

Telling Stories through Multi-User Dialogue by Modeling Character Relations

Wai Man Si

Prithviraj Ammanabrolu

Mark O. Riedl

School of Interactive Computing

Georgia Institute of Technology

{wssi, raj.ammanabrolu, riedl}@gatech.edu

Abstract

This paper explores *character-driven story continuation*, in which the story emerges through characters' first- and second-person narration as well as dialogue—requiring models to select language that is consistent with a character's persona and their relationships with other characters while following and advancing the story. We hypothesize that a multi-task model that trains on character dialogue plus character relationship information improves transformer-based story continuation. To this end, we extend the Critical Role Dungeons and Dragons Dataset (Rameshkumar and Bailey, 2020)—consisting of dialogue transcripts of people collaboratively telling a story while playing the role-playing game Dungeons and Dragons—with automatically extracted relationships between each pair of interacting characters as well as their personas. A series of ablations lend evidence to our hypothesis, showing that our multi-task model using character relationships improves story continuation accuracy over strong baselines.

1 Introduction

Automated storytelling can be thought of as creative, long-form text generation and understanding—requiring explicit long-term memory, consistency, and creativity among other pre-requisites. Most modern (neural) automated storytellers are *plot-driven* and frame the task in terms of sequentially generating plot points that narrate the story in third-person (Kiros et al., 2015; Mostafazadeh et al., 2016; Martin et al., 2018; Fan et al., 2018). This approach does not generally place much weight on individual characters or their interactions—information known to be critical for creating stories (Riedl and Young, 2010).

We are inspired by the idea of *character-driven and emergent storytelling* wherein narrative emerges through characters' interactions as seen

| | |
|------------------|---|
| Relations | { Scanlan, neutral, Vexahlia }, { Keyleth, positive, Scanlan }, { Grog, negative, Vexahlia }, { Scanlan, positive, Vaxildan } ... |
| Summary | They wake up in the morning, preparing for the coming battle. Scanlan turns them all into Ravenites with light clothing. The sleet storm is starting. ... |
| Vexahlia: | Bundle up! |
| Scanlan: | Okay. How will we know when it's time for me to release? We have to wait for Tooma to go report. |
| Vexahlia: | Is Vorugal back? He's back. |
| Scanlan: | I assume. |
| Vexahlia: | Do we see Larkin around? |
| DM: | No, you do not see Larkin around. |
| Scanlan: | Vax , do you want to go look? |
| Vaxildan: | For Larkin? No Larkin. <i>I attempt to see see if Tooma is coming.</i> I don't want to release this thing before Tooma is there reporting to Vorugal. |
| Scanlan: | |
| Vaxildan: | (Grog voice) Six. It said six. |

Table 1: A sample from CRD3 extended, showing: pairwise character relationships; historical context via the summary; and current character interactions in the form of dialogue, *first-person* (green), and *second-person* (blue) narration. DM refers to the Dungeon Master who provides arbitration and additional context to players.

in Table 1. In addition to the challenges faced by automated storytellers, a character-driven storytelling system must produce language while simultaneously: (1) keeping each character's personas consistent while acting; (2) keeping track of relationships between characters that will affect their interactions; and (3) follow and logically advance the plot of the story.

To better explore how to give automated systems these two abilities, we focus on the task of *story continuation* solely through dialogue—i.e. picking the next character response that best continues a story. The task and data are seen in Table 1. We build off the Critical Role Dungeons and Drag-

ons Dataset or CRD3 (Rameshkumar and Bailey, 2020), a unique dataset that contains dialogue transcripts of a small group of around six players role-playing various characters while playing the table top role-playing game Dungeons and Dragons—their adventures and interactions forming a narrative that stretches hundreds of chapters, with each chapter forming a subplot. The original dataset was intended to be used for abstractive summarization and contains ground-truth summaries for each chapter. To better suit our purpose of studying *character-driven storytelling*, we automatically augment the dataset with information regarding character persona as well as relationship types between pairs of characters (friends, enemies, etc.) by clustering crowdsourced descriptions of character interactions from the Critical Role Wiki.¹

This extended dataset lets us break down the problem of *character-driven story continuation* into two sub-tasks corresponding to the three challenges mentioned earlier in terms of interacting within the confines of a story while staying consistent with respect to character personas and relationships. We show that training a system to optimize for both of these sub-tasks significantly improves story continuation accuracy.

Our work’s two primary contributions are thus: (1) the extension to CRD3 enabling a study of *character-driven storytelling* and the corresponding methodology used; and (2) a multi-task learning system that leverages character relation and persona information to better complete stories.

2 Related Work and Background

Storytelling. Storytelling systems that use symbolic planning (Lebowitz, 1987; Gervás et al., 2005; Porteous and Cavazza, 2009; Riedl and Young, 2010; Ware and Young, 2011) focused on ensuring coherence and consistency of plot through explicitly listed rules in the form of pre- and post-conditions, often requiring extensive knowledge engineering. Modern neural language-model based approaches generally attempt to learn to tell *plot-driven stories* from a corpus of stories via learning objectives that optimize reconstructing the story itself (Kiros et al., 2015; Roemmele and Gordon, 2018; Khalifa et al., 2017; Fan et al., 2018). In particular, a two-step process in which the high level plot is first generated, followed by filling in rest of

¹https://criticalrole.fandom.com/wiki/Critical_Role_Wiki

the story constrained to the plot has emerged (Martin et al., 2017, 2018; Ammanabrolu et al., 2020; Tambwekar et al., 2019; Yao et al., 2019; Ippolito et al., 2019). Ammanabrolu et al. (2021) look at plot generation from a *character-driven* perspective using commonsense knowledge, though do not model character interactions at all. Closely related to the spirit of our task is the Story Cloze test (Mostafazadeh et al., 2016), which measures the ability of a model to correctly predict the end of a story. Like the other works mentioned here, however, this task does not require dialogue or other forms of character interactions.

Dialogue. Contemporary neural dialogue retrieval systems, both chit-chat and goal-oriented, more explicitly model agent interactions than most storytelling systems (Henderson et al., 2014; El Asri et al., 2017). Particularly relevant to our work are dialogue systems that attempt to model and stay consistent with an agent’s persona, such as Persona Chat (Zhang et al., 2018), or using further contextual information such as setting in addition to character personas using a crowd-sourced fantasy text-game such as LIGHT (Urbanek et al., 2019). None of these works, however, have any notion of story or plot, often using significantly less long-term context than most storytelling systems.

3 Character-Driven Storytelling

This section first describes the automated extensions to the CRD3 dataset, specifically information on character relationships, followed by the multi-task learning setup and transformer architecture that leverage the new data for story continuation.

| | Train | Valid | Test |
|-------------------------------|-------|-------|-------|
| Avg. no. of turns in a chunk | 38.37 | 61.17 | 62.18 |
| Avg. no. of char.s in a chunk | 4.06 | 4.07 | 4.36 |
| No. of chunks | 11400 | 815 | 761 |

Table 2: CRD3 (extended) dataset statistics.

3.1 CRD3 Automated Dataset Extension

CRD3, as originally seen in Rameshkumar and Bailey (2020), contains two seasons of 159 transcribed Critical Role episodes, consisting of 398,682 turns in total. It further contains 34,243 ground truth human-written summary dialogue chunks that abstractively summarize dialogue chunks. The chunks themselves consist of a sequence of dialogue and first- and second-person narration turns

that form a semantically cohesive unit—with the end of a chunk signifying the completion of a subplot or change in location. Table 2 provides statistics for the number of chunks in the train, as well as the average number of character turns and number of characters within a chunk.

To enable a more effective study of *character-driven storytelling* using this dataset, we automatically extend CRD3 by adding descriptions of character relations from the Critical Role Wiki. These descriptions are free form text and often summarize character emotions during their interactions with another character. To condense them down, we cluster the character relation descriptions in an unsupervised fashion by calculating the vectorized TF-IDF representation of the description and applying the K-means algorithm. Varying the number of clusters changes the qualitative information conveyed by the cluster. For example, if we set the number of clusters to three, we can then also use the popular sentiment analysis tool VADER (Hutto and Gilbert, 2014) to provide human interperable relationship labels for each of the three clusters—positive, negative, or neutral as seen in Table 1. We specifically focus on incorporating these 3 relation types into our models. These relationship labels are attached to every dialogue chunk based on the characters appearing in that chunk. Further information regarding clusters is found in Appendix A.3.

3.2 Multi-task Learning

Based on the hypothesis that modeling character interaction information is critical for our overall task of *character-driven storytelling*, our system optimizes for two sub-tasks: next character prediction and story continuation. The next character prediction task can be summarized as: given current context, predict the next character who will act or speak—providing a proxy for judging who is most likely to respond to the current character in a multi-character setting. Similarly, the story continuation task refers to predicting the next character response that continues the story given the same context. The context itself contains information regarding: (1) a summary of the story so far using the dialogue summary chunks provided in CRD3 and described in Section 3.1, (2) pairwise relationship cluster labels between all characters within the dialogue chunk, and (3) the last n -turns of character interactions.

Our model’s architecture is shown in Figure 1.

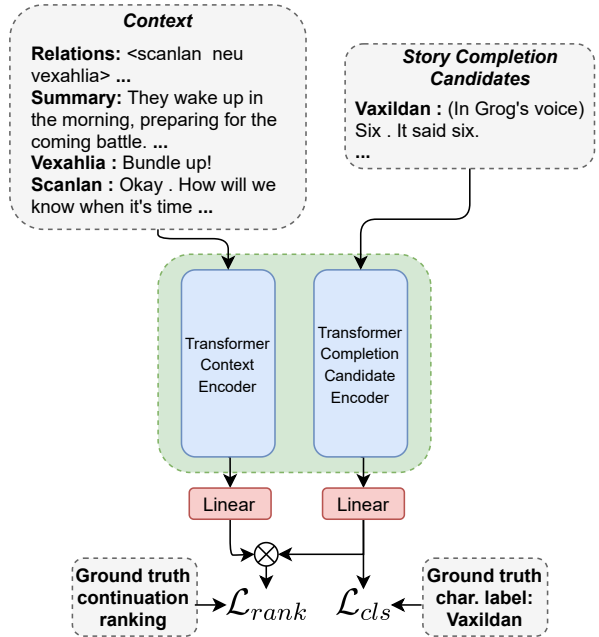


Figure 1: Multi-task learning overall architecture. The red shaded linear layers are task-specific and blue transformer blocks are pre-trained. Both transformer blocks share parameters across tasks.

It is inspired by the bi-encoder featured in Urbanek et al. (2019). In this model, two separate transformers are used to produce vector representations for the input context and each candidate utterance for the response retrieval task. All candidates are scored by via dot product between their vector representations and the context representation and trained using a ranking loss \mathcal{L}_{rank} . For the task of next character prediction, we use the same vector representation for the context and pass it through an additional linear layer with softmax layer to predict the correct character from the list of all possible characters. This sub-task uses a cross entropy loss \mathcal{L}_{cls} . The entire system is trained jointly by optimizing $\mathcal{L} = \lambda_1 \mathcal{L}_{rank} + \lambda_2 \mathcal{L}_{cls}$ for some hyperparameters λ_i . By virtue of the architecture, network parameters are shared between the tasks.

4 Evaluation

We conduct two ablations studies that analyze: (1) the complexity of performing *character-driven story continuation* on the dataset; and (2) the effectiveness of imbuing the model with relation information via input context and multi-task learning.

Our base transformer model that we build off of in each of these is the bi-encoder ranker described in Section 3.2. The transformer encoder is a similar architecture as BERT (Devlin et al., 2019), with

| Eval Task | Character Prediction | | Story Continuation | | | |
|------------------------|----------------------|-------------|--------------------|-------------|-------------|-------------|
| Metric | Weighted F1 | | Hits@1/10 | | Hits@5/10 | |
| Training Task Type | Single | Multi | Single | Multi | Single | Multi |
| Base | 47.3 | 47.6 | 18.0 | 18.3 | 70.6 | 73.9 |
| Base+Summary | 48.4 | 49.0 | 18.0 | 20.4 | 71.7 | 74.3 |
| Base+Relations | 49.0 | 48.8 | 17.6 | 20.2 | 70.6 | 74.0 |
| Base+Summary+Relations | 48.8 | 48.8 | 18.0 | 21.3 | 72.9 | 74.6 |

Table 3: Multi-task ablations.

| Eval Task | Char. Pred. | Story Continuation |
|-----------|-------------|--------------------|
| Metric | Weighted F1 | Hits@1/10 |
| 1 | 24.2 | 17.0 |
| 2 | 42.6 | 18.8 |
| 5 | 47.2 | 18.2 |
| 10 | 47.6 | 20.5 |

Table 4: Historical context ablations.

256 million total parameters, and is pre-trained using the Reddit dataset extracted and made available on pushshift.io (Baumgartner et al., 2020) seen in Roller et al. (2020). This dataset has been shown to result in an improved understanding of conversational natural language (Yang et al., 2018; Mazaré et al., 2018). Further hyperparameter and training details are shown in Appendix A.2.

For story continuation report standard retrieval metrics of Hits@ N , where we measure the ability of the model to output the gold standard dialogue candidate in the top- N of the given candidates. For character prediction, we report F1 weighted by the number of instances of each character type.

4.1 Historical Context Ablations

The first set of ablations measures performance on each of the two sub-tasks as a function of historical context required in an attempt to assess the complexity of the CRD3 extended dataset and its suitability for exploring *character-driven storytelling*. Recall that the CRD3 dataset provides summaries for each separate dialogue chunk. In Table 4, we vary the number of prior chunks of such summaries used as input context to the model and measure performance on each of the sub-tasks after training the model jointly on both sub-tasks.

The trends shown in Table 4 are quite clear—indicating that, overall, the CRD3 dataset requires very long contexts to ensure effective performance. On average, across both evaluation tasks performance gain between using a single historical context chunk and using two is greater than the corresponding differences when using even more chunks. Additionally, performance continues to rise with

added historical context up to the maximum context length we tested of 10. We note that this is a significantly greater amount of context than generally required for state-of-the-art chit-chat dialogue datasets (Roller et al., 2020) as well as prior story completion datasets such as ROC Stories (Mostafazadeh et al., 2016), reinforcing our hypothesis that the CRD3 dataset is well suited to enabling *character-driven storytelling* by focusing on interactions requiring long-term memory.

4.2 Multi-task Ablations

These ablations focus on analyzing the effects of our methods to imbue the agent with relationship and character information, specifically including the relationship cluster labels in the input and multi-task training. Table 3 outlines these results when evaluated on both the character prediction and story continuation sub-tasks with different: (1) inputs types—with base referring to only character interactions and additional information as seen in Figure 1; and (2) training methods—single referring to training on only the evaluation task and multi to jointly training on both tasks.

We would first like to note that we use the same relationship labels for characters through the entire story—i.e. across all the dialogue chunks. Our approach intuitively averages the relationship type between characters through time—e.g. characters that are friends at first and then become enemies will have a neutral label throughout all the story. While more fine grained relationship labels that do not perform such averaging might perform better, they would also require extensive additional human annotations to track relationships through time.

For character prediction, the Base+Summary multi-task and Base+Relations single-task models perform best though are closely comparable to the Base+Summary+Relations multi-task model. For story continuation, the Base+Summary+Relations multi-task model outperforms all others. In all story continuation experiments, multi-task trained models outperform their counterpart single-task trained

model. Through these results, we can infer that imbuing character relationship information through *both* input relationship cluster information as well as next character prediction helps models continue stories—while staying consistent with a particular character’s persona—more accurately.

5 Conclusions

We hypothesized that injecting models with information on relationships between characters would improve their ability to complete *character-driven stories*. A series of ablation studies support this, with a key insight being that a particularly efficient way of giving story continuation models this information is by multi-task training them on both character dialogue and relationship information automatically extracted from online sources.

6 Broader Impacts

Our work on *character-driven storytelling* has potential implications extending to the creation of learning agents that communicate using natural language, especially those requiring an agent to stay consistent with a character or persona throughout an interaction. As our system is trained entirely using a dataset collected from character interactions of a set of players role-playing in a fantasy Dungeons and Dragons world, we are prone to echoing biases found in the data. Some of these biases are necessary for effective story continuation, enabling a reader to identify the genre and conveying thematic information. Others may potentially involve non-normative language usage—acceptable in a fantasy world but inappropriate in the real world. Restricting our system to *story continuation* through a retrieval mechanism as opposed to generating text mitigates, though do not eliminate some of these biases. We urge future researchers and application developers that use automated storytelling techniques to similarly clarify the origins and methodology behind the creation of delivered story content.

References

- Prithviraj Ammanabrolu, Wesley Cheung, William Broniec, and Mark O Riedl. 2021. Automated storytelling via causal, commonsense plot ordering. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- Prithviraj Ammanabrolu, Ethan Tien, Wesley Cheung, Zhaochen Luo, William Ma, Lara J Martin, and Mark O Riedl. 2020. Story realization: Expanding plot events into sentences. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Jason Baumgartner, Savvas Zannettou, Brian Keegan, Megan Squire, and Jeremy Blackburn. 2020. The pushshift reddit dataset. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 830–839.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. Frames: a corpus for adding memory to goal-oriented dialogue systems. In *Proceedings of the 18th Annual SIGDial Meeting on Discourse and Dialogue*, pages 207–219, Saarbrücken, Germany. 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*, pages 889–898.
- Pablo Gervás, Belén Díaz-Agudo, Federico Peinado, and Raquel Hervás. 2005. Story plot generation based on CBR. *Knowledge-Based Systems*, 18(4-5):235–242.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014. The second dialog state tracking challenge. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 263–272.
- C. Hutto and E. Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *ICWSM*.
- Daphne Ippolito, David Grangier, Chris Callison-Burch, and Douglas Eck. 2019. **Unsupervised hierarchical story infilling**. In *Proceedings of the First Workshop on Narrative Understanding*, pages 37–43, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ahmed Khalifa, Gabriella A. B. Barros, and Julian Togelius. 2017. **DeepTingle**. In *International Conference on Computational Creativity*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.

- Michael Lebowitz. 1987. Planning Stories. In *Proceedings of the 9th Annual Conference of the Cognitive Science Society*, pages 234–242.
- Lara J. Martin, Prithviraj Ammanabrolu, Xinyu Wang, William Hancock, Shruti Singh, Brent Harrison, and Mark O. Riedl. 2018. Event Representations for Automated Story Generation with Deep Neural Nets. In *Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 868–875, New Orleans, Louisiana.
- Lara J. Martin, Prithviraj Ammanabrolu, Xinyu Wang, Shruti Singh, Brent Harrison, Murtaza Dhuliawala, Pradyumna Tambwekar, Animesh Mehta, Richa Arora, Nathan Dass, Chris Purdy, and Mark O. Riedl. 2017. **Improvisational Storytelling Agents**. In *Workshop on Machine Learning for Creativity and Design (NeurIPS 2017)*, Long Beach, CA.
- Pierre-Emmanuel Mazaré, Samuel Humeau, Martin Raison, and Antoine Bordes. 2018. **Training millions of personalized dialogue agents**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2775–2779, Brussels, Belgium. Association for Computational Linguistics.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. **A Corpus and Evaluation Framework for Deeper Understanding of Commonsense Stories**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849.
- Julie Porteous and Marc Cavazza. 2009. **Controlling narrative generation with planning trajectories: The role of constraints**. In *Joint International Conference on Interactive Digital Storytelling*, volume 5915 LNCS, pages 234–245. Springer.
- 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.
- Mark O Riedl and R Michael Young. 2010. **Narrative Planning: Balancing Plot and Character**. *Journal of Artificial Intelligence Research*, 39:217–267.
- Melissa Roemmele and Andrew S Gordon. 2018. **An Encoder-decoder Approach to Predicting Causal Relations in Stories**. In *Proceedings of the First Workshop on Storytelling*, pages 50–59, New Orleans, Louisiana. Association for Computational Linguistics.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M Smith, et al. 2020. **Recipes for building an open-domain chatbot**. *arXiv preprint arXiv:2004.13637*.
- Pradyumna Tambwekar, Murtaza Dhuliawala, Lara J. Martin, Animesh Mehta, Brent Harrison, and Mark O. Riedl. 2019. **Controllable Neural Story Plot Generation via Reward Shaping**. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*.
- Jack Urbanek, Angela Fan, Siddharth Karamcheti, Saachi Jain, Samuel Humeau, Emily Dinan, Tim Rocktäschel, Douwe Kiela, Arthur Szlam, and Jason Weston. 2019. Learning to speak and act in a fantasy text adventure game. In *EMNLP*.
- Stephen Ware and R. Michael Young. 2011. **Cpocl: A narrative planner supporting conflict**. In *Proceedings of the 7th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.
- Yinfei Yang, Steve Yuan, Daniel Cer, Sheng-Yi Kong, Noah Constant, Petr Pilar, Heming Ge, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. Learning semantic textual similarity from conversations. *arXiv preprint arXiv:1804.07754*.
- Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. **Plan-And-Write: Towards Better Automatic Storytelling**. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. **Personalizing dialogue agents: I have a dog, do you have pets too?** In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

A Appendices

A.1 CRD3 Extended Examples

| | |
|------------------|--|
| Relations | <Scanlan, neutral, Vexahlia>, <Grog, neutral, Scanlan>, ... |
| Summary | Scanlan deceives the clasp leader with a blue gem that can grant one wish if they say the password while holding the gem. He gives the leader the gem and promises to give him the password if they can visit riskel. The leader reveals the clasp helped riskel prepare for his escape. ... |
| Keyleth: | okay ! |
| DM: | He looks over at the gentleman who inspected it earlier and nods his head. "accepted." and they continue walking forward. |
| Grog: | Lucky fucking druid. |
| DM: | It is the piece you put in the actual– |
| Scanlan: | It's a blue shard that we found in–long, long ago– it's real crystal and it's real magic. |
| DM: | Yes. I know what that is. |
| Scanlan: | Because I don't. |
| DM: | Well, it was sufficient upon inspection for this. |
| Scanlan: | Okay. |
| Vexahlia: | Whoa, I think it opens a portal to another plane. |
| Scanlan: | I don't know what it is, but it's magic. |

Table 5: Randomly selected CRD3 extended examples

| | |
|------------------|--|
| Relations | <Grog, neutral, Vexahlia>, <Keyleth, positive, Scanlan>, ... |
| Summary | Rejoining the party, Vex wonders aloud why desmond is still in the cell. Percy responds that it was originally for his own protection, but that since the problem has been taken care of, it is a precaution that is no longer needed. ... |
| Vexahlia: | Are there days of the week? what is a weekend? |
| Keyleth: | Yeah, There's days of the week . |
| Scanlan: | What is this world? How does time work here? |
| DM: | There are days of the week, I'm not gon na go into the specifics of it because I'm working on it. This question hasn't really arisen before and I probably should figure that out. It 's the equivalent of a thursday. |
| Scanlan: | It's always thursday. |

Table 6: Randomly selected CRD3 extended examples