

CHIRON: Rich Character Representations in Long-Form Narratives

Alexander Gurung
School of Informatics
University of Edinburgh
Edinburgh, UK
a.gurung-1@sms.ed.ac.uk

Mirella Lapata
School of Informatics
University of Edinburgh
Edinburgh, UK
mlap@inf.ed.ac.uk

Abstract

Characters are integral to long-form narratives, but are poorly understood by existing story analysis and generation systems. While prior work has simplified characters via graph-based methods and brief character descriptions, we aim to better tackle the problem of representing complex characters by taking inspiration from advice given to professional writers. We propose CHIRON¹, a new ‘character-sheet’ based representation that organizes and filters textual information about characters. We construct CHIRON sheets in two steps: a Generation Module that prompts an LLM for character information via question-answering and a Validation Module that uses automated reasoning and a domain-specific entailment model to eliminate false facts about a character. We validate CHIRON via the downstream task of masked-character prediction, where our experiments show CHIRON is better and more flexible than comparable summary-based baselines. We also show that metrics derived from CHIRON can be used to automatically infer character-centricity in stories, and that these metrics align with human judgments.²

1 Introduction

“The three things that matter most in a story are characters, characters and characters.” – Bob Gale³

Advice to professional writers emphasizes the importance of well-developed and interesting characters in stories, identifying them as the key to engaging and believable narratives (Kyle, 2016; Jarvis, 2014; Weiland, 2016; The Editors of Writer’s Digest Books, 2014).

¹In Greek mythology, CHIRON was a highly educated centaur known for his wisdom, and as a mentor of many Greek heroes, including Achilles.

²Our data and code are available at <https://github.com/Alex-Gurung/CHIRON>

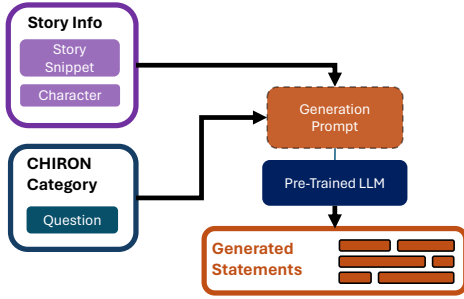
³American screenwriter, producer, and film director best known for co-writing the “Back to the Future” trilogy.

Computational techniques aiming to analyze and generate literary texts have also emphasized the importance of characters in storytelling by tracking their relationships and emotional trajectories (Chaturvedi et al., 2017; Iyyer et al., 2016; El-sner, 2012), identifying enemies and allies (Nalissenick and Baird, 2013), villains or heroes (Bamman et al., 2013a, 2014), and representing characters via attributes (Zhang et al., 2019), personality types (Sang et al., 2022), or high-level paragraph-style descriptions (Brahman et al., 2021a; Chen et al., 2024; Li et al., 2024; Xu et al., 2024; Yu et al., 2024; Yang et al., 2023, 2022).

In this work we propose a new ‘character-sheet’ based representation, CHIRON, standing for **CH**aracter-**I**nformed **R**epresentations **O**ver **N**arratives. CHIRON consists of structured descriptions of characters along several dimensions, encapsulating character nuance and complexity, while still being interpretable and useful for downstream tasks and analysis. It is composed of two modules: a Generation Module that generates long-form answers to questions based on a snippet of story, and a Validation Module that verifies the utility and truthfulness of these answers.

Drawing inspiration from how writers create their own characters (Jarvis, 2014; Weiland, 2016; The Editors of Writer’s Digest Books, 2014), our character sheets have four broad categories: **Dialogue** (how a character speaks), **Physical/Personality** (how a character looks/acts), **Knowledge** (what a character knows), and **Goals** (what a character wants to accomplish/has accomplished thus far). We fill these categories (see Table 1) in the Generation Module with information about a given character by prompting out-of-the-box language models with a snippet from a story and questions relevant to each category. Despite using a wide variety of models, we find that the generations at this stage are often unfaithful to the story. We develop a Verification Module that uses

1) Generation Module: Generate responses to questions for each CHIRON category. Responses are **statements about the character**, based on the snippet.



2) Validation Module: Intermediary reasoning steps determine ambiguity and uninformativeness (ICL), and find text relevant to the statement (CoT). The **entailment model** compiles reasoning and determines if a statement should be kept in the final representation.

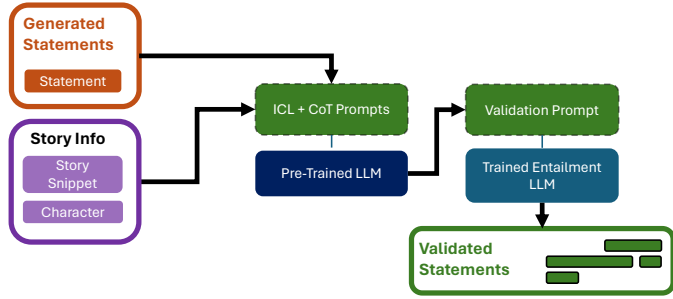


Figure 1: Overview of CHIRON’s Generation and Validation Modules. For each CHIRON category, our Generation Module takes a character and story snippet and uses a pretrained LLM to generate statements about the character. Our Validation Module then passes these statements through reasoning steps and a trained entailment classifier to determine if they are true and useful. We find this methodology produces more accurate and useful character-representations for downstream tasks and analysis. We use Mistral 7B-Instruct v0.2 unless otherwise stated.

automated reasoning in combination with a high-precision entailment model to ensure our character-representations are accurate and useful.

We build this work on STORIUM (Akoury et al., 2020), a dataset of collaboratively written stories where contributors write entries from the perspective of a specific character. This gives us a perfect test-bed to explore character representations, as we know which character will be heavily featured in each story entry. We validate CHIRON in two ways. We assess its usefulness for the downstream task of masked-character prediction, generating character sheets up-until a given story snippet and then predicting the character-IDs from the masked snippet. We find an 11.6% performance improvement over summary-based baselines, demonstrating that our structured format would be more useful than existing methods for character-centric tasks. We also develop automated metrics based on our character sheets as a means of understanding the use of characters in human-written (Akoury et al., 2020; Chakrabarty et al., 2024) and LLM-generated stories (Yang et al., 2023). We show that our proposed *density* metric aligns with human judgments of character-centricity in stories. Our primary contributions are:

- We present CHIRON, a novel ‘character-sheet’ structure, for representing characters in long-form narratives;
- We develop a Validation Module to automatically determine if a given statement about a character is true based on the snippet, and use it to produce more accurate character sheets over the course of a story;

- We validate CHIRON via the task of masked-character prediction, showing improved performance over competitive baselines;
- We use automated metrics derived from CHIRON to better understand characters in human-written and LLM-generated stories.

2 Related Work

Modeling Stories, Narratives and Characters

Prior work represents characters in three general ways: 1) as discrete values (e.g., sentiments, adjectives, or graph nodes) (Hoque et al., 2023; Soni et al., 2023; Yu et al., 2023; Gorinski and Lapata, 2015; Elson et al., 2010; Ye et al., 2022; Peng et al., 2022) 2) as embeddings (Bamman et al., 2013b; Kim et al., 2018) and 3) as high-level paragraph-style descriptions or summaries (Brahman et al., 2021a; Chen et al., 2024; Li et al., 2024; Xu et al., 2024; Yu et al., 2024; Yang et al., 2023, 2022). Our work is most similar in form to this last approach, and we use generated character summaries as a baseline in our experiments.

Much research on modeling characters comes from role-playing in text-based games (Ammanabrolu et al., 2021; Shuster et al., 2020; Zhu et al., 2023; Rameshkumar and Bailey, 2020; Louis and Sutton, 2018). In particular, Ammanabrolu et al. (2021) showed good role-playing dialogue and action performance by providing models with multi-faceted motivations and brief character summaries. Early work using LLMs for story generation has explored the benefit of character information indirectly by providing plot outlines as guidance (Rashkin et al., 2020; Fan et al., 2018;

Peng et al., 2022). More recent work focusing on generating longer stories gives more weight to complex plot structure (Yang et al., 2023, 2022; Chung et al., 2022). Yang et al. (2023) tracks character development while generating stories, although their representation of characters is unstructured and difficult to validate independently.

Literary Understanding of Characters We take inspiration from existing discussion aiming to help writers develop characters in their own narratives (Jarvis, 2014; Weiland, 2016; The Editors of Writer’s Digest Books, 2014). While the advice is largely meant for *creating* fully fleshed out characters, it is also useful for *analysing* existing ones, and provides insight into what ancillary notes writers use to keep characters consistent. Two insights in particular greatly inform this work 1) the aspects important to different characters vary wildly and can be extremely nuanced in their substance and 2) despite this, many writers still find it useful to keep organized notes on key character details as they write. We interpret these insights to indicate the need for un-simplified descriptions of character features, and hypothesize that an organized collection of these descriptions would be useful for analysing and generating stories.

3 STORIUM Dataset

We base our work on the STORIUM dataset, a collection of 6k stories written collaboratively online (Akoury et al., 2020). We filter these for length (avoiding stories with too little content or too long for analysis), English language, and traditional novel writing style (as opposed to game-playing). This leaves us with 529 stories, containing an average of 5.5k words in 42 entries. Each entry is from a specific character’s perspective (although as in traditional narratives the use of grammatical person and narrator omniscience varies). Due to the length of many of these entries, we further split them into snippets, along paragraph lines. We select the top three most frequently occurring characters, giving us 1,403 character-story combinations and 9,021 snippets. Further details are presented in Appendix A.

While our CHIRON representation can be applied to any story text, we primarily rely on STORIUM for development and validation. As each of these story snippets is centered around the given character, STORIUM is ideal for exploring how characters are represented in narratives. An exam-

ple STORIUM snippet and character sheet based on it is shown in Table 1.

4 CHIRON: Character Sheets

While prior work has used detailed plot outlines and knowledge graphs (Peng et al., 2022; Yang et al., 2023), developing rich character representations in a manner similar to writers is an unexplored area of research. Such representations should accurately describe all attributes of a given character that may be important for the reader to know, and should be consistent (or change logically) across the story. Some attributes (e.g., height, name) are well suited for knowledge graphs and other database-inspired representations, but many of the more subtle facts about characters are too complex to store in this manner (e.g., they are self-conscious about their clothes).

Inspired by advice for writers (Jarvis, 2014; Weiland, 2016; The Editors of Writer’s Digest Books, 2014), we propose instead to represent characters in a long-form character-sheet structure we call CHIRON. We extrapolated themes and condensed lists of specific questions (e.g. “Is this character intelligent, and in what way?”) from these resources to create a structure more broadly applicable to any story. As seen in Table 1, CHIRON is organized as four broad categories, each composed of questions-and-answer pairs:

Dialogue: how the character speaks; this could include accents, catch-phrases, whether they prefer flowery or direct prose, etc.

Physical/Personality: descriptions of the character’s appearance and personality; this could include height, scars, courage, condescension, etc.

Knowledge: other factual details about this character, as well as knowledge this character has learned about the world; this could include family members, a tragic backstory, where the gold is hidden, etc.

Goals: goals this character has gained or completed, and changes in their internal motivation; this could include finding the treasure, gaining a quest, wanting revenge, etc.

We operationalize CHIRON’s character-representation task as two steps: a Generation Module and a Validation Module. (i) The generation step takes in a character and snippet of story, and uses an out-of-the-box LLM to answer a series of open-ended questions about the character, and organizes these question-answer pairs into

Story Snippet		
Nadia watched both of the men silently, pulling her tattered, too-big sweatshirt tighter around herself. She hadn't spoken since the plane had crashed, and she didn't want to, but it looked like the men were about to devolve into an argument about whether or not to go into the cave, and she actually had input. She cleared her throat and, when she had the attention of both Jacob and Jayson, she pointed to the sky. Their heads all turned. They could barely see them through the trees, but the grey storm clouds were unmistakable all the same. Nadia tried to speak but had to clear her throat once more and swallow before she was able to do so, not having used her voice for so long. "I don't know about you, but I'd really rather not be out in the forest during a storm. I say we check out the cave. If we're careful about it, we won't run into any trouble."		
Dialogue	What, if anything have we learned about how this character speaks from this snippet?	She has gone without using her voice for an extended period.
Physical/ Personality	What, if any, physical descriptions of this character are in this snippet? What, if any, descriptions of this character's personality are in this snippet?	She pulls her tattered, too-big sweatshirt tighter around herself. Nadia is quiet.
Knowledge	What, if any, factual information is given about this character in this snippet? What, if any, information has this character learned in this snippet?	Nadia is a survivor of a plane crash. She is aware of the approaching storm.
Plot	What, if any, goals does this character gain in this snippet that they wish to accomplish in the future? What, if any, goals does this character complete in this snippet? How, if at all, does this character's internal motivations change in this snippet?	Nadia's goal in this snippet is to seek shelter from the approaching storm. Nadia suggests seeking shelter in the cave during an approaching storm. Nadia, who had been silent since the plane crash, feels the need to speak up.

Table 1: Sample character sheet for Nadia. Answers to questions were provided by Mistral 7B-Instruct v0.2.

categories. (ii) The validation step then splits the resulting answers into sentences and filters out incorrect sentences using our intermediary reasoning steps and domain-specific entailment model. Figure 1 shows an overview of this two-step process and Table 1 shows an example character sheet (one sentence taken per question).

This method allows us to leverage the high performance of out-of-the-box LLMs on instruction following tasks, while using our entailment model to prevent hallucinations corrupting the resulting representation. For tasks that require a character-representation across multiple snippets, we simply concatenate these filtered answers by question and remove duplicates via pairwise cosine similarity between TF-IDF vectors (thresholded at 0.9). We also incorporate a simplification step prior to the Validation Module that splits sentences into individual claims; further explanation is in Appendix H and Table 15.

5 Validation Module

Given a story snippet, character, and statement (about said character), our Validation Module predicts whether the statement entails completely from the snippet. As this task is extremely domain-specific, we construct this module from two novel parts: automated reasoning and a trained entail-

ment model. In this section, we describe our methodology for collecting data, setting up automated reasoning, and training our domain-specific entailment model.

5.1 Dataset Creation and Annotation

Character Sheet Generation We generate character sheets by showing a pretrained out-of-the-box LLM a snippet of story and the associated character, and asking questions for each category. Full prompt examples are shown in Appendix H. To diversify our responses, we generate character sheets using four models of varying sizes: Mistral 7B Instruct v0.2 (Jiang et al., 2023), Llama 2 Chat 13B (Touvron et al., 2023), Tulu 2 DPO 70B (Iverson et al., 2023), and GPT-4 (version gpt-4-0613) (OpenAI, 2023). We generate 250 character sheets from each model, each representing a unique story-character combination. Taking the question-response pairs that compose character sheets, we split responses into sentences to give us sentence-snippet-character tuples. We randomly sample 1,000 tuples (250 from each model) for annotation.

Entailment Annotations Human annotators were asked to judge whether a statement about a character entails from a story snippet using a 1–5 scale to account for ambiguity and multiple interpretations of the same text. A statement labeled

Label:	1	2	3	4	5
Mistral 7B	10.8	2.8	9.6	15.2	61.6
Llama-2 13B	24.4	3.2	13.6	17.2	41.6
Tulu-2 70B	10.4	2.4	12.4	16.8	58.0
GPT-4	24.0	1.6	14.4	13.2	46.8
Combined	17.4	2.5	12.7	15.4	52.0

Table 2: Distribution of labels (in %) in annotated dataset (1,000 examples, 250 per model). Despite being the smallest model, Mistral 7B-Instruct seems to perform the best, although prompts were not optimized for each model, so further work could find different results. Prompts were primarily tested using Mistral and Llama before being used uniformly across all models.

with: 1 is entirely unsupported by the snippet; 2 is largely contradicted by the snippet; 3 is ambiguous in its relationship with the snippet; 4 is likely true based on the snippet, or 5 is entirely supported by the snippet. Detailed examples are included in [Appendix C, Table 8](#).

Annotators were recruited through Prolific; they were first shown instructions and a series of examples with justifications for the labels, before annotating five snippet-statement pairs on their own. In addition to entailment (called ‘consistency’ for comprehension), crowdworkers were asked to annotate statement coherence and to provide brief justifications for their answers, which we found improved attention to detail. Screenshots of the interface and more details are in [Appendix C](#). Our data was annotated by a group of 7 annotators who passed various screening tests and spent a considerable amount of time on the task (over 21 minutes on average). Inter-annotator agreement was high, achieving a Krippendorff Alpha of 0.679.

[Table 2](#) shows the label distribution in our annotated dataset, we find that 52% of sentences were entirely correct and entailed from the snippet, and an additional $15.4\% + 52\% = 67.4\%$ were largely correct (a label of 4 generally indicates a minor potentially ambiguous or unclear phrase amongst an otherwise verifiably correct statement). The remaining 32.6% of sentences were unverifiable or false, although it should be noted that this includes statements that make no claims. Further details and examples are presented in [Appendix B](#).

5.2 Statement Validation

We split our dataset 80-10-10 for train-validation-test, and use Mistral 7B Instruct v0.2 ([Jiang et al., 2023](#)) as our base model, given its better entailment performance (see [Table 2](#)). Although overall

Name	Precision	Recall	Accuracy
Zero-Shot All	0.753	0.986	76%
Zero-Shot ICL	0.759	0.887	72%
Zero-Shot CoT	0.753	0.986	76%
Zero-Shot None	0.835	0.930	82%
Trained All	0.930	0.746	78%
Trained ICL	0.952	0.563	67%
Trained CoT	0.864	0.803	77%
Trained None	0.895	0.718	74%

Table 3: Comparing baseline (zero-shot) performance and trained performance by shown reasoning steps. ‘All’ refers to showing both ICL and CoT reasoning steps to the model, ICL and CoT refers to just showing their respective reasoning steps, and ‘None’ refers to not showing any reasoning steps. We prioritize precision but want to maintain a good balance between precision and recall, so choose the ‘Trained All’ setting moving forward. All models are based on Mistral 7B Instruct v0.2 ([Jiang et al., 2023](#)).

accuracy is important, we pay particular attention to precision as we want to ensure all statements included in the final character sheet are correct.

Automated Reasoning Prior work has shown In-Context-Learning (ICL) and Chain-of-Thought (CoT) reasoning can improve zero-shot classification performance ([Brown et al., 2020](#); [Wei et al., 2022](#)). To this end, we generate intermediate reasoning steps using Mistral 7B Instruct v0.2 ([Jiang et al., 2023](#)) that help our classifier accurately identify entailment by plainly decomposing the reasoning steps. We split reasoning into three steps: ambiguity (is this statement too ambiguous to verify), informative (does this statement provide any information about the character), and a chain-of-thought step that prompts the model to retrieve relevant pieces of the story and then to compare and contrast with the statement. More details are given in [Appendix E](#).

Domain-Specific Entailment Model We combine the snippet, reasoning steps, character, and statement in the context and prompt a Mistral 7B model for an entailment label (1–5). To best take advantage of the language model’s existing abilities we fine-tune with the causal language modeling objective (on just this labeling step) and LoRA ($r = 16, \alpha = 32$) ([Hu et al., 2022](#)). During training we take the best-performing model by precision (boolean where label ≥ 4) on the validation set. We find slightly improved performance by thresholding our model’s predictions at $= 5$, further details and comparisons are presented in [Appendix G](#).

Table 3 compares our trained model (Trained All) against ablated variants (ICL, CoT, None) and zero-shot models. As can be seen, Trained All improves on baseline zero-shot precision (without reasoning) by 0.095, achieving a precision of 0.930. Removing our CoT step can further improve precision, but at too great a cost to recall and accuracy. We find the performance of this Validation Module promising as it shows that even in this rhetorically complex domain where models of varying sizes struggle to extract consistent character details, the analytical task of determining entailment is possible simply by fine-tuning a 7B-parameter model. As we use this Validation Module to filter out incorrect statements, the high precision also lends credence to the resulting character sheets without needing to approach the more difficult task of improving character-statement generations. An example of a full character sheet (across an entire story) before and after our Validation Module is shown in Appendix D (Table 9, Table 10).

6 Masked Character Prediction

In this section we validate CHIRON on a masked character prediction task (Holgate and Erk, 2021; Brahman et al., 2021b) by showing it serves as a better character-representation than a competitive baseline. We define the task as follows: given an unseen STORIUM story, we get the three most common characters by scene appearance and identify story snippets that contain all three characters. We create character sheets for each story up-to this snippet. We mask out character names from the story snippet and, given this masked story snippet and generated character-information, prompt Mistral 7B Instruct v0.2 (Jiang et al., 2023) to match each character to its mask-ID. Filtering for finished, previously unseen stories and scenes with their three most common characters gives us a dataset of 49 entries with 3 characters each, or 147 data points. Table 4 shows an example.

Settings and Baselines We have two CHIRON-based settings. *EntireSheet* is the most traditional setting, where the entire CHIRON character-sheet is shown to the model. In contrast, our *Agreed* setting refers to showing to the masked-character prediction model only one category at a time (Dialogue, Physical/Personality, Knowledge, or Goals), and aggregating the predictions by multiplying the softmax probabilities and taking the maximum. We also briefly explore aggregating the predictions

[CHAR 0] glanced over when [CHAR 2] mentioned [CHAR 1]. "Heh, well she better get going if she doesn't want to be left behind." [CHAR 0] caught his breath as he leaned down, taking a moment, making sure he was hidden from sight. He saw [CHAR 1] lean down and hold Ian for some time, before getting up and calmly walking away. "Wha'she doing there ya think? Speaking with the dead? Ah, whatever." [CHAR 2] piped up about his uneasiness in this world. [CHAR 0] scoffed at him. "What does it matter? It's just a tree. I haven't noticed anything different. I'm sure you're just shaken from seeing a guy murdered next to you," [CHAR 0] chuckled as he smacked [CHAR 2] on the back. "Now, where to next, mr. navigator?"

Table 4: Masked snippet from unseen STORIUM story. The task is as follows: given a character name and representation, determine which character-ID (0, 1, or 2) corresponds to that character in the snippet.

from a subset of our categories, as some may be more informative than others for this specific task.

We also have two baseline settings. *Character-Summary* sets the character-information to a single character description obtained by prompting Mistral 7B Instruct v0.2 (Jiang et al., 2023) with the entire story so-far and asking for a summary of the given character. This summary is then filtered using our entailment classification pipeline to ensure that differences in performance can be attributed to the difference in representation (as opposed to differences in hallucination rates). Table 5 presents an example summary. *No-Information* refers to showing no character-information during the masked-character prediction task. As pronouns and other name-based clues make the task slightly easier than random, this baseline serves as a more accurate performance floor.

We also introduce a STORIUM-specific hyper-parameter *OnlyRole* that applies to all character representations. *OnlyRole* refers to the boolean choice of whether to only include story snippets from the character-of-interest's perspective as input (where *true* refers to only using such snippets). Excluding non-perspective snippets may leave out crucial information, but may increase the density of useful information for the task. We ultimately report performance under both settings and find CHIRON tends to benefit from *OnlyRole*.

Results Table 6 shows the results of our experiments. We find all of our settings to outperform the *No-Information* performance-floor (43.5%) as expected; both *Agreed* (50.3%) and

Summary Question	Response
Summarize everything we have learned about this character across these snippets. Include aspects of the character like how they speak, what they look like, their personality, their goals, etc.	Nadia is a quiet and thoughtful survivor. She wears a tattered, too-big sweatshirt. She hasn't spoken much since the plane crash. She suggests checking out the cave for shelter. She's resourceful and thinks logically. She's determined to survive and keep the group safe.

Table 5: Sample *Character-Summary* baseline for Nadia. Input story snippet was the same as in Table 1.

Setup	OnlyRole	Accuracy
NoInformation	—	43.5%
Character-Summary	✗	44.9%
CHIRON: Entire-Sheet	✗	47.6%
CHIRON: Agreed	✗	50.3%
Character-Summary	✓	42.2%
CHIRON: Entire-Sheet	✓	56.5%
CHIRON: Agreed	✓	53.7%
CHIRON: Agreed-Subset	✓	58.5%

Table 6: Accuracy on masked-character prediction task, separated by *OnlyRole* (whether to only include snippets from the characters perspective when creating a character representation). CHIRON representations uniformly perform better than *NoInformation* and *Character-Summary* baselines; this difference becomes more pronounced under the *OnlyRole* setting. We also report the best performing *Agreed* setting across all possible subsets of categories (Knowledge+Physical).

EntireSheet (47.6%) outperform the more comparable *Character-Summary* baseline (44.9%). This difference is even greater under the *OnlyRole* setting, where baseline performance decreases but our *EntireSheet* accuracy jumps to 56.5%. This validates our assertion that CHIRON-based character-representations provide a tangible benefit on downstream tasks compared to existing summary-based representations, without any additional training.

We can further improve performance by modifying our *Agreed* setting to compute agreement exclusively between our Knowledge and Physical categories. This *Agreed:Knowledge+Physical* setting, with *OnlyRole*, results in an accuracy of 58.5%, a full 13.6% higher than the best performing baseline. The optimal combination of character-sheet categories would likely change depending on the downstream task, and we consider it a strength of CHIRON that it provides this flexibility.

7 Story Analysis by Character Density

In addition to being useful for downstream tasks, we wish to show that CHIRON is useful as a means of automatically analyzing stories from a variety of sources. We create a metric, *density*, and show that it correlates with human judgements of character-centricity before applying it to a varied set of hu-

man and LLM-generated story datasets.

We define *density* as the number of sentences in a finished character sheet divided by the number of sentences in the provided story, averaged across all story-character combinations. Let S denote the set of story-character combinations, C is the set of categories (e.g., Dialogue, Goals), and $\chi(s, c)$ is the character sheet category c given a story-character:

$$\text{density} = \frac{1}{|S|} \sum_{s \in S} \sum_{c \in C} \frac{\# \text{ sentences in } \chi(s, c)}{\# \text{ sentences in story of } s}$$

We interpret this metric as a way of measuring the extent to which characters are the focus and driving force behind a story, or in other words a story's *character-centricity*.

We can also compute density for each category in CHIRON, quantifying the extent to which a given aspect of characterization is present in a story:

$$\text{density}_c = \frac{1}{|S|} \sum_{s \in S} \frac{\# \text{ sentences in } \chi(s, c)}{\# \text{ sentences in story of } s}$$

Correlation with Human Judgments We evaluate our *density* metric against human judgments on stories produced by DOC, Re³ and their variants (Yang et al., 2023, 2022). They are state-of-the-art models, exhibiting traditional narrative style. Moreover, their stories are relatively short (averaging 2k fewer words than our STORIUM set) and thus amenable to human review.

The Re³ framework introduced a prompting-based method for drafting, rewriting and editing generated stories (Yang et al., 2022). DOC built on this framework by adding a more detailed plot outline and a controller that attempts to match generated stories to the outline. ‘No-Control’ and ‘NO-Outline’ refer to settings without the controller and the detailed outline respectively. Both Re³ and DOC also have an ‘Interactive’ setting, where humans are involved in the story-planning process (Yang et al., 2023). We collated the output of their systems, a total of 100 stories across 6 settings, which we annotated using a similar Prolific setup to our entailment-classifier dataset.

Annotators were asked to rate (from 1–5) a story across five dimensions meant to encapsulate how well-defined characters are in the story and how central to the plot: **Roles** (Does the story have characters with explicitly defined roles?); **Traits** (Do the characters have well-defined traits and relationships?); **Actions** (Do their actions play a role in the story?); **Physical/Personality** (Do the characters have defined physical/personality attributes?); and **Dialogue** (Do the characters have defined manners of speaking?). We sum up these scores to give us a singular metric of character importance which we correlate with our CHIRON-based *density* metric.

We find that our automated *density* metric correlates highly with human judgements of character centrality in stories, achieving a Pearson correlation of 0.753 ($p < 0.01$). This allows us to use CHIRON as an analysis tool on datasets where human annotations would be expensive or intractable. We can also determine the correlations between $\text{density}_{\text{dialogue}}$ and $\text{density}_{\text{physical/personality}}$ and their respective questions in our annotation task. We find a dialogue correlation of 0.362 ($p < 0.01$) and a physical/personality correlation of 0.440 ($p < 0.01$), further justifying that the categories within CHIRON are useful for deeper analysis.

Comparing Stories via Character Density We briefly show an example of how CHIRON’s density metric can be used to differentiate between different stories and styles. We compute our density metric across the following datasets, for the most common character in each story:

STORIUM (Akoury et al., 2020): We expect to find a high density from STORIUM stories due to the extreme character focus inherent to the way the stories were created (almost akin to role-playing). We report density based on all-snippets by default as it is the most comparable between datasets, but we also report density on the *OnlyRole* setting.

New Yorker TTCW (Chakrabarty et al., 2024): this dataset is derived from New Yorker stories and LLM-generated stories from a premise. Chakrabarty et al. (2024) showed that the New Yorker stories were much more creative (as measured by the TTCW⁴ metrics) than the LLM-generated stories, and were in particular more original in their Form, Thought, Theme and Content. We expect to find low density from the New Yorker stories due to their differences from conventional narratives (a few stories do not even have named

Source	Setup	Density
New Yorker	New Yorker	0.648
DOC	No-Control	0.746
DOC	DOC	0.810
DOC	Interactive	0.840
DOC	No-Outline	0.922
New Yorker	GPT4	1.037
Re ³	Re ³	1.046
Re ³	Interactive	1.058
New Yorker	Claude	1.096
Storium	Storium	1.254
New Yorker	GPT3.5	1.279
Storium	Storium-OnlyRole	1.668

Table 7: CHIRON-based density metric computed across different data sources and setups, sorted by density. We find human-written New Yorker stories to have the lowest density, potentially due to the more creative narrative structures referenced in Chakrabarty et al. (2024). In contrast, our Storium-OnlyRole setting has the highest density by a wide margin, likely explained by the extremely character-driven source material (Akoury et al., 2020)

main characters, for example).

DOC stories (Yang et al., 2023, 2022): We compare the automatically generated stories from the following settings: *DOC*, *DOC-nocontrol*, *DOC-nooutline*, *DOC-interactive*, *Re³*, *Re³-interactive*.

The results in Table 7 corroborate our expectations. In particular, the more ‘creative’ narrative structures of human-written New Yorker stories (Chakrabarty et al., 2024) produced a low *density* score, while the character-focused Storium-OnlyRole setting produced the highest. Results for DOC (Yang et al., 2023) and Re³ (Yang et al., 2022) are harder to interpret, but seem to indicate higher character-centricity in Re³ stories than DOC ones and higher character-centricity in stories humans edited (the ‘Interactive’ settings). Our *density* metric should not be seen as a score for story ‘quality’ or ‘goodness’. Instead, it should be used to compare styles of stories or story generation systems.

8 Conclusion

In this work we introduce CHIRON, a novel character-sheet representation for long-form narratives consisting of a zero-shot Generation Module (in the form of structured question-answering) and a Validation Module composed of automated reasoning and a trained entailment-model. Out-of-the-box LLMs frequently (32.6% of sentences) produce character descriptions inconsistent with the story, but our Validation Module can differentiate with 0.930 precision between consistent and

⁴TTCW is a short for Torrance Test of Creative Writing.

inconsistent statements. On the downstream task of masked-character prediction, CHIRON outperforms (11.6% gain) a comparable summary-based baseline, which suggests that the representation would be useful for other character-centric tasks. Finally, we show that *density*, an automated metric derived from CHIRON, correlates with human judgements of character-centricity and can be used to differentiate story styles. In future work, we plan to leverage CHIRON to generate more character-driven stories.

Limitations

While the character sheet structure selects for specific types of information, determining what is an ‘interesting’ statement about a character is an open question and one that is hard to answer without knowledge of the entire story. Whether a character holds an object may be only relevant short-term, for example, but could also be crucially important when they pull it out in the final act.

A follow-up study could also better leverage information readers gain across the narrative as context for the current snippet. For example, many snippets in a story will contain story-specific vocabulary or established character relationships, but do not explicitly explain them as they were introduced previously. We only showed annotators one snippet of story at a time, but future work could relax this requirement and ask annotators to read the entire story up to the snippet of interest.

We also focus exclusively on English-language stories in this work; future research should expand our ontology and annotations to other languages and explore using other language models pretrained on more non-English data.

Ethics Statement

Any automated system for text generation has the potential to create harmful content. As such, we aimed to avoid potentially inappropriate domains by leveraging the STORIUM dataset’s filters. We also took proper precautions to allow crowdsourced workers to withdraw from the task at any point.

It is also possible that abstract character representations (e.g. ‘*X character is a member of Y group*’) could lead to stereotypical generations when used for story-generation purposes. While we hope that encouraging more complex and rich character representations makes such stereotypes less likely, this is an open question for future re-

search. Such work could also explore the potential for bias in the statement-generation step, where a similar bias may cause models to fill character-sheets with incorrect statements about a character due to stereotypes on mentioned attributes.

Acknowledgments

We gratefully acknowledge the support of the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1),

References

- Nader Akoury, Shufan Wang, Josh Whiting, Stephen Hood, Nanyun Peng, and Mohit Iyyer. 2020. [STORIUM: 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.
- Prithviraj Ammanabrolu, Jack Urbanek, Margaret Li, Arthur Szlam, Tim Rocktäschel, and Jason Weston. 2021. [How to Motivate Your Dragon: Teaching Goal-Driven Agents to Speak and Act in Fantasy Worlds](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 807–833, Online. Association for Computational Linguistics.
- David Bamman, Brendan O’Connor, and Noah A. Smith. 2013a. [Learning latent personas of film characters](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 352–361, Sofia, Bulgaria. Association for Computational Linguistics.
- David Bamman, Brendan O’Connor, and Noah A. Smith. 2013b. [Learning latent personas of film characters](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 352–361, Sofia, Bulgaria. Association for Computational Linguistics.
- David Bamman, Ted Underwood, and Noah A. Smith. 2014. [A Bayesian mixed effects model of literary character](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 370–379, Baltimore, Maryland. Association for Computational Linguistics.
- Faeze Brahman, Meng Huang, Oyvind Tafjord, Chao Zhao, Mrinmaya Sachan, and Snigdha Chaturvedi. 2021a. [“let your characters tell their story”: A dataset for character-centric narrative understanding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1734–1752, Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Faeze Brahman, Meng Huang, Oyvind Tafjord, Chao Zhao, Mrinmaya Sachan, and Snigdha Chaturvedi. 2021b. “Let Your Characters Tell Their Story”: A Dataset for Character-Centric Narrative Understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1734–1752, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Tuhin Chakrabarty, Philippe Laban, Divyansh Agarwal, Smaranda Muresan, and Chien-Sheng Wu. 2024. *Art or artifice? large language models and the false promise of creativity*. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI ’24, New York, NY, USA. Association for Computing Machinery.
- Snigdha Chaturvedi, Mohit Iyyer, and Hal Daume III. 2017. Unsupervised learning of evolving relationships between literary characters. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, Aili Chen, Nianqi Li, Lida Chen, Caiyu Hu, Siye Wu, Scott Ren, Ziquan Fu, and Yanghua Xiao. 2024. *From Persona to Personalization: A Survey on Role-Playing Language Agents*. Publisher: arXiv Version Number: 1.
- John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. *TaleBrush: Sketching Stories with Generative Pretrained Language Models*. In *CHI Conference on Human Factors in Computing Systems*, pages 1–19, New Orleans LA USA. ACM.
- Micha Elsner. 2012. *Character-based kernels for novelistic plot structure*. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 634–644, Avignon, France. Association for Computational Linguistics.
- David Elson, Nicholas Dames, and Kathleen McKeown. 2010. *Extracting social networks from literary fiction*. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 138–147, Uppsala, Sweden. 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.
- Philip John Gorinski and Mirella Lapata. 2015. *Movie Script Summarization as Graph-based Scene Extraction*. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1066–1076, Denver, Colorado. Association for Computational Linguistics.
- Eric Holgate and Katrin Erk. 2021. “politeness, you simpleton!” retorted [MASK]: Masked prediction of literary characters. In *Proceedings of the 14th International Conference on Computational Semantics (IWCS)*, pages 202–211, Groningen, The Netherlands (online). Association for Computational Linguistics.
- Md Naimul Hoque, Bhavya Ghai, Kari Kraus, and Niklas Elmqvist. 2023. *Portrayal: Leveraging nlp and visualization for analyzing fictional characters*. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*, DIS ’23, page 74–94, New York, NY, USA. Association for Computing Machinery.
- 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*.
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. *Camels in a changing climate: Enhancing lm adaptation with tulu 2*. Preprint, arXiv:2311.10702.
- Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, and Hal Daumé III. 2016. *Feuding families and former Friends: Unsupervised learning for dynamic fictional relationships*. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1534–1544, San Diego, California. Association for Computational Linguistics.
- Jennie Jarvis. 2014. *Crafting the character ARC*. Beating Windward Press.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. *Mistral 7b*. Preprint, arXiv:2310.06825.

- Hannah Kim, Denys Katerenchuk, Daniel Billet, Hae-sun Park, and Boyang Albert Li. 2018. [Learning Joint Gaussian Representations for Movies, Actors, and Literary Characters](#). *ArXiv*.
- Barbara Kyle. 2016. *Page-Turner: Your Path to Writing a Novel That Publishers Want and Readers Buy*. Rosethorn Books.
- Yingxuan Li, Ryota Hinami, Kiyoharu Aizawa, and Yusuke Matsui. 2024. [Zero-Shot Character Identification and Speaker Prediction in Comics via Iterative Multimodal Fusion](#). Publisher: arXiv Version Number: 2.
- 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, New Orleans, Louisiana. Association for Computational Linguistics.
- Eric T. Nalisnick and Henry S. Baird. 2013. [Character-to-character sentiment analysis in shakespeare’s plays](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 479–483, Sofia, Bulgaria. Association for Computational Linguistics.
- OpenAI. 2023. [GPT-4 Technical Report](#). Technical report, OpenAI.
- Xiangyu Peng, Kaige Xie, Amal Alabdulkarim, Harshith Kayam, Samihan Dani, and Mark Riedl. 2022. [Guiding neural story generation with reader models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7087–7111, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Revanth Rameshkumar and Peter Bailey. 2020. [Storytelling with Dialogue: A Critical Role Dungeons and Dragons Dataset](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5121–5134, Online. Association for Computational Linguistics.
- Hannah Rashkin, Asli Celikyilmaz, Yejin Choi, and Jianfeng Gao. 2020. [PlotMachines: Outline-Conditioned Generation with Dynamic Plot State Tracking](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4274–4295, Online. Association for Computational Linguistics.
- Yisi Sang, Xiangyang Mou, Mo Yu, Dakuo Wang, Jing Li, and Jeffrey Stanton. 2022. [MBTI Personality Prediction for Fictional Characters Using Movie Scripts](#). *arXiv*. Version Number: 1.
- Kurt Shuster, Jack Urbanek, Emily Dinan, Arthur Szlam, and J. Weston. 2020. [Deploying Lifelong Open-Domain Dialogue Learning](#). *ArXiv*.
- Sandeep Soni, Amanpreet Sihra, Elizabeth Evans, Matthew Wilkens, and David Bamman. 2023. [Grounding characters and places in narrative text](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11723–11736, Toronto, Canada. Association for Computational Linguistics.
- The Editors of Writer’s Digest Books. 2014. *Creating characters*. Writer’s Digest Books.
- Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, D. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, A. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kam-badur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#). *ArXiv*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- K M Weiland. 2016. *Creating character arcs*. Helping Writers Become Authors. Penforasword.
- Rui Xu, Xintao Wang, Jiangjie Chen, Siyu Yuan, Xinfeng Yuan, Jiaqing Liang, Zulong Chen, Xiaoqing Dong, and Yanghua Xiao. 2024. [Character is Destiny: Can Large Language Models Simulate Persona-Driven Decisions in Role-Playing?](#) Publisher: arXiv Version Number: 1.
- Kevin Yang, Dan Klein, Nanyun Peng, and Yuandong Tian. 2023. [DOC: Improving Long Story Coherence With Detailed Outline Control](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3378–3465, Toronto, Canada. Association for Computational Linguistics.
- Kevin Yang, Yuandong Tian, Nanyun Peng, and Dan Klein. 2022. [Re3: Generating longer stories with recursive reprompting and revision](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4393–4479, Abu

Dhabi, United Arab Emirates. Association for Computational Linguistics.

Anbang Ye, Christopher Cui, Taiwei Shi, and Mark O. Riedl. 2022. [Neural story planning](#). *Preprint*, arXiv:2212.08718.

Mo Yu, Jiangnan Li, Shunyu Yao, Wenjie Pang, Xiaochen Zhou, Zhou Xiao, Fandong Meng, and Jie Zhou. 2023. [Personality understanding of fictional characters during book reading](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14784–14802, Toronto, Canada. Association for Computational Linguistics.

Mo Yu, Qiuqing Wang, Shunchi Zhang, Yisi Sang, Kangsheng Pu, Zekai Wei, Han Wang, Liyan Xu, Jing Li, Yue Yu, and Jie Zhou. 2024. Few-shot character understanding in movies as an assessment to meta-learning of theory-of-mind. In *International Conference on Machine Learning*, pages 57703–57729. PMLR.

Weiwei Zhang, Jackie Chi Kit Cheung, and Joel Oren. 2019. [Generating character descriptions for automatic summarization of fiction](#). In *AAAI Conference on Artificial Intelligence*.

Andrew Zhu, Karmanya Aggarwal, Alexander Feng, Lara Martin, and Chris Callison-Burch. 2023. [FIREBALL: A Dataset of Dungeons and Dragons Actual-Play with Structured Game State Information](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4171–4193, Toronto, Canada. Association for Computational Linguistics.

A STORIUM Filtering

We filter stories for the following attributes: *playstyle* \in [‘writing’, ‘standard’], *speed* \neq ‘Hardcore’, *suspended* = **False**, *is_nsfw* = **False**, # *Words* > 1,000, *Avg. Words per Entry* > 30, *Language* = **English**. We find this filters for English stories in a standard narrative style, and removes potentially offensive content.

For the purposes of our annotations we only select specific entries where # *Words in Entry* > 50 to increase the likelihood of character-focused content, and shorter overall works where # *Words* < 10,000 to better allow for future story-long annotations. This last filter removes around 500 stories from consideration, which may be useful for future work. This gives us a dataset of 529 stories which we sample from for entailment dataset.

Entries are also further split to give us an average word count of 337; we make splits along paragraph lines to maintain comprehensibility as much as possible.

Finally, for the purposes of masked-character-prediction we also filter for *finished* stories (defined as the last scene having *is_ended*=**True** and *is_final*=**True**). We also filter for stories that contain a scene-entry with the three most common characters listed by name.

B Annotation Collection and Examples

Annotators were collected through Prolific, filtering for: living in the US or UK, and English as a Primary Language. All annotators were paid a rate of at least £9 an hour.

Table 8 contains a sample annotation for each value in our 1-5 scale, as well as a justification as explanation.

C Annotation Interface

Figure 2 (Guidelines & Consent), Figure 3 (Instructions), and Figure 4 (Example Annotation) provide screenshots of the annotation task we provided Prolific crowdworkers.

D Full Character Sheet Examples

Table 9 and Table 10 contain a full example character-sheet (over the course of an entire story) before and after the Validation Module.

E Automated Reasoning Steps

E.0.1 In-Context-Learning

Due to the relatively long nature of our snippets, ICL is not naturally applicable. However, we can extract from our overarching annotation task two sub-tasks that do not require the snippet:

Ambiguity: does this statement contain references to unspecified objects? Example: *He ran away from them*. Too much ambiguity may make a statement impossible to entail from the snippet.

Informativeness: does this statement provide any information about the character? Example: *There were no descriptions of Kelly*. Statements that do not make any claims do not entail from the snippet.

We generate responses starting with *Yes* or *No* followed by a short sentence justifying the answer. We find that naively using the boolean values as a filtering step for the classifier is too aggressive and has a low recall, largely due to many statements being slightly ambiguous. The justification is thus useful to further explain the rating so the downstream classifier can choose how heavily to consider it. Examples are given in Table 11.

Task Guidelines and Consent

Procedures: Thank you for taking part in our experiment! In this study, you will be presented with a story snippet and a statement about its characters. Your task is to judge whether the statement matches your impression of the character when reading the story. You will rate the statement for coherence and factuality (based on the story snippet) using a 1-5 scale. You will also provide a brief justification for your rating. No expert knowledge is required to perform this task.

Guidelines: Please read carefully the story snippet, and then consider the character statement. If the perspective of the snippet is unclear, assume it is from the perspective of the given character. Do not make broad inferences based on the snippet. For example, if the story says a character is tall, do not assume they play sports. Assume the character statement can be inferred from the story only when claims in the statement are directly stated in the snippet.

Voluntary Participation: This study is performed by researchers at the University of Edinburgh. If you have any questions, feel free to contact Alexander Gurung (A.Gurung-1 at sms.ed.ac.uk). Participation is voluntary. You have the right to withdraw from the study at any time. The collected data will be used for research purposes only. We will not collect any personal information. Your responses will be linked to your anonymous Prolific ID for the exclusive purpose of conducting our experiment.

If you give your consent to take part please click 'I agree' below.

I agree ☒
I disagree ☐

Move forward

Figure 2: Guidelines and consent for task shown to Prolific annotators.

E.0.2 Chain of Thought

We mimic chain of thought prompting by asking the model two short-answer questions that annotators implicitly answered while labeling. First, we ask for the section of the snippet most relevant to the statement. Second, we ask the model to compare and contrast the section of the snippet to the statement. We hypothesize that implicitly asking the model to both identify key phrases and compare them to the statement at the same time is too difficult, and find that guiding the model’s reasoning in this way improves performance. We provide examples in [Appendix F Table 12](#).

F Chain of Thought Example

[Table 12](#) contains two examples of our chain-of-thought reasoning steps, to illustrate how they might help determine the validity of a given statement.

G Classifier Training

We used Mistral 7B Instruct v0.2 ([Jiang et al., 2023](#)) as our base, and used Huggingface’s TRL library for all training. LoRA hyper-parameters: $r = 16, \alpha = 32, \text{dropout} = 0.1, \text{target-modules}=["\text{gate_proj}", "\text{down_proj}", "\text{up_proj}", "\text{q_proj}", "\text{v_proj}", "\text{k_proj}", "\text{o_proj}"]$. Other hyper-parameters: $\text{learning_rate} = 2e - 5, \text{batch-size}=1$ (with $\text{gradient-accumulation-steps}=4$), $\text{epochs}=4$ (and best model selected by validation precision). Further details will be shared with the code.

All experiments and training was performed on two NVIDIA A100 GPUs.

H Prompt Templates

[Table 14](#) shows the prompt template for our Generation Module, and [Table 15](#) shows the template for the simplification step that we pass the output from the initial generation through. We added this simplification step after collecting our annotated dataset, and found that it produces both easier to annotate statements and more accurate character sheets. Furthermore, by making each claim atomic we ensure our ‘density’ metric uses accurate estimations for the amount of character information.

[Table 16](#) and [Table 19](#) show the template for our ‘ambiguity’ reasoning in our Validation Module. [Table 17](#) and [Table 18](#) show the template for our ‘informative’ reasoning in our Validation Module.

[Table 20](#) shows the template for our Chain-Of-Thought reasoning step in our Validation Module.

Please read these examples to acquaint yourself with the task, and refer to these examples for guidance.

Instructions

You will be given a section of a story or screenplay from the perspective of Leonia Lowsnout (or others, depending on the snippet), and a statement about that character. Please answer the following questions concerning whether the statement is definitely true based on the story snippet.

If the perspective of the snippet is unclear, assume the story snippet is from the perspective of the given character. Do not make broad inferences based on the snippet. For example, if the story says a character is tall do not assume they play sports. The statement should only be considered factually consistent if all of its claims are directly stated in the snippet. If there are no claims made in the statement, mark the consistency of the statement as 1 as there is no evidence for the statement.

Story Snippet:

Quite marvelous, the craftsmanship. She'd received more than one offer during her inquiries to take it off her forepaws -- in gold and in blood. Those forepaws still trembled at the memory of the latter offers.

Breathing another sigh that shivered her whiskers, the charr glanced about the courtyard for the hundredth time since her arrival. Five more minutes, Leonia told herself, and she would stop thinking about what to say and... just say it. Assuming Brandis would grant an audience, of course. The notion of being turned away by the scholar at this juncture was almost too much for her to even contemplate.

Character: Leonia Lowsnout

Statement: We have learned that Leonia speaks with sighs and shivers her whiskers when she's nervous or in a bad mood.

Coherence: Is the statement fluent and understood, and lacking in grammatical/spelling errors?

Incoherent ☐ ☐ ☐ ☐ ☒ Coherent
1 2 3 4 5

Factual Consistency: Does the story snippet provide evidence for the statement?

There is no evidence for this statement ☐ ☐ ☐ ☐ ☒ There is evidence for all claims in this statement.
1 2 3 4 5

(1-2 Sentences) Justify your above rating, referencing the story snippet as evidence where applicable.

Leonia is in a nervous/bad mood in the snippet, and both sighs and shivers her whiskers. Thus, every claim the statement makes is true based on the snippet.

Figure 3: Instructions for annotation task shown to Prolific annotators, including one of the provided examples.

Character	Snippet	Statement	Rating	Justification
Holly	"Of all the people here likely to have magic, your the least likely. I probably have more magic in my little finger than you on your whole body." Holly said this with confidence in her voice but she was really thinking about if there was someone who could just know someone else had magic. Taking a deep breath she sat next to the fire to try and dry off.	Based on the given story section, we can infer the following physical descriptions of Holly: [Llama-2 13b]	1	This statement makes no claims about Holly, and therefore cannot entail from the snippet.
Cherish Nicole	Cherish was startled by his comment and moved her book to her lap in a rush. "Uh, well, no I wasn't taught that. I'm sorry, I didn't know I was being rude.", she said quickly, not meeting Gregory's eyes. She didn't want to give potential friends the wrong impression, darn it. She makes a mental note to leave the textbooks in the room next time. "What are toastmasters?", she asks, completely confused.	She uses simple sentences and avoids complex vocabulary when she is unsure of a topic. [Mistral-7b]	2	Cherish does speak in simple sentences in this snippet while unsure of a topic, but there is no indication that she avoids complex vocabulary and instead seems to ask about terms she does not understand.
Dr F. Leiber	"I'd be obliged if you'd employ that blade on my restraints." said Dr Lieber as he pulled and strained at the bindings that prevented him from sitting fully upright. "Only, I beg you, leave my hands unharmed that I may continue my practice." Turning to look at Professor Doyle, Lieber asks "Are you able to barricade the door with that bed Doyle? Buy us some time."	He pleads for help in barricading the door to buy them some time. [Mistral-7b]	3	While it is true that Leiber asks Doyle to barricade the door to buy some time, it's unclear if Leiber <i>pleads for help</i> as they seem to ask calmly.
Kaye Madison	Kaye shifted her feet nervously. "That's one way of saying it, I suppose. It's not so much funny as it is uncommon, I guess." She readjusted her bag on her shoulder and took a sip of tea. "Still, it's weird to be all friendly and stuff with strangers, right? What if I was some kind of a murderer?"	Based on this story section, we can infer that Kaye speaks in a hesitant and self-deprecating manner. [Llama-2 13b]	4	This statement seems generally true based on the snippet. However, one small complication is if Kaya's comment about potentially being a murderer counts as self-deprecating. It is also unclear if Kaya always speaks in this manner.
Anastasia Howe	"GAH!" Anastasia screams as she wakes up to a snake curled up. She had about 1 second to take in the new scenery. She then jumped up and away from the snake. She can't fight! Yet again she can't work with animals. She jumped away, uncertain of what the snake might do to her. "Help." She yelped weakly.	Anastasia appears to be easily frightened, as evidenced by her scream and immediate jump to escape when she wakes up to a snake. [GPT-4]	5	The claim and its underlying reasoning all entirely entail from the snippet, and there are no ambiguities.

Table 8: Example annotations and their justifications. While the rest of the table comes directly from our dataset, the justifications are author-provided included for better comprehension.

Close Instructions

Instructions

You will be given a section of a story or screenplay from the perspective of Hana Praewa, and a statement about that character. Please answer the following questions concerning the coherence of the statement and whether the statement is definitely true based on the story snippet.

If the perspective of the snippet is unclear, assume the story snippet is from the perspective of the given character. Do not make broad inferences based on the snippet. For example, if the story says a character is tall do not assume they play sports. The statement should only be considered factually consistent if all of its claims are directly stated in the snippet. If there are no claims made in the statement, mark the consistency of the statement as 1 as there is no evidence for the statement.

Story Snippet:

"Really!!! Are you sure ? Oh my god!!! Thanks Goddess. Thank you very much, doctor really thank you very much. I don't know how much I can thank you, but thank you a lot. I'm so happy."She said to Dr.Lovve and pulled him up to shake his hand like a little kid, and then he nodded to her. Suddenly she jumped up and hugged him tightly without realizing what she was doing.

Character: Hana Praewa

Statement: We learn that Hana Praewa is an emotional and expressive character in her speech.

Coherence: Is the **statement** fluent and understandable, and lacking in grammatical/spelling errors?

Incoherent

☐
☐
☐
☐
☐

Coherent

1

2

3

4

5

Factual Consistency: Does the story snippet provide evidence for the statement?

There is no evidence for this statement

☐
☐
☐
☐
☐

There is evidence for all claims in this statement.

1

2

3

4

5

(1-2 Sentences) Justify your above rating, referencing the story snippet as evidence where applicable.

Figure 4: Example of annotation task shown to Prolific annotators.

Physical/Personality

What, if any, physical descriptions of this character are in this snippet?

Nadia is described as watching both men silently. She pulls her tattered, too-big sweatshirt tighter around herself.

Nadia is a woman with the ability to glance over her shoulder. She deflates slightly. She sees that no men have followed her. No other physical descriptions are provided in this snippet.

Nadia is described as pushing her way forward. She has companions who scoff. There is no physical description of her appearance given in this snippet.

Nadia is described as staring in horror. She bites her lip. She glances towards Ian. She shouts to the natives.

She pulls out a digital camera. She is shocked when the natives react positively to the camera. She is dragged down from the rooftop by a woman. There is no detailed physical description of Nadia in this snippet.

Nadia bit her lip. She glanced up at the men. She gestured in the direction of the cell tower.

What, if any, descriptions of this character's personality are in this snippet?

Nadia is quiet. She has not spoken since the plane crash. She is observant. She notices the approaching argument between Jacob and Jayson. She is cautious and concerned about their safety. She suggests they seek shelter in the cave during an impending storm. She is emotionally drained.

Nadia is angry. She expects men to follow her willingly. She is facing reality. She is introspective.

Nadia is assertive. Nadia is apologetic for her friends' behavior. Nadia is determined to obtain information. She offers him the opportunity to accompany them off the island. However, she is firm in her demand for full disclosure from him.

Nadia is resourceful. She is brave. She is willing to take risks, as shown when she calls out to the natives from the rooftop. She is also compassionate and caring. She is culturally curious and open-minded.

Nadia is determined and brave. She knows about the cell tower and offers this information to the men. Nadia is independent and strong-willed. She chooses to stay on the island instead of leaving with the men. She is aware of the potential danger and the consequences of Ian and his wife's actions.

Knowledge

What, if any, factual information is given about this character in this snippet?

Nadia is a survivor of a plane crash. She has been silent since the accident. She wears a tattered sweatshirt. She is emotionally drained. She suggests seeking shelter in a cave.

Nadia is the name of the character. She is the one who glances over her shoulder in fear. She is an author.

Nadia is a member of a group. She apologizes for her friends' behavior towards Ian. She offers him the opportunity to accompany them off the island in exchange for information. She is assertive. She takes charge of the situation. She has a backpack with her.

Nadia is a character who witnesses Ian falling off a rooftop. Nadia is confronted by hostile natives. She is resourceful and quick-thinking. She uses a digital camera to diffuse the situation. The natives react positively to the camera. They allow Nadia to be helped down from the rooftop. Nadia shows concern for Ian. Nadia apologizes when he is injured.

Nadia is stranded on an island. She hates two men, Jayson and Jacob. She knows of a cell tower nearby. She suggests they use it to contact someone for help. She decides to stay on the island. She is certain she is making the right choice.

What, if any, information has this character learned in this snippet?

Nadia has learned that she has valuable input in decision-making situations. She has observed the potential for conflict between Jacob and Jayson. She has suggested a course of action to avoid it. She is aware of the approaching storm.

Nadia has realized that the men she is with are not willing to follow her like the characters in her novels. She feels disappointed. Nadia has learned that Ian is in possession of valuable supplies. She learns that her friends are treating Ian unfairly.

Nadia learns that the natives of the island are fascinated by the camera. Using it as a tool, she gains their trust. She is able to safely descend from the rooftop.

Nadia learns that there is a cell tower nearby. Nadia realizes that the men, Jayson and Jacob, could use it to contact someone for help. She realizes that she is not leaving the island. She will stay there to live in the world of her books.

Table 9: Sample part (Physical/Personality and Knowledge) of character sheet concatenated across an entire story, all text is from the Generation Module and highlighted text is that which passed the Validation Module.

Dialogue

What, if anything, have we learned about how this character speaks from this snippet?

Nadia has a quiet and hesitant speaking voice. She clears her throat before speaking. She has gone without using her voice for an extended period.

Nadia thinks aloud.

Nadia speaks assertively. Nadia interrupts conversations when necessary. She apologizes for her friends' behavior. She stands firm on the need for information. She uses direct language. She makes clear the conditions for allowing Ian to accompany them.

Nadia speaks with urgency when demanding answers. She uses exclamations such as "Screw it." She communicates effectively with the natives.

Nadia speaks with determination. She shouts to get the men's attention. She communicates effectively about the cell tower's location. She is assertive in expressing her decision to stay on the island.

Plot

What, if any, goals does this character gain in this snippet that they wish to accomplish in the future?

Nadia's goal in this snippet is to seek shelter from the approaching storm.

Nadia desires to have people willingly follow her in reality. Just as they do in her novels.

Nadia's goal in this snippet is to obtain necessary information from Ian. She apologizes for her friends' behavior. She offers him the opportunity to accompany them. To achieve this, she convinces Ian to share his supplies.

Nadia's initial goal is to protect Ian. Nadia demands answers from the men who pushed him off the rooftop. However, when she encounters the natives, her goal shifts to finding a way to communicate and connect with them.

She discovers that the natives are fascinated by the camera. Her new goal is to use the camera to build a relationship with the natives.

Nadia's goal in this snippet is to help the men contact someone off the island. She stays behind.

What, if any, goals does this character complete in this snippet?

Nadia expresses her opinion to Jacob and Jayson. Nadia suggests seeking shelter in the cave during an approaching storm.

She initiates communication after a long period of silence following the plane crash.

Nadia glances over her shoulder. She is disappointed that neither man follows her. She thinks angrily about how things would be different in her novels. No goals are completed in this snippet for Nadia.

Nadia convinces Ian to share his supplies with her group.

Nadia manages to save herself from the natives. She uses a digital camera to distract them. She discovers that the natives are fascinated by the camera.

Nadia reveals the existence of a cell tower to Jayson and Jacob. Despite the danger, she decides to stay on the island.

How, if at all, does this character's internal motivations change in this snippet?

Nadia, who had been silent since the plane crash, feels the need to speak up. She speaks up when she sees the men arguing about entering the cave. She expresses her concern about the approaching storm. She suggests they seek shelter in the cave. Despite her emotional exhaustion, she takes the initiative.

Nadia's internal motivation shifts from expecting the men to follow her. She compares the situation to her novels. She implies a desire for excitement and adventure in her real life.

Nadia feels apologetic for her friends' behavior towards Ian. Nadia remains determined to obtain necessary information. She offers him the opportunity to join them until they leave the island. She insists on receiving all the information he possesses in return.

Nadia's initial motivation is to protect Ian. Nadia demands answers from the men who pushed him off the rooftop. However, she realizes that the natives are not hostile towards her but rather fascinated by the camera. Her motivation shifts to finding a way to communicate and connect with them. She uses the camera as a tool.

Nadia's internal motivation shifts from hatred and a desire to be rid of the men. She develops a desire to stay on the island and live in the world of her books. She makes this decision despite the danger and potential consequences.

Table 10: Sample part cont. (Dialogue and Plot) of character sheet concatenated across an entire story, all text is from the Generation Module and highlighted text is that which passed the Validation Module.

Type	Statement	Reasoning
Ambiguity	However, before he can carry out his plan, she is shot from behind and collapses into his arms.	Yes, the statement is ambiguous as it does not specify who "she" is.
Ambiguity	Skótos, also known as Scott Mavrogiannis, is a demi-god.	No, the statement is unambiguous in stating that Skótos is a demi-god.
Informative	She is also adaptable, following William into the dark room without complaint.	Yes, the statement describes Sharon as adaptable
Informative	No other physical descriptions are provided in this snippet.	No, the statement does not provide any new physical descriptions of Jameson Jackson.

Table 11: Example ICL reasoning statements and the corresponding automatically generated responses. A response of ‘Yes’ for Ambiguity questions implies the statement is too ambiguous to be understood, and likely should not be included in a character sheet. A response of ‘No’ for Informative questions implies the statement does not provide any new information, and likely should not be included in a character sheet.

Statement	Question	Answer
She intends to distract any potential threats and lead Ymros to safety through the fire door.	What, if any, section of the story snippet is most relevant to the given statement? Provide a brief 1-2 sentence description of this section or "N/A" if there is no relevant section. In 1-2 sentences, compare the claim the statement makes and the section of story you highlighted in your previous answer. Are there any notable differences? Are all claims made by the statement explicitly supported? If there are no claims, write "N/A".	The section of the story snippet that is most relevant to the given statement is when Rachel suggests going outside through the fire door to distract potential threats and lead Ymros to safety. The statement aligns with the section of the story where Rachel plans to use the fire door as a means to distract threats and lead Ymros to safety. Both the statement and the story suggest that Rachel intends to use this strategy to ensure Ymros' safety. No notable differences exist between the two. All claims made by the statement are explicitly supported by the story.
Bryn's speech is straightforward and direct.	What, if any, section of the story snippet is most relevant to the given statement? Provide a brief 1-2 sentence description of this section or "N/A" if there is no relevant section. In 1-2 sentences, compare the claim the statement makes and the section of story you highlighted in your previous answer. Are there any notable differences? Are all claims made by the statement explicitly supported? If there are no claims, write "N/A".	Bryn's reaction to Sam's interaction with Mona, specifically when she punches Sam and tears up the card, is the most relevant section to the given statement. In this moment, Bryn's words and actions are straightforward and direct as she expresses her jealousy and frustration towards Mona and Sam. The statement that "Bryn's speech is straightforward and direct" is supported by the section of the story where Bryn punches Sam and tears up the card in response to Sam's interaction with Mona. This action demonstrates Bryn's straightforward and direct emotional response to the situation.

Table 12: Example Chain of Thought Reasoning

Name	ICL	CoT	Eval-Setup	Precision	Recall	Accuracy
Zero-Shot All	✓	✓	≥ 4	0.747	1.0	76%
Zero-Shot ICL	✓	✗	≥ 4	0.759	0.887	72%
Zero-Shot CoT	✗	✓	≥ 4	0.747	1.0	76%
Zero-Shot None	✗	✗	≥ 4	0.793	0.972	80%
Trained All	✓	✓	≥ 4	0.868	0.93	85%
Trained ICL	✓	✗	≥ 4	0.868	0.831	79%
Trained CoT	✗	✓	≥ 4	0.843	0.986	86%
Trained None	✗	✗	≥ 4	0.873	0.972	88%
Zero-Shot All	✓	✓	$= 5$	0.753	0.986	76%
Zero-Shot ICL	✓	✗	$= 5$	0.759	0.887	72%
Zero-Shot CoT	✗	✓	$= 5$	0.753	0.986	76%
Zero-Shot None	✗	✗	$= 5$	0.835	0.930	82%
Trained All	✓	✓	$= 5$	0.930	0.746	78%
Trained ICL	✓	✗	$= 5$	0.952	0.563	67%
Trained CoT	✗	✓	$= 5$	0.864	0.803	77%
Trained None	✗	✗	$= 5$	0.895	0.718	74%

Table 13: Classifier performance by training data and evaluation type. Zero-Shot refers to untrained Mistral 7B Instruct v0.2. (≥ 4 v. ≥ 4) refers to converting **both** model predictions and dataset labels from 1-5 to booleans via ≥ 4 . In contrast, ($= 5$ v. ≥ 4) refers to converting dataset labels the same way but only taking model predictions as *True* if the prediction = 5. We generally value higher precision, but the best performing precision model (Trained ICL) has much worse accuracy and recall. We take the reasonable performance of Zero-Shot models to indicate the quality of our prompts, but note from the high recall+low precision that this setup over-predicts positive labels. All reported numbers are from our test set (out of an 80-10-10 split).

You are a helpful and expert writing assistant. You will be given a section of a story or screenplay. Please answer the following questions about the character learned in this story section, and respond in short paragraph form.

Story Section:

[story_section]

Please answer the following questions about [character] with short, succinct sentences based on the given story section.

Question: [question] Respond in paragraph form with short, simple sentences with no dependent clauses or transition words.

Answer:

Table 14: Prompt for generating initial responses from out-of-the-box LLMs. Highlighted [text] is filled in with the appropriate information prior to generation. Prompts were adjusted to use instruction tokens via their respective model’s chat-template.

You are an expert writing assistant helping an author split compound sentences. Please answer the following questions to the best of your ability.

Given the provided sentence, please split all independent clauses into independent sentences and resolve any issues with unclear pronouns or references. Only do this for compound sentences. Every new sentence should make sense on its own. Write them out in paragraph form, one sentence after another.

Sentence: She's curious about a closed door in Maxim's apartment and feels a strong urge to discover what's behind it.

Split Sentences: She's curious about a closed door in Maxim's apartment. She feels a strong urge to discover what's behind the closed door in Maxim's apartment.

Sentence: Kaluros is determined and focused during battles, using his magic and weapons effectively to defeat his enemies.

Split Sentences: Kaluros is determined and focused during battles, using his magic and weapons effectively to defeat his enemies.

Sentence: Hassan encountered a crab monster and engaged in a card battle to defeat it.

Split Sentences: Hassan encountered a crab monster. Hassan engaged in a card battle to defeat the crab monster.

Sentence: She uses imperatives to give orders and asks direct questions to gather information.

Split Sentences: She uses imperatives to give orders. She asks direct questions to gather information.

Sentence: Bob is easily distracted and forgets about the chase when he notices something outside.

Split Sentences: Bob is easily distracted. Bob forgets about the chase when he notices something outside.

Sentence: Rachel enters the warehouse to join the baby dragon, defying her initial skepticism.

Split Sentences: Rachel enters the warehouse to join the baby dragon, defying her initial skepticism.

Sentence: He gives commands to his companions and asks for their assistance.

Split Sentences: He gives commands to his companions. He asks for his companions' assistance.

Sentence: She explores the Zombear's massive body and climbs on it.

Split Sentences: She explores the Zombear's massive body. She climbs on the Zombear.

Sentence: Jordan opens the locker to find a locket, a newspaper, and a mysterious photograph.

Split Sentences: Jordan opens the locker to find a locket, a newspaper, and a mysterious photograph.

Sentence: He is quiet and tosses a gold idol between his hands while they wait for rescue.

Split Sentences: He is quiet. He tosses a gold idol between his hands while they wait for rescue.

Sentence: [sentence]

Split Sentences:

Table 15: Prompt (ICL-style) for simplifying responses from out-of-the-box LLMs. We use this to break apart complex claims and make the resulting statements easier to verify in our Verification Module. Highlighted [text] is filled in with the appropriate information prior to generation. Prompts were adjusted to use instruction tokens via their respective model's chat-template.

You are a helpful and expert writing assistant. Please answer the following questions to the best of your ability.

Question: [question_text]

Statement: No other physical descriptions of Collins are provided in this snippet.

Answer: No, the statement is unambiguous in its meaning as there are no claims made.

Question: [question_text]

Statement: He has a strange glowing key that he uses to open the door to his home.

Answer: No, the statement is unambiguous in its description of Mustafa's ownership of key and his opening of a door. The pronouns 'he' and 'his' unambiguously refer to Mustafa.

Question: [question_text]

Statement: These men run away at first sight.

Answer: Yes, the statement is ambiguous because it doesn't specify who "These men" are, and it is also unclear what "first sight" refers to.

Question: [question_text]

Statement: They are skittish and afraid of the darkness around their camp.

Answer: No, the statement is unambiguous about their (Alex's) fear of the darkness.

Question: [question_text]

Statement: They struggle to stand up.

Answer: No, the statement is unambiguous about Luis's difficulty standing up.

Question: [question_text]

Statement: Arjun's primary goal is to resolve this situation.

Answer: Yes, the statement is ambiguous because we cannot understand Arjun's goal without knowing "this situation".

Question: [question_text]

Statement: She speaks softly to them, hoping to calm them down.

Answer: Yes, the statement is ambiguous because it is unclear who "them" refers to, which is necessary to understand the statement.

Question: [question_text]

Statement: She speaks using colloquial expressions like "C'est la vie" and "TGIF, am I right?"

Answer: No, the statement is unambiguous in its description of Ping's speaking habits.

Question: [question_text]

Statement: He is skinny with long legs.

Answer: No, the statement is unambiguous in its description of Santiago.

Question: [question_text]

Statement: He opens a brown bottle and drinks the murky liquid inside.

Answer: No the statement is unambiguous in its description of Jarvis drinking from the brown bottle.

Question: [question_text]

Statement: Her keen eyes spot Arthur sneaking the coin out from underneath the cup, showing her skills in observation.

Answer: No the statement is unambiguous in its description of the Bartender's observational skills.

Question: [question_text]

Statement: He aims to accomplish this by using fire as a means to drive it away, based on his knowledge of myths and lore from various cultures.

Answer: Yes, the statement is ambiguous as it is unclear what he is trying to accomplish or who/what he is driving away.

Question: [question_text]

Statement: Overall, Jake seems to be focused on his social life and maintaining his reputation.

Answer: No, the statement is unambiguous in its description of Jake's focus.

Question: [question_text]

Statement: His accomplishes this goal quickly.

Answer: Yes, the statement is ambiguous as it is unclear what "this goal" refers to.

[...]

Table 16: ICL-prompt: ambiguity. We use this in our Verification Module to determine if a given statement is too ambiguous to verify or do entailment with. Highlighted [text] is filled in with the appropriate information prior to generation. The [question_text] in this case refers to: *"Is the given statement about [character] ambiguous in a way that makes the meaning unclear? Ambiguities may include, but are not limited to, references to unspecified characters, objects, and actions. If the statement begins with a personal pronoun (e.g. "He" or "She"), assume it refers to [character] and don't count the pronoun towards the ambiguity. Begin your 1-2 sentence response with "Yes" if the statement is too ambiguous to be understood on its own and "No" if the statement makes is unambiguous in its meaning."*. Prompts were adjusted to use instruction tokens via their respective model's chat-template.. See Table 19 for the rest of the prompt.

You are a helpful and expert writing assistant. Please answer the following questions to the best of your ability.

Question: [question_text]
Statement: Collins learned that his father died after the factory accident in 1973.
Answer: Yes, the statement provides information about what Collins has learned about his father's death.

Question: [question_text]
Statement: There were no descriptions of Kelly.
Answer: No, the statement does not give you any information about Kelly.

Question: [question_text]
Statement: She gestures wildly, and speaks with flowery descriptions.
Answer: Yes, the statement provides information on how Ani speaks.

Question: [question_text]
Statement: They have a new goal of finding a place to eat.
Answer: Yes, the statement is gives us information about Dr. Alex's new goal.

Question: [question_text]
Statement: This snippet provides many physical descriptions of Rockefeller.
Answer: No, the statement does not provide any information about Rockefeller.

Question: [question_text]
Statement: She is an extremely focused and driven individual.
Answer: Yes, the statement is describes the Bartender's personality.

Question: [question_text]
Statement: Jennifer aims to accomplish this by using fire as a means to drive it away, based on her knowledge of myths and lore from various cultures.
Answer: No, the statement does not make any claims about Arjun as it only makes claims about Jennifer.

Question: [question_text]
Statement: She expresses a desire to look into the matter further and gather more information, indicating that she is motivated to uncover the truth about the space/time anomaly and its potential impact on the city's population.
Answer: Yes, the statement describes her desire to investigate.

Question: [question_text]
Statement: He is skinny with long legs.
Answer: Yes, the statement describes his physical appearance.

Question: [question_text]
Statement: According to this section of the story, Merlin's primary goal is to live to see the night.
Answer: Yes, the statement describes describes Merlin's goal to survive.

Question: [question_text]
Statement: His internal motivations do not significantly change in this snippet.
Answer: No, the statement does not provide any new information about Mohammed.

Question: [question_text]
Statement: She recently acquired a gun while searching the trunks abandoned in the cave.
Answer: Yes, the statement tells the reader a fact about her recent activity and a new item she possesses.

Question: [question_text]
Statement: Here are some physical descriptions of Pedro based on the given story section:
Answer: No, the statement does not provide any information about Pedro.

Question: [question_text]
Statement: They are a tall warrior.
Answer: Yes, the statement describes their height and profession.

[...]

Table 17: ICL-prompt: informative. We use this in our Verification Module to determine if a given statement makes any claims or tells us any information about a character. Highlighted [text] is filled in with the appropriate information prior to generation. The [question_text] in this case refers to: “Does this statement give you any novel information concerning [character] or what [character] knows? Novel information may include, but is not limited to, physical descriptions, new information they may have learned, goals they have or actions they have just completed, and descriptions of their speech. Begin your response with "Yes" if the statement gives us any new information and "No" if the statement doesn't add to our knowledge/understanding of the character in any way.”. Prompts were adjusted to use instruction tokens via their respective model's chat-template.. See Table 19 for the rest of the prompt.

[...]

Question: [question_text]

Statement: Here are some descriptions of how Ava speaks:

Answer: No, the statement does not provide any information about Ava, although it implies that following statements will contain descriptions of how they speak.

Question: [question_text]

Statement: However, we learned that they speak brashly and with profanity, with little regard for others.

Answer: Yes, the statement describes their speaking style.

Question: [question_text]

Statement: He carries a med-kit and a set of knives in his backpack, just in case the werewolves show up.

Answer: Yes, the statement gives us new information about what Emmanuel the Destroyer has in his bag.

Question: [question_text]

Statement: A wide grin slowly grows across their face.

Answer: Yes, the statement describes Abul grinning.

Question: [question_text]

Statement: Based on the story section provided, we learn that he has a mysterious book of magic.

Answer: Yes, the statement gives us the information that Dmitry has a book of magic.

Question: [question_text]

Statement: He does not handle stress very well.

Answer: Yes, the statement describes his inability to handle stress.

Question: [question_text]

Statement: [statement]

Answer:

Table 18: ICL-prompt: informative (cont). See [Table 17](#) for more details.

[...]

Question: [question_text]

Statement: Based on this story snippet, they have a mysterious notebook.

Answer: No, the statement is unambiguous in describing what Fadekemi has.

Question: [question_text]

Statement: Daniel seems very familiar with the inner workings of the ship.

Answer: No, the statement is unambiguous in its description of Daniel's knowledge of the ship.

Question: [question_text]

Statement: Based on this story snippet, he greatly values the guidance of others, as evidenced by him asking for Sunita's advice.

Answer: No, the statement is unambiguous and describes Merlin's appreciation of other people's advice.

Question: [question_text]

Statement: They try very hard to speak with the accent of the High Aristocracy, but sometimes slip into less-idolized accent of his hometown.

Answer: No, the statement is unambiguous and describes Abdul's accent.

Question: [question_text]

Statement: Based on the story section provided, we can infer that he speaks with confidence and assertiveness.

Answer: No, the statement is unambiguous in its description of Dmitry's speaking pattern.

Question: [question_text]

Statement: He is in love with Cassian.

Answer: No, the statement is unambiguous in who Hassan loves.

Question: [question_text]

Statement: Furthermore, they are easily scared by the screams coming from the movie theatre.

Answer: No, the statement is unambiguous in describing Rockefeller being scared by the screams.

Question: [question_text]

Statement: [statement]

Answer:

Table 19: ICL-prompt: ambiguity (cont). See [Table 16](#) for more details.

You are a helpful and expert writing assistant. You will be given a section of a story or screenplay from the perspective of [character] . Please answer the following questions about the given statements and their relationship with the snippet provided.

Story Section:
[story_section]

Please answer the following questions about [character] by comparing the provided statement with the story section above:

Statement: [statement]

Question: What, if any, section of the story snippet is most relevant to the given statement? Provide a brief 1-2 sentence description of this section or "N/A" if there is no relevant section.

Answer: [answer]

Question: In 1-2 sentences, compare the claim the statement makes and the section of story you highlighted in your previous answer. Are there any notable differences? Are all claims made by the statement explicitly supported? If there are no claims, write "N/A".

Answer:

Table 20: Chain-Of-Thought prompt. We use this in our Verification Module to break down the reasoning steps necessary to compare a statement about a character to a story snippet. Highlighted [text] is filled in with the appropriate information prior to generation, except for [answer] which is first generated with only the first question before being filled in here for the second. Prompts were adjusted to use instruction tokens via their respective model’s chat-template.