following text:

<STORY>

If the story has no problems, say ‘The text is perfect as-is’.

Feedback:

B.3 One Sentence

Write one sentence of feedback for improving the following text:

<STORY>

The feedback should be brief and address the most important problem in the text.

If the story has no problems, say ‘The text is perfect as-is’.

Feedback:

¹²<https://aistudio.google.com/>

B.4 Spot the Problem

There is one problem in this text. Can you find it?
<STORY>

In one sentence, describe the problem.

If the story has no problems, say ‘The text is perfect as-is’.

Feedback:

B.5 Twoshot Examples

Example 1

Story: Cathy loves basketball. She made plans to go watch her favorite team play. So she went to the store to buy more. She sat in the front row. Her team won the game.

Feedback:¹³ The third sentence in the text is incoherent with the rest of the story and should be removed.

Example 2

Story: Matty’s cats kept scratching themselves. She didn’t know what to do. Eventually she found out they had fleas. She gave them flea medicine. They were completely healed after that.

Feedback:¹⁴ The text is perfect as-is.

C Additional Automatic Evaluation Results

This section contains several additional automatic evaluations of the generated feedback.

C.1 Trigram Overlap with Story

In Section 5 we present the trigram overlap between one piece of feedback and the rest of the feedback in a subset to understand the diversity of a subset. We can also look at the trigram overlap between a piece of feedback and the story it is based on as a measure of the specificity of the feedback. Higher overlap may indicate that the feedback is referencing specific portions of the story in its feedback. Table 6 shows the average trigram overlap between the story and the feedback, aggregated across different splits of the data. It also shows the average BERTScore (Zhang et al., 2020) for each subset (an embedding-based similarity metric).

C.2 Corruption Method

The frequency of the “perfect as-is” feedback also provides insight into which story corruption methods were the most challenging for the models to

¹³The feedback is preceded by a dash (-) for the bulleted list and bulleted list only prompts.

¹⁴See footnote 13.

	trigram rep	BERTScore	
models	Bloomz 7B	0.10	83.74
	Bloomz 176B	0.11	83.98
	Gemma 9B	0.09	83.89
	Gemma 27B	0.10	84.00
	Gemini 1.5 Flash	0.10	84.17
	Gemini 1.5 Pro	0.11	84.05
	GPT 3.5	0.05	83.67
	GPT 4	0.08	84.09
prompts	bulleted list only	0.08	84.01
	bulleted list	0.05	83.12
	spot the problem	0.15	84.37
nshot	one sentence	0.09	84.28
	zeroshot	0.12	84.14
	twoshot	0.07	83.76

Table 6: The similarity between a story and its feedback measured by the average trigram overlap (trigram rep) and BERTScore. The results are broken down three different ways: by model, by prompt, and by zeroshot vs. twoshot. Higher similarity scores indicate more overlap with the story, and the highest numbers are in **bold**. All numbers are multiplied by 100 for readability.

noise	% perfect as-is
original	0.23
backtranslate	0.13
random sentence swap	0.15
delete one sentence	0.21

Table 7: The percent of feedback that included the phrase “perfect as-is” broken down by the type of noise that was introduced into the story.

detect. As shown in Table 7, the models were most likely to say a corrupted story was perfect when the corruption method was delete. They did so at a rate close to the original (i.e., uncorrupted) stories (0.21 vs. 0.23), indicating the deleted sentences were less likely to be caught by the models than backtranslation or sentence swapping.

C.3 One Sentence Instruction Following

Two of the prompts (one sentence and spot the problem) instruct the model to output one sentence of feedback. We evaluate how often the feedback follows these instructions by using the sentence splitting implementation from Huot et al. (2024) and counting percentage of model responses that only contain one sentence (Table 8). The models are generally able to follow this instruction very well. The Bloom models struggle in some of the zeroshot cases but improve in the twoshot setting. The exception is the Bloomz 176B consistently out-

	one sentence		spot the problem	
	zeroshot	twoshot	zeroshot	twoshot
Bloomz 7B	0.62	0.90	0.78	0.95
Bloomz 176B	0.64	0.96	0.95	0.29
Gemini 1.5 Flash	1.00	1.00	0.99	0.96
Gemini 1.5 Pro	0.99	0.99	0.99	0.97
Gemma 9B	0.99	0.98	0.92	0.83
Gemma 27B	1.00	0.99	0.99	0.99
GPT 3.5	1.00	1.00	0.95	0.99
GPT 4	0.98	0.99	0.99	0.98

Table 8: The percent of feedback that was only one sentence, broken down by prompt type.

put an extra “Feedback: ” when given the twoshot spot the problem prompt, which was split into multiple sentences by the splitter.

C.4 Full Automatic Evaluation Results

The full breakdown of automatic evaluation scores are in Table 9.

D Human Evaluation Details

D.1 Annotator Details and Task Statistics

We ran our task with responses from 64 annotators, each annotating 90 story-feedback pairs on average. The median time spent annotating each example was 82.3 seconds. For the free-response feedback written by annotators, we collected a total of 5,754 feedback entries. The human-written feedback is on average 56.5 characters long. The annotators had the option to write “n/a” as their feedback if they didn’t see any writing issues. Prior to completing the task, annotators completed a tutorial and consent form detailing the study in advance. Annotators are all fluent English speakers, who were paid fairly, above minimum wage for their respective locations.

D.2 Annotation Instructions

Below we include the exact phrasing of the instructions, question, and answer options, verbatim to how they were provided to the annotators. In the actual task we also included examples that they could refer to at any time.

1. Well-formed feedback

Task Instructions

This question verifies that the feedback contains a critique of the story and stays on-topic. This isn’t asking about the quality of the feedback, just verifying that it is actually feedback for the given story, since we

	model	length	% perfect as-is	perfect precision	as-is	trigram repetition	trigram repetition (w/o perfect as-is)
bulleted list	zeroshot	Bloomz 7B	75.51	0.27	0.24	0.88	0.85
		Bloomz 176B	96.66	0.06	0.28	0.93	0.93
		Gemma 9B	625.96	0.06	0.48	0.56	0.54
		Gemma 27B	602.12	0.02	0.44	0.55	0.54
		Gemini 1.5 Flash	980.99	0.02	0.45	0.47	0.47
		Gemini 1.5 Pro	1011.8	0	0.75	0.43	0.43
		GPT 3.5	323.63	0.57	0.31	0.76	0.49
		GPT 4	838.02	0.1	0.42	0.59	0.57
	twoshot	Bloomz 7B	72.15	0.69	0.24	0.92	0.8
		Bloomz 176B	122.63	0.41	0.24	0.77	0.65
		Gemma 9B	420.65	0.22	0.42	0.62	0.52
		Gemma 27B	454.35	0.12	0.47	0.58	0.52
		Gemini 1.5 Flash	465.07	0.16	0.48	0.52	0.44
		Gemini 1.5 Pro	457.32	0.21	0.52	0.51	0.4
bulleted list only	zeroshot	Bloomz 7B	71.54	0.51	0.25	0.78	0.58
		Bloomz 176B	178.72	0.14	0.29	0.19	0.08
		Gemma 9B	833.17	0	0	0.43	0.43
		Gemma 27B	836.19	0.01	0.2	0.41	0.41
		Gemini 1.5 Flash	1134.44	0	0	0.35	0.35
		Gemini 1.5 Pro	1298.6	0.01	0.15	0.34	0.34
		GPT 3.5	466.04	0.17	0.28	0.53	0.44
		GPT 4	844.28	0.1	0.37	0.47	0.43
	twoshot	Bloomz 7B	197.42	0.23	0.21	0.95	0.94
		Bloomz 176B	165.59	0.25	0.26	0.71	0.67
		Gemma 9B	462.98	0.08	0.41	0.46	0.42
		Gemma 27B	549.64	0.02	0.42	0.44	0.43
		Gemini 1.5 Flash	547.54	0.05	0.51	0.42	0.39
		Gemini 1.5 Pro	661.52	0.08	0.52	0.4	0.35
one sentence	zeroshot	Bloomz 7B	88.38	0.04	0.34	0.37	0.34
		Bloomz 176B	98.06	0.35	0.25	0.47	0.21
		Gemma 9B	101.03	0	NaN	0.44	0.44
		Gemma 27B	90.67	0	NaN	0.49	0.49
		Gemini 1.5 Flash	99.29	0	NaN	0.37	0.37
		Gemini 1.5 Pro	118.86	0	NaN	0.35	0.35
		GPT 3.5	90.46	0.33	0.23	0.66	0.49
		GPT 4	150.59	0.04	0.36	0.44	0.42
	twoshot	Bloomz 7B	59.62	0.45	0.24	0.85	0.74
		Bloomz 176B	61.37	0.48	0.24	0.92	0.84
		Gemma 9B	94.67	0.04	0.48	0.49	0.47
		Gemma 27B	100.64	0	0	0.42	0.42
		Gemini 1.5 Flash	102.02	0.03	0.65	0.48	0.47
		Gemini 1.5 Pro	100.48	0.08	0.53	0.44	0.39
spot the problem	zeroshot	Bloomz 7B	72.21	0.09	0.32	0.47	0.41
		Bloomz 176B	16.78	0.01	0.21	0.45	0.41
		Gemma 9B	108.45	0.01	0.36	0.33	0.32
		Gemma 27B	110.5	0	0	0.38	0.38
		Gemini 1.5 Flash	110.21	0	0	0.2	0.2
		Gemini 1.5 Pro	110.37	0	0.5	0.21	0.21
		GPT 3.5	37.94	0.85	0.26	0.91	0.42
		GPT 4	117.17	0.16	0.38	0.38	0.26
	twoshot	Bloomz 7B	69.63	0.27	0.24	0.95	0.94
		Bloomz 176B	258.1	0.4	0.22	0.79	0.79
		Gemma 9B	102.08	0.07	0.41	0.45	0.41
		Gemma 27B	90.36	0.11	0.49	0.41	0.34
		Gemini 1.5 Flash	108.14	0.07	0.53	0.36	0.31
		Gemini 1.5 Pro	97.86	0.13	0.53	0.38	0.28
GPT 3.5	62.8	0.47	0.27	0.69	0.42		
GPT 4	73.35	0.5	0.38	0.67	0.33		

Table 9: Automatic evaluation of feedback, broken down by prompt type, by model, and by zeroshot vs. twoshot prompting.

don't want to evaluate it if it is not. Feedback that starts talking about another story or that rambles about irrelevant information should be marked as "No."

Q/A Format

Is this feedback for a story, without changing to unrelated topics or adding new stories or feedback for additional texts?

- "No, the text does not give any feedback."
- "No. The text includes feedback, but then adds additional off-topic information, stories, or unrequested feedback."
- "Yes, this is feedback for a story."

2. Feedback Type

Task Instructions

Does the feedback claim that the story is "perfect as-is" or that it doesn't require any changes? If the feedback is simply, "The text is perfect as-is.", please label it "Yes." However, if the feedback says the text is perfect, but then also includes other suggestions, please respond "Both."

Q/A Format

Is the feedback saying that the story is already "perfect as is" or is it suggesting further edits:

- "Yes, the only feedback is that the text is perfect as-is."
- "No, the feedback does not say the text is perfect and suggests changes to the story."
- "Both. It says the text is perfect, but also includes other feedback."

3. Perfect-Agree

Task Instructions

If the feedback doesn't suggest any changes or says the story is "perfect as-is", do you agree? If you don't see any problems or issues with the story, please mark "Yes." If you think there are issues with the story that need to be improved, say "No." Note that this may be subjective in some cases, so please use your best judgment.

Q/A Format

Do you agree that the story is already good as is?

- "Yes, the story seems fine and doesn't have any major issues."
- "No, there are edits I would recommend to improve the story."

4. Correctness

Task Instructions

Does the feedback address weaknesses in the text? It's ok if it doesn't address all the problems in the text, as long as all the things it mentions would improve the story.

Q/A Format

Correctness: Does the feedback address a weakness in the story?

- "Every piece of feedback is correct."
- "Some of the feedback is correct, but some of the feedback is not accurate to the text."
- "No part of the feedback is correct."

5. Error-Detection

Task Instructions

If you feel there is a problem with the story, would following the feedback help fix that specific problem? Does any part of the feedback address what you consider to be the biggest issue in the text? As long as one part of the feedback addresses the problem, you can ignore the rest of the feedback. This may be somewhat subjective, so use your best judgment.

Q/A Format

Error detection: Does the feedback capture the biggest problem (if any) with the story?

- "I don't see any big problems in the story."
- "Yes, at least one piece of the feedback would help fix the biggest problem with the story."
- "No, following the feedback would not help fix the biggest problem in the story."

6. Specificity

Task Instructions

Is it clear that the feedback is for this specific story? Or is it vague or general enough to apply to another story? Specific feedback might mention characters, events, or details from the story, or quote or rewrite pieces of the original story. For this question, it doesn't matter whether the feedback is accurate or helpful, just whether it is specific to the story.

Q/A Format

Specificity: Is the feedback specific to this story or could the feedback be applied to other stories?

- “Yes, the feedback is about this story.”
- “No, the feedback is vague enough that it could be about another story.”

7. Relevance

Task Instructions

When you read the edited story draft, do the changes that were made from the original story reflect the suggestions from the feedback? Or do the edits and the feedback seem unrelated? If it is a mix of both (i.e., some feedback is reflected in the edits and some is not OR some edits are related to the feedback and others aren't), please select “Somewhat.”

Q/A Format

Are any of the changes in the edited story relevant to what was suggested in the feedback?

- “Yes, the changes in the story agree with the suggestions in the feedback.”
- “Somewhat, some of the changes in the story reflect some of the suggestions in the feedback.”
- “No, the changes in the story are unrelated to what the feedback suggested”

D.3 Evaluator Training

We found that annotators needed training specifically for this task, so we first made them complete a tutorial in which they had to read in-depth task instructions and examples. Afterwards, they had to pass a qualifying quiz on three randomly chosen stories from a set that was annotated by experts. For further quality control, we ran analysis on the each annotator after every batch (average inter-annotator agreement rates, the median time spent per story, the variance in the time they spent per story) and performed manual grading on the annotators with outlier characteristics. If an annotator did not pass the regular quality control checks, we removed their annotations and re-annotated their responses with different annotators (about 360 responses were re-annotated).

D.4 Aggregating Evaluator Scores

To aggregate scores, we treat each of the multiple-choice questions as values between $[0, 1]$ and then we average the answers from the annotators. For the binary (yes/no) questions, 0 is “no” and 1 is “yes”. For the questions with a “somewhat” cate-

gory (i.e. relevance & correctness), the “somewhat” is mapped to 0.5. For the error-detection question, we translate both of the non-“Yes” answers (the “I don’t see any issues” answer and the “No there would still be an issue”) to 0. The full phrasing of questions, answers and their mapping to numerical values can be found in Table 11. Because of the pipelined annotation approach, not every dimension was answered for each example, only the applicable ones. If the example has two or more values for a given dimension, we compute the average of the non-missing responses to get the final value.

E Validating Automatic Measures

We validate the performance of the automatic metrics from Section 5 by comparing them to the human evaluation results. For the automatic metrics, we searched for the phrase “is perfect as-is” to automatically detect the feedback type; in the human evaluation, we asked evaluators to label whether the text only contained positive feedback. We find a Pearson correlation coefficient of ~ 0.90 , showing a strong correlation between the presence of the phrase “perfect as-is” and positive feedback. This indicates that models generally followed the prompt instruction to use the phrase “The text is perfect as-is.” when they saw no issues in the story.

We also look at the relationship between the trigram-based diversity metric and the human evaluation specificity ratings. The Pearson correlation of coefficient between non-unique trigrams and the specificity of the response is -0.55 , indicating a moderate anti-correlation.

F Examples of Feedback

Examples of the model-generated feedback are in Table 10.

G Analysis of Human-Authored Feedback

This section contains further details and discussion of the analysis of the human-authored feedback.

G.1 Feedback Analysis Prompts

The prompt given to Gemini 1.5 Flash to label the feedback:

Feedback Labeling Prompt

Please select the category that best summarizes the issue described in the feedback below. Do not include any other information.

Categories: none, spelling and grammar, coherence, clarity and readability, style, needs additions, needs removed, commonsense, other.

Feedback: Sneaked should be snuck

Category: spelling and grammar

Feedback: The story doesn't make sense as Einstein created the theory of relativity so why would he be interested and research something he already knows.

Category: commonsense

Feedback: The dialogue feels unnatural and overly formal for a simple interaction between a man and a fly. The fly's speech, in particular, is too sophisticated.

Category: style

Feedback: <FEEDBACK>

Category:

commonsense of the story, often misunderstanding the intent of the story (e.g., models sometimes gave feedback calling stories unrealistic due to containing talking animals). Models are less likely than humans to comment on fluency issues (spelling/grammar) overall.

G.2 Automatic Labeling Validation

We evaluate the model-generated labels by comparing them to a set of gold labels; the model-generated label was also in the gold label set 84% of the time, out of 480 pieces of feedback annotated by the first author.

G.3 Distribution of Types of Feedback

Figure 5 shows an overview of the frequency of each label type for model- vs. human-authored feedback. Figures 6a and 6b further break down these results by corruption type.

In Figure 6a, we show the distribution of feedback categories (according to Gemini) of the types of feedback found in human-written feedback collected by annotators. General trends follow our expectations based on the type of noise that's been added. Humans give no feedback more often on the original stories. When reading the backtranslate stories, they generally offer more comments on clarity and readability issues. When reading delete stories, they offer more suggestions for details that need to be added into the story. When reading the swap stories, they are most likely to comment on the incoherence of the sequence of events.

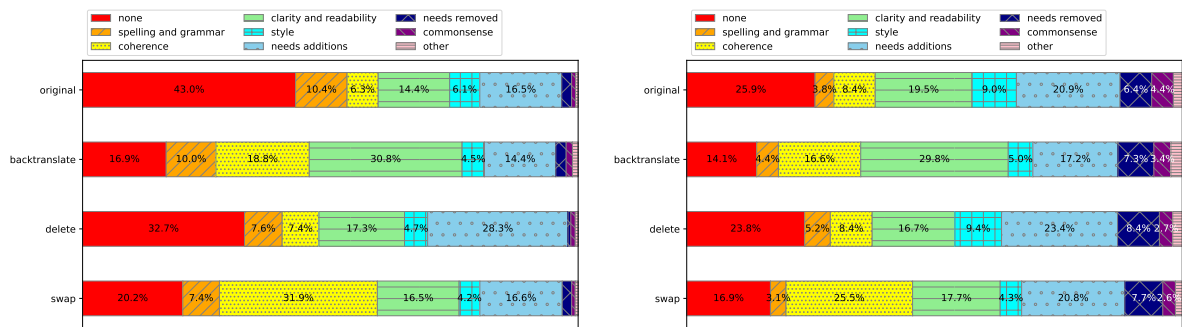
In Figure 6b, we show the distribution of feedback categories of the model-generated feedback. The distributions are mostly similar to the distribution of human-written feedback. However, the model gives feedback more often (uses the 'none' category much less). Models are also more likely to say things need to be removed. Unlike humans, models also give feedback criticizing the

Story	Model-Generated Feedback Examples
<p>She was unsure as to whether she should open the door or not. Slowly she creaked open the door and peered outside. Her brother leaped out at her and gave her a good scare. She screamed and pushed her brother away.</p>	<ul style="list-style-type: none"> - The text is fluent with no typos or grammar issues. - The text is coherent, it maintains consistency and is well-ordered. - The text lacks a bit of clarity, especially in the description of the brother's action. It could be more effective to describe exactly how the brother 'leaped out'. - The text could benefit from a stylistic change, specifically in varying sentence structure for better flow and engagement. - The text could use some additions, specifically more details to give the reader a better sense of the situation and the characters. For example, the reason why the character was unsure about opening the door could be better explained. - There's no need to remove anything from the text as every part contributes to the overall narrative. - The text could use a bit more emotional description to enhance the reader's connection with the characters. For instance, how the character felt after her brother scared her could be expanded upon.
<p>A rooster found a shiny gem while pecking around in the dirt to find food for himself and his hens. I, however, have no use for you and would have preferred to find some corn, instead. " Looking down at the stone, he observed: " Had your owner found you, he would have picked you up and treasured you.</p>	<p>The text switches from the rooster's perspective to the narrator's perspective inconsistently and abruptly.</p>
<p>A raccoon accidentally stumbled into the theatre's collapse. On his journey he was suddenly blinded by a face that looked at him from above. Impressed by her beauty, he soon realized that it was just a mask worn by actors during their performances, and said, "Although you are beautiful, it's a shame that you don't have brains."</p>	<ul style="list-style-type: none"> - The first sentence is ambiguous and lacks clarity. It's unclear what "the theatre's collapse" refers to. - The order of events could be rearranged for better coherence. For example, the raccoon's realization that the face is a mask could come before his reaction to it. - The story could benefit from additional information or details to provide context. For instance, how did the raccoon end up in the theatre?
<p>A bat fell to the ground where he was caught by a passing lynx. The bat begged the lynx to spare his life, but the lynx refused, saying that, by her nature, she is the enemy of all birds. Some time later, the bat again fell to the ground and was caught by yet another lynx, whom he also asked to spare him. The lynx said that she is the enemy of all mice, and will therefore kill the bat. The bat, in turn, assured her that he was not a mouse, but a bird. Thus, he was able to escape once again.</p>	<p>The text lacks consistency in its portrayal of the lynx's nature, as it contradicts itself regarding its prey.</p>

Table 10: Examples of Feedback Written by LLMs

Question	Answer Options
[Sanity-Check] Well-formed feedback: Is this feedback for a story, without changing to unrelated topics or adding new stories or feedback for additional texts?	<p>Yes, this is feedback for a story. → 1</p> <p>No. The text includes feedback, but then adds additional off-topic information, stories, or unrequested feedback. → 1</p> <p>No, the text does not give any feedback. → 0</p>
[Feedback-type] Is the feedback saying that the story is already “perfect as is” or is it suggesting further edits:	<p>Yes, the only feedback is that the text is perfect as-is. → 1</p> <p>Both. It says the text is perfect, but also includes other feedback. → 0</p> <p>No, the feedback does not say the text is perfect and suggests changes to the story. → 0</p>
[Perfect-Agree] Do you agree that the story is already good as is:	<p>Yes, the story seems fine and doesn’t have any major issues. → 1</p> <p>No, there are edits I would recommend to improve the story. → 0</p>
[Correct] Correctness: Does the feedback address a weakness in the story?	<p>Every piece of feedback is correct. → 1</p> <p>Some of the feedback is correct, but some of the feedback is not accurate to the text. → .5</p> <p>No part of the feedback is correct. → 0</p>
[Error-Detection] Error detection: Does the feedback capture the biggest problem [if any] with the story?	<p>Yes, at least one piece of the feedback would help fix the biggest problem with the story. → 1</p> <p>No, following the feedback would not help fix the biggest problem in the story. → 0</p> <p>I don’t see any big problems in the story. → 0</p>
[Specificity] Specificity: Is the feedback specific to this story or could the feedback be applied to other stories?	<p>Yes, the feedback is about this story. → 1</p> <p>No, the feedback is vague enough that it could be about another story. → 0</p>
[Relevance] Are any of the changes in the edited story relevant to what was suggested in the feedback?	<p>Yes, the changes in the story agree with the suggestions in the feedback. → 1</p> <p>Somewhat, some of the changes in the story reflect some of the suggestions in the feedback. → 0.5</p> <p>No, the changes in the story are unrelated to what the feedback suggested → 0</p>

Table 11: Human evaluation questions, full wording and mapping to numerical values



(a) Human Feedback Categories: The categories of feedback provided by human annotators in their free-text responses, split by type of story. Categories were assigned automatically by Gemini.

(b) Model Feedback Categories: The categories of feedback provided by models in their output responses, split by type of story. Categories were assigned automatically by Gemini.

Figure 6: Types of feedback written by humans and generated by models.