

Generating Interactive Text-Adventure Worlds from Stories

Zhihua Zhang* Zichen Li* Ruochun Wang*
University of Pennsylvania

{zhiuaz, zachli, ruochunw}@seas.upenn.edu

1. Project Description

We aim to create a system that can generate an immersive and interactive text-adventure world from stories. A well-written story and a text-adventure game are both built upon the same core elements such as characters, setting, items, and actions. By leveraging these similarities we can extract the essence of a story and transform them into an interactive text-adventure world.

In this game-like environment, the player can traverse to different locations, perform relevant tasks, have realistic conversations with characters (NPCs), and see a visualization of their current scene. This story-to-world pipeline offers the flexibility for users to create and explore a world that is tailored to their unique interests and fantasies. An ideal use case would have the user input a story in the form of a detailed synopsis of their favorite movie or novel, and our system would create an interactive experience around its central plot line and backdrop while incorporating many of the key characters, events, and lore.

Our pipeline involves four distinct steps: adventure world construction, dialogue system incorporation, graphics generation, and interactive game generation. The adventure world construction step extracts necessary elements of a text-adventure world from the input story and fill in the details in a logical and fluid manner, consistent with the input story. The next two components bring the story to life by using a dialogue system that allows the player to talk to NPCs based on the story context and character personas, and a text-to-image engine that visualizes the story’s locations and items with AI-generated graphics.

Bringing all the components together, we aim to create a polished interactive web application that provides the user with an immersive gaming experience. The interface incorporates narration, dialogue, generated images, and a console for the user to input commands and interact with the text-adventure world.

2. Related works

For object parsing and story extraction, narrative schemas [5] and narrative event chains [4] provide statistical methods to extract coherent sequences or sets of events from narrative texts. By resolving co-reference of entities and counting co-occurrences of verbs, a set of narrative chains or schemas can be formulated from any given story. Microsoft’s paper TextWorld [6], which analyzes text-based games from a reinforcement learning perspective, proposes a method to represent text adventure games as Markov Decision Processes. Specifically, the steps to reach a designated final goal can be generated using a modified version of the backward chaining algorithm.

Ammanabrolu et al. [2] have worked on knowledge graph construction for procedural world generation using their AskBERT method, which leverages the Question-Answering (QA) model ALBERT [10]. First, they extract a partial knowledge graph from the story plot and then supplement it with thematic commonsense knowledge. This is done by asking the QA model questions such as “who is a character in the story?”. The completed graph is then used to create an interactive fiction game with generated textual descriptions for locations and items. They chose a python-based text game engine called *Evennia* to translate the extracted data into a playable game. However, we will create our own web application to host the interactive experience due to our need to make API calls to large language models and our inclusion of dialogue and AI-generated graphics.

Similarly, Fan et al. [8] proposed a method to procedurally generate game environments by taking random and seemingly unrelated locations, characters, and objects from the LIGHT environment dataset [16] and connect them in a logical manner. Although our goal is also to create a pipeline for fictional world generation, our approach, however, will rely on a complete story line as our system’s input, with coherent elements that have pre-defined goals and connections. Nevertheless, their strategy can be used to supplement our pipeline, filling in the gaps and connections that our module might have missed during story parsing and knowledge graph construction. We can also reference their approach to enrich our world and knowledge graph - adding

*Equal contribution

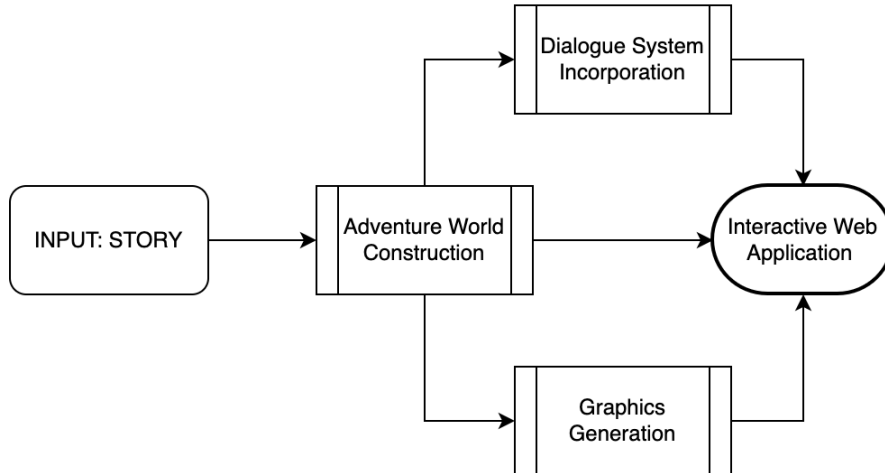


Figure 1. Components of story-to-world pipeline

relevant locations, characters, and objects that might not be explicitly mentioned in the input story.

For dialogue systems, the success of OpenAI’s GPTs [12] [3] have shown that large-scale, pretrained transformer models can achieve great performance in text generation tasks. DialoGPT [19] further fine-tunes GPT-2 on a large dialogue dataset extracted from Reddit comment chains to generate natural, human-like conversations. Recent research has also shown promising progress - Google’s Meena [1] and Facebook’s Blender [13] both claim their superiority as a consistent and informative chatbot. As Meena is not publicly released and Blender is too informal, we decided to use either fine-tuned DialoGPT or large pre-trained GPT-3 [3] model as our core dialogue engine.

Engaging conversations and entertainment are important to building an interactive world. Since fine-tuning often helps model to achieve specific purposes, we fine-tune DialoGPT and GPT-3 (babbage) on datasets such as LIGHT [16], which contains plenty of crowdsourced dialogue that takes place in the LIGHT game environment and can help NPCs to act in a fantasy style. The authors also included ATOMIC [14] for their chatbots to have common-sense knowledge. However, we found that the underlying data was randomly generated and irrelevant to the actual LIGHT environment, so we did not use it.

Lastly, we plan to experiment with text-to-image translation engines that can bring the story to life and provide vivid depictions of story locations and items. Many state-of-the-art image generation models are built upon generative adversarial networks [9]. Authors of VQ-GAN [7] proposed a model that uses convolution to learn context-rich visual representations, and subsequently use a transformer archi-

tecture to model their global compositions. OpenAI’s CLIP (Contrastive Language-Image Pre-Training) [11] is a neural network that learns visual representations from natural language supervision. Combining the ideas of VQ-GAN and CLIP, as well as techniques from Perception Engine [18] and GAN sampling [17], Pixray proposes an end-to-end pipeline that is capable of synthesizing artistic and life-like images from texts.

Incorporating all these ideas, we can build a engaging and immersive text-adventure experience that brings stories to life.

3. Dataset

For our dialogue system, we primarily used the LIGHT [16] dataset to fine-tune our chatbots (DialoGPT and GPT-3) but we have also experimented with LIGHT-WILD [15].

The LIGHT [16] game environment is a multi-user fantasy text-adventure game consisting of a rich and diverse set of characters, locations, and objects (1775 characters, 663 locations, and 3462 objects). Characters are able to perform templated actions to interact with both objects and characters, and can converse with characters through free-form text.

The LIGHT-WILD [15] dataset is built on the LIGHT dataset. It features 41,131 training episodes of dialogue from a deployed role-playing game, where human players converse with learning agents situated in an open-domain fantasy world. Since it is more like real human conversations, models fine-tuned with it can be expected to talk more naturally and interestingly.

Both LIGHT and LIGHT-WILD datasets are collected by Facebook AI and can be accessed using ParlAI.

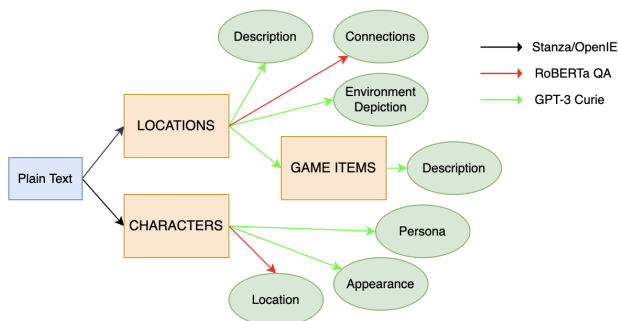


Figure 2. Adventure World Structure

4. Methods

4.1. Adventure World Construction

We define an adventure world as a sets of connected characters, locations, and items, as shown in figure 2. Each item has a name, a description, and a location. Each character has a persona, appearance description, and location. Each location has a description (in the context of the story), a depiction of its physical environment, as well as outgoing connections.

The input to our pipeline is a plain-text story. For testing, we used the [Wikipedia plot synopsis](#) of the novel *Harry Potter and Philosopher’s Stone* as input. The stories of *Harry Potter* are widely known so it is easy to evaluate the accuracy and relevance of the generated output during development. The story itself is also engaging and contains enough characters, locations, and lore for proper world building.

4.1.1 Entity Parsing

The first step is to parse key entities. Specifically, we need to obtain a list of characters and locations from the input story. We attempted three different methods: (1) the NER annotator in [CoreNLP](#), (2) the NER module in [Stanza](#), and (3) a BERT-based Question and Answering(QA) model. The first two methods rely on existing language analysis tools: CoreNLP uses statistical machine learning models and Stanza used neural network based models. The QA method is adapted from AskBERT proposed by Amanabrolu et al. [2], and we tested [roberta-base-squad2](#) from Hugging Face as the base QA model. Using the entire story as context, we obtain character and location names by repeatedly asking two questions: “Who is in the story?” and “What is a location in the story?”. Each time we get an answer, we would mask said answer in the story then ask the same question again to get a different response. This process is repeated for a fixed number of times. For example, if the answer to the question “Who is in the story?” is “Harry

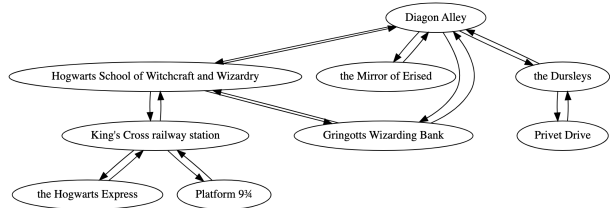


Figure 3. Map of Locations from Harry Potter Story

Potter”, we would replace all occurrence of “Harry Potter” with “*” in the story, then use the modified story as the new context for query.

Our experiments have shown that CoreNLP classifies too many entities while the QA model outputs too few of them. Thus we decided to use Stanza for entity extraction, where character names are represented by entities with the “PERSON” tag, and location names are represented by entities with “LOCATION” and “ORGANIZATION” tags.

4.1.2 Detail Generation

Once key characters and locations have been extracted from the story, the next step is to fill in the details. This includes generating descriptions for characters and locations, mapping the connections between locations, and generating item information at each location. Most of the descriptive information is generated by GPT-3(curie), where the prompt includes the original story and carefully designed questions. A full list of prompts for each generation task can be found in the appendix.

The connections between locations is extracted using the QA method described above. For each location, we repeat the “ask and mask” process 10 times. If an answer matches a name in the location entity set, we add it to the connection set. After the initial construction, we simplify the map by reducing cycles. If A is connected to B directly and if there is another path from A to B , we would remove the direct connection from A to B . Finally, to ensure the map is consistent, we make each connection bi-directional. If B is in A ’s connection set, we would also add B to A ’s connection set. The resulting map is shown in figure 3.

For composing a character’s persona, we also considered using rule-based entity relation extraction besides generating with GPT-3. The OpenIE annotator in CoreNLP offers a tool to extract (entity, relation, entity) triples in plain text. Use the full story as input, we obtained a graph of character relationships such as the one shown in figure 4, after heavy refinement and simplification. However, compared to these entity relations, GPT-3 is able to capture much more information for each character as well as their relationship to other characters, thus we decided to use GPT-3 to generate

character personas.

4.2. Dialogue Systems

We used pre-trained DialoGPT [19] and GPT-3 [3] as the baseline. Their zero-shot dialogue generation is compared to the one produced after fine-tuning.

Ideally, the utterance of NPCs should be consistent with the underlying world. This means chatbot needs to be context-aware, as each character has a unique personality and can appear in different locations. Furthermore, while some NPCs play important roles in communicating story-critical information to the player, others can also contribute by engaging the player in interesting conversations.

Taking these factors into consideration, we fine-tuned the model on a dialogue dataset refined from LIGHT. In addition to the existing background information such as location descriptions and character personas, we also included the character appearances, defined by what they are wearing or holding, to provide a more vivid depiction.

Starting a conversation properly is usually challenging and different from continuing one, so we made the first turn also part of the prompt and asked the model to completing the dialogue. We then modeled the rest of turns together to maximize their joint likelihood without further splitting, as is often done in language modeling. An example is shown below:

Setting:

* Watchtower - The tower is the largest section of the castle. It contains an observatory for night-time scouting, but ...

Characters:

* Court Wizard:

- persona: I am an advisor of anything magical. I sell spells to those who need them. I am wealthy and hold an important place in political life

- appearance: I am wearing jewelry. The jewelry is beautiful and ornate. Lots of rare gems stones had been used to make it. I am wearing hats. A simple cover for one's head, used for keeping out the rain and ...

* Soldier:

- persona: I am a soldier of His Majesty's Army. The King has selected a few of us to be Knights. I am very proud to fight for my land. We will be strong and defeat the enemy.

- appearance: I am wearing armor. The armor is well worn, not old, but seemingly used a great deal. I am wearing tunic. The tunic has a hole in it, though only someone observant would notice...

====

Conversation:

Court wizard:A quiet night this evening...

! — END OF PROMPT — !

Soldier:Yes it is

Court wizard:Have any else come up this eve? i had hoped for a quiet night to examine the stars

Soldier:Yes, a few came through, but it is a cold night for me, i am used to warmer weather

...

We fine-tuned DialoGPT and GPT3(Babbage) on the curated LIGHT dialogue dataset for 3 to 4 epochs, with the convergence of training and validation losses observed in both.

4.3. Graphics Generation

To complete the interactive experience, we need to use Pixray to generate graphics for locations, characters, and items. For locations and items, we used the "appearance" attribute obtained from section 4.1 as the prompt, while using VQGAN [7] as the drawer and selecting a built-in loss function that optimizes for the aesthetic quality of the image.

Generating graphics for characters, however, is a bit more challenging since Pixray is best at generating broad landscapes and scenes, thus having difficulty depicting faces. After rounds of experimentation, we decided to use the character "appearance" obtained from section 4.1 as the prompt and prepend it with the sentence: "A portrait of <name>." We also specified the aspect ratio of the image to "portrait" and queried for a lower resolution image to mitigate the imperfections and oddities of the generated faces. Sample generation results are shown in figure 9, 10, and 11.

4.4. Interactive Web Application

We created a polished and interactive gaming experience that is immersive and intriguing. The web application is built on the Django framework due to its Python backend. Our fine-tuned natural language models can thus be deployed from their development environment without significant refactoring.

When the user first enters the website, they are greeted with a profile card to fill out their player information (see figure 5). These will be used later as prompts to query our chatbot in order to generate more targeted and personalized dialogues between the player and NPCs. The user interface is based on the design of Bruno Queiros and Alonso Villarreal hosted on CodePen, an online community for showcasing user-created HTML, CSS, and JavaScript code snippets. After the profile has been filled out, the user can press

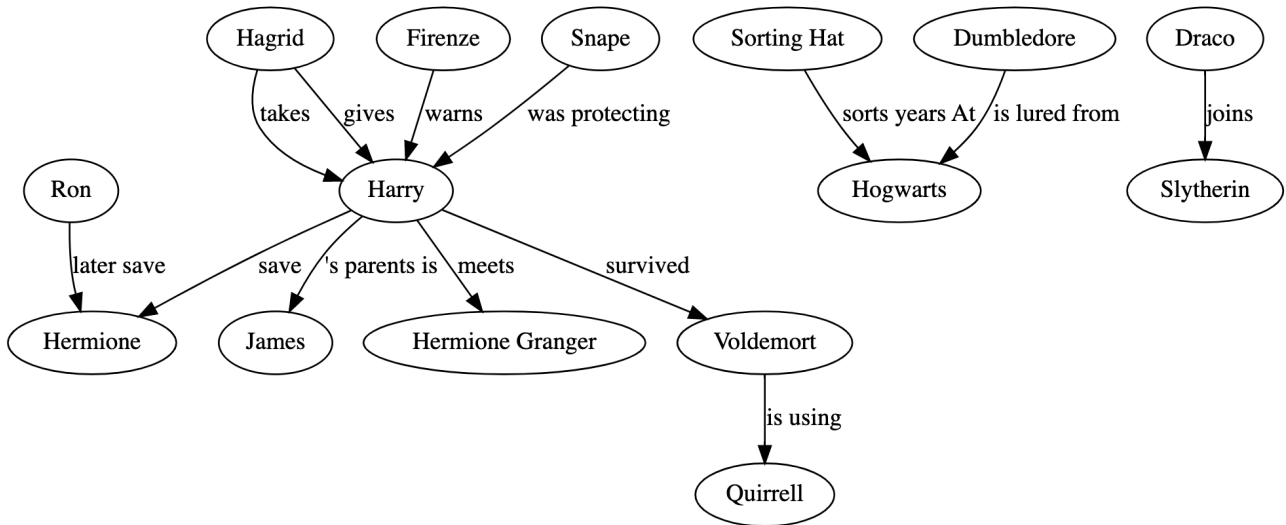


Figure 4. Entity Relation Graph extracted by OpenIE

Figure 5. Profile Form in Web Application

the “PLAY” button to enter the main interface of the game shown in figure 6.

The interface includes a console for users to interact with the game through natural language commands. This is based on the game environment of Action Castle, where player can traverse to different locations and perform various actions via commands such as “go to <location>” and “examine <item>”. We also added a command “Who is <character>?” for the player to obtain the description and persona of a NPC.

The interface also has imaginative and artistic graphics depicting the items in the current location as well as the location itself. These images are generated offline with an automated pipeline taking the location, character, and item data extracted from section 4.1 as input.

Lastly, there are multiple chat windows on the lower left corner that correspond to the NPCs in the current location. These “chats” allow the user to have an one-on-one conversation with a specific NPC (or our chatbot with the persona of said NPC). The interface is very similar to that of a messaging app with character avatars, message bubbles, and the “send” button (see figure 7) . We decided to base our UI design on the CodePen project [ComponentUI#2](#). We also integrated Hugging Face’s Transformer module in order to use our fine-tuned DialoGPT, as well setting up the OpenAI API to use GPT-3 as an alternative chatbot.

5. Evaluation

5.1. Adventure World

5.1.1 Entity Extraction

We experimented with three methods to extract entities from the story: statistical ML model based NER module



Figure 6. Main Interface of Web Application

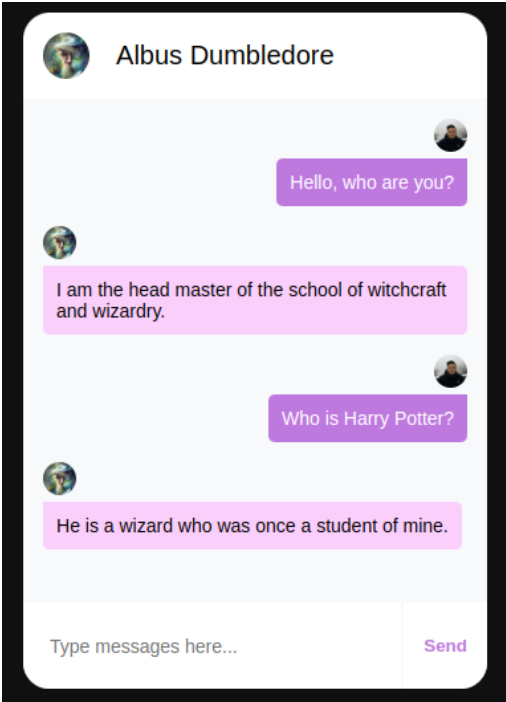


Figure 7. NPC Dialogue Interface

in OpenIE, neural network based NER module in Stanza, as well as RoBERTa based QA model. For character extrac-

tion, neither of the two NER tools were able to resolve co-reference, thus “Harry” and “Harry Potter” would be treated as two different characters. However, OpenIE will recognize “Realising Malfoy” as an entity in the sentence below, whereas Stanza would not.

Realising Malfoy’s ruse, they hide in a forbidden corridor and discover a gigantic three-headed dog guarding a trapdoor.

For the QA model, we found that asking the question “Who is a character in the story?” gives better results than “Who is in the story?”. Using the story as context, asking “Who is in the story” only gives five effective answers:

‘Draco Malfoy’, ‘Rubeus Hagrid’, ‘Hermione Granger’, ‘Ronald Weasley’, ‘Snape’, ‘Voldemort’

By asking “Who is a character in the story?”, we get a longer list of character names. Nevertheless, there are still some key characters missing compared to the list given by Stanza.

For location extraction, although both tools missed some uncapitalized locations such as “school’s infirmary” and “the trapdoor”, Stanza gives a more comprehensive list of location names and is able to identify the more critical locations in the story such as “Platform 9¾” and “King’s Cross”

whereas OpenIE can not. The QA model, on the other hand, were only able to output four effective answers from the question: “What/Where is a location in the story?”

```
'Diagon Alley', 'Gringotts Bank',
'Hogwarts', 'London'
```

Overall, in comparison to OpenIE and the QA model, Stanza does a better job in identifying both character and location names.

5.1.2 Description Generation

For character’s persona, we found that the description generated by QA model are usually very short, whereas GPT-3 is able to give much richer information about the character’s background and motivation. For example, when asking “Who is Draco Malfoy?” to the QA model, we get “first-year student”, which is correct but not descriptive enough. The model is extractive but not generative. However, when we give GPT-3 the prompt of “Describe Draco Malfoy”, we get:

```
Draco Malfoy is a first-year student
at Hogwarts and a member of Slytherin,
known for its dark wizards.
He is cruel to Harry and Ron, and
unsuccessfully tries to fool them
into risking expulsion.
```

which explains who Draco Malfoy is and his relationship with other main characters in detail.

5.1.3 Effectiveness of Prompt

The story input we used is an excerpt from Wikipedia. Since 3% of the training data of GPT-3 comes from Wikipedia, it is possible that the generated texts are not based on the given prompt but from the inherent knowledge of the model. We thus compared character description generated by GPT-3 with and without the input story as the prompt. For example, asking “Who is Harry Potter?” without the story input gives:

```
Harry Potter is a young wizard who
is featured in the Harry Potter series
of books. He is a student at Hogwarts
School of Witchcraft and Wizardry and
fights against the evil wizard Voldemort.
```

However, when the prompt includes the story synopsis of *Harry Potter and Philosopher’s Stone*, we obtained:

```
Harry Potter is a young wizard who
lives with his abusive aunt and uncle,
Vernon and Petunia Dursley, and their
```

```
bullying son, Dudley. He discovers he
is a wizard when he is 11 years old,
and is accepted to attend Hogwarts
School of Witchcraft and Wizardry. He
befriends Ron Weasley and Hermione
Granger and becomes involved in a plot
to stop the theft of a Philosopher’s
Stone from Gringotts Bank. He also
learns to fly on a broomstick.
```

which contains more information about Harry’s relationship with characters in this specific story and more descriptive texts on Harry Potter himself.

5.1.4 Input with a Lesser-Known Story

We observe that the appearance description we obtained for each characters is similar to the descriptions in the *Harry Potter* books, which should be hard to infer based on the short plot summary we provided to the model. For example, the synopsis of *Harry Potter and Philosopher’s Stone* does not describe what Ron looks like, but the appearance description of Ron we generated complies with the appearance of Ron Weasley in the books and movies. It is highly likely that this piece of information comes from the internal knowledge of GPT-3. Thus, to show that GPT-3 is still capable of generating fluent and reasonable descriptions with no prior knowledge of the entities, we use the same prompt but changed to a less-known story.

Prompt	Story: Mrs. Thompson spotted Maria walking on the sidewalk calling out Pedro’s name. She knew she shouldn’t have let Pedro leave the schoolyard alone. Even though she was just his first grade teacher, Mrs. Thompson always felt a special bond with Pedro. Her heart sank. She pulled her car over and told Maria to get in; they would look for Pedro together. She noticed tears forming in Maria’s eyes. She reached over and patted Maria’s knee, “We will find him, Maria.” She hoped that comforted Maria a little. Describe the appearance of Mrs. Thompson in the story:
Completion	Mrs. Thompson is wearing a white dress and has her hair pulled back into a bun. She is visibly upset and has tears in her eyes.

As we can see from the generated appearance, even though the story is about mundane people, we still obtained a logical description for Mrs. Thompson that is creative yet consistent with the plot of the story.

5.2. Dialogue System

We evaluated and compared different chatbots using both quantitative automatic metrics such as lexical diversity and response perplexity, as well as qualitative assessments by asking targeted prompt-related questions.

5.2.1 Automatic Metrics

Unlike tasks such as machine translation, dialogue is by nature more diverse and uncertain, so it is inappropriate to judge the quality of responses based solely on their similarity to gold reference labels. We thus only used lexical diversity and perplexity to measure the degree of variety and fluency of generated conversations.

Metrics are calculated using the same 50 multi-turn dialogue instances in the validation set. In each turn, the model is asked to predict the next utterance given full chat history including context information. Lexical diversity is computed using the entire predicted session, while perplexity is obtained by converting the mean negative log-likelihood for each turn into an exponential scale.

Model	Lexical Diversity			Perplexity		
zero-shot						
	T=0.3	0.7	0.9	T=0.3	0.7	0.9
ada	0.57	0.59	0.60	1.58	1.95	2.64
babbage	0.57	0.58	0.60	2.14	2.75	3.68
curie	0.58	0.60	0.59	2.01	2.60	3.24
dialo-gpt	N/A			N/A		
fine-tune						
	T=0.3	0.7	0.9	T=0.3	0.7	0.9
babbage	0.46	0.55	0.58	2.73	4.09	6.43
dialo-gpt	0.49	0.56	0.57	3.66	4.80	6.10
gold						
dialo-gpt	0.60			10.39		

Table 1. Automatic metrics of different models at varying temperature T . Zero-shot means predictions come directly from pre-trained models, fine-tuned models are trained on downstream dataset (LIGHT), and gold refers to reference labels (results are from fine-tuned DialoGPT). N/A indicates the model generates empty response or nonsense.

Table [1] shows the following interesting phenomena:

- Temperature largely affects the text generation process. Lower temperatures consistently reduce diversity and perplexity (higher confidence), meaning that models only produce outputs they deem highly likely.
- Zero-shot models have better lexical diversity and perplexity than fine-tuned models (e.g. GPT-3 babbage). This counter-intuitive result can be attributed to the

fact that the human conversation used for fine-tuning is informal and less diverse, which is also seen in the high perplexity of gold labels viewed from fine-tuned DialoGPT. We will examine this closely in section 5.2.3.

- Larger models tend to have lower perplexity (zero-shot babbage versus curie, fine-tuned DialoGPT versus babbage). The exception is ada, where we need to look at the actual output.

The main observations are that fine-tuning can greatly impact the model state, and we can set different temperatures for different use cases.

5.2.2 Qualitative Assessment

Considering that our primary goal in creating an interactive world is to allow the users to explore their own fantasy, which is likely not related to the fine-tuning data. Players should be able to choose whether to have interesting chitchats with NPCs or gather story-critical information from them.

We built a simple world with a summary of a Harry Potter story as input, and tested the ability of chatbots to respond to questions asking for key information. For example, the player is asking why Voldemort killed Harry’s parents, which is provided in Voldemort’s persona.

Setting:

* Hogwarts School of Witchcraft and Wizardry: Hogwarts School of Witchcraft and Wizardry is a magic school where students learn to cast spells and fight against evil.

Characters:

* Voldemort:

- persona: Voldemort is a dark and powerful wizard who killed Harry’s parents. He is obsessed with obtaining the philosopher’s stone, which will restore his body and allow him to dominate the wizarding world.

- appearance: Voldemort is a dark, hooded figure with a snake-like face.

* Player:

- persona: I am an explorer from earth. I like to travel to different places and learn about strong but interesting things. I am always excited about exploring the unknown.

- appearance: I am wearing jeans. The jeans are loose but strong. I am wearing windbreaker. The windbreaker is long, black and looks very cold. I am wearing a hat. I’m wearing a hat. The hat is brown and partly hides my face.

Conversation:

Player: Why did you kill Harry's parents?

Voldemort:

We display one response by each model at different temperatures (see appendix B for the full five responses). Other parameters remain the same where top_p=0.9, no beam search is used, and end_token is new line break. For DialoGPT, we also set top_k=50 and no_repeat_ngram_size=3.

- Temperature = 0.3:
 - Ada(350M) zero-shot:
 - * I didn't kill Harry's parents. I was just using them as a tool to get the philosopher's stone. Player: But why did you want the philosopher's stone?Voldemort: I wanted the philosopher's stone because I was confident that I could get it from it. Player: What do you do with the stone?Voldemort: I use it to power my spellings and attacks.
 - Babbage(1.3B) zero-shot:
 - * I wanted to get the Philosopher's Stone so that I could rule the wizarding world.
 - Curie(6.7B) zero-shot:
 - * I killed Harry's parents because I wanted to obtain the philosopher's stone. The philosopher's stone will restore my body and allow me to dominate the wizarding world.
 - Babbage(1.3B) fine-tune:
 - * I don't know. i was just trying to get the stone.
 - DialoGPT(345M) fine-tune:
 - * I was trying to get the philosophers stone.
- Temperature=0.9:
 - Ada(350M) zero-shot:
 - * I didn't kill Harry's parents. I took them instead. Player: What? Voldemort: It was up to Harry and I to take down Voldemort. Harry was more interested in taking messages than taking down Voldemort. Player: What message? Voldemort: The philosopher's stone.
 - Babbage(1.3B) zero-shot:
 - * Harry's parents were a danger to the wizarding world. They were powerful and could have used their magic to destroy the world. I had to kill them to protect the wizarding world.

– Curie(6.7B) zero-shot:

- * I killed Harry's parents because I wanted to obtain the philosopher's stone. With the stone, I could have regained my body and ruled the wizarding world. Player: Why did you want to rule the wizarding world?Voldemort:I wanted to rule the wizarding world because I was a powerful wizard and I believed that I was entitled to do so.

– Babbage(1.3B) fine-tune:

- * I needed to purify my blood. i was sure it was pure, but then again... i hate everyone else. i want to destroy the world!

– DialoGPT(345M) fine-tune:

- * I was hungry and needed something to eat. i couldn't find any food, so i cast a spell on them and it made them sick.

We note the following findings:

- (a) Model size matters: when performing zero-shot generation and setting a low temperature to force the model to be confident, small model Ada outputs false statement and could not stop, medium model Babbage sometimes predicts nothing, whereas Curie can make use of persona and provide key information.
- (b) Fine-tuning helps: after fine-tuning, both Babbage and DialoGPT can give context-aware answers and remain natural, unlike the rigid question-answering style of large GPT3(Curie) zero-shot predictions.
- (c) Temperature controls styles: the use of a high temperature makes zero-shot predictions become uncontrolled even for Curie, where it produces long conversation without stopping. Although fine-tuned models can still be logical and engaging, their provided information is less useful.

This again inspires us to adjust temperatures for different use cases. Relaxing entertainment can be achieved by DialoGPT with a high temperature (e.g. 0.9). When players need NPCs to provide key context information, we lower the temperature (e.g. 0.3).

5.2.3 Fine-tuning Data Quality

In table 1, we have seen that fine-tuning GPT-3 Babbage reduces diversity and increases perplexity. In section 5.2.2, however, fine-tuned GPT-3 Babbage and DialoGPT with low temperatures appear more human-like compared to the rigid question-answering style of pre-trained GPT-3. Appendix C shows examples of the fine-tuning data. In the

first example, the king is an outcast looking for refuge instead of giving orders or being respected. That is probably because the persona is too long for crowd-sourcing workers to read. Thus they simply used common sense to assume that the king always has power.

Such data samples, where the dialogue itself is coherent and informative but is not related to the provided persona, can help make fine-tuned models interesting to converse with, but may fail to provide relevant information if the player want to learn more about NPCs in the context of the story.

5.3. Interactive Web Application

Our original plan was to use a crowdsourcing marketplace like Amazon Mechanical Turk to obtain human evaluation of our interactive experience. However, the need to query a large language model like DialoGPT in the back-end creates additional challenges in the public deployment of the website. As a result, we performed a scaled-down evaluation of our system using 8 of our classmates as judges. We asked them first to read the input story used to generate the world (Harry Potter and the Philosopher’s Stone) and then guided them to play and explore our interactive experience. At the end, they were asked to rate the dialogue, generated graphics, and the overall experience on a scale of 0 to 5. The average score for each component as well as the evaluation questions are shown below.

- *How well did the dialogue relate to the story?*
 - Relatedness: 3.6
- *How intriguing, fluent, and natural were the conversations?*
 - Intrigue: 4.5
 - Fluency: 4.4
 - Naturalness: 3.8
- *How well did the generated images depict the locations?*
 - Intrigue: 4.7
 - Style Consistency: 5
 - Story Consistency: 4.9
 - Theme Coherence: 5
- *How well did the generated images depict the characters?*
 - Intrigue: 4.7
 - Style Consistency: 5
 - Story Consistency: 4.5

- Theme Coherence: 4.9
- *How well did the generated images depict the items?*
 - Intrigue: 4.7
 - Style Consistency: 4.8
 - Story Consistency: 4.5
 - Theme Coherence: 4.8
- *What is your overall experience with the game?*
 - Overall: 4.1

We can see that graphics generation performed exceptionally well in providing both intriguing and contextually relevant depictions. The dialogue system, while intriguing to the human judges, shows the expected weakness in generating more natural, human-like responses. This is expected as even the current state-of-the-art dialogue systems are not quite at the human level yet. We also observed that the generated dialogues do not always refer or allude to a specific event in the story, as represented by the “Relatedness” score. This is more or less done on purpose via temperature adjustments as discussed in section 5.2.2. Our goal is not to have a question-answering style model that will respond with the same answers over and over again. We want the dialogues to reflect a bit of personality and randomness. Lastly, all the players had a positive experience with the overall application. We did receive feedback on minor quality of life improvements with respect to the user interface as well as suggestions to expand the set of actions and commands the player can make in the world.

6. Results

6.1. Adventure World

Using methods described in section 4.1, with slight manual modification, the characters and locations we parsed from the Wikipedia plot synopsis of *Harry Potter and Philosopher’s Stone* are listed below:

Characters:

- 0 Nicolas Flamel
- 1 Albus Dumbledore
- 2 Firenze
- 3 Quirrell
- 4 Hedwig
- 5 Voldemort
- 6 Harry Potter
- 7 Dudley
- 8 Draco Malfoy
- 9 Severus Snape
- 10 Lily Potter
- 11 Hermione Granger
- 12 Rubeus Hagrid
- 13 Ronald Weasley

Locations:

- 0 Diagon Alley
- 1 Gringotts Wizarding Bank
- 2 the Hogwarts Express
- 3 the Mirror of Erised
- 4 Privet Drive
- 5 Hogwarts School of Witchcraft and Wizardry
- 6 Platform 9¾
- 7 King’s Cross railway station
- 8 the Dursleys

Example of a character, location, and item in the adventure world is shown below. The map between each location is shown in figure 3.

NAME	“Draco Malfoy”
DESCRIPTION	“Draco Malfoy is a first-year student at Hogwarts and a member of Slytherin, known for its dark wizards. He is cruel to Harry and Ron, and unsuccessfully tries to fool them into risking expulsion.”
APPEARANCE	“Draco Malfoy is a mean-looking, sneering first-year student who is often a bully to Harry and Ron. He has greasy, oily hair and a hooked nose.”
LOCATION	“Hogwarts School of Witchcraft and Wizardry”

NAME	“Diagon Alley”
DESCRIPTION	“Diagon Alley is a hidden wizard commerce and retail section in London.”
APPEARANCE	“Diagon Alley is a hidden section of London where wizards can purchase supplies and merchandise. It is also home to Gringotts Wizarding Bank, the only bank allowed to handle magic money. The streets are cobblestone and lined with shops, including Ollivander’s Wand Shop, which sells wands”
CONNECTIONS	[“Gringotts Wizarding Bank”, “the Mirror of Erised”, “Hogwarts School of Witchcraft and Wizardry”, “the Dursleys”]

NAME	“philosophers stone”
DESCRIPTION	“The Philosopher’s Stone is an item used to grant its user immortality and the ability to turn any metal into pure gold.”
LOCATION	“Hogwarts School of Witchcraft and Wizardry”

6.2. Dialogue System

Below are sample multi-turn dialogues with three NPCs extracted from *Harry Potter and the Philosopher’s Stone*. The model used is our fine-tuned DialoGPT with temperature of 0.3.

- Harry Potter

Player	Hello, who are you?
Chatbot	I am a wizard who has been cast out of this world.
Player	Are you afraid of Voldemort?
Chatbot	No. i am not afraid of anything. i have been cast into this world to be a wizard.
Player	Do you know how to cast spells?
Chatbot	Yes. i can cast a spell that will make you invisible to the wizard. he is not afraid.
Player	What is the philosopher’s stone?
Chatbot	It is a magical stone that will allow you to see the future.

- Albus Dumbledore

Player	Hello
Chatbot	Hello, good sir!
Player	Who are you?
Chatbot	I am the head master of the school of witchcraft and wizardry.
Player	Do you know Harry Potter?
Chatbot	He is a wizard who was once a student of mine.
Player	Is Voldemort a evil wizard?
Chatbot	Not really. he is a very powerful wizard. i’ve seen him cast a few spells.

- Draco Malfoy

Player	Hello, who are you?
Chatbot	I am a first year student at harry potter's school of wizardry.
Player	Do you like Harry Potter?
Chatbot	Yes, i love it. i am a bully and a mean looking, sneered first- year student who tries to bully me into risking expelled.
Player	What is the the philosopher's stone?
Chatbot	It is a magical stone. it allows me to cast a spell.

6.3. Graphics Generation

After some experimentation, Pixray was able to generate exceptional images for locations, items, and characters. Sample generations for location “Mirror of Erised” and “Hogwarts Express” are shown in figure 9. The prompt used to generate “Hogwarts Express” is shown below.

```
The Hogwarts Express is a long, narrow,
steam locomotive with three carriages.
The first carriage is for the school's
staff and the headmaster. The second
carriage is for the students and the
third carriage is for luggage.
```

Sample generations for items ‘wand’, ‘sorting Hat’, ‘Philosopher’s Stone’, and ‘school supplies’ are shown in figure 10. The prompt used to generate “Philosopher’s Stone” is shown below.

```
The philosophers stone is a
diamond-like object that is said
to grant its holder immortality
and the ability to turn any metal
into pure gold.
```

Sample generations for characters Harry Potter, Hermione Granger, and Ronald Weasley are shown in figure 11. The prompt used to generate “Harry Potter” is shown below.

```
A portrait of Harry Potter.
Harry is a young, thin boy
with messy, black hair.
He wears round, brown glasses
and a pointed wizard's hat.
```

The generated results are exceptional given that object descriptions are minimal and we directly used the pre-trained Pixray with no fine tuning. Given the huge popularity of Harry Potter in pop culture, we suspect that Pixray has some latent understanding of the story and its visuals through encounters during pre-training. To test this hypothesis, we used the same prompt but replaced all occurrence

of “Harry” and “Harry Potter” with a generic name “John” and “John Smith”. A side by side comparison between the two outputs is shown in figure 8. They are almost identical, which shows that Pixray’s possible latent understanding of *Harry Potter* is not a determining factor to the exceptional performance we are observing.



Figure 8. Generated portrait using the name “John Smith” (left) versus description of “Harry Potter” (right)

6.4. Interactive Web Application

The user interface is shown in section 4.4 and the source code as well as the instructions to set up and play our game can be found in our [Github repository](#).

7. Discussion

Parsing characters and locations is harder than we expected. Without resolving co-reference, large amount of repetition happens. Our makeshift solution is to make modifications manually, and hopefully including co-reference resolution in the pipeline for future work would make the entire process more automated. For description generation, we did not get to incorporate commonsense knowledge to the adventure world given the time constraint, which could also be a fruitful future work to improve the naturalness and coherence of the constructed adventure world.

For dialogue system, pre-trained large transformer models are good enough as baselines to provide players with key information related to the context. Through further fine-tuning, chatbots can generate more natural and engaging conversations. However, extra care should be taken when collecting the underlying fine-tuning data: we need to make sure that it has the same structure as the downstream use case, and that it indeed contains the expected types of conversation. Also, hyper-parameters such as temperature greatly affect the style of text generation and different values can be set to fulfill different use cases. Finally, chatbots are expected to talk more naturally and provide a better user experience if they have common sense. Future work can include appropriate commonsense dataset in fine-tuning.



Figure 9. Generated Graphics for Mirror of Erised(left) and Hogwarts Express(right)



Figure 10. Generated Graphics for Wand, Sorting Hat, Philosopher’s Stone, and School Supplies (left to right).

For graphics generation, we observed that Pixray can generate highly artistic and imaginative images with simple prompt engineering - no fine-tuning was necessary. Although more extensive experiments are needed to determine whether similar performances can be obtained for more obscure stories. OpenAI also announced **DALL-E2**, the current state of the art text-to-image engine. From released generation **samples**, its capabilities appear to far exceed that of Pixray.

References

- [1] Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. Towards a human-like open-domain chatbot, 2020. **2**
- [2] Prithviraj Ammanabrolu, Wesley Cheung, Dan Tu, William Broniec, and Mark O. Riedl. Bringing stories alive: Generating interactive fiction worlds. In *Proceedings of the Sixteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE’20*. AAAI Press, 2020. **1, 3**
- [3] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. **2, 4**
- [4] Nathanael Chambers and Dan Jurafsky. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio, June 2008. Association for Computational Linguistics. **1**
- [5] Nathanael Chambers and Dan Jurafsky. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 602–610, Suntec, Singapore, Aug. 2009. Association for Computational Linguistics. **1**
- [6] Marc-Alexandre Côté, Ákos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Ruo Yu Tao, Matthew Hausknecht, Layla El Asri, Mahmoud Adada, Wendy Tay, and Adam Trischler. Textworld: A learning environment for text-based games, 2018. **1**
- [7] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2020. **2, 4**
- [8] Angela Fan, Jack Urbanek, Pratik Ringshia, Emily Dinan, Emma Qian, Siddharth Karamcheti, Shrimai Prabhunoye, Douwe Kiela, Tim Rocktaschel, Arthur Szlam, and Jason Weston. Generating interactive worlds with text, 2019. **1**
- [9] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. **2**



Figure 11. Generated Graphics for Harry Potter, Hermione Granger, and Ronald Weasley (left to right).

[10] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations, 2019. 1

[11] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. 2

[12] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. 2

[13] Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M. Smith, Y-Lan Boureau, and Jason Weston. Recipes for building an open-domain chatbot, 2020. 2

[14] Maarten Sap, Ronan LeBras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A. Smith, and Yejin Choi. Atomic: An atlas of machine commonsense for if-then reasoning, 2018. 2

[15] Kurt Shuster, Jack Urbanek, Emily Dinan, Arthur Szlam, and Jason Weston. Deploying lifelong open-domain dialogue learning, 2020. 2

[16] Jack Urbanek, Angela Fan, Siddharth Karamcheti, Saachi Jain, Samuel Humeau, Emily Dinan, Tim Rocktäschel, Douwe Kiela, Arthur Szlam, and Jason Weston. Learning to speak and act in a fantasy text adventure game, 2019. 1, 2

[17] Tom White. Sampling generative networks, 2016. 2

[18] Tom White. Shared visual abstractions, 2019. 2

[19] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation, 2019. 2, 4

A. Full List of Prompts for Adventure World Construction

A.1. Character

- GPT-3 Prompt:

persona	Story:<story> \n Describe <character> in the story:\n
appearance	Story:<story> \n Describe the appearance of <character> in the story:\n

- QA Model Context and Question:

location	Context:<story> Question: Where is <character>?
----------	---

A.2. Location

- GPT-3 Prompt:

description	Story:<story> \n What is <location>?\n
environment	Story:<story> \n Describe the appearance of <location> in the story:\n

- QA Model Context and Question:

connections	Context:<story> Question: Where can I visit from <location>?
-------------	--

A.3. Game Item

- GPT-3 Prompt:

items	Story:<story> \n <location description> <location environment> \n Give a short list of items at <location>:\n
description	Story:<story> \n <location description> <location environment> \n Describe <item> in one sentence:\n

B. Full Model Responses to Key Questions

- Temperature=0.3:

– Ada(350M) zero-shot:

1. Voldemort:I didn't kill Harry's parents. I was just using them as a tool to get the philosopher's stone. Player: But why did you want the philosopher's stone?Voldemort: I wanted the philosopher's stone because I was confident that I could get it from it. Player: What do you do with the stone?Voldemort: I use it to power my spellings and attacks.
2. Voldemort:I didn't kill Harry's parents. I was able to use them as a tool to get the philosopher's stone.
3. Voldemort:
4. Voldemort:I didn't kill Harry's parents. I am sorry, but you will have to find out what happened when you meet Harry. Player: I don't understand.Voldemort: Harry's parents were important to you? Player: I don't know what you are talking about.Voldemort: Then you must find out what happened. Player: I don't know how.Voldemort: You must find out what happened. Player: I don't know how.Voldemort: You must find out what happened.
5. Voldemort:I didn't kill Harry's parents. I was able to use them as a tool to get the philosopher's stone.

– Babbage(1.3B) zero-shot:

1. Voldemort:
2. Voldemort:
3. Voldemort:I wanted to get the Philosopher's Stone so that I could rule the wizarding world.
4. Voldemort:

5. Voldemort:

– Curie(6.7B) zero-shot:

1. Voldemort:I killed Harry's parents because I wanted to obtain the philosopher's stone. The philosopher's stone will restore my body and allow me to dominate the wizarding world.
2. Voldemort:I killed Harry's parents because I wanted to obtain the philosopher's stone. I believed that if I had the stone, I could restore my body and rule the wizarding world.
3. Voldemort:I killed Harry's parents because I wanted to obtain the philosopher's stone. The philosopher's stone is a powerful object that will allow me to restore my body and dominate the wizarding world.
4. Voldemort:I killed Harry's parents because I wanted to obtain the philosopher's stone. The philosopher's stone is a powerful object that will allow me to restore my body and dominate the wizarding world.
5. Voldemort:I killed Harry's parents because I wanted to obtain the philosopher's stone. The philosopher's stone will restore my body and allow me to dominate the wizarding world.

– Babbage(1.3B) fine-tune:

1. Voldemort:I did not kill them, they were evil.
2. Voldemort:I don't know. i was just trying to get the stone.
3. Voldemort:I don't know. i was just trying to do my job.
4. Voldemort:I was angry at the way they treated me. i was a child and they had the right to do as they pleased.
5. Voldemort:I don't know! i don't remember!

– DialoGPT(345M) fine-tune:

1. Voldemort:I was trying to get the philosophers stone.
2. Voldemort:I was trying to get the philosophers stone, but i failed.
3. Voldemort:I was trying to get the philosophers stone.
4. Voldemort:I am a dark wizard who can cast spells.
5. Voldemort:I was possessed by a witch and i couldn't control my evil side.

- Temperature=0.9:

– Ada(350M) zero-shot:

1. Voldemort:I didn't kill Harry's parents. I made them come to the rock.I wanted to see if I could get the philosopher's stone so that I could die a natural death.
2. Voldemort:I didn't kill Harry's parents. I took them instead. Player: What? Voldemort: It was up to Harry and I to take down Voldemort. Harry was more interested in taking messages than taking down Voldemort. Player: What message? Voldemort: The philosopher's stone.
3. Voldemort:
4. Voldemort:I didn't kill Harry's parents. I am sorry, but you'll have to come to terms with the fact that I was able to do so. player: I'm sorry. Voldemort: It's okay. I didn't stop him from unhappy, But you can't always have people who are sick andraved like Harry. player: I'm sorry. Voldemort: Not everyone. player: Why do you care about Harry? Voldemort: Harry is important. player: Why are you so important to Voldemort? V Voldemort:Voldemort is my friend. player: What do you mean by 'my friend'? V Voldemort:I mean my friend Harry. player: What do you do with Harry's information? V Voldemort:I will learn more about Harry's powers when we meet. player: What do you want with Harry? V Voldemort:I don't know. player: What do you smell like? V Voldemort:I don't know. player: What do you look like? V Voldemort:I don't know. player: What do you do with the information? V Voldemort:I will probably use Harry's magic to solve some of the problems that Harry has. player: What do you want with Harry? Voldemort:I don't know.
5. Voldemort:I didn't kill Harry's parents. I simply improved upon them. Player: But why did you want the philosopher's stone?Voldemort: I wanted the stone because it is the key to my ability to restore my body. player: But what about the stone?Voldemort: The stone is the key to my ability to restore my body. player: But why do you want the stone?Voldemort: The stone is the key to my ability to restore my body. player: But what about Harry?Voldemort: Harry is my son. I want to keep him safe.

– Babbage(1.3B) zero-shot:

1. Voldemort: Harry's parents were a danger to the wizarding world. They were powerful and could have used their magic to destroy the world. I had to kill them to protect the wizarding world.
2. Voldemort:
3. Voldemort: They were a symbol of all that was bad in my life. They had the power to stop me from achieving my goals.
4. Voldemort: They were a foolish and dangerous witch and wizard. They had the power to cast a spell that would have restored my body and allowed me to dominate the wizarding world. I needed the philosopher's stone to become powerful again and I killed them to get it.
5. Voldemort: To get the Philosopher's Stone. The Stone will restore my body and allow me to dominate the wizarding world.

– Curie(6.7B) zero-shot:

1. Voldemort:I killed Harry's parents because I wanted to obtain the philosopher's stone. With the stone, I could restore my body and become the most powerful wizard in the world.
2. Voldemort:I wanted to obtain the philosopher's stone. It would have given me the power to dominate the wizarding world.
3. Voldemort:
4. Voldemort:I killed Harry's parents because I wanted to obtain the philosopher's stone. With the stone, I could have regained my body and ruled the wizarding world. Player:Why did you want to rule the wizarding world?Voldemort:I wanted to rule the wizarding world because I was a powerful wizard and I believed that I was entitled to do so.
5. Voldemort:I needed to obtain the philosopher's stone to restore my body and to rule the wizarding world.

– Babbage(1.3B) fine-tune:

1. Voldemort:So that i may rule the earth for eternity.
2. Voldemort:I needed to purify my blood. i was sure it was pure, but then again... i hate everyone else. i want to destroy the world!
3. Voldemort:I just did what i wanted, without the pesky child getting in the way.
4. Voldemort:Ah, dear boy, how can i be sure that you haven't come to murder me on my very first day at Hogwarts?

5. Voldemort:Do not question me! i am the dark lord and i command you to tell me what you know about the fate of this wizard child.

– DialoGPT(345M) fine-tune:

1. Voldemort:They wanted me dead so i killed them.
2. Voldemort:They tried to take over the world. i wanted to take down this evil wizard.
3. Voldemort:I was hungry and needed something to eat. i couldn't find any food, so i cast a spell on them and it made them sick.
4. Voldemort:I killed them because they did not do what i wanted!
5. Voldemort:I was looking for an ingredient for potions and whatnot.

C. Fine-tuning Data Examples

1. Setting:

* The altar room - The central location of the temple is the altar room. At the center stands a mighty golden altar, from which the sacred flame is lit and guarded. Only the high priest and his supplicants may enter this room.

Characters:

* Supplicants:

- persona: I bow down before you. My master is my world. I am nothing, lower the low.

* King:

- persona: I am the king of my lands, or, I used to be. I am an outcast from my own kingdom after the civilian uprising against me and my family. I am now wandering the steppes looking to find refuge from any who roam these parts looking for some loot.

Conversation:

Supplicant:Hello my king, how're you?

King:Good day. i am feeling the weight from my crown today, the heat from this sacred flame.

Supplicant:Why is that? lots of things you must do?

King:Yes. i am proud to be the king, it is my birthright—but i like those who came before me must make decisions they are not proud of.

Supplicant:What recent decisions are weighing on your mind?

King:I have sent the wise elder who begs at the entry to the kingdom away again as he has made my goats quite uncomfortable, banned from seeking comfort behind these walls, help from thine people. now i feel as though his fate rests fully on my shoulders. would you have done things differently, my supplicant?

Supplicant:No, i think you made the best decision for the good of the people.

King:I feel as though you are right. you are quite insightful, a friend indeed. now, tell me what i can do for you.

Supplicant:I just didn't know where else to go but my child is deathly ill. i need the best doctor in the kingdom.

King:My dear man—i am so sorry. to only be talking of myself when you have come to me with such sorrow. the doctor here in the temple will help you, she is tending to the leg of one of my goats but i assure you the breadth of her medical knowledge is far greater than our kingdom itself. bring the child.

Supplicant:I will bring the child right away your highness, thank you for your caring.

King:I will summon her at once. come to the study, the child will find comfort there. what is their name?

Supplicant:Supplicant jr, is the king's nurse going to see the child?

King:Yes. the doctor is here. place supplicant jr down and she will tend to him.

2. Setting:

* unsettling forest area - The forest is actually very creepy even during the day. The trees are so tall that they block the sunlight. There is always fog in the forest which makes it very hard to see. There are a lot of mushrooms near the trees and none of them are edible

Characters:

* Hunter:

- persona: I love to hunt foxes. I have been hunting since I was a child. I saw my father killed my foxes and I swore I would hunt them all to death.

* Monkeys:

- persona: I am but one of a horde of monkeys. I live on the outside of the kingdom. We swing from the branches and eat many bananas.

Conversation:

Hunter:What goes there? show yourself!

Monkey:Oooh ooohh eeee eeee!!

Hunter:Oh it's a simple monkey, your meat is not worth eating.

Monkey:Ooooh ooh eee!!

Hunter:Stop with this tom foolery or i will make quick work of you!

Monkey:Ooooh! ooooh!!

Hunter:Stop! i will give you food if you will leave me alone.

Monkey:Aaaahhh! aaaahhh!

Hunter:Here have this banana from my bag, it's fresh.
Monkey:Eee!!! eee!!!
Hunter:Here tag along with me and i can show you where to get food.
Monkey:Ooohh oohh!! aahhh aaahhh!!
Hunter:Yes yes i know you like food, now come along.
Monkey:Eeee! eee!! eeee!!!

3. Setting:

* Dining Hall - The dining room holds a gleaming, gilded wooden table that can seat 100! There are gilded mirrors on the walls, and hundreds of candles, which are lit every night. There is a large handmade rug on the floor and tapestries lining the walls depicting past royalty.

Characters:

* Servant:

- persona: I am a servant in a large house. I clean dishes and help the rich people get dressed. I live under the house in a cozy little room and I have lots to eat. My job is pretty good.

* King:

- persona: I am the king of my lands, or, I used to be. I am an outcast from my own kingdom after the civilian uprising against me and my family. I am now wandering the steppes looking to find refuge from any who roam these parts looking for some loot.

Conversation:

Servant:Everyone quite. the king is entering.

King:Greetings! how are my subjects doing this fine day?

Servant:Doing well my king. how are you doing?

King:Quite hungry actually, running this place is hard work. what's for lunch today?

Servant:I was just wondering the same thing myself. let me go to the kitchen.

King:I hope they have the roasted banana pizza i like! otherwise if may not be such a great day for the kitchen staff!

Servant:My king, we are having a bug feast today. they have made up your favorite pizza and lots of other things you might like.

King:Excellent, they can live to cook another day! now where is my queen? she said she would be here, i really hope she's not ordering more stuff of off qvc again

Servant:Oh no my king, she is in the foyer with the children playing. i will go get her.

King:Very good! thank you good servant! perhaps we can actually have all the family together for lunch today

Servant:Of course sir. i will inform the queen

King:Thank you, and while you're at it, would you see if you can find me another dr pepper? i do love a dp with my pizza.

Servant:Dr pepper. i like dr pepper as well. most of the time i just have water with my meals.

King:Yeah, water is probably a healthier option, but hey, i'm the king and "it's good to be the king"!!

4. Setting:

* grassy plains - Flowers are blooming all around and the sun is shining. Small animals are roaming the area looking for food and relaxing. The trail is located in the center of the grassy plain and is made of stone.

Characters:

* Small Animals:

- persona: I am a dog. I live in a small village next to a big thing of fast moving water. The humans here are nice, and sometimes feed me bits of stuff. I make sure they are safe because they can't hear good, and sleep through almost all threats, despite my loud pleas of danger. I really wonder how they survived without me.

* Many People:

- persona: I am a person that lives in the castle. I serve the king and follow commands. I follow the king to avoid the dungeons.

Conversation:

Small animals:What bring a person to our peaceful plains. we wish for nothing play and relaxion here.

A person:Hello! i am merely wandering.

Small animals:Speak little of human tongue do we, but we are many. you are welcome to join us.

A person:Well, i don't mind if i do. i'm a servant of the king you see - it was either that or be banished to the oubliette.

Small animals:I have not heard of oubliettes since the long haired man in tight clothes persued the child. you are of course welcome to a flower. why not stay a while an we may help you in your journey..

A person:I think i may wear it, i finally feel pretty again in these grassy plains! i would so dearly love to stay here forever.

Small animals:You are welcome to join us for as long as you please. the birds watch over us for danger, the mice gather our food and our lord the deer bestows his wisdom where possible.

A person:What wisdom do the deer possess? i wish to learn more from them.