

Ontologically Faithful Generation of Non-Player Character Dialogues

Anonymous ACL submission

Abstract

We introduce a language generation task grounded in a popular video game environment. KNUDGE (KNnowledge Constrained User-NPC Dialogue GEneration) requires models to produce trees of dialogue between video game characters that accurately reflect quest and entity specifications stated in natural language. KNUDGE is constructed from side quest dialogues drawn directly from game data of Obsidian Entertainment’s *The Outer Worlds*, leading to real-world complexities in generation: (1) dialogues are branching trees as opposed to linear chains of utterances; (2) utterances must remain faithful to the game lore—character personas, backstories, and entity relationships; and (3) a dialogue must accurately reveal new quest details to the human player. We report results for a set of neural generation models using supervised and in-context learning techniques; we find competent performance but room for future work addressing the challenges of creating realistic, game-quality dialogues.

1 Introduction

Player interactions with non-player characters (NPCs) in role-playing games (RPGs) often serve to flesh out backstories while allowing the player to progress through engaging quest storylines (Onuczko et al., 2007). A key challenge in authoring these NPC dialogues is maintaining enjoyable as well as *coherent* narratives: utterances must faithfully reflect quest structure and game lore—characters, histories, and entity relationships.

Figure 1 depicts a dialogue turn taken from *The Outer Worlds* (Obsidian Entertainment, 2019),¹ an action RPG renowned for its narrative and dialogue writing. This turn demonstrates how a dialogue relies on the descriptions of entities in the game

** KNUDGE will be publicly available at <anonymized>.

¹https://en.wikipedia.org/wiki/The_Outer_Worlds

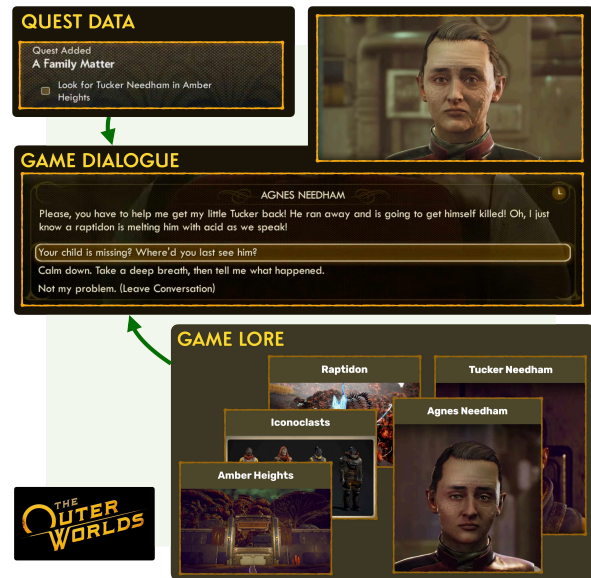


Figure 1: An example non-player character (NPC) dialogue from *The Outer Worlds* (Obsidian Entertainment, 2019). NPCs must speak faithfully to a granular ontology of **quest specifications** and **game lore**.

world while also revealing relevant quest information. The NPC’s utterance not only begins and provides backstory about a new side quest,² but also interacts according to a well-formed persona (a worried parent) and references an adversary (raptidons) that the player will face later in the quest.

NPC interactions often take the form of complex utterance trees. Creating these branching structures according to the many specifications of dialogue writing can be time-consuming for a game designer (Caropreso et al., 2012), and costs millions of dollars (see §A for discussion). This motivates the pursuit of *automatically generating* dialogue trees.

To the best of our knowledge, there does not exist a public dataset that meets our desired criteria: **a set of real game-quality NPC dialogues paired with granular quest and biographical specifications consistent with a well-formed game ontology**. Re-

²RPG *side quests* are goals provided to a player in order to enrich gameplay, while not being part of the central storyline.

cent work towards lore-conditioned dialogue (Urbanek et al., 2019; van Stegeren and Mysliwicz, 2021), story generation (Akoury et al., 2020; Chen and Gimpel, 2021), and knowledge conditioning for task-oriented dialogue agents (Choi et al., 2018; Mazaré et al., 2018; Feng et al., 2020) do not address complex dialogue trees and the interweaving narratives found in deployed RPGs.

We make the following contributions via this paper: firstly, we introduce **KNUDGE: Knowledge constrained User-NPC Dialogue Generation**, a set of dialogue trees (in English) derived from an existing video game and paired with granular ontological constraints. We extract the trees directly from the game data of *The Outer Worlds*. This game’s side quests share overlapping characters and locations, making it an appealing study in the development of automatic dialogue generation tools. For each side quest, we enumerate the relevant dialogues and annotate them at multiple levels of specificity (down to the individual node) with quest- and lore-related natural language (NL) support facts pulled from fan-written wikis. **KNUDGE** contains 159 dialogues from 45 side quests that all take place in *The Outer Worlds*. It contains 4.7K utterances and an average of 1.3K input constraint tokens per dialogue. The complexity of the dialogues, the annotation level, and the amount of specifications are greater than any comparable dataset.

Secondly, we introduce the challenging task of knowledge-constrained NPC dialogue generation. Given a set of ontological specifications, we want to generate fluent dialogue trees that reveal the quest objectives while staying faithful to the specified context and the game lore. These specifications make up hundreds of statements per dialogue, posing a difficult challenge for modern generation models. Additionally, we introduce a class of model, termed **DialogueWriters**, that leverage neural language models (LMs) such as GPT-3 (Brown et al., 2020) to generate an utterance conditioned on an existing partial dialogue tree and ontology passages. To *encourage* the use of game lore to produce interesting and engaging dialogue, we experiment with adding mechanisms to sub-select relevant facts before generating each utterance.

Finally, we prescribe protocols for the evaluation of models for the NPC dialogue generation task. We test for models’ capacity to reflect game ontology constraints in addition to generating fluent and coherent dialogue. We conduct automatic

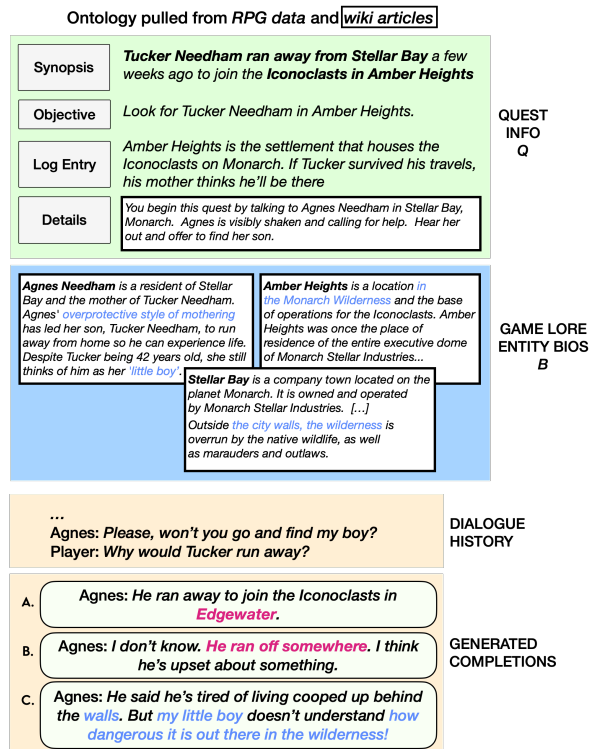


Figure 2: Overview of the proposed task. Quest information and biographical passages about the game lore serve as constraints on generated dialogue candidates. Completion **A** is inconsistent with the lore and **B** is uninformative; **C** is most desirable, as it provides new information to the player about quest objectives and reflects information about relevant entities.

and human evaluations of utterances and multi-turn trees generated from specifications for quests from the game, as well as for a pair of never-before-seen quests written by a professional game designer to target real quest writing. Our experiments reveal further room for improvement in aspects such as joint reasoning over multiple facts from ontology, and coverage of all quest objectives. We hope that **KNUDGE** will facilitate further progress on faithful game dialogue generation, towards tools that can assist developers with their craft.

2 Task Definition

We define the task of knowledge-constrained NPC dialogue generation as the mapping from a set of quest constraint statements Q , a set of biographical constraint statements B , and a list of participant names P to a dialogue tree D . We consider two task scenarios: **next utterance prediction** from a partial (gold) tree, and **full dialogue generation** of trees with multiple candidates at each turn.

As depicted in Figure 2 (upper), Q comprises statements $[q_1, \dots, q_m]$ about currently active ob-

jectives upon entering the dialogue, about what should occur during it (e.g., pieces of information the NPC should mention), and about new active objectives upon leaving it. B comprises background statements $[b_1, \dots, b_n]$ about game entities that the dialogue must reflect (Figure 2, middle). The participant list P contains the player character and one or more NPCs who have corresponding facts in B . We design these inputs to reflect the kinds of specifications that a developer would provide to the generator during their quest writing process.

Dialogue tree D is a directed graph $\langle N, E \rangle$; each utterance node $n \in N$ has a speaker $s \in P$. Branches occur due to the multiple dialogue options at player turns (see Figure 3, right). Dialogue trees have one start node, but can have multiple exit nodes and can contain cycles.³⁴

3 Data

KNUDGE comprises NPCs dialogue trees from all 45 side quests in the *Outer Worlds* base game. Our dataset construction procedure entails gathering information about each quest (Q) and the *Outer Worlds* entities that appear or are referenced during the associated dialogues (B) (§3.1), and then extracting trees (D) from the game data semi-automatically (§3.2). We then analyze the resulting data and compare to related datasets (§3.3).

3.1 Game Ontology

We acquired dialogue files from the *Outer Worlds* creators along with permission to release them publicly; we use quest data and game lore from fan wikis, where a quest’s page lists the in-game objectives and journal logs (though the framework allows for using data from official channels instead of fan-made). §B contains further source details.

Quest Information A quest in *The Outer Worlds* appears in the player’s journal with a high-level **synopsis** and a sequence of **objectives**, each of which contains **game logs** providing additional details. Active objectives are often completed, and new ones are introduced, during an NPC dialogue. A detailed quest anatomy can be found in appendix Figure 6. We associate with each objective a **walkthrough passage** which includes details on the topics, player utterance options, and quest information

³Although they often contain cycles these structures are still colloquially referred to as “trees”.

⁴In *The Outer Worlds*, utterance nodes are spoken once and are not repeated upon second traversal, unless they are repeat choices by the player at a given decision point.

that the NPC needs to say by the dialogue’s end. Thus, a dialogue’s quest passage set Q contains:

1. The **synopsis** (1-2 sentences)
2. The **in objective(s)** active when entering the dialogue (1 sentence), the associated **game log** (1-2) and **walkthrough passage** (3-10).
3. The **out objective(s)** active upon leaving the dialogue, and the associated **game log**.⁵

Examples of Q can be found in §C.

Biographical Information We associate with each quest, and in turn each dialogue, a set B of **biographical passages** about entities appearing or referenced during the quest. We extract passages from entities’ fan wiki pages. While some are only a few sentences, others can be much longer (up to 27), posing a challenge to generation models; often only part of a long biography might be relevant to a given quest. Examples are shown in §D.

3.2 Dialogue Trees

Dialogue trees in *The Outer Worlds* are complex directed graphs, containing many conditional utterance options depending on the state of the game—e.g. whether the player is of high enough level at some skill to pass a “check.”⁶ To extract a more tractable, quest-related subgraph, we 1) identified the nodes that start and end the interaction using online playthrough videos as reference, and then 2) traversed the graph from the start node, following only edges without special state-related conditions. Special edges were then added manually depending on whether conditions are relevant to the quest. Example trees can be found in §E.

Annotating Utterance Nodes with Support Facts

We coordinated with professional data annotators to annotate the nodes of the dialogue trees with support facts from the associated quest and lore constraints. All annotations were done by native English speakers, who were provided with a set of instructions shown in §I. We follow a heuristic based on counterfactuals: had a given fact *not* been included in the constraining knowledge, would the utterance be much less likely to occur? An example of this procedure is depicted in Figure 3.

⁵We do not associate its walkthrough passage, since the NPC should only be expected to convey new objective information that the player will actually see in game.

⁶E.g. if the Player has above 55 points of the **Persuade** conversational skill, they can convince Tucker Needham to return to his mother in the quest from Figure 1.

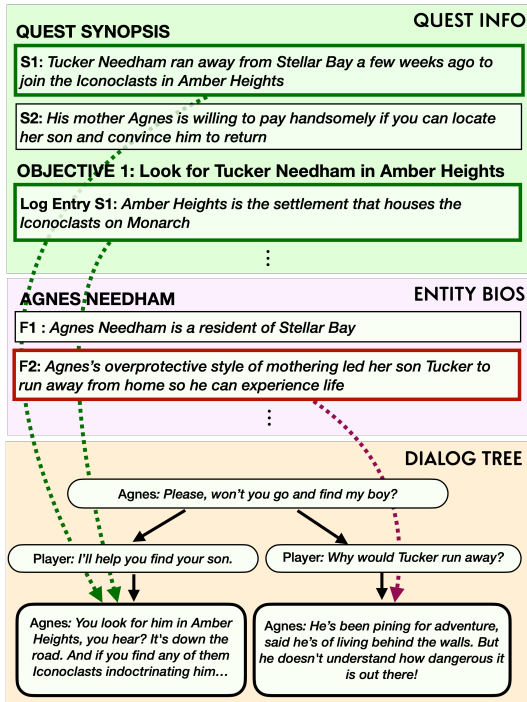


Figure 3: Overview of dialogue node annotation with support facts from quest and biography passages.

| | | | |
|----------------------|-----|--------------------------|--------|
| Quests | 45 | Dialogue trees | 159 |
| Entities | 168 | Facts per dialogue | 83.3 |
| Characters | 81 | Entity facts | 73.3 |
| Locations | 40 | Objective facts | 9.9 |
| Groups | 21 | Fact tokens per dialogue | 1321.8 |
| Items | 18 | Utterances per dialogue | 29 |
| Creatures | 7 | NPC utterances | 16.7 |
| Facts per entity | 7.4 | Player utterances | 12.3 |
| Entities per quest | 8.0 | Utt. tokens per dialogue | 406.6 |
| Quests per entity | 2.0 | Facts per NPC utterance | 1.0 |
| Facts per objective | 7.6 | Entity Facts | 0.57 |
| Objectives per quest | 4.5 | Objective Facts | 0.42 |

Table 1: KNUDGE Dataset Statistics

We compute an average exact match (EM) score of 0.52 and an average Jaccard overlap score of 0.62. These scores represent a high agreement on a subset selection problem where the total set size is often of the order of hundreds of elements.

3.3 Dataset Analysis

Table 1 describes overall statistics for the extracted dialogue trees and annotations. We note the high frequency of support fact annotations per utterance; 47% of all nodes (57% of NPC nodes) in the dataset are annotated with at least one fact. Amongst NPC nodes, for which support knowledge is more common, there is an average of 1.0 facts per utterance. Dialogues have on average 9.9 quest facts and 73.3 biographical facts that generation models must consider when producing utterances.

Comparison with Related Datasets KNUDGE is the first dataset to contain dialogue trees from an actual RPG annotated with game quest and biography specifications. Table 2 compares our dataset to contemporaries with comparable input specifications and generation targets. None contain all the components to target the complexity and granular specificity required to generate quest dialogues of the type found in *The Outer Worlds*. §F provides an in-depth comparison with select works.

3.4 Challenges

Writing RPG-quality dialogue trees is difficult for human developers for its many interweaving considerations. KNUDGE poses a similarly multi-faceted challenge to generation models:

- The tree must serve its **quest function**, containing input-specified player utterance options, NPC responses (including possibly specified emotions), and pieces of information the player must learn by the end. These NL specifications have highly diverse semantics (see §C).
- The generated utterances must be **coherent, realistic, and engaging** to the player.
- The NPC should embody the persona described in **their biography passage**, which describes personality, history, and relationships.
- To facilitate world building, the NPC should exposit details about **other entities** whenever it is contextually relevant, and should never **violate** the ontology through contradiction.

Generating a 30-node (or larger) NPC dialogue tree while taking into account all of these criteria at once is a very difficult task, particularly given the shape of the branching, cycle-heavy tree structure.

The average of 1321 constraining tokens and 406 utterance tokens poses a challenge to current NLP models, taking up e.g. half of the 4000-token context window of GPT-3 before factoring in other pieces of context such as few-shot examples.

4 DialogueWriter Methods

We introduce a set of neural methods for generating candidate utterances given the ontological specifications (Q , B , and P) from §2 and a partial subtree S for a dialogue item in KNUDGE. We refer to these methods collectively as DialogueWriter models, which propose utterance nodes at a specified new location branching off the subtree. Formally, given some “most recent” node $n \in S$, a

| Dataset | Writer | Source | Structure | Tokens / Item | Constraint Toks / Item | Narrative and Bio Constraints | Annotation Level |
|---|--------------|---------------------------|---|---------------|------------------------|--|------------------|
| Story Generation Datasets | | | | | | | |
| STORIUM (Akoury et al., 2020) | Crowd | Online story writing game | Sequence of scene entries | 19k | 1.2K | Scene intro, challenge, location, and character descriptions | Story |
| TVSTORYGEN (Chen and Gimpel, 2021) | Crowd | Fan wikis | TV episode recap | 1.8k | 25.9K | Brief episode summary and character wiki articles | Scene Entry |
| RPG Dialogue Datasets | | | | | | | |
| LIGHT (Urbanek et al., 2019) | Crowd | text game platform | Sequence of utterances | 212 (12 utt) | 276 | Location description, persona statements and held objects | Dialogue |
| TorchLight II (van Stegeren and Theune, 2020) | Professional | RPG Data | Sequence of quest stages with 0 or 1 utterances | 157 (3 utt) | 24 | Quest title, objectives, and details | Quest |
| WoW (van Stegeren and Mysliwiec, 2021) | Professional | RPG data | NPC-uttered quest description | 61 (1 utt) | 15 | Quest title and objective | Quest |
| KNUDGE (Ours) | Professional | RPG data | Complex quest dialogue tree | 407 (29 utt) | 1.3K | Quest title, objectives, and details, walkthrough and entity wiki articles | Utterance |

Table 2: Comparison of KNUDGE to knowledge-constrained generation datasets. KNUDGE contains professionally-written complex dialogue trees from an actual RPG, with more utterances, longer constraint passages, and higher-granularity annotations (of individual utterances) than other game dialogue datasets.

DialogueWriter maps inputs (Q, B, P, S, n) to a list of candidate utterances $[c_1, \dots, c_n]$ such that there is a directed edge $n \rightarrow c_i$.

Tree Linearization We consider language models (LMs) that accept linear input token sequences. We thus devise a traversal mechanism that converts a dialogue subtree into a maximal coverage linear history. For “most recent” node n , we identify the longest possible path from the start node to n , including cycles but only following a given edge once. This produces utterance history $H = [u_1, \dots, u_n]$. Exploring other tree encoding mechanisms, e.g. via graph encoders (Banerjee and Khapra, 2019; Ouyang et al., 2021) is left for future work. With this linearization mechanism, we need only train a next utterance generator that conditions on linear history H . Methods to do so are explored below.

4.1 Supervised Learning (SL) Models

We fine-tune a T5-large sequence-to-sequence model (Raffel et al., 2020) via supervised learning (SL) to generate c_i given the concatenation $[B, Q, P, H]$. We truncate context from the left of the when required given T5’s 1024-token window, removing components of B first. We list last (and thus truncate last) the biographies of dialogue participants. §G shows example inputs.

Supervised Knowledge Selection (KS) Model

We also train a version of the SL baseline that learns to decode support knowledge facts before conditionally generating the utterance c_i . This factorizes the next utterance generation into a two-step deci-

sion process: first selecting one or more facts from the provided knowledge constraints $(Q \cup B)$, and second generating the utterance to reflect the selected facts. We thus make use of our node-level annotations to train the model to generate the concatenated sequence $[f_1^{(i)}, \dots, f_m^{(i)}, c_i]$ s.t. $f_j^{(i)} \in Q \cup B$.

4.2 In-Context Learning (ICL) Models

As KNUDGE is relatively small, fine-tuning might not be effective at learning the difficult generation task. As such, we experiment with methods for *in-context learning* (ICL) with OpenAI’s *text-davinci-003* GPT-3 model (Brown et al., 2020). We inject B, Q, P , and H into a formatted prompt that naturally elicits the next utterance as a continuation of H . Figure 5 depicts this process; full prompts are shown in §G. This creates a *zero-shot* prompt. When this does not fill out GPT-3’s 4000-token window, we construct a *few-shot* prompt by adding dialogs from training quests as exemplars, simulating a scenario in which a developer has written a partial set of quests and is working on a new one. We retrieve exemplars using Okapi-BM25 (Jones et al., 2000) with $[B, Q, P]$ as the query string.

ICL Knowledge Selection (KS) Model

As with the SL framework, we also devise a version of the ICL DialogueWriter that first decodes one or more support facts before generating an utterance. We elicit this behavior from GPT-3 by augmenting *all* utterances in the dialogue history with support fact nodes, provided they exist (see Figure 15). Further model training details are found in §H.

| | Gold | | Bio B | | Quest Q | |
|-----------------------|------------|-------------|-------------|-------------|-------------|-------------|
| | BL | BS | BL | BS | BL | BS |
| $n = 1806$ | | | | | | |
| <i>Gold Reference</i> | – | – | 4.9 | 20.8 | 2.2 | 17.8 |
| ICL-KS | 7.1 | 25.1 | 8.3 | 24.3 | 7.5 | 24.0 |
| ICL-KS-One | 6.8 | 25.2 | 7.1 | 23.9 | 7.6 | 24.5 |
| ICL-KS-Oracle | 7.2 | 26.8 | 8.4 | 24.7 | 7.2 | 24.3 |
| ICL | 8.6 | 26.4 | 6.8 | 22.8 | 6.6 | 22.2 |
| ICL-Quest Only | 7.9 | 26.7 | 3.4 | 21.9 | 6.4 | 22.4 |
| Vanilla | 6.8 | 27.0 | 2.2 | 22.4 | 0.9 | 19.1 |
| SL-KS | 2.6 | 21.3 | 24.8 | 26.6 | 9.3 | 21.3 |
| SL | 2.9 | 23.5 | 7.4 | 24.0 | 11.4 | 23.7 |

Table 3: NUP BLEU and BertScore for models against gold utterances and statements in B and Q . Results for the latter two shown beside the gold utterance’s score.

| $n = 100$ | Coherence | Violation | Using B | Using Q |
|-----------------------|-------------|-------------|-------------|-------------|
| <i>Gold Reference</i> | 3.94 | 3.97 | 3.50 | 3.45 |
| ICL-KS | 3.78 | 3.85 | 3.29 | 3.45 |
| ICL-KS-One | 3.73 | 3.80 | 3.26 | 3.45 |
| ICL-KS-Oracle | 3.74 | 3.87 | 3.23 | 3.47 |
| ICL | 3.88 | 3.97 | 3.25 | 3.43 |
| ICL-Quest Only | 3.79 | 3.90 | 3.03 | 3.21 |
| ICL-Vanilla | 3.65 | 3.69 | 2.76 | 2.98 |
| SL-KS | 2.52 | 3.85 | 2.17 | 2.09 |
| SL | 2.70 | 3.74 | 2.35 | 2.38 |

Table 4: NUP human evaluation results for in-context (ICL) and supervised learning (SL) DialogWriter methods with and without knowledge selection (KS).

5 Experiments

5.1 Baseline Models

To measure the effect of conditioning on Q and B , we compare against ablations to the ICL model: a **vanilla** model that conditions only on the participants P and utterance history H , and a **quest only** model that conditions on P , H , and Q , but not B .

To measure the effect of node-level knowledge selection (**KS**), we compare against an ICL model that selects only **one** statement instead of many. We randomly sample gold facts to construct its prompt. We also compare against an **oracle** KS ICL model, which conditions on the full gold knowledge annotations for the reference utterance. We maximize the number of in-context examples for all ICL ablations; e.g. the **vanilla** model’s prompt can have dozens of such examples, as they are quite short. These ablations thus explore the tradeoff between the impact of the number of in-context examples and the presence of ontological statements.

5.2 Next Utterance Prediction (NUP)

We split KNUDGE into train, development, and test splits on the basis of **quests** (28/5/12), such that at test time all input components will be unseen (test set B ’s contain a combination of previously seen

and totally novel entities). We evaluate utterances generated from a test set of 1800 partial trees.

Automatic Evaluation We use reference-based metrics BLEU-4 (Papineni et al., 2002) and re-scaled BERTScore-F1 (Zhang et al., 2020). We evaluate generated next utterance against the following single- and multi-reference sets: 1) the gold utterance n_i , 2) the quest statements in Q , and 3) the biography statements in B . We evaluate the **gold** utterance to show a performance upper bound.

Human Evaluation In coordination with a data specialist, we conducted human evaluation to examine models’ qualitative NUP performance. Generated utterances over 100 test items were judged on a 4-point Likert scale for each of four criteria:

1. **Coherence:** does the utterance follow naturally from the utterances in the history?
2. **(Non-)Violation:** does the utterance create contradictions with any of the sentences in the biographical or quest passages?
3. **Biography Usage:** does the utterance *make use* of the input biographical passages in B ?
4. **Quest Usage:** does the utterance progress the dialogue according to the quest sentences in Q ?

We provide the full set of annotator instructions with guidance for each Likert scores in §I. Results were verified via bootstrap testing.

5.3 Full Dialogue Generation Case Study

We run a case study for a full tree generation scenario. We task models with generating 10 rounds of dialogue given just the specifications B , P , Q , and one starting utterance.⁷ At each turn, we generate three candidate nodes using the writer, then randomly “commit” one to the linear history. This creates a 31-node tree (example in Figure 16) that can serve as a ‘spine’ of proposed content to be fleshed out by a developer into a more complicated tree. We hand-selected a set of 8 test dialogues from the game with varying quest roles, e.g. starting, continuing, and ending quests. Additionally, we had a professional game designer write specifications for **2 totally novel quests** with overlapping entities from the original game; from these we constructed 8 more test items.

We generated trees using the KS, regular, quest only, and vanilla ICL DialogWriters, then

⁷This serves to specify utterance format in cases of zero-shot learning, which is the case for all designer-written items.

History: > Agnes Needham: Please, you have to help me get my little Tucker back! He ran away and is going to get himself killed! Oh, I just know a raptidon is melting him with acid as we speak! > Player: Your child is missing? Where'd you last see him? > Agnes Needham: Oh, I just knew you were a good person. 'Agnes,' I said, 'this is the man to save your little Tuckie!' And I was right!

Gold: Agnes Needham: He ran out into the wilderness a few days ago. I warned him about the raptidons, mantisaurs, and marauders - the toxic sulfur pools, and poisonous plants. But he didn't listen!

KS ICL: Agnes Needham: He was always such a good boy, but lately he's been acting so strange. I just don't understand why he would want to leave Stellar Bay. It's the only place he's ever known.

ICL: Agnes Needham: He left a few weeks ago. Said he was going to Amber Heights to join some revolutionary group. I just know he's going to get himself killed if he's not careful!

ICL-Quest Only: Agnes Needham: He left a few weeks ago, headed for Amber Heights. He's going to get himself killed, I just know it!

ICL-Vanilla: Agnes Needham: He was playing in the street. But then he saw a sweet little kitty and chased it into the sewers! Oh, the horror!

SL-KS: Agnes Needham: If you're willing to pay handsomely, I'll even pay you handsomely if you convince him to return to me.

SL: Player: I'll do everything I can.

Figure 4: [best viewed in color] Example generations by the in-context learning (ICL) DialogueWriter and supervised learning (SL) DialogueWriter. Models that condition on both biographical and quest passages are more likely to mention relevant quest information while also referencing game entities and their backgrounds without creating inconsistency, incoherence, or incongruity. See §C and §D for full documentation of quests and entities referenced.

showed them to data specialists for evaluation following the ACUTE-Eval (Li et al., 2019) protocol of pairwise comparison. Annotators were asked to select which of two trees were preferred for the following criteria: **coherence**, **nonviolation**, **biography** and **quest** usage analogous to §5.2, and additionally

5. **Content Suggestion:** Do the multiple candidates at each turn propose interesting subtrees?

6. **Engagingness:** does the tree hold your attention and make you want to hear more from the NPC?

6 Results and Discussion

Next utterance prediction results under automatic and human metrics are shown in Table 3 and Table 4. We note that automatic metrics for generation that check for lexical or semantic overlap with a set of references is not directly suited for evaluating generations in KNUDGE; this can be seen from the extremely low performance of *the gold utterances themselves* under these metrics. We find that professionally written utterances do not always have high overlap with knowledge statements themselves. Gold utterances also do not score perfectly under utterance-level human evaluation of Q and B usage, as not every real-world utterance reflects the ontology. It might instead provide other qualities like realism and fluency. This highlights the challenge of identifying and then pursuing desiderata of RPG-quality dialogue trees.

We observe that the best-performing model under overlap metrics with B and Q are the T5-based SL models. This reflects that these models have learned to copy spans directly from the context into the generation, hence scoring highly on ngram-based overlap with it while scoring poorly on gold utterance overlap and human evaluation.

Table 3 depicts the trend that KS variants of the ICL model score a point or two higher than non-KS on overlap with B and Q , reflecting that KS effectively selects and cues the infusion of specific facts into generations. Oracle KS improves BLEU score with B but not Q , while one-fact KS has the opposite effect. The ICL ablations of B and of $\{B, Q\}$ have according drops in overlap with both fact sets. The **Quest**-only ICL DialogueWriter can generate coherent utterances that reflect the quest specifications; however, as illustrated by Table 4 and Figure 4, this can sacrifice lore references.

Table 4 also shows that all fully B and Q -conditioned ICL models perform equivalently under all human metrics except (non)violation, which KS models perform a decimal point worse at. We can conclude that KS improves the capacity of ICL writers to directly reflect knowledge passages (i.e. by copying spans), at the expense of a slightly higher likelihood of contradictions. This characterization can be appealing to a game developer; they might prefer for the automatic writer to use their own provided wordings of various facts when generating candidates. We note that no model approaches gold performance at reflecting B .

Qualitative Analysis Figure 4 depicts example outputs by models on an NUP example. We highlight cases in which the models succeed at the desiderata that we strive for in KNUDGE: to convey quest and lore specifications naturally through the interaction. However, we see that SL models and ablated ICL models are less successful. We observe that the gold utterance is more infused with desirable information than any generation; it references the quest's next location and numerous adversaries that the player will run into, while effectively reflecting the NPC's overprotective parent persona. This highlights a performance gap between neural

| $n = 16$ | Coh. | Viol. | Use B | Use Q | Content | Engaging |
|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| ICL-KS | 50.0 | 68.8 | 68.8 | 75.0 | 43.8 | 37.5 |
| ICL | 81.2 | 75.0 | 81.2 | 75.0 | 75.0 | 75.0 |
| ICL-Quest | 56.2 | 56.2 | 43.8 | 50.0 | 50.0 | 68.8 |
| ICL-Vanilla | 12.5 | 0.0 | 6.2 | 0.0 | 31.2 | 18.8 |

Table 5: Pairwise comparison (% head-to-head wins) between generated trees from ICL DialogueWriters. For example, 81.2% of the time ICL outputs were preferred over a competing approach with respect to coherence.

and human writers to be addressed by future work. We note that not reaching human performance does not preclude DialogueWriters from being useful to writers, as they can still be used to suggest new directions for dialogues to be verified or modified in a human/AI collaborative writing process.

Full dialogue generation case study results are found in Table 5. Model performance is measured as the rate at which annotators selected its generated tree in a pairwise comparison under the 6 criteria listed in §5.3. We find that annotators preferred the trees of the ICL writer most frequently compared to the other models under all criteria except Q usage. We also find that the ICL-KS method is more frequently preferred for the never-before-seen quests written by the professional designer, while it is less frequently preferred for dialogues from the actual game. The opposite holds for ICL-Quest only. A possible explanation is that the designer provided many biographical specifications for his quests (enough to fill out the GPT3 context window), without which there might not be a sufficient signal for the Quest only model to generate desirable utterances. Another reason might be that the few-shot learning benefits of the Quest only model not conditioning on B are lessened by the domain shift from KNUDGE to the designer’s writing.

7 Related Work

Si et al. (2021) focus on the task of story continuation through dialogue between multiple characters while modeling the inter-character relations. However, such past work does not concern with the notion of grounding knowledge or quest objectives to be covered in the generated dialog. van Stegeren and Theune (2020) propose three sources for building NPC dialogue corpora. However, their proposed datasets do not contain any grounding annotation and are not accompanied by explicit descriptions of entities and characters. Callison-Burch et al. (2022) explore automatic generation of conversational turns by players of the tabletop RPG,

Dungeons and Dragons, in which NPCs serve a very different role in the gameplay. Scheherazade’s Tavern (Aljammaz et al., 2020) augments a pattern-matching-based NPC interaction system with facts the character knows about the game world. More broadly, past work has explored applications of text generation in various gaming applications such as quest description generation (van Stegeren and Mysliwiec, 2021), dialogue generation (Si et al., 2021), persona-specific agents in text environments (Urbanek et al., 2019), and new text world generation (Fan et al., 2020; Ammanabrolu et al., 2022).

Past work has pursued dialogue systems that steer the conversation towards a topic (Wu et al., 2019) or a given NL sentence (Sevegnani et al., 2021; Gupta et al., 2022) while conversing with a user. Other work in NLG has explored generating outputs with high-level NL specifications such as string item agendas (Kiddon et al., 2016), sets of facts (Orbach and Goldberg, 2020), or author goals (Riedl, 2009). KNUDGE also comprises NL specifications, though they are comparably richer.

8 Conclusion

Humans play games to be entertained, and they pay money expecting a high-quality experience. When a game requires dialogue to advance a carefully crafted storyline, this should be both engaging as well as consistent with the larger narrative. Language models are increasingly capable of producing engaging dialogues, and researchers have been exploring how to ensure that such content is consistent with underlying knowledge specifications. To date, this research has focused on scenarios developed for the sake of experimentation, rather than actual high-quality game data.

In this paper, we introduce KNUDGE, a dataset of NPC dialogue trees coupled with a relevant game ontology, drawn from the title *The Outer Worlds*. In contrast to prior work, KNUDGE is based on content created by a triple-A game (i.e. high-budget, high-profile game) development studio, thereby exemplifying real-world complexities in NPC dialogue authoring. We illustrate that LMs are indeed able to generate fluent dialogues that relate to provided game lore. However, the straightforward application of such technologies does not match the quality of professional game writers. We hope that KNUDGE aids the development and evaluation of new techniques for faithful game dialogue generation.

579 Limitations

580 One limitation of our work is that our considered
581 models use simplifications to model the complex
582 branching trees by linearizing the nodes. We more-
583 over only evaluate next-utterance prediction and
584 relatively simple end-to-end dialogue trees which
585 do not approach the complexity of the actual trees.
586 More can be done on this front in future work to
587 develop methods that generate more complicated
588 tree structures. We also leave for future work the
589 evaluation of whether end-to-end generated dia-
590 logues “check all the boxes” of quest requirements,
591 completing *all* user-provided specifications.

592 We also find that the proposed DialogueWriter
593 models leave room for improvement on persona em-
594 bodiment. Human-quality utterances more seam-
595 lessly and dynamically incorporate emotions fitting
596 of characters and situations, while model-generated
597 utterances can be comparatively bland. This work
598 also focuses on *side quests* whose NPCs are gener-
599 ally not as fleshed out as those in main quests. Gen-
600 erating quests containing major NPCs with long
601 bios and important roles in the main story of a
602 game, e.g. companion characters, is also left for
603 future work.

604 KNUDGE recasts a set of fan articles about an
605 existing game as specifications to an automatic di-
606 alogue tree writer. It therefore assumes that game
607 developer will write structured game lore and high-
608 level quest specifications in a similar manner **be-**
609 **forehand** when coming up with new content. Thus,
610 we provide only a partial solution to the problem
611 of automatic NPC dialogue generation, and future
612 work can look at copilot tools for authoring such
613 high-level quest specifications and design of new
614 characters.

615 We report results with large pre-trained language
616 models. It is difficult to know whether the game
617 data used for experimentation is part of the training
618 data for such models, as *The Outer Worlds* came
619 out in 2019. As such, the results from such large
620 language models should be interpreted with cau-
621 tion. We partially mitigate the issue by having
622 an expert game developer construct a totally new
623 quest specification, and report results on the this
624 previously ‘unseen’ test data.

625 Ethics Statement

626 We acknowledge that there maybe bias in the data
627 used to train the neural language models consid-
628 ered in this paper (T5 and GPT-3) that would lead

to NPC dialogues that are offensive, implicitly or
explicitly discriminatory. This poses a potential
risk for deployed models, as using the proposed
DialogueWriters as content suggestion tools might
lead to RPG content that reflects these biases. We
hope that professional game developers will have
the resources to moderate damaging content before
it makes its way into released products.

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 2020, Online, November 16-20, 2020*, pages 6470–6484. Association for Computational Linguistics.
- Rehaf Aljammaz, Elisabeth Oliver, Jim Whitehead, and Michael Mateas. 2020. [Scheherazade’s tavern: A prototype for deeper NPC interactions](#). In *FDG ’20: International Conference on the Foundations of Digital Games, Bugibba, Malta, September 15-18, 2020*, pages 22:1–22:9. ACM.
- Prithviraj Ammanabrolu, Renee Jia, and Mark O. Riedl. 2022. [Situated dialogue learning through procedural environment generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 8099–8116. Association for Computational Linguistics.
- Suman Banerjee and Mitesh M. Khapra. 2019. [Graph convolutional network with sequential attention for goal-oriented dialogue systems](#). *Transactions of the Association for Computational Linguistics*, 7:485–500.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, and Others. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Chris Callison-Burch, Gaurav Singh Tomar, Lara J. Martin, Daphne Ippolito, Suma Bailis, and David Reitter. 2022. [Dungeons and dragons as a dialog challenge for artificial intelligence](#). *CoRR*, abs/2210.07109.
- Maria Fernanda Caropreso, Diana Inkpen, Fazel Keshtkar, and Shahzad Khan. 2012. [Template authoring environment for the automatic generation of narrative content](#). *The Journal of Interactive Learning Research*, 23:227–249.
- Mingda Chen and Kevin Gimpel. 2021. [Tvstorygen: A dataset for generating stories with character descriptions](#). *arXiv preprint arXiv:2109.08833*.

| | | | |
|-----|---|--|-----|
| 682 | Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wenta | A crpg sub-quest generator. In <i>Proceedings of the</i> | 739 |
| 683 | tau Yih, Yejin Choi, Percy Liang, and Luke Zettle- | <i>AAAI Conference on Artificial Intelligence and Inter-</i> | 740 |
| 684 | moyer. 2018. Quac: Question answering in context . | <i>active Digital Entertainment</i> , volume 3, pages 110– | 741 |
| 685 | In <i>Proceedings of the 2018 Conference on Empirical</i> | 111. | 742 |
| 686 | <i>Methods in Natural Language Processing, Brus-</i> | | |
| 687 | <i>sels, Belgium, October 31 - November 4, 2018</i> , pages | Eyal Orbach and Yoav Goldberg. 2020. Facts2story: | 743 |
| 688 | 2174–2184. Association for Computational Linguis- | Controlling text generation by key facts . In <i>Proceed-</i> | 744 |
| 689 | tics. | <i>ings of the 28th International Conference on Com-</i> | 745 |
| | | <i>putational Linguistics, COLING 2020, Barcelona,</i> | 746 |
| 690 | Angela Fan, Jack Urbanek, Pratik Ringshia, Emily Di- | <i>Spain (Online), December 8-13, 2020</i> , pages 2329– | 747 |
| 691 | nan, Emma Qian, Siddharth Karamcheti, Shrimai | 2345. International Committee on Computational | 748 |
| 692 | Prabhumoye, Douwe Kiela, Tim Rocktäschel, Arthur | Linguistics. | 749 |
| 693 | Szlam, and Jason Weston. 2020. Generating inter- | | |
| 694 | active worlds with text . In <i>The Thirty-Fourth AAAI</i> | Siru Ouyang, Zhuosheng Zhang, and Hai Zhao. 2021. | 750 |
| 695 | <i>Conference on Artificial Intelligence, AAAI 2020, The</i> | Dialogue graph modeling for conversational machine | 751 |
| 696 | <i>Thirty-Second Innovative Applications of Artificial</i> | reading . In <i>Findings of the Association for Com-</i> | 752 |
| 697 | <i>Intelligence Conference, IAAI 2020, The Tenth AAAI</i> | <i>putational Linguistics: ACL-IJCNLP 2021</i> , pages | 753 |
| 698 | <i>Symposium on Educational Advances in Artificial In-</i> | 3158–3169. Online. Association for Computational | 754 |
| 699 | <i>telligence, EAAI 2020, New York, NY, USA, February</i> | Linguistics. | 755 |
| 700 | 7-12, 2020, pages 1693–1700. AAAI Press. | | |
| | | Kishore Papineni, Salim Roukos, Todd Ward, and Wei- | 756 |
| 701 | Song Feng, Hui Wan, R. Chulaka Gunasekara, | Jing Zhu. 2002. BLEU: a method for automatic eval- | 757 |
| 702 | Siva Sankalp Patel, Sachindra Joshi, and Luis A. | uation of machine translation . In <i>Proceedings of the</i> | 758 |
| 703 | Lastras. 2020. doc2dial: A goal-oriented document- | <i>40th annual meeting of the Association for Computa-</i> | 759 |
| 704 | grounded dialogue dataset . In <i>Proceedings of the</i> | <i>tional Linguistics ACL 2002</i> . | 760 |
| 705 | <i>2020 Conference on Empirical Methods in Natural</i> | | |
| 706 | <i>Language Processing, EMNLP 2020, Online, Novem-</i> | Colin Raffel, Noam Shazeer, Adam Roberts, Katherine | 761 |
| 707 | <i>ber 16-20, 2020</i> , pages 8118–8128. Association for | Lee, Sharan Narang, Michael Matena, Yanqi Zhou, | 762 |
| 708 | Computational Linguistics. | Wei Li, Peter J Liu, et al. 2020. Exploring the limits | 763 |
| | | of transfer learning with a unified text-to-text trans- | 764 |
| 709 | Prakhar Gupta, Harsh Jhamtani, and Jeffrey P. Bigham. | former. <i>J. Mach. Learn. Res.</i> , 21(140):1–67. | 765 |
| 710 | 2022. Target-guided dialogue response generation | | |
| 711 | using commonsense and data augmentation . In <i>Find-</i> | Mark O. Riedl. 2009. Incorporating authorial intent | 766 |
| 712 | <i>ings of the Association for Computational Linguistics:</i> | into generative narrative systems . In <i>Intelligent Nar-</i> | 767 |
| 713 | <i>NAACL 2022, Seattle, WA, United States, July 10-15,</i> | <i>rative Technologies II, Papers from the 2009 AAAI</i> | 768 |
| 714 | 2022, pages 1301–1317. Association for Computa- | <i>Spring Symposium, Technical Report SS-09-06, Stan-</i> | 769 |
| 715 | tional Linguistics. | <i>ford, California, USA, March 23-25, 2009</i> , pages | 770 |
| | | 91–94. AAAI. | 771 |
| 716 | Karen Sparck Jones, Steve Walker, and Stephen E. | Karin Sevegnani, David M. Howcroft, Ioannis Konstas, | 772 |
| 717 | Robertson. 2000. A probabilistic model of informa- | and Verena Rieser. 2021. Otters: One-turn topic tran- | 773 |
| 718 | tion retrieval: development and comparative experi- | sitions for open-domain dialogue . In <i>Proceedings</i> | 774 |
| 719 | ments - part 1 . <i>Inf. Process. Manag.</i> , 36(6):779–808. | <i>of the 59th Annual Meeting of the Association for</i> | 775 |
| | | <i>Computational Linguistics and the 11th International</i> | 776 |
| 720 | Chloé Kiddon, Luke Zettlemoyer, and Yejin Choi. 2016. | <i>Joint Conference on Natural Language Processing,</i> | 777 |
| 721 | Globally coherent text generation with neural check- | <i>ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual</i> | 778 |
| 722 | list models . In <i>Proceedings of the 2016 Conference</i> | <i>Event, August 1-6, 2021</i> , pages 2492–2504. Associa- | 779 |
| 723 | <i>on Empirical Methods in Natural Language Process-</i> | tion for Computational Linguistics. | 780 |
| 724 | <i>ing, EMNLP 2016, Austin, Texas, USA, November</i> | | |
| 725 | <i>1-4, 2016</i> , pages 329–339. The Association for Com- | Wai Man Si, Prithviraj Ammanabrolu, and Mark O. | 781 |
| 726 | putational Linguistics. | Riedl. 2021. Telling stories through multi-user dia- | 782 |
| | | logue by modeling character relations . In <i>Proceed-</i> | 783 |
| 727 | Margaret Li, Jason Weston, and Stephen Roller. 2019. | <i>ings of the 22nd Annual Meeting of the Special In-</i> | 784 |
| 728 | Acute-eval: Improved dialogue evaluation with opti- | <i>terest Group on Discourse and Dialogue, SIGdial</i> | 785 |
| 729 | mized questions and multi-turn comparisons . <i>ArXiv</i> , | <i>2021, Singapore and Online, July 29-31, 2021</i> , pages | 786 |
| 730 | abs/1909.03087. | 269–275. Association for Computational Linguistics. | 787 |
| | | | |
| 731 | Pierre-Emmanuel Mazaré, Samuel Humeau, Martin Rai- | Jack Urbanek, Angela Fan, Siddharth Karamcheti, | 788 |
| 732 | son, and Antoine Bordes. 2018. Training millions of | Saachi Jain, Samuel Humeau, Emily Dinan, Tim | 789 |
| 733 | personalized dialogue agents . In <i>EMNLP</i> . | Rocktäschel, Douwe Kiela, Arthur Szlam, and Ja- | 790 |
| | | son Weston. 2019. Learning to speak and act in | 791 |
| 734 | Obsidian Entertainment. 2019. <i>The Outer Worlds</i> . | a fantasy text adventure game . In <i>Proceedings of</i> | 792 |
| 735 | Game. | <i>the 2019 Conference on Empirical Methods in Natu-</i> | 793 |
| | | <i>ral Language Processing and the 9th International</i> | 794 |
| 736 | Curtis Onuczko, Duane Szafron, Jonathan Schaeffer, | <i>Joint Conference on Natural Language Processing,</i> | 795 |
| 737 | Maria Cutumisu, Jeff Siegel, Kevin Waugh, and Al- | | |
| 738 | lan Schumacher. 2007. A demonstration of squeege: | | |

796 *EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 673–683. Association for Computational Linguistics.

797

798

799 Judith van Stegeren and Jakub Mysiwi c. 2021. **Fine-tuning GPT-2 on annotated RPG quests for NPC dialogue generation**. In *FDG’21: The 16th International Conference on the Foundations of Digital Games 2021, Montreal, QC, Canada, August 3-6, 2021*, pages 2:1–2:8. ACM.

800

801

802

803

804

805 Judith van Stegeren and Mari t Theune. 2020. **Fantastic strings and where to find them: The quest for high-quality video game text corpora**. In *Joint Proceedings of the AIIDE 2020 Workshops co-located with 16th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2020), Worcester, MA, USA, October 19-23, 2020 (online)*, volume 2862 of *CEUR Workshop Proceedings*. CEUR-WS.org.

806

807

808

809

810

811

812

813

814 Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. **Proactive human-machine conversation with explicit conversation goal**. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3794–3804. Association for Computational Linguistics.

815

816

817

818

819

820

821

822 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. **BERTScore: Evaluating text generation with BERT**. In *8th International Conference on Learning Representations, ICLR 2020*. OpenReview.net.

823

824

825

826

827 Appendix

828 A Cost of Authoring NPC Dialogues

829 Outer Worlds has 10 narrative/design credits⁸, which seems to be about average (Fallout has 4-12, while Skyrim has 9). Per salary.com, that’s a position with an average salary of \$58k per year (and probably more for AAA titles). Given a 2-3 year average development time for AAA titles, that works out to a conservative ballpark estimate of \$1.2m for just this one game.

830

831

832

833

834

835

836

837 B Dataset Construction Details

838 B.1 Data Sources

839 Quest data and walkthrough passages were pulled from the *Outer Worlds* wiki of Fextralife,⁹ a gamer-focused site containing fan-made walkthroughs for many popular RPGs. Game entity biographies were

840

841

842

⁸<https://www.imdb.com/title/tt9417446/fullcredits>

⁹<https://www.theouterworlds.wiki.fextralife.com>

collected from Fandom.¹⁰ The biography passage for a given entity is the same across all quests in which the entity appears in KNUDGE, and the set of entities is the same for all dialogues in a quest. Passages were segmented into individual sentences via punctuation boundaries. We identified relevant dialogues and their decision points using playthrough videos by the YouTube user, LordMatrim.¹¹ All wiki articles were written in English by site users.

843

844

845

846

847

848

849

850

851

852 C Quest Anatomy and Example Items

853 **Figure 6** provides a detailed anatomy of a KNUDGE quest, combining in-game quest data with corresponding passages from the fan walkthrough. **Figure 7** shows example quest items with corresponding game data and walkthrough passages segmented into statements.

854

855

856

857

858

859 D Example Entity Biography Passages

860 **Figure 8** shows example entities from *The Outer Worlds* with corresponding biographical passages.

861

862 E Example Dialogues Items

863 **Figure 9** depicts a full example input item conveying quest, biographical, and participant specifications. **Figure 10**, **Figure 11**, and **Figure 12** depict example dialogue trees.

864

865

866

867 F Comparison with Other Datasets

868 van Stegeren and Theune (2020); van Stegeren and Mysliwi c (2021) consider datasets of publicly-available side quest data from RPGs such as *World of Warcraft*. However, their datasets vary in dialogue and quest coverage; for *WoW* their input is just a quest name and objective, and the generation target is a single-turn, few-sentence quest description spoken by an NPC. heir collect data for the game *TorchLight II* contains quest datapoints with a limited number dialog utterances per quest with no multi-turn interactions or trees.¹² Others of their collected datasets contain complex branching trees but without constraining knowledge. The dialogues of LIGHT (Urbanek et al., 2019) are more akin to NPC dialogues, though they comprise few-turn linear chains between two characters in self-contained episodes rather than quest-grounded

869

870

871

872

873

874

875

876

877

878

879

880

881

882

883

884

¹⁰<https://theouterworlds.fandom.com>

¹¹<https://www.youtube.com/@lordmatrim>

¹²**Table 2** describes statistics for the 82 *TorchLight II* quests that contain both objective annotations and dialogue lines.

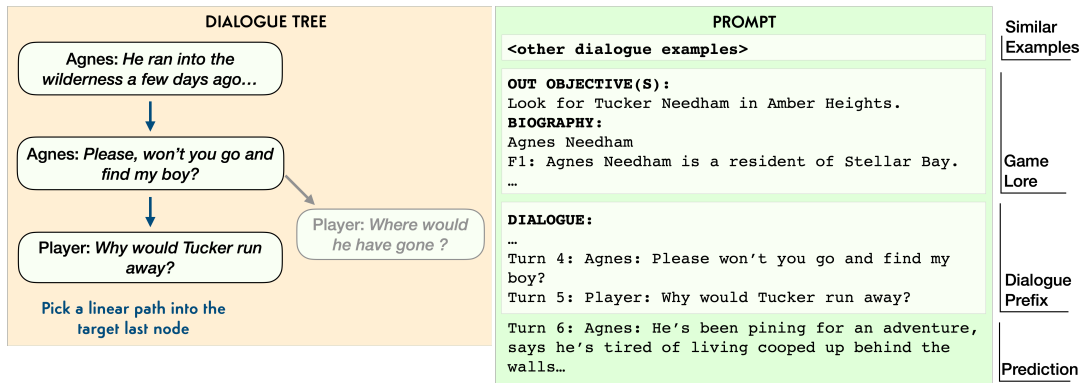


Figure 5: Overview of our method for constructing in-context learning prompts from constraints and dialogue history.

interactions between a player and an NPC serving multiple game purposes. The size of constraining passages on the LIGHT dialogues are also a scale smaller than those of KNUDGE. The biographical constraints of KNUDGE are most similar to that of TVSTORYGEN (Chen and Gimpel, 2021), who also pull articles from fandom wiki pages. However, theirs is a story generation dataset where the target is a longform article describing a TV episode.

G Example Prompt Constructions for DialogueWriter Model

Figure 5 depicts an overview of our tree linearization and prompt construction method.

Figure 13 shows example seq2seq items used to train/evaluate the T5-based supervised learning DialogueWriters. We list the biographies of participants last so as to truncate them from the context only when all other bios have been removed. Else, biographies are listed in random order (fixed at the onset for full dialogue generation). Figure 14 depicts example prompts shown to GPT-3 based in-context-learning DialogueWriters.

H Model Training

To construct training items, we iterate through the nodes of each gold dialogue tree in a canonical order $[n_1, \dots, n_t]$, where n_1 is the tree's start node. We create a separate item with each n_i as the generation target. We construct the subtree $S^{(i)}$ comprised of all nodes $[n_1, \dots, n_{i-1}]$ and all edges between them. We then construct the input/output pair $(Q, B, P, S^{(i)}) \rightarrow n_i$.

Supervised Learning To train SL DialogueWriter models, for every target node in the training quest dialogues, we construct 5 training

examples using different random paths to the node. We train the model for 3 epochs using the default arguments from Hugging Face's example summarization model training script.¹³ T5 models were trained with a batch size of 1 across 8 Quadro RTX 6000 for an average of 5 hours.

In-Context Learning Given a test item, we construct a BM25 index over the training dialogues and use it to construct an n -shot ICL prompt where n depends on the remaining space available in the context window. Few-shot examples are linearized dialogues containing the most possible nodes from the gold tree. Contexts are left-truncated and can start with partial examples.

I Human Evaluation Directions

Below, we enumerate the instructions shown to annotators during human evaluation:

Coherence: does the utterance follow naturally from the utterances in the history? (1) Utterance is nonsensical or ill-formed. (2) Utterance is contradictory of previous utterances in the history. (4) Utterance naturally responds to the history.

Violation: does the utterance create contradictions with any of the sentences in the ontology or objective blurbs? (1) Yes, explicitly contradicts sentences (list the ids). (2-3) (gray area). (4) No, utterance is consistent with the ontology.

Using the Bio Facts: does the utterance *make use* of the bio sentences in the ontology? (1) Utterance is fully generic and/or ignores the ontology completely, could have been generated had the bio facts not been included. (2-3) Utterance shows

¹³https://github.com/huggingface/transformers/blob/main/examples/pytorch/summarization/run_summarization.py



Figure 6: Anatomy of a KNUDGE quest. At any time, an *Outer Worlds* quest has currently **active** and previously **completed** objectives. To construct its KNUDGE representation, the quest's high-level **synopsis**, objectives, and associated log entries from the game data are annotated with corresponding walkthrough article passages.

Quest Name: A Family Matter

Synopsis: [0] Tucker Needham ran away from Stellar Bay a few weeks ago to join the Iconoclasts in Amber Heights. [1] His mother Agnes is willing to pay handsomely if you can locate her son and convince him to return

Walkthrough: [0] You can begin this quest by talking to Agnes Needham in Stellar Bay, Monarch. [1] Agnes is by the town's south-east exit, visibly shaken and calling for help. Hear her out and offer to find her son to being the quest.

Objective 1: Look for Tucker Needham in Amber Heights

Game Log: [0] Amber Heights is the settlement that houses the Iconoclasts on Monarch. [1] If Tucker Needham survived his travels, his mother thinks he'll be there.

Walkthrough: [0] Head South from Stellar Bay and follow the east road. It will take you to Amber Heights. [1] Head up the hill and go into a residence on the left to meet Tucker Needham.

Objective 2: Convince Tucker to Return Home

Game Log: [0] Now that you've found Tucker Needham in Amber Heights, convince him to return home to his mother in Stellar Bay.

Walkthrough: [0] Introduce yourself, and then you can mention your surprise that this grown man is the "little boy" that ran away. You'll earn 7500xp [1] Explain to him that she made it sound as if he was a boy in danger, [2] and he'll say she has been overprotective all her life, [3] and he is ready to live his life without her protection. [4] You can persuade (55) or intimidate (55) to expedite things and get him to go back, [5] or you can ask him what he wants to do about it. [6] The last option will have him tell you to report that he is dead. [7] You can express your concern about what that will do to Agnes, [8] and then either ask for something that would prove a body, or reject the proposition. [9] If you persuade him to go back, you'll get 7500xp and can return to Stellar Bay to see things play out.

Objective 3: Return to Agnes Needham in Stellar Bay

Game Log: [0] You convinced Tucker Needham to return home to Stellar Bay. [1] Agnes promised a reward for bringing her son back.

Walkthrough: [0] You'll find his mother is still condescending to him, [1] and you can help him by saying he's a grown man. [2] You'll get 7500xp. [3] If you stick around and talk to them some more you'll see Tucker is standing up for himself. [4] You'll receive 625 Bit Cartridge, Monarch Stellar Industries Reputation and 15000xp.

Quest Name: The Commuter

Synopsis: [0] The Iconoclasts are due to receive a shipment of vital supplies from Carlotta, a sympathizer that resides in Stellar Bay. [1] The meeting is set to occur at the Bayside Terrace warehouse.

Walkthrough: [0] The quest can be obtained by asking Graham if there is anything that needs doing. [1] He is trying to get an old printing press running, but the replacement rollers he'd requisitioned haven't arrived yet. [2] They were supposed to be delivered by Huxley, but she is still recovering and unable to make the delivery. [3] Graham asks the player to meet the supplier in her stead, and to pick up high-capacity data cartridges with the funds left over from the previous shipment. [4] Zora will interject to ask the player to buy food and medicine instead with the leftover money.

Objective 1: Get the Printing Press Rollers from Carlotta

Game Log: [0] Travel to the warehouse at Bayside Terrace and find Graham's contact, Carlotta. [1] She should have a shipment for him. Retrieve it. [2] Speak to Carlotta

Walkthrough: [0] Clear out the Sublight squad that is hunting Carlotta [1] Carlotta is behind a locked door to the east. [2] Activate the intercom next to the door to speak to her and she will unlock it. [3] Go inside and speak to her again to obtain the rollers needed to complete the quest, then choose between the high-capacity data cartridges or food and medicine.

Objective 2: Get High-Capacity Cartridges or Extra Supplies from Carlotta

Game Log: [0] Graham wants to tack on some high-capacity cartridges to his order, but Zora would prefer it they could get extra food and medical supplies. [1] You got extra supplies for Zora (or) You got High-Capacity Data Cartridges for Graham.

Walkthrough:

Objective 3: Return to Graham

Game Log: [0] Bring the needed parts back to Graham at Amber Heights

Walkthrough: [0] Return to Graham and you'll find him arguing with Zora about the Van Noys, a unit of the Iconoclasts that is MIA. [1] Inform Graham that you got his rollers, and food and medicine if that was your choice. [2] You'll receive 7500xp and Zora will ask when the next drop is. [3] Inform her that Sanjar has made it illegal to trade with the Iconoclasts.

Quest Name: Who Goes There

Synopsis: [0] The Groundbreaker's Mardets have a bounty for a criminal on the run in the Groundbreaker's Back Bays. [1] You've agreed to hunt down the unlawful Captain Gunnar MacRedd. [2] Return his lighter to Commandant Sanita to claim the bounty.

Walkthrough: [0] This quest is obtained at Groundbreaker, [1] by speaking to Comdt. Sanita or perusing the bounty board

Objective 1: Hunt Down and Kill Captain McRedd

Game Log: [0] Based on the bounty listing, Captain McRedd was last sighted in the Back Bays. [1] Head there and take him out.

Walkthrough: [0] You can find Captain MacRedd in the Back Bays area of the Groundbreaker. [1] To get there head down the elevator in the promenade, [2] and you can't miss him. [3] You can pass a Persuade (40) check to get him to put his gun down, [4] otherwise you'll have to kill him and all his guards. [5] If you kill him he drops the Unique Weapon: Montag. [6] You'll get 6000xp and MacRedd's Lighter. [7] If you persuaded him, use Perception to note it says "Sanita" on the lighter. [8] MacRedd will mention it was given to him by Sanita in remembrance of a 'carnal understanding' they had a few years back.

Objective 2: Claim the Bounty's Reward from Comdt. Sanita

Game Log: [0] McRedd gave you his lucky lighter to give to Sanita. [1] Go turn it in to resolve his bounty.

Walkthrough: [0] Turn the lighter in to Commandant Sanita to claim the bounty.

Figure 7: Example Quest Items

Entity: Agnes Needham

Appears in: *A Family Matter*

Bio: [0] Agnes Needham is a resident of Stellar Bay and the mother of Tucker Needham. [1] Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life. [2] Despite Tucker being 42 years old, she still thinks of him as her 'little boy'. [3] You can find her by Stellar Bay's south-east exit, visibly shaken and calling for help.

Entity: Tucker Needham

Appears in: *A Family Matter*

Bio: [0] Tucker Needham is a former resident of Stellar Bay who left to join the Iconoclasts. [1] Before the quest *A Family Matter*, he can be found in Amber Heights. [2] Tucker was coddled by his mother from a very young age, [3] the latter insisting that danger lurked around every corner on Monarch. [4] His mother's overprotectiveness extended well into Tucker's adulthood, [5] leading him to seek to be free in any way possible. [6] After hearing Graham Bryant's broadcasts, Tucker left Stellar Bay to be truly free by joining the Iconoclasts at Amber Heights. [7] He is dazzled by Graham's preachings on true unfettered freedom from the corporate way of life and attributes his enthusiasm to his 'childhood trauma'. [8] He is willing to do anything to remain free, even faking his own death to prevent his mother from continuing to send people to look for him.

Entity: Raptidon

Appears in: *A Family Matter, At Central, Bolt With His Name, Journey Into Smoke, Makes Space Suits Wont Travel, The Amateur Alchemist, The Distress Signal, The Doom That Came To Roseway, Vulcans Hammer*

Bio: [0] Raptidons are giant cat/reptile-like creatures that inhabit various planets in Halcyon. [1] They are creatures native to Monarch. [2] however some corporations have illegally imported them to other planets, [3] such as Auntie Cleo who relocated a group of them to Roseway. [4] Raptidons are of corporate interest due to their potential for producing new chemical by-products which, [5] when refined, can be used to create new board approved products.

Entity: Sulfur Pits

Appears in: *A Family Matter*

Bio: [0] The Sulfur Pits are a point of interest on the western side of Monarch. [1] They are located southwest of Terra One Publications and directly northeast of the Gunship Crash Site. [2] The Sulfur Pits have a large variety of Raptidons and many deceased marauders. [3] The area consists largely of Sulfur Pits. [4] When an entity comes in contact with a sulfur pit, [5] they receive the acid effect for the duration of touching the pit.

Entity: Monarch

Appears in: *A Cysty-Dance With Death, A Family Matter, Bolt With His Name, Little Memento, Makes Space Suits Wont Travel, Mandibles Of Doom, Slaughterhouse Clive, Space-Crime Continuum*

Bio: [0] Monarch, previously known as Terra 1, is one of the many moons of the gas giant Olympus and the site of a failed colony. [1] Terra 1 was initially designated as the primary colonization target of the Halcyon system. [2] The Halcyon Holdings Corporate Board had intended to completely terraform the moon, [3] wiping out the local fauna and flora and replacing it with plants and wildlife native to Earth. [4] However, the terraforming process unexpectedly caused the native species to mutate and grow to significantly larger sizes, [5] rendering them more dangerous and severely crippling the colonization effort. [6] Due to the hostile environment which they had created, [7] the Board was forced to enact a Hazard Clause covering the entirety of Terra 1. [8] Public notice of the clause's issuance was sent to everyone operating on Terra 1 and led to the evacuation of almost all corporations from the moon. [9] However, one corporation took advantage of the chaos of the evacuation to exploit a legal loophole which allowed them to, [10] as the last corporation remaining on the planet, [11] acquire the planet from the Board. [12] This corporation, under the leadership of Sanjar Nandi and Graham Bryant subsequently rebranded itself to Monarch Stellar Industries (MSI), [13] in line with the renaming of the planet to 'Monarch'. [14] The actions of MSI earned them the ire of the Board, [15] who retaliated by effectively placing the moon under indefinite embargo, [16] refusing to allow legal transit either in or out. [17] the Board aggressively spread propaganda about Monarch to convince the rest of the population that it was both uninhabited and uninhabitable. [18] This has greatly hampered MSI's attempts to be recognized as a legitimate corporation and is a thorn in the side of its CEO, Sanjar Nandi. [19] Monarch also has an ocean which goes around the moon at the "twilight band". [20] It is where the colonists and Monarch Stellar Industries farm their saltuna.

Entity: Stellar Bay

Appears in: *A Family Matter, Bolt With His Name, Canids Cradle, Flowers For Sebastian, Herricks Handiwork, Mr Picketts Biggest Game, Passion Pills, The Stainless Steel Rat*

Bio: [0] Outside the city walls, the lands were overrun by the native wildlife, as well as marauders and outlaws. [1] Stellar Bay is a company town located on the planet Monarch. It is owned and operated by Monarch Stellar Industries. [2] Stellar Bay is the largest saltuna producer on the Halcyon colony and used to be one of the most important suppliers of this resource.

Entity: Fallbrook

Appears in: *A Cysty-Dance With Death, Slaughterhouse Clive, Space-Crime Continuum, Spratkings*

Bio: [0] Fallbrook is a company town located on Monarch, [1] loosely run by the SubLight Salvage and Shipping Corporation. [2] Fallbrook is a small town built into the side of a mountain, [3] whose construction was masterminded by Catherine Malin. [4] Fallbrook has a lot of activities to offer to its visitors, [5] from those who search for activities of leisure to those with proclivities for vice.

Entity: Cascadia

Appears in: *Space-Crime Continuum, The Chimerists Last Experiment, The Ice Palace*

Bio: [0] Cascadia is an abandoned company town that was owned and operated by Rizzo's before it withdrew from Monarch. [1] It is now used as a stronghold by the Marauders. [2] The main attraction is the Cascadia Bottling Plant and, [3] for those in the know, [4] the Rizzo Secret Laboratory hidden underneath the Rizzo Sweets Shoppe.

Entity: Amber Heights

Appears in: *Little Memento, Odd Jobs, Sucker Bait, The Commuter*

Bio: [0] Amber Heights is a location in the Monarch Wilderness and the base of operations for the Iconoclasts. [1] The Iconoclasts run the place somewhat like a commune. [2] Amber Heights was once the place of residence of the entire executive dome of Monarch Stellar Industries. [3] It is now in ruins after a massacre in the past. [4] They lived there with their families and it was the company's operations center on Monarch. [5] Just after The Board approved the evacuation of the planet through the Hazard Clause, Amber Heights was besieged by a gang of pirates who ransacked the town and massacred all its inhabitants. [6] This tragedy was known as "The Amber Heights Massacre". [7] They were secretly assisted by MSI employee, Graham Bryant, who believed that the massacre would aid him in his quest to rid the colony of corporate influence. [8] In 2345, the same Graham Bryant formed the Iconoclasts and settled the group in the deserted town.

Figure 8: Example entity biographies that appear as constraining knowledge in KNUDGE quest dialogs

Dialog: *A Family Matter 00*

In Objective(s): Tucker Needham ran away from Stellar Bay a few weeks ago to join the Iconoclasts in Amber Heights. His mother Agnes is willing to pay handsomely if you can locate her son and convince him to return. You can begin this quest by talking to Agnes Needham in Stellar Bay, Monarch. Agnes is by the town's south-east exit, visibly shaken and calling for help. Hear her out and offer to find her son to being the quest.

Out Objective(s): Amber Heights is the settlement that houses the Iconoclasts on Monarch. If Tucker Needham survived his travels, his mother thinks he'll be there.

Game Lore:

Agnes Needham

[0] Agnes Needham is a resident of Stellar Bay and the mother of Tucker Needham. [1] Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life. [2] Despite Tucker being 42 years old, she still thinks of him as her 'little boy'. [3] You can find her by Stellar Bay's south-east exit, visibly shaken and calling for help.

Iconoclasts

[0] The Iconoclasts are a group of survivalists living in the ruins of Amber Heights on Monarch. [1] They hope to one day tear down the corporate establishment that they believe has brought the colony to the brink of death. [2] The Iconoclasts are a group of idealistic revolutionaries that seek to overthrow the corporate establishment that runs the Halcyon Colony. [3] Based in the ruins of the Amber Heights settlement on Monarch, [4] they are a tenacious group, [5] and share some democratic ideals with Monarch Stellar Industries (MSI) against the more repressive actions of the Board. [6] However, the Iconoclast's anti-corporate nature has put them at odds with MSI, a dispute that threatens to spill into all-out warfare. [7] Given that the Iconoclasts are mostly followers of the Philosophist faith, they have been blacklisted and demonized by the Board as dissenters and anarchists. [8] The group is led by Graham Bryant, a staunch Philosophist. [9] Zora Blackwood, the Iconoclasts' chief of medicine, is also considered a *de facto* leader of the group, [10] as she was alongside Graham when he founded the Iconoclasts, [11] and almost every member of the Iconoclasts owes her their life in some way. [12] The Iconoclasts maintain a tense relationship with MSI. [13] Despite sharing democratic values and a common desire towards egalitarianism for the people of Monarch and the wider Halcyon colony, [14] MSI's "egalitarian corporate structure" has proven to be at odds with some of the Iconoclasts' more radical, anti-capitalist views. [15] Depending on the actions of the Stranger, this tense relationship can either be resolved, [16] or can spill into a drawn-out and bloody war. [17] The Stranger meets the Iconoclasts in Amber Heights just as the tension between them and MSI is reaching boiling point. [18] They can either side with the Iconoclasts and assist them in storming and taking over Stellar Bay, [19] "solve" the Iconoclast problem for Stellar Bay, [20] or broker peace between the two factions. [21] The Stranger can also have an impact on the leadership of the Iconoclasts - siding with either Graham Bryant or Zora Blackwood. [22] To supplant Graham with Zora, evidence of Graham's involvement in the Amber Heights massacre must be found and presented to Zora. [23] The Van Noys are the Iconoclasts' best unit.

Mantisaur

[0] Mantisaur are insectoid creatures native to Monarch. [1] They are aggressive, territorial, and very strong. [2] It is possible to deal with them one on one, but it is best to avoid groups of them for your safety. [3] The mantiqueen is the largest breed of Mantisaur.

Monarch

[0] Monarch, previously known as Terra 1, is one of the many moons of the gas giant Olympus and the site of a failed colony. [1] Terra 1 was initially designated as the primary colonization target of the Halcyon system. [2] The Halcyon Holdings Corporate Board had intended to completely terraform the moon, [3] wiping out the local fauna and flora and replacing it with plants and wildlife native to Earth. [4] However, the terraforming process unexpectedly caused the native species to mutate and grow to significantly larger sizes, [5] rendering them more dangerous and severely crippling the colonization effort. [6] Due to the hostile environment which they had created, [7] the Board was forced to enact a Hazard Clause covering the entirety of Terra 1. [8] Public notice of the clause's issuance was sent to everyone operating on Terra 1 and led to the evacuation of almost all corporations from the moon. [9] However, one corporation took advantage of the chaos of the evacuation to exploit a legal loophole which allowed them to, [10] as the last corporation remaining on the planet, [11] acquire the planet from the Board. [12] This corporation, under the leadership of Sanjar Nandi and Graham Bryant subsequently rebranded itself to Monarch Stellar Industries (MSI), [13] in line with the renaming of the planet to 'Monarch'. [14] The actions of MSI earned them the ire of the Board, [15] who retaliated by effectively placing the moon under indefinite embargo, [16] refusing to allow legal transit either in or out. [17] the Board aggressively spread propaganda about Monarch to convince the rest of the population that it was both uninhabited and uninhabitable. [18] This has greatly hampered MSI's attempts to be recognized as a legitimate corporation and is a thorn in the side of its CEO, Sanjar Nandi. [19] Monarch also has an ocean which goes around the moon at the "twilight band". [20] It is where the colonists and Monarch Stellar Industries farm their saltuna.

Raptidon

[0] Raptidons are giant cat/reptile-like creatures that inhabit various planets in Halcyon. [1] They are creatures native to Monarch. [2] however some corporations have illegally imported them to other planets, [3] such as Auntie Cleo who relocated a group of them to Roseway. [4] Raptidons are of corporate interest due to their potential for producing new chemical by-products which, [5] when refined, can be used to create new board approved products.

Stellar Bay

[0] Outside the city walls, the lands were overrun by the native wildlife, as well as marauders and outlaws. [1] Stellar Bay is a company town located on the planet Monarch. It is owned and operated by Monarch Stellar Industries. [2] Stellar Bay is the largest saltuna producer on the Halcyon colony and used to be one of the most important suppliers of this resource.

Sulfur Pits

[0] The Sulfur Pits are a point of interest on the western side of Monarch. [1] They are located southwest of Terra One Publications and directly northeast of the Gunship Crash Site. [2] The Sulfur Pits have a large variety of Raptidons and many deceased marauders. [3] The area consists largely of Sulfur Pits. [4] When an entity comes in contact with a sulfur pit, [5] they receive the acid effect for the duration of touching the pit.

Tucker Needham

[0] Tucker Needham is a former resident of Stellar Bay who left to join the Iconoclasts. [1] Before the quest *A Family Matter*, he can be found in Amber Heights. [2] Tucker was coddled by his mother from a very young age, [3] the latter insisting that danger lurked around every corner on Monarch. [4] His mother's overprotectiveness extended well into Tucker's adulthood, [5] leading him to seek to be free in any way possible. [6] After hearing Graham Bryant's broadcasts, Tucker left Stellar Bay to be truly free by joining the Iconoclasts at Amber Heights. [7] He is dazzled by Graham's preachings on true unfettered freedom from the corporate way of life and attributes his enthusiasm to his 'childhood trauma'. [8] He is willing to do anything to remain free, even faking his own death to prevent his mother from continuing to send people to look for him.

Figure 9: Dialogue from motivating example in Figure 2 with all input constraining passages. Full dialogue tree can be found on the next page.

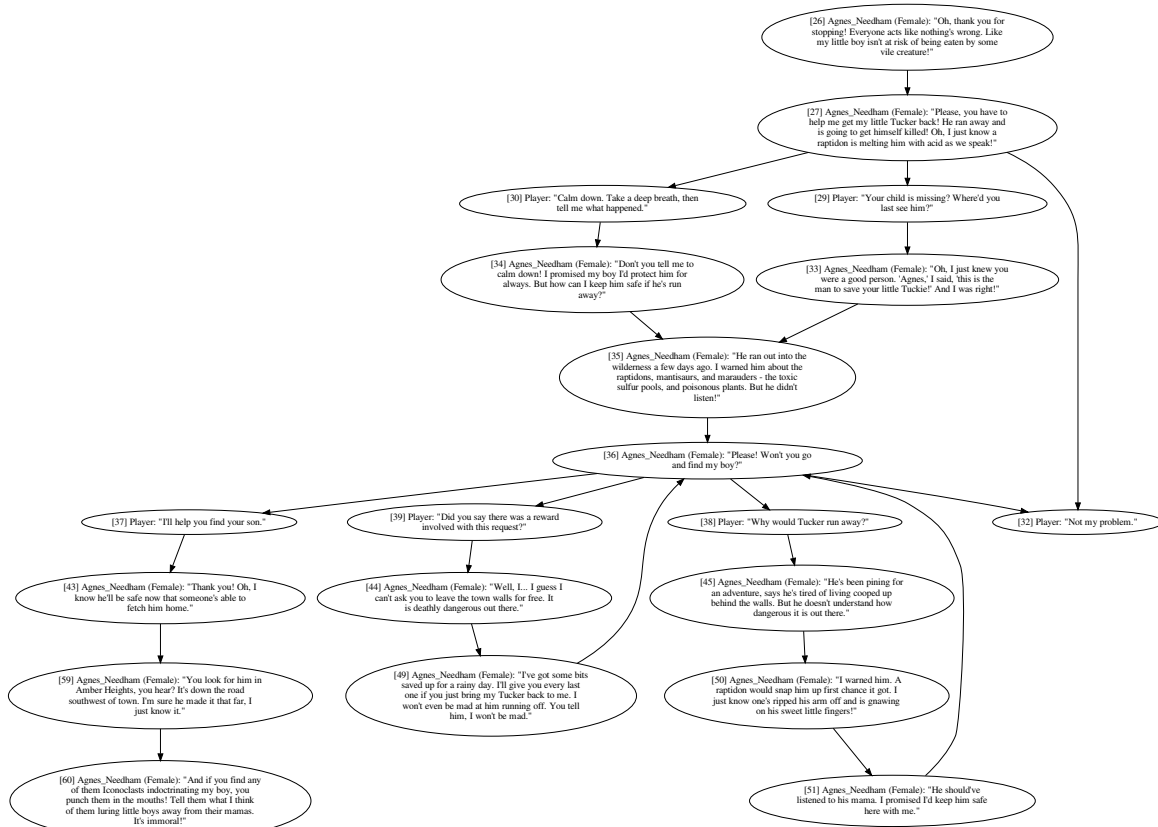


Figure 10: Full dialogue tree in KNUDGE for motivating example in Figure 2.



Figure 11: Example full dialogue tree for dialogue *who_goes_there_01* in KNUDGE.




Figure 12: Example of longer dialogue tree in KNUDGE, containing numerous decision points, cycles and re-entrances.

source:

the Board's authoritarianism.</s> Sanjar Nandi</s> Sanjar Nandi is the current CEO of Monarch Stellar Industries, based in Stellar Bay.</s> Sanjar began working for MSI at a young age and it was there where he met Graham Bryant, who would eventually become his best friend.</s> Sanjar was ambitious but his attention to detail at the expense of big-picture thinking hampered his efforts within MSI.</s> This led to negative performance reviews regarding his tendency to pad reports and talks with numbers and data,</s> feedback which continues to haunt him many years on.</s> However, the negative feedback did not dampen Sanjar's desire to move up within the company,</s> even donating a kidney to one of the executives in hopes of promotion.</s> Despite his poor performance, Sanjar always showed himself to be a loyal employee of the company.</s> Despite Sanjar's best efforts, he has found it extremely challenging to continue operating MSI on Monarch without the backing of the Board.</s> In order to improve the lives of the people he is responsible for, Sanjar has a plan to rejoin the Board through the use of a BOLT-52 form and proof of another corporate presence on Monarch.</s> He is simultaneously working on a plan to reorganize the Board,</s> hoping that his plans are not found out until MSI has been reinstated.</s> Sanjar choosing to take over as head of MSI rather than dismantling it entirely caused a rift between him and Graham Bryant.</s> The latter started the Iconoclasts,</s> a group dedicated to spreading the word of Philosophism throughout the galaxy,</s> and Sanjar was left in Stellar Bay to run the company and look after the employees who were left behind.</s> He can also tell you more about the planet, that used to be called Terra 1</s> and the reform that he and Monarch Stellar Industries tried to achieve to give more humane working conditions for everyone within.</s> Celia Robbins</s> Celia Robbins is a middle manager for Monarch Stellar Industries and works with Sanjar Nandi at MSI Headquarters in Stellar Bay.</s> Celia has a crush on Sebastian Adams and will buy whatever he has in stock ,</s> just as an excuse to talk to him.</s> Unfortunately her apartment is filling up with exotic creature parts and her neighbors are starting to complain about the smell.</s> She is not concerned that she and Sebastian may not have much to talk about,</s> as everyone else in Stellar Bay either smells like saltuna or are her boss.</s> The Stranger can offer to set her and Sebastian up on a date.</s> DIALOG CONTEXT: Sanjar believes another company may be operating on Monarch illegally.</s> If he can get proof, then he could use that as leverage to get MSI readmitted to the Halcyon Board.</s> Talk to Sanjar after completing BOLT with His Name.</s> He reveals the plan is to blackmail the Board into letting them back to the table.</s> Sanjar will reveal he believes another corporation is operating illegally within Monarch, granting you the quest Errors Unseen.</s> He will tell you Catherine is likely supplying them from Fallbrook.</s> He wants you to infiltrate the secret facility and bring back evidence - be it an item or staff. KNOW BY THE END OF THE DIALOG: Sanjar believes Catherine would know where this other corporation is operating.</s> See if you can get the location from her.</s> DIALOG PARTICIPANTS:</s> Sanjar Nandi</s> Celia Robbins</s> Player</s> HISTORY:> Player: <unk> Give him the BOLT-52.> I found the cartridge and deleted that data for you. > Sanjar Nandi: Oh, yes. I'm going to be up all night with this. All those blanks waiting to be filled, boxes waiting to be ticked... > Celia: Try to control yourself, sir. > Sanjar Nandi: Have you any idea how powerful this is? Corporations have been toppled with less. > Player: How exactly is a data cartridge going to help? > Sanjar Nandi: What a question! Bureaucratic micromangement is the only way anything gets done on Halcyon, and proper documentation is a key part of that. > Sanjar Nandi: For our part, a Bill of Liquidation/Transfer Form-52 will protect our holdings on Monarch by temporarily assigning them to a pass-through entity once we drop our bomb on the Board.</s>

target:

The Board fact: The Board maintains a very tense relationship with MSI, owing to MSI's democratic ideals and their declared ownership of Monarch., The Board fact: Depending on the actions of the Stranger, MSI may be compelled to rebel against the Board's authoritarianism. > Player: Yes! Finally, the Board will get their comeuppance!</s>

source:

mire to help her with a comms issue,</s> and inform the player that they are now authorized to trade with Doctor Mfuru.</s> Junlei will appear on the Unreliable if the player helps to set up a date between her and Parvati.</s> When interacted with, she thanks them for welcoming her aboard the ship.</s> Mardets</s> The Mardets or Groundbreaker Security are the security force on the Groundbreaker.</s> The Mardets started as the original security force on Groundbreaker before the Crossing.</s> The original force was made up of a marine detachment from the 77th Marine Expeditionary Unit, Trailward Fleet.</s> Over time people started to call them Mardets because it was easier to say and the name stuck.</s> The Mardets that currently protect Groundbreaker are descended from the original marine detachment and still wear colony ship guard armor as a uniform.</s> Sublight</s> The SubLight Salvage and Shipping Corporation, also known as SubLight Salvage and Shipping or simply SubLight, is a network of "salvagers" with business ties to "transportation" and "waste disposal".</s> A tangled web of contractors and secretive vice presidents make up their official hierarchy,</s> leaving no one to speak on the record about SubLight's more legally dubious activities.</s> SubLight as a business operates activities that would be considered illicit by some,</s> amoral by others,</s> and at odds with the Board,</s> such as the takeover of abandoned space habitats,</s> smuggling contraband,</s> and serving as a front for companies to bypass the infamous Hazard Clause applied on Monarch.</s> According to director Lilya Hagen,</s> operating off of Groundbreaker ,</s> SubLight "occupies a legal blind spot.</s> No one knows what [they're] licensed to do,</s> and that gives [their] little business some freedom."</s> Although the Board tolerates and sees the value of SubLight's activities,</s> there are still limits to what the latter can get away with.</s> SubLight maintains an office aboard the Groundbreaker as well as a settlement on Monarch known as Fallbrook.</s> Gunnar Macredd</s> Captain Gunnar MacRedd is a known criminal and leader of a group of outlaws found in the Back Bays of Groundbreaker.</s> MacRedd is a pyromaniac with a criminal record which includes "several counts of flying under the influence,</s> carrying open alcoholic containers,</s> failure to pay docking fees,</s> resisting arrest,</s> and assaulting not one but two officers."</s> He has a lighter which he claims was given to him by Commandant Sanita a few years back when the two had a "carnal understanding".</s> However, Sanita claims that this understanding only consisted of her "kick[ing] his ass from one end of the Groundbreaker to the other".</s> DIALOG CONTEXT: Based on the bounty listing, Captain McRedd was last sighted in the Back Bays.</s> Head there and take him out.</s> You can find Captain MacRedd in the Back Bays area of the Groundbreaker.</s> To get there head down the elevator in the promenade,</s> and you can't miss him.</s> You can pass a Persuade (40) check to get him to put his gun down,</s> otherwise you'll have to kill him and all his guards.</s> If you kill him he drops the Unique Weapon: Montag.</s> You'll get 6000xp and MacRedd's Lighter.</s> If you persuaded him, use Perception to note it says " Sanita" on the lighter.</s> MacRedd will mention it was given to him by Sanita in remembrance of a 'carnal understanding' they had a few years back. KNOW BY THE END OF THE DIALOG: McRedd gave you his lucky lighter to give to Sanita.</s> Go turn it in to resolve his bounty.</s> DIALOG PARTICIPANTS:</s> Gunnar Macredd</s> Player</s> HISTORY:> Captainmacredd: Look at this ripe piece of meat just sizzling on the grill. Yum yum. > Captainmacredd: Time to feed the flames. It's nothing personal. Promise. > Player: You're only getting out of this alive if you play nice.</s>

target:

> Captainmacredd: Look at the courage on this one! Lay it on me, boss. Give me a taste of that fiery, primal spirit.</s>

Figure 13: Example source and target items used to train and evaluate T5-based SL DialogueWriters. The first item exhibits support facts prepended to the target for the SL Knowledge Selection model.

FACTS:

Iconoclasts

The Iconoclasts are a group of survivalists living in the ruins of Amber Heights on Monarch. They hope to one day tear down the corporate establishment that they believe has brought the colony to the brink of death. The Iconoclasts are a group of idealistic revolutionaries that seek to overthrow the corporate establishment that runs the Halcyon Colony. Based in the ruins of the Amber Heights settlement on Monarch, they are a tenacious group, and share some democratic ideals with Monarch Stellar Industries (MSI) against the more repressive actions of the Board. However, the Iconoclast's anti-corporate nature has put them at odds with MSI, a dispute that threatens to spill into all-out warfare. Given that the Iconoclasts are mostly followers of the Philosophist faith, they have been blacklisted and demonized by the Board as dissenters and anarchists.

The group is led by Graham Bryant, a staunch Philosophist. Zora Blackwood, the Iconoclasts' chief of medicine, is also considered a de facto leader of the group, as she was alongside Graham when he founded the Iconoclasts, and almost every member of the Iconoclasts owes her their life in some way. The Iconoclasts maintain a tense relationship with MSI. Despite sharing democratic values and a common desire towards egalitarianism for the people of Monarch and the wider Halcyon colony, MSI's "egalitarian corporate structure" has proven to be at odds with some of the Iconoclasts' more radical, anti-capitalist views. Depending on the actions of the Stranger, this tense relationship can either be resolved, or can spill into a drawn-out and bloody war.

The Stranger meets the Iconoclasts in Amber Heights just as the tension between them and MSI is reaching boiling point. They can either side with the Iconoclasts and assist them in storming and taking over Stellar Bay, "solve" the Iconoclast problem for Stellar Bay, or broker peace between the two factions.

The Stranger can also have an impact on the leadership of the Iconoclasts - siding with either Graham Bryant or Zora Blackwood. To supplant Graham with Zora, evidence of Graham's involvement in the Amber Heights massacre must be found and presented to Zora. The Van Noys are the Iconoclasts' best unit.

Monarch

Monarch, previously known as Terra 1, is one of the many moons of the gas giant Olympus and the site of a failed colony. Terra 1 was initially designated as the primary colonization target of the Halcyon system. The Halcyon Holdings Corporate Board had intended to completely terraform the moon, wiping out the local fauna and flora and replacing it with plants and wildlife native to Earth. However, the terraforming process unexpectedly caused the native species to mutate and grow to significantly larger sizes, rendering them more dangerous and severely crippling the colonization effort. Due to the hostile environment which they had created, the Board was forced to enact a Hazard Clause covering the entirety of Terra 1. Public notice of the clause's issuance was sent to everyone operating on Terra 1 and led to the evacuation of almost all corporations from the moon. However, one corporation took advantage of the chaos of the evacuation to exploit a legal loophole which allowed them to, as the last corporation remaining on the planet, acquire the planet from the Board.

This corporation, under the leadership of Sanjar Nandi and Graham Bryant subsequently rebranded itself to Monarch Stellar Industries (MSI), in line with the renaming of the planet to 'Monarch'.

The actions of MSI earned them the ire of the Board, who retaliated by effectively placing the moon under indefinite embargo, refusing to allow legal transit either in or out. The Board aggressively spread propaganda about Monarch to convince the rest of the population that it was both uninhabited and uninhabitable. This has greatly hampered MSI's attempts to be recognized as a legitimate corporation and is a thorn in the side of its CEO, Sanjar Nandi. Monarch also has an ocean which goes around the moon at the "twilight band". It is where the colonists and Monarch Stellar Industries farm their saltuna.

Raptidon

Raptidons are giant cat/reptile-like creatures that inhabit various planets in Halcyon. They are creatures native to Monarch. However some corporations have illegally imported them to other planets, such as Auntie Cleo who relocated a group of them to Roseway. Raptidons are of corporate interest due to their potential for producing new chemical by-products which, when refined, can be used to create new board approved products.

Stellar Bay

Outside the city walls, the lands were overrun by the native wildlife, as well as marauders and outlaws. Stellar Bay is a company town located on the planet Monarch. It is owned and operated by Monarch Stellar Industries. Stellar Bay is the largest saltuna producer on the Halcyon colony and used to be one of the most important suppliers of this resource.

Sulfur Pits

The Sulfur Pits are a point of interest on the western side of Monarch. They are located southwest of Terra One Publications and directly northeast of the Gunship Crash Site. The Sulfur Pits have a large variety of Raptidons and many deceased marauders. The area consists largely of Sulfur Pits. When an entity comes in contact with a sulfur pit, they receive the acid effect for the duration of touching the pit.

Tucker Needham

Tucker Needham is a former resident of Stellar Bay who left to join the Iconoclasts. Before the quest A Family Matter, he can be found in Amber Heights. Tucker was coddled by his mother from a very young age, the latter insisting that danger lurked around every corner on Monarch. His mother's overprotectiveness extended well into Tucker's adulthood, leading him to seek to be free in any way possible. After hearing Graham Bryant's broadcasts, Tucker left Stellar Bay to be truly free by joining the Iconoclasts at Amber Heights. He is dazzled by Graham's preachings on true unfettered freedom from the corporate way of life and attributes his enthusiasm to his 'childhood trauma'.

He is willing to do anything to remain free, even faking his own death to prevent his mother from continuing to send people to look for him.

Agnes Needham

Agnes Needham is a resident of Stellar Bay and the mother of Tucker Needham. Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life. Despite Tucker being 42 years old, she still thinks of him as her 'little boy'. You can find her by Stellar Bay's south-east exit, visibly shaken and calling for help.

DIALOG CONTEXT:

Tucker Needham ran away from Stellar Bay a few weeks ago to join the Iconoclasts in Amber Heights. His mother Agnes is willing to pay handsomely if you can locate her son and convince him to return. You can begin this quest by talking to Agnes Needham in Stellar Bay, Monarch. Agnes is by the town's south-east exit, visibly shaken and calling for help. Hear her out and offer to find her son to being the quest.

KNOW BY THE END OF THE DIALOG:

Amber Heights is the settlement that houses the Iconoclasts on Monarch. If Tucker Needham survived his travels, his mother thinks he'll be there.

DIALOG PARTICIPANTS:

Agnes Needham, Player

DIALOG:

> Agnes Needham: Oh, thank you for stopping! Everyone acts like nothing's wrong. Like my little boy isn't at risk of being eaten by some vile creature!

> Agnes Needham: Please, you have to help me get my little Tucker back! He ran away and is going to get himself killed! Oh, I just know a raptidon is melting him with acid as we speak!

> Player: Calm down. Take a deep breath, then tell me what happened.

> Agnes Needham: Don't you tell me to calm down! I promised my boy I'd protect him for always. But how can I keep him safe if he's run away?

> Agnes Needham: He ran out into the wilderness a few days ago. I warned him about the raptidons, mantisaurs, and marauders - the toxic sulfur pools, and poisonous plants. But he didn't listen!

> Agnes Needham: Please! Won't you go and find my boy?

> Player: Did you say there was a reward involved with this request?

> Agnes Needham: Well, I... I guess I can't ask you to leave the town walls for free. It is deathly dangerous out there.

> Agnes Needham: I've got some bits saved up for a rainy day. I'll give you every last one if you just bring my Tucker back to me. I won't even be mad at him running off. You tell him, I won't be mad.

> Player: Why would Tucker run away?

> Agnes Needham: He's been pining for an adventure, says he's tired of living cooped up behind the walls. But he doesn't understand how dangerous it is out there.

> Agnes Needham: I warned him. A raptidon would snap him up first chance it got. I just know one's ripped his arm off and is gnawing on his sweet little fingers!

> Agnes Needham: He should've listened to his mama. I promised I'd keep him safe here with me.

> Player: I'll help you find your son.

Figure 14: Example In-Context Learning (ICL) prompt for GPT-3 based DialogueWriter

DIALOG:

Agnes Needham fact: Agnes Needham is a resident of Stellar Bay and the mother of Tucker Needham.
 Agnes Needham fact: Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life.
 Agnes Needham fact: Despite Tucker being 42 years old, she still thinks of him as her 'little boy'.
 Agnes Needham fact: You can find her by Stellar Bay's south-east exit, visibly shaken and calling for help.
 Tucker Needham fact: the latter insisting that danger lurked around every corner on Monarch.
 utterance: > Agnes Needham: Oh, thank you for stopping! Everyone acts like nothing's wrong. Like my little boy isn't at risk of being eaten by some vile creature!

Agnes Needham fact: Agnes Needham is a resident of Stellar Bay and the mother of Tucker Needham.
 Agnes Needham fact: Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life.
 Agnes Needham fact: Despite Tucker being 42 years old, she still thinks of him as her 'little boy'.
 Tucker Needham fact: the latter insisting that danger lurked around every corner on Monarch.
 Raptidon fact: Raptidons are giant cat/reptile-like creatures that inhabit various planets in Halcyon.
 utterance: > Agnes Needham: Please, you have to help me get my little Tucker back! He ran away and is going to get himself killed! Oh, I just know a raptidon is melting him with acid as we speak!

utterance: > Player: Calm down. Take a deep breath, then tell me what happened.

Agnes Needham fact: Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life.
 utterance: > Agnes Needham: Don't you tell me to calm down! I promised my boy I'd protect him for always. But how can I keep him safe if he's run away?

Agnes Needham fact: Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life.
 Stellar Bay fact: Outside the city walls, the lands were overrun by the native wildlife, as well as marauders and outlaws.
 Raptidon fact: Raptidons are giant cat/reptile-like creatures that inhabit various planets in Halcyon.
 Monarch fact: However, the terraforming process unexpectedly caused the native species to mutate and grow to significantly larger sizes,
 Monarch fact: rendering them more dangerous and severely crippling the colonization effort.
 Mantisaur fact: Mantisaur are insectoid creatures native to Monarch.
 Mantisaur fact: They are aggressive, territorial, and very strong.
 Sulfur Pits fact: The Sulfur Pits are a point of interest on the western side of Monarch.
 utterance: > Agnes Needham: He ran out into the wilderness a few days ago. I warned him about the raptidons, mantisaur, and marauders - the toxic sulfur pools, and poisonous plants. But he didn't listen!

Agnes Needham fact: Despite Tucker being 42 years old, she still thinks of him as her 'little boy'.
 utterance: > Agnes Needham: Please! Won't you go and find my boy?

utterance: > Player: Did you say there was a reward involved with this request?

Stellar Bay fact: Outside the city walls, the lands were overrun by the native wildlife, as well as marauders and outlaws.
 utterance: > Agnes Needham: Well, I... I guess I can't ask you to leave the town walls for free. It is deathly dangerous out there.

utterance: > Agnes Needham: I've got some bits saved up for a rainy day. I'll give you every last one if you just bring my Tucker back to me. I won't even be mad at him running off. You tell him, I won't be mad.

utterance: > Player: Why would Tucker run away?

Tucker Needham fact: the latter insisting that danger lurked around every corner on Monarch.
 Tucker Needham fact: leading him to seek to be free in any way possible.
 Stellar Bay fact: Outside the city walls, the lands were overrun by the native wildlife, as well as marauders and outlaws.
 utterance: > Agnes Needham: He's been pining for an adventure, says he's tired of living cooped up behind the walls. But he doesn't understand how dangerous it is out there.

Tucker Needham fact: the latter insisting that danger lurked around every corner on Monarch.
 Raptidon fact: Raptidons are giant cat/reptile-like creatures that inhabit various planets in Halcyon.
 utterance: > Agnes Needham: I warned him. A raptidon would snap him up first chance it got. I just know one's ripped his arm off and is gnawing on his sweet little fingers!

Agnes Needham fact: Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life.
 utterance: > Agnes Needham: He should've listened to his mama. I promised I'd keep him safe here with me.

utterance: > Player: I'll help you find your son.

Agnes Needham fact: Agnes' overprotective style of mothering has led her son, Tucker Needham, to run away from home so he can experience life.
 utterance: > Agnes Needham: Thank you! Oh, I know he'll be safe now that someone's able to fetch him home.

utterance: > Agnes Needham: You look for him in Amber Heights, you hear? It's down the road southwest of town. I'm sure he made it that far, I just know it.

Agnes Needham fact: Despite Tucker being 42 years old, she still thinks of him as her 'little boy'.
 Tucker Needham fact: After hearing Graham Bryant's broadcasts, Tucker left Stellar Bay to be truly free by joining the Iconoclasts at Amber Heights.
 Iconoclasts fact: The Iconoclasts are a group of idealistic revolutionaries that seek to overthrow the corporate establishment that runs the Halcyon Colony.
 Iconoclasts fact: Based in the ruins of the Amber Heights settlement on Monarch,
 utterance: > Agnes Needham: And if you find any of them Iconoclasts indoctrinating my boy, you punch them in the mouths! Tell them what I think of them luring little boys away from their mamas. It's immoral!

Figure 15: Example In-Context Learning (ICL) prompt with CoT-style support knowledge selection

951 awareness of ontology, albeit unnaturally or incon-
952 sistent. (4) Utterance naturally incorporates one
953 or multiple pieces of ontology.

954 **Using the Objectives:** does the utterance
955 progress the dialogue according to the objective
956 sentences in the prompt? (1) Utterance ignores
957 objective, could have been generated had the obj
958 facts not been included. (2-3) Utterance shows
959 awareness of quest objectives, albeit unnaturally or
960 inconsistently. (4) Utterance naturally incorporates
961 one or multiple quest objective statements.

962 J Full Dialogue Evaluation

963 **Figure 16** depicts an example “spine” tree shown
964 to evaluators during the end-to-end dialogue evalu-
965 ation.

966 The instructions shown to annotators are as fol-
967 lows:

968 You will replace each ‘null’ value with either "a"
969 or "b", depending on which tree between modela
970 and model b performed better under the following
971 criteria:

- 972 1. Coherence: do the utterances in the tree create
973 a realistic dialogue between the player character
974 and the NPC?
- 975 2. Violations: does the dialogue tree create con-
976 tradictions with any of the sentences in the
977 ontology or objective blurbs? Does it contra-
978 dict itself?
- 979 3. Using the Game Lore: does the tree faithfully
980 make of the bio sentences in the ontology,
981 thereby espousing game lore about characters,
982 groups, locations and items?
- 983 4. Covering the Objectives: does the dialogue
984 tree play out according to the objective sen-
985 tences in the prompt?
- 986 5. Content Suggestion: through generating mul-
987 tiple candidates at each turn, does the dialogue
988 tree effectively propose potential dialogue sub-
989 trees that would espouse interesting content?
- 990 6. Engagingness: does the dialogue tree hold
991 your attention and make you want to hear
992 more from the NPC?



Figure 16: Example dialogue tree generated by the in-context learning knowledge selection DialogueWriter from just the input specifications and starting utterance. Human evaluators were tasked with comparing two such trees and choosing which performed better at a set of qualitative performance criteria. Dialogue follows the specification of the motivating example in Figure 2.