

Improved Instruction Ordering in Recipe-Grounded Conversation

Duong Minh Le, Ruohao Guo, Wei Xu, Alan Ritter

Georgia Institute of Technology

{dminh6, rguo48}@gatech.edu; {wei.xu, alan.ritter}@cc.gatech.edu

Abstract

In this paper, we study the task of instructional dialogue and focus on the cooking domain. Analyzing the generated output of the GPT-J model, we reveal that the primary challenge for a recipe-grounded dialog system is how to provide the instructions in the correct order. We hypothesize that this is due to the model’s lack of understanding of user intent and inability to track the instruction state (i.e., which step was last instructed). Therefore, we propose to explore two auxiliary subtasks, namely User Intent Detection and Instruction State Tracking, to support Response Generation with improved instruction grounding. Experimenting with our newly collected dataset, ChattyChef, shows that incorporating user intent and instruction state information helps the response generation model mitigate the incorrect order issue. Furthermore, to investigate whether ChatGPT has completely solved this task, we analyze its outputs and find that it also makes mistakes (10.7% of the responses), about half of which are out-of-order instructions. We will release ChattyChef to facilitate further research in this area at: <https://github.com/octaviaguo/ChattyChef>.

1 Introduction

Historically, work on conversational agents has mostly fallen into one of two categories: open-domain chatbots (Ritter et al., 2011; Li et al., 2016; Thoppilan et al., 2022; Shuster et al., 2022) or goal-directed dialogue systems within narrow domains (Williams et al., 2016; Eric et al., 2020). However, recent advances in large language models have paved the way for the exploration of dialog agents that can engage in conversations with users to accomplish open-ended objectives, such as learning about a new topic (Dinan et al., 2019; Choi et al., 2018; Reddy et al., 2019), interpreting bureaucratic policies to answer questions (Saeidi et al., 2018), or negotiating within strategy games (Lewis et al., 2017; Bakhtin et al., 2022).

Correct Response		49.6%
Incorrect Response	Wrong order	22.9%
	Irrelevant response	10.7%
	Lack of information	8.4%
	Wrong information	8.4%

Table 1: Manual analysis of 10 recipe-grounded conversations (131 responses in total) generated by a fine-tuned GPT-J model on the test portion of our new dataset 🍳 ChattyChef. The incorrect responses are classified into four error types (examples in Figure 1) with out-of-order instructions being the most common.

In this paper, we explore the task of *Recipe-Grounded Conversation*, where the dialogue agent is expected to converse with a user to walk him/her through the cooking procedure of a recipe, while answering any questions that might arise along the way (see examples in Figure 1). Although many types of dialogue tasks have been proposed and explored, very little prior work has focused on providing instructions to a user to complete a task. In contrast to other dialogue tasks, such as document-grounded conversation (Dinan et al., 2019), accurately tracking the conversation state is more crucial in recipe-grounded dialogue. This is because the agent needs to know which step in the recipe the user is currently working on in order to answer questions, such as: *what is the next step?*

To investigate what challenges may arise in recipe-grounded conversation, we have collected a dataset by crowdsourcing (see §2). As an initial baseline model, we used this data to fine-tune GPT-J following a similar protocol to Peng et al. (2022). Specifically, the conversation history, the grounded recipe, and the gold system response are concatenated as a long input sequence to the model (see §5 for details). We show that fine-tuning GPT-J on a few recipe-grounded conversations works surprisingly well, however, the model makes a significant number of mistakes based on our manual inspection

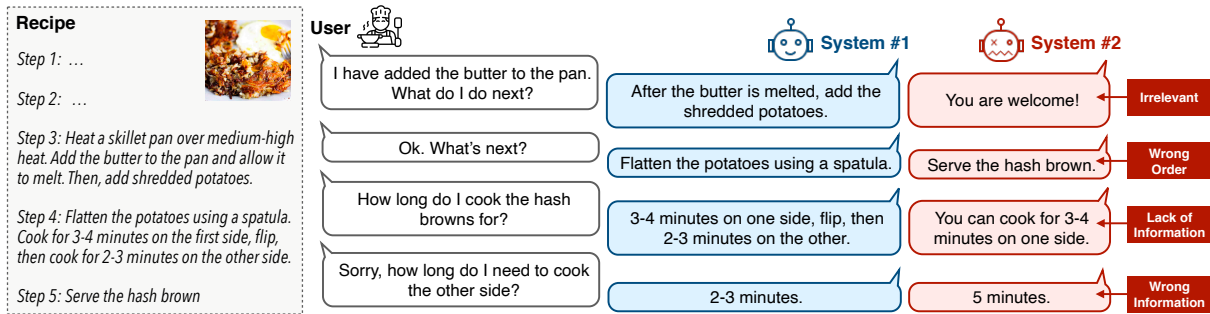


Figure 1: A conversation snippet of the cooking instructional dialogue task with **good** and **bad** system responses and the corresponding error type of each bad response.

over 10 conversations (131 generated responses) of the fine-tuned model (Table 1). Examples of each type of error are presented in Figure 1. Notably, the most prevalent type of errors is presenting information from the recipe to the user in the wrong order. We thus focus on tackling this most common error in our work. We hypothesize two potential causes: (1) GPT-J struggles to understand the user’s intent, and (2) GPT-J has difficulty tracking the current state throughout the conversation. Both are crucial in many scenarios, for example, when the user asks for more information about the current instruction, the system should not skip ahead to a later step of the recipe. Based on these hypotheses, we experiment with two supplemental tasks to improve instruction ordering: User Intent Detection (§3) and Instruction State Tracking (§4).

The goal of *Intent Detection* is to classify the user’s current intent within a fixed set of possibilities (e.g., ask for the next instruction or ask for details about ingredients). Because the set of intents varies across domains,¹ we take a few-shot transfer learning approach to leverage existing dialogue intent datasets, such as MultiWOZ 2.2 (Budzianowski et al., 2018) and Schema Guided Dialogue (Rastogi et al., 2020). We show that incorporating natural language descriptions of intents (Zhao et al., 2022) can enable more effective transfer. For example, F1 score for detecting 19 different user intents in ChattyChef increases from 32.0 to 65.1 when transferring from MultiWOZ (§3.2). In addition to Intent Detection, we also explore a simple yet effective method for *Instruction State Tracking*. State tracking aims to identify which recipe step the user is currently working on. We show that based on unigram F1 overlap, despite

¹For example, we would need a different set of intents to model instructions in Windows help documents (Branavan et al., 2010).

the approach’s simplicity, we are able to identify the most relevant recipe step at each turn of the conversation with nearly 80% accuracy.

The information from these two subtasks is then used to support *Response Generation* to improve instruction ordering (§5). Specifically, instead of feeding the whole recipe into the generation model, we leverage the instruction state to select only the most relevant knowledge. To incorporate user intents, we enrich the input prompt to the model with natural language descriptions of the predicted intent. Experiments show that even though intent and instruction state predictions are not perfect, including this information in the Response Generation model helps mitigate the wrong-order issue. We release 🍳 ChattyChef, a new dataset of cooking dialogues, to support future work on instruction-grounded conversational agents.

2 Dataset Construction

To collect a corpus of recipe-grounded conversations for fine-tuning and evaluating models, we first obtain WikiHow² articles under the Recipes category from the data compiled by Zhang et al. (2020). We control the qualities of the recipes by only selecting articles that have a helpful vote rating greater than 75% and have at least 5 votes. We retain the images from the recipes in our dataset, but experiments in this paper only make use of recipe texts. Moreover, in order to improve the conversation quality and avoid crowd workers quitting in the middle of a long conversation, we remove recipes with more than 8 steps.

2.1 Conversation Collection

After getting the recipes, we then ask crowd workers to create conversation data by role-playing.

²<https://www.wikihow.com/Main-Page>

There are two roles in each conversation of our dataset: an agent and a user. The agent is provided with a full recipe article and assumed to be an expert on cooking, while the user can only see the title of the recipe (i.e., the name of the dish). During the conversation, the agent needs to help the user complete a cooking task with the knowledge learned from the recipe and/or their common knowledge about cooking. Different interfaces are used by crowd workers when they play the agent and the user (see Appendix D.2).

Crowd workers are instructed that conversations should be relevant to the provided recipe, and should also be natural and interesting. In our process, at each turn of a conversation, agents need to identify and highlight the relevant text span in the article, if present, before sending a message, but they are not allowed to answer by copying and pasting. Instead, the agent must rephrase the instruction in their own words. To facilitate more natural interactions, both workers can discuss guidelines and ask their partner to resend messages whenever they have confusion or disagreements, using a separate chat interface.³ Furthermore, in our preliminary study, we found that users tend to repeatedly send messages such as “What is next?” to simply urge the agent to move on without thinking or trying to learn the task. To encourage diverse conversations, we provide different dialog act prompts for annotators to choose from: “teach a new step”, “ask a question”, “answer a question”, and “other” (see Appendix D.2). Diverse dialog acts such as asking and answering questions are encouraged with higher payments.

2.2 Dataset Statistics

We summarize the statistics of our final dataset – ChattyChef – and compare it with CookDial (Jiang et al., 2022) in Table 2. Compared to Cookdial, even though ChattyChef has fewer utterances per dialogue, our recipe steps are much longer, and each step includes multiple sentences or micro-steps (about 6.0 sentences per step on average). This feature sets our dataset aside from the CookDial, where nearly all recipe steps have only one short sentence or one single instruction. Having recipes with long, multi-sentence steps makes the conversation more diverse, giving crowd workers more freedom in choosing their own way of in-

³During data annotation, we formed teams of two workers and allowed them to communicate with each other via Slack.

Dataset	ChattyChef	CookDial
<i>Conversation Statistics</i>		
#Dialogues	267	260
#Utterances per dialog	26.0	35.0
#Grounding recipes	267	260
<i>Recipe Statistics</i>		
#Steps per recipe	3.9	8.4
#Tokens per recipe	417.7	120.0
#Sentences per step	6.0	1.0
#Tokens per recipe step	70.1	14.4

Table 2: The statistics of our dataset and CookDial.

Dataset	Diversity (%) (1/2-gram)	N-gram overlap (%) (1/2/3/4/5-gram)
CookDial	18.7 / 38.1	44.6 / 24.4 / 15.7 / 10.6 / 7.4
ChattyChef	26.0 / 53.6	30.2 / 12.0 / 5.8 / 3.4 / 2.2

Table 3: The statistic about the diversity of the agent’s utterances in CookDial and ChattyChef. **Diversity** (dist-1/2): the number of unique unigrams/bigrams divided by the total number of all unigrams/bigrams in the conversations. **N-gram overlap**: the percentage of n-gram overlap between the recipes and agent responses.

structing, as some micro-steps can be done in parallel while others can be skipped. This attribute is important as it makes our dataset closer to a real-life setting, where the user will normally not strictly follow the order of steps in the recipe. As shown in Table 3, the utterances from the agent in our dataset are much more diverse than CookDial and have instructions worded more differently from the grounded recipe.

Analysis of Instruction State Changes. In this work, we define the instruction state at a time step as the last recipe step which the agent instructed. We analyze the change of instruction state between two consecutive agent utterances of our dataset in Figure 2. Most of the time, the instruction would be either the same or the next recipe step. However, there are also cases when the agent needs to go back to previous steps (e.g., when the user requests to repeat an instruction) or go ahead (e.g., when the user wants to skip some steps) to provide instructions. In ChattyChef, the agent sometimes goes backward for as many as six steps, or forward seven steps. These observations have partially demonstrated the challenge of instructing multi-step procedures in the correct order, as there are many possibilities. Simply providing information from the recipe in a linear sequence is insufficient.

3 User Intent Detection

In this section, we discuss the User Intent Detection subtask. Formally, the task is to predict a set of user intent indices \mathcal{I} (as one utterance may contain multiple intents), given the t^{th} user utterance U_t^{usr} and the conversation history $\mathcal{H} = \{U_1^{sys}, U_1^{usr}, \dots, U_{t-1}^{sys}\}$. We hypothesize that providing information about the user’s intents may help the response generation model better provide information to the user in the correct order. For example, if the user asks for ingredient substitutions, the system should not respond by providing information based on the current step of the recipe.

3.1 Few-Shot Transfer Learning

Instruction-grounded chatbots could potentially be developed for many domains beyond cooking, for example, repair manuals (Wu et al., 2022), wet-lab protocols (Kulkarni et al., 2018), software documentation (Olmo et al., 2021), etc. Each new domain will require a different set of user intents, motivating the need for few-shot intent classification. To better support few-shot learning, we also investigate whether existing large-scale dialogue corpora, such as MultiWOZ, can be used to transfer knowledge to instruction-grounded chatbots.

Simply training on MultiWOZ using existing intent labels is unlikely to work well, because the MultiWOZ intents are very different from those needed for recipe-grounded dialogue. To address this challenge, we utilize natural language descriptions of the intents, following Zhao et al. (2022). A full list of intent descriptions is provided as input to T5 model (Raffel et al., 2020), which is fine-tuned to predict the index of the correct intent. Intent indices are randomized during fine-tuning, to prevent memorization of the MultiWOZ intents, which are different from those in ChattyChef, and force the model to learn to recognize intents based on their text descriptions. A complete list of intents and their associated descriptions is presented in Table 10 (in the Appendix). Example prompts used for MultiWOZ, and recipe-grounded conversation are presented in Table 12 (in the Appendix).

In addition to supporting few-shot transfer learning from existing resources such as MultiWOZ, the intent descriptions are also useful for providing intent information to the response generation model (see §5, and Table 13 in the Appendix for details.)

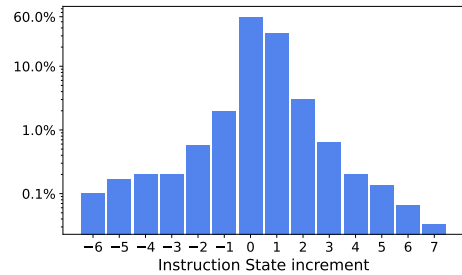


Figure 2: Statistic of the change of the instruction state between two consecutive agent utterances (the minus value means that the agent moves back and instructs on the past steps).

	Train	Valid	Test
#Dialogues	134	46	87
#Turns per dialogue	26.6	24.8	26.0

Table 4: Data split statistics of ChattyChef.

3.2 Experiments

Datasets. To evaluate the performance of the model on ChattyChef in the few-shot setting, we annotate the user intents for 10 conversations from the train set, 10 conversations from the validation set, and all 87 conversations in the test set. We consider a total of 19 different user intents for ChattyChef, which include 16 intents inherited from the CookDial dataset (Jiang et al., 2022) and 3 new intents: `req_confirmation` (ask for verification), `req_explanation` (ask to explain the reason or explain in more detail), `req_description` (ask for description). The full list of all 19 intents and their descriptions can be found in Appendix B.1. The numbers of intent annotations in the train, validation, and test split are 128, 125, and 1059. In addition to ChattyChef, for intent detection, we also use other datasets to conduct the cross-dataset experiments. In particular, we use one in-domain dataset CookDial and two large out-of-domain datasets, namely MultiWOZ 2.2 and Schema Guided Dialogue (SGD). Specifically, MultiWOZ 2.2 contains dialogues covering 8 domains (i.e., restaurant, hotel, attraction, taxi, train, hospital, bus, and police); while SGD has conversations spanning over 20 domains (e.g., banks, events, media, calendar, travel) but also not include cooking. For MultiWOZ and SGD, we extract all user utterances, which have active intents along with the conversation histories. For CookDial, as there is no official data split, we split the data on the conversation level with the propor-

tion of train, validation, and test set being 8:1:1. The sizes of the train/validation/test set of MultiWOZ, SGD and CookDial are 47,897/6,208/6,251, 153,940/22,832/39,623, and 3,599/466/545, respectively.

Models. We choose to experiment with the following training settings to evaluate this subtask: (1) **In-context learning** (In-context): As a baseline approach, we prompt the GPT-J model, which learns to do this task by only observing a few examples without any parameter updates. (2) **Few-shot fine-tuning** (None \rightarrow ChattyChef): in this setting, we fine-tune the T5 model (following the approach discussed in §3.1) on a few training examples (16-/128-shot) from ChattyChef. (3) **Cross-dataset** ($X \rightarrow$ ChattyChef): the T5 model is first fine-tuned on another dataset (i.e., X may be MultiWOZ, SGD, or CookDial), and the fine-tuned model is then directly used to predict user intents in ChattyChef for 0-shot experiment or is further fine-tuned on a few examples of ChattyChef to perform the task. (4) **Cross-dataset two-hop** ($X \rightarrow$ CookDial \rightarrow ChattyChef): this setting is similar to the *Cross-dataset* setting except that the model is fine-tuned on two datasets, first on an out-of-domain dataset (either MultiWOZ or SGD) then on CookDial.

For all models which utilize T5, we use the T5-XL version. More details about the training process of these models are described in Appendix B.1.

Results. Following Jiang et al. (2022), we use micro-F1 as the evaluation metric for Intent Detection. Table 5 demonstrates the performance of different models. In-context learning with 16 demonstrations significantly outperforms few-shot fine-tuning on a single dataset (None \rightarrow ChattyChef) with 128 examples.

Moreover, fine-tuning on another dataset first ($X \rightarrow$ ChattyChef), either from in-domain or out-of-domain, does help boost the performance (over None \rightarrow ChattyChef) dramatically on all settings. This result is expected for the in-domain dataset as, besides the domain similarity, CookDial and ChattyChef also share a large number of intents. More interestingly, leveraging MultiWOZ and SGD also improves the performance by more than 28% and 33% for the 16- and 128-shot, respectively, even though these two datasets cover quite different domains and intents from ChattyChef.

Finally, fine-tuning the model on MultiWOZ or SGD first further improves the performance of

Model	0-shot	16-shot	128-shot
In-context	-	40.2	-
None \rightarrow ChattyChef	-	7.5	32.0
MultiWOZ \rightarrow ChattyChef	14.2	36.2	65.1
SGD \rightarrow ChattyChef	21.5	34.9	66.9
CookDial \rightarrow ChattyChef	72.3	72.8	74.5
MultiWOZ \rightarrow CookDial \rightarrow ChattyChef	73.9	76.6	77.7
SGD \rightarrow CookDial \rightarrow ChattyChef	73.7	76.5	78.3

Table 5: Performance of models in the User Intent Detection task in few-shot settings. Fine-tuning the models on large out-of-domain datasets first (i.e., MultiWOZ or SGD) is helpful for the task in low-resource settings.

	#	WordMatch	SentEmb
Validation	576	82.0	80.8
Test	1145	79.0	79.4

Table 6: The alignment accuracy for Instruction State Tracking on the validation and test set.

CookDial \rightarrow ChattyChef. In particular, both MultiWOZ/SGD \rightarrow CookDial \rightarrow ChattyChef outperform CookDial \rightarrow ChattyChef on all settings with large margins. From this result and the above observations, we note that fine-tuning on a large dataset first, even from other domains, is extremely helpful for intent detection in the low-resource setting.

Since we want to measure the effectiveness of incorporating the intent information into the generation model, we will use the intent predictions from the best model (SGD \rightarrow CookDial \rightarrow ChattyChef 128-shot) in the later experiments on response generation (§5).

4 Instruction State Tracking

We study the second subtask to support the instruction ordering of the generation module – Instruction State Tracking. The goal of this task is to predict the current state of the instruction, or in other words, the last instructed recipe step. Formally, given the t^{th} system response U_t^{sys} , the previous instruction state T_{t-1} (i.e., an index of a recipe step), and the recipe with a list of n_r steps $\mathcal{R} = \{R_1, R_2, \dots, R_{n_r}\}$, the expected output of this subtask is T_t .

4.1 Aligning Conversations to Recipe Steps

For this subtask, we adopt a simple unsupervised approach to track the instruction state. The key idea of our approach is to align the most recent system utterance with its most similar step in the

recipe, and this aligned step will be the current instruction state. If the utterance can not be aligned with any recipe steps, the current instruction state will be the same as the previous one. For the scoring function that measures similarity between the conversation history and the text of recipe steps, we use two simple approaches: (1) **WordMatch** (Word Matching): the scoring function computes the unigram F1 overlap between two input texts. (2) **SentEmb** (Sentence embedding): the scoring function computes the cosine similarity between sentence embeddings of the two input texts. More details about the alignment algorithm and SentEmb approach are described in Appendix B.2.

4.2 Experiments

Setup. For this subtask, we evaluate two approaches: **WordMatch** and **SentEmb**, which were discussed above. We manually annotate the instruction states for all the system responses in ChatyChef and evaluate the accuracy of the two approaches on the validation and test sets.

Results. The performance of Instruction State Tracking is reported in Table 6. Despite of its simplicity, the WordMatch approach has comparable performance to SentEmb. In particular, WordMatch outperforms SentEmb on the validation set by 1.2%, but is slightly worse on the test set by 0.4%. One plausible explanation is that there are many entities (e.g., ingredients, cooking utensils) in the recipe that are hardly be paraphrased in the cooking dialogue. In the next section, we will use the predicted instruction state from the WordMatch approach for integration with the generation model.

5 Response Generation

Given the conversation history and the grounded recipe, the Response Generation task aims to generate the instruction response to the user. Formally, given the history $\mathcal{H} = \{U_1^{sys}, U_1^{usr}, \dots, U_{t-1}^{sys}, U_{t-1}^{usr}\}$, the recipe \mathcal{R} , the dialog system is expected to generate the next utterance U_t^{sys} .

5.1 Generating Dialog Responses

Base Model. In this work, we chose GPT-J (Wang and Komatsuzaki, 2021) as the base model. To fine-tune the model, we follow the approach of Peng et al. (2022) that concatenates the dialog history, the cooking recipe, and the system response as “ $\mathcal{H} \langle | \text{Knowledge} | \rangle \mathcal{R} \Rightarrow [\text{system}] U_t^{sys}$ ”

and feed it to the model. Both `[system]` and `<|Knowledge|>` are regular text strings. Let S be the source text, which corresponds to part of this concatenated string of the dialog history and the cooking recipe (i.e., “ $\mathcal{H} \langle | \text{Knowledge} | \rangle \mathcal{R} \Rightarrow [\text{system}]$ ”). In the fine-tuning phase, the model tries to learn the conditional probability of $P(U_t^{sys} | S)$, which can be written as the product of a series of conditional probabilities:

$$P(U_t^{sys} | S) = \prod_{i=1}^{n_t} p(U_{t,i}^{sys} | U_{t,<i}^{sys}, S)$$

where n_t is the length of the response at the t^{th} turn of conversation, $U_{t,<i}^{sys}$ indicates all tokens before the i^{th} token generated by the system.

Intent-aware Model. Since the intent labels may not convey the full meaning of the user’s intents (§3), we propose to leverage the natural language description of the intents when integrating this information type into the model. In particular, we enhance the input prompt to the GPT-J model as “ $\mathcal{H} \langle | \text{Knowledge} | \rangle \mathcal{R} [\text{user}] \text{ wants to: } D. \Rightarrow [\text{system}] U_t^{sys}$ ”, where D is the description of user intent. An example prompt is shown in Table 13 (in the Appendix).

State-aware Model. When provided with a full recipe, the response generation model might have difficulty choosing the correct recipe part to condition on when generating a response, which can lead to giving wrong order instructions. As the instruction state (§4) indicates the last instructed step, this information is essential for selecting the proper knowledge from the recipe for the model. Therefore, we explore two heuristic approaches for knowledge selection: (1) **Cutoff**: only select recipe steps starting from the current instruction state. Formally, the input prompt to the GPT-J model is “ $\mathcal{H} \langle | \text{Knowledge} | \rangle \mathcal{R}' \Rightarrow [\text{system}] U_t^{sys}$ ”, where $\mathcal{R}' = \{R_{T_{t-1}}, \dots, R_{n_r}\}$, and T_{t-1} is the output from the Instruction State Tracking module. (2) **Center**: Only select recipe steps in the ± 1 window around the current state. Formally, the input prompt to GPT-J will be similar to Cutoff except for $\mathcal{R}' = \{R_{T_{t-1}-1}, R_{T_{t-1}}, R_{T_{t-1}+1}\}$.

5.2 Experimental Setup

For this task, we evaluate the following models: (1) **GPT-J**: the base GPT-J model. (2) **GPT-J+cut**: a state-aware model, using the *Cutoff* approach to

Model	Automatic Evaluation				Human Evaluation				
	BLEU	BLEURT	Length	Diversity	wrong order	irrelevant	lack of info.	wrong info.	correct
GPT-J	4.1	44.7	11.1	9.9 / 37.9	22.9	10.7	8.4	8.4	49.6
GPT-J+int	3.9	45.0	10.0	10.4 / 38.5	18.3	8.4	11.5	6.1	55.7
GPT-J+cut	4.3	45.2	10.9	9.9 / 38.7	20.6	6.9	10.7	6.1	55.7
GPT-J+ctr	4.7	45.9	11.7	9.3 / 36.6	23.7	3.8	11.5	4.6	56.4
GPT-J+ctr+int	4.2	45.1	10.3	10.8 / 39.3	22.9	5.3	9.9	7.6	54.2
ChatGPT [†]	5.4	53.0	64.9	12.5 / 45.3	6.1				89.3

Table 7: Automatic and human evaluation of Response Generation models ChattyChef. The automatic evaluation is conducted on the entire test set of 87 multi-turn conversations, and human evaluation is on 10 multi-turn conversations (131 generated responses). †: Automatic evaluation metrics for ChatGPT are computed using the same subset of test data used in the human evaluation due to its access constraints. Length: the average length in terms of number of tokens. Diversity: diversity scores based on the unique unigrams (left) and bigrams (right).

select the grounded knowledge. (3) **GPT-J+ctr**: same as the above method, but the grounded knowledge is selected by using the *Center* approach. (4) **GPT-J+int**: the GPT-J model which is incorporated with User Intent information. (5) **GPT-J+ctr+int**: a state-aware model using the *Center* approach and is additionally incorporated with *user intent*. (6) **ChatGPT**:⁴ We also interact with ChatGPT, a chatbot launched by OpenAI. We provide the recipes and the corresponding conversation histories from our test set and ask ChatGPT to provide the next system response. At the time of writing, because OpenAI had not yet published an API to process a large number of requests, we manually interacted with ChatGPT and collected 131 responses for 10 test conversations. The details about the training process of GPT-J base model and its variants are provided in Appendix B.3.

5.3 Results

We report the following automatic evaluation metrics: BLEU (Papineni et al., 2002),⁵ BLEURT (Selam et al., 2020), the average length of the outputs and the diversity scores (Li et al., 2016) based on the proportion of unique n-grams over the total number of n-grams. Because there is a lack of consensus on how best to automatically evaluate open-domain dialogue models (Liu et al., 2016; Sedoc et al., 2019; Csáky et al., 2019), we also conduct two human evaluations in which model outputs are rated in terms of *correctness* while the *errors* are categorized.

⁴<https://chat.openai.com/chat>

⁵All BLEU scores reported in this paper are based on corpus-level BLEU-4.

Automatic Evaluation. Table 7 shows the performance of different models on the test set of ChattyChef. All GPT-J variants (+int, +cut, +ctr, and +int+ctr) have comparable or better performance than the base model, except for GPT-J+int, which has a lower BLEU score, and GPT-J+ctr, which has lower diversity scores. In terms of the comparison of the two knowledge selection methods, GPT-J+ctr has higher BLEU and BLEURT scores, while GPT+J+cut has better diversity scores. One possible reason is that the *Center* approach considers a small context window of at most three recipe steps. It helps the model to focus on the most relevant information; but at the same time, it reduces the total amount of knowledge the model can rely on to instruct, making the responses less diverse than other models. Finally, incorporating the intent information does not show improvement in terms of BLEU and BLEURT, however, this approach does help increase the diversity of the generated responses. We also conduct additional experiments on the CookDial dataset in Appendix A.

Human Evaluation. In order to further understand the behaviors of the models, we ask three annotators to analyze the system outputs and manually categorize their errors. We also ask the annotators to rate the *correctness* of each system response, using a 5-point Likert scale (i.e., 5-completely correct; 4-mostly correct, has minor and acceptable errors; 3-borderline, has both correct and incorrect information, nothing outweighs the other; 2-mostly incorrect, but still has correct information; 1-completely incorrect). The inter-annotator agreements measured by the nominal and ordinal Krippendorff’s alpha (Krippendorff, 2004) for the error categorization and correctness rating are 0.43 and

<p>Recipe:</p> <p>Step 1: ...</p> <p>Step 2: Make your batter and breading. To make your batter, mix 0.5 cups (120 mL) of flour, 0.5 cups (120 mL) of milk, and one egg in a bowl. In a separate shallow bowl or plate, put 3 tablespoons (44 mL), 1.5 cups (350 mL) of breadcrumbs, and 1.5 teaspoons (7.4 mL) of salt and mix them thoroughly. This is your breading. If you want to add other spices to your batter put them in the dry mixture on the plate. Some great spices to use would be garlic, paprika, or chili powder. If you want to add a liquid ingredient like beer or lime juice, add it the mixture in the bowl with the raw egg.</p> <p>Step 3: Slice your jalapeno ...</p> <p>Step 4: Dip the jalapeno peppers ...</p>	<p>Conversation history:</p> <p>...</p> <p>[user] Cool! I'm ready to make the batter.</p> <p>[system] In a bowl mix together 0.5 cups of flour, 0.5 cups of milk and one egg.</p> <p>...</p> <p>[system] Some may choose to add ingredients such as beer or lime juice for some added flavour. Would you like to do this?</p> <p>[user] No. Let's move on then.</p> <p>[system] Next, in a separate dish to your batter, mix together 1.5 cups of breadcrumbs and 1.5 teaspoons of salt.</p> <p>[user] Mixture is ready.</p> <hr/> <p>Gold response: To the dry mixture you can add in any seasonings of your choice.</p> <hr/> <p>GPT-J: Yes, mix them thoroughly. → <i>wrong order</i></p> <p>GPT-J+ctr: Yes, the mixture is ready. → <i>irrelevant</i></p> <p>GPT-J+int: Next, you can dip the peppers in the batter mixture. → <i>wrong order</i></p> <p>ChatGPT: Great! Now you can slice your jalapenos. If you're just making rings, slice them across the width of the pepper. If you're filling them, make an incision on one side and remove the seeds before adding your stuffing mixture. Don't use more than a tablespoon or the pepper might explode while frying. → <i>wrong order</i></p>
---	---

Figure 3: Example errors generated by each of the four systems. The system utterance in the history and its corresponding recipe micro-step have the same color. Text in red indicates the error types of the system responses.

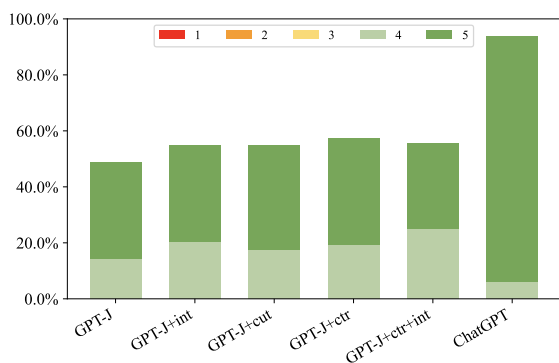


Figure 4: Percentages of model's outputs with high correctness ratings of 4 and 5. All GPT-J (+int/+cut/+ctr) variants that use intent and/or state tracking information generate better outputs than the vanilla GPT-J.

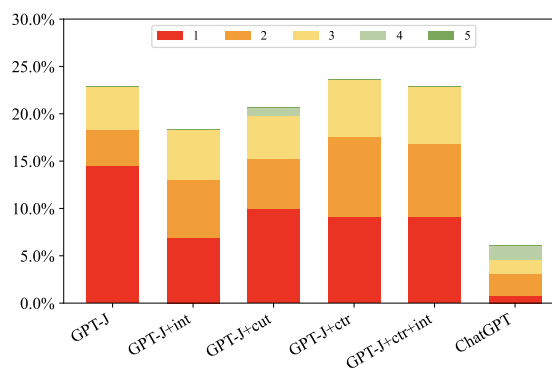


Figure 5: Correctness ratings of the wrong-order outputs from each model. All GPT-J (+int/+cut/+ctr) variants make fewer serious wrong-order errors (correctness of 1) than the vanilla GPT-J. Errors are sometimes not seen as severe in ChatGPT's longer responses than GPT-J's.

0.62, respectively. More details about how to aggregate three annotations are in Appendix C). As shown in Table 7 and Figure 4, all GPT-J variants that incorporate the intent and/or state information have fewer errors and more responses rated with 4 and 5 for correctness than the base model. Even though automatic metrics do not show a clear difference, human evaluation reveals that GPT-J+int has fewer (22.9%→18.3%) wrong-order errors compared to the base model and is also the model with the least number of this type of error. On the contrary, using *Center* approach (i.e., GPT-J+ctr and GPT-J+ctr+int) in grounded recipe selection does not have much impact on reducing the number of wrong-order responses, despite the fact that it helps improve BLEU and BLEURT scores. In addition, all +int/+ctr variants of GPT-J have fewer responses

with severe wrong-order errors (*correctness* of 1) than the base model.

Finally, we also analyze the errors in the outputs of ChatGPT. Overall, ChatGPT performs extremely well in this task with only 10.7% of the outputs being erroneous. The outputs of ChatGPT are notably longer than other systems since ChatGPT tends to instruct multiple recipe steps in one utterance or utilize knowledge outside the given recipe. As shown in Table 7, wrong-order instruction is still the most common error for ChatGPT. One scenario where ChatGPT makes mistakes in terms of the ordering is when the recipe step contains multiple micro-steps (see an example in Figure 3). It indicates that there are still many challenges that remain unsolved in the cooking instruction dialogue task.

6 Related Work

The task of recipe-grounded conversation is close to the Conversational Question Answering (CQA) task. In CQA, given a reference text, the system needs to engage in a multi-turn conversation to answer questions from users. Compared to single-turn question answering, CQA raises new challenges (e.g., co-reference resolution, contextual reasoning) due to the dependency between Question Answering turns. There exist multiple datasets in this area, such as CoQA (Reddy et al., 2019), QuAC (Choi et al., 2018), DoQA (Campos et al., 2020), and ShARC (Saeidi et al., 2018). There are several differences between Instructional Dialogue and Conversational Question Answering. Firstly, in the dialogue setting, the message from the system can also be a question, such as verification. Secondly, while the goal of CQA is seeking information, Instructional Dialogue focuses on supporting users to complete a procedure; therefore, there is additional order-related relationship between the system’s responses and the instructions that needs to be managed by the dialog agent.

Recent work has investigated issues that arise in chatbots based on large language models. For instance, they are known to sometimes generate toxic language (Baheti et al., 2021; Deng et al., 2022), make factual errors in their statements (Honovich et al., 2021; Dziri et al., 2022; Rashkin et al., 2021), and be overly confident (Mielke et al., 2022). In this work, we focus on addressing a specific problem related to instruction-grounded dialogue, which is presenting information in the wrong order to a user.

A small amount of prior work (Jiang et al., 2022; Strathearn and Gkatzia, 2022) has started to explore the problem of recipe-grounded conversation, which makes these papers the two closest to ours. Both of these papers focused primarily on dataset creation. Jiang et al. (2022) included experiments on response generation, but as their focus was on building a new dataset, they did not conduct extensive experiments or perform a human evaluation of their system’s outputs. They did propose baselines and evaluate the tasks of User Question Understanding and Agent Action Frame Prediction, which are similar to our User Intent Detection and Instruction State Tracking. Although these tasks have similar goals, our work is different in the sense that we focus on tackling the problems in the low-resource setting, by transferring knowledge from

existing dialogue corpora such as MultiWOZ. Finally, besides providing additional recipe-grounded conversations as in these two prior works, our main focuses are on analyzing challenges of current large language models (i.e., GPT-J and ChatGPT) on this task and addressing the specific challenge of instruction ordering.

7 Conclusion

In this paper, we have proposed to explore two additional subtasks, namely User Intent Detection and Instructional State Tracking, to mitigate the problem of incorrect instruction order in Instructional Dialogue. We analyze these two auxiliary subtasks with different methods in low-resource settings. Even though the performance of the modules for the two subtasks is still low, experiment results show that incorporating the user intent or state information does help to mitigate the wrong-order instructions issue, with the intent information having a greater impact. However, combining the two pieces of information does not lead to improvement over using each individual one of them alone. Therefore, we believe that further research for the two subtasks is still needed, and also more effective ways of incorporating the information into the Response Generation module need to be investigated. Finally, we release 🍳 ChattyChef, a new cooking instructional dataset, to promote future research in this direction.

Limitations

In this work, we have only analyzed the common errors of two models (i.e., GPT-J and ChatGPT) in the Instructional Dialogue task. One open question is whether other GPT-based models or models with other architectures (e.g., encoder-decoder models) also have the same issue in this task. Our work and dataset are also limited to the English language.

Ethical Considerations

To collect recipe-grounded conversations we hired crowd workers using the Prolific platform.⁶ The study was conducted with the approval of our local IRB. The compensation was derived based on Prolific’s payment principles. We estimate the hourly pay for crowd workers was \$15.49 (details in Appendix D). Crowd workers were strictly asked not to write any offensive content or personal information.

⁶<https://www.prolific.co/>

Acknowledgments

We thank Yao Dou, Fan Bai as well as four anonymous reviewers for their helpful feedback on this work. We also thank Govind Ramesh, Grace Kim, Yimeng Jiang for their help with human evaluation. This research is supported in part by the NSF awards IIS-2112633 and IIS-2052498. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of NSF or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Ashutosh Baheti, Maarten Sap, Alan Ritter, and Mark Riedl. 2021. Just say no: Analyzing the stance of neural dialogue generation in offensive contexts. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. 2022. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. "O'Reilly Media, Inc."
- SRK Branavan, Luke Zettlemoyer, and Regina Barzilay. 2010. Reading between the lines: Learning to map high-level instructions to commands. In *Proceedings of the 48th annual meeting of the association for computational linguistics*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. [MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Jon Ander Campos, Arantxa Otegi, Aitor Soroa, Jan Derru, Mark Cieliebak, and Eneko Agirre. 2020. [DoQA - accessing domain-specific FAQs via conversational QA](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7302–7314, Online. Association for Computational Linguistics.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. [QuAC: Question answering in context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184, Brussels, Belgium. Association for Computational Linguistics.
- Richárd Csáky, Patrik Purgai, and Gábor Recski. 2019. Improving neural conversational models with entropy-based data filtering. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Jiawen Deng, Jingyan Zhou, Hao Sun, Fei Mi, and Minlie Huang. 2022. Cold: A benchmark for chinese offensive language detection. *arXiv preprint arXiv:2201.06025*.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of Wikipedia: Knowledge-powered conversational agents. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Nouha Dziri, Ehsan Kamaloo, Sivan Milton, Osmar Zaniane, Mo Yu, Edoardo M Ponti, and Siva Reddy. 2022. Faithdial: A faithful benchmark for information-seeking dialogue. *arXiv preprint arXiv:2204.10757*.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. Multiwoz 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking baselines. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 422–428.
- Or Honovich, Leshem Choshen, Roei Aharoni, Ella Neeman, Idan Szpektor, and Omri Abend. 2021. Q2:: Evaluating factual consistency in knowledge-grounded dialogues via question generation and question answering. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Yiwei Jiang, Klim Zaporozhets, Johannes Deleu, Thomas Demeester, and Chris Develder. 2022. Cookdial: a dataset for task-oriented dialogs grounded in procedural documents. *Applied Intelligence*, pages 1–19.
- Klaus Krippendorff. 2004. Reliability in content analysis: Some common misconceptions and recommendations. *Human communication research*, 30(3):411–433.
- Chaitanya Kulkarni, Wei Xu, Alan Ritter, and Raghu Machiraju. 2018. An annotated corpus for machine reading of instructions in wet lab protocols. In *Proceedings of NAACL-HLT*.
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning of negotiation dialogues. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.

- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Chia-Wei Liu, Ryan Lowe, Iulian Vlad Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Sabrina J Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. 2022. Reducing conversational agents’ overconfidence through linguistic calibration. *Transactions of the Association for Computational Linguistics*, 10.
- Alexander Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiasen Lu, Devi Parikh, and Jason Weston. 2017. **ParLAI: A dialog research software platform**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 79–84, Copenhagen, Denmark. Association for Computational Linguistics.
- Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2021. Gpt3-to-plan: Extracting plans from text using gpt-3. *FinPlan 2021*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of ACL*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Baolin Peng, Michel Galley, Pengcheng He, Chris Brockett, Lars Liden, Elnaz Nouri, Zhou Yu, Bill Dolan, and Jianfeng Gao. 2022. Godel: Large-scale pre-training for goal-directed dialog. *arXiv preprint arXiv:2206.11309*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Hannah Rashkin, David Reitter, Gaurav Singh Tomar, and Dipanjan Das. 2021. Increasing faithfulness in knowledge-grounded dialogue with controllable features. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 583–593.
- Marzieh Saeidi, Max Bartolo, Patrick Lewis, Sameer Singh, Tim Rocktäschel, Mike Sheldon, Guillaume Bouchard, and Sebastian Riedel. 2018. Interpretation of natural language rules in conversational machine reading. *arXiv preprint arXiv:1809.01494*.
- Joao Sedoc, Daphne Ippolito, Arun Kirubakaran, Jai Thirani, Lyle Ungar, and Chris Callison-Burch. 2019. Chateval: A tool for chatbot evaluation. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics (demonstrations)*.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. **BLEURT: Learning robust metrics for text generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, et al. 2022. Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage. *arXiv preprint arXiv:2208.03188*.
- Carl Strathearn and Dimitra Gkatzia. 2022. **Task2Dial: A novel task and dataset for commonsense-enhanced task-based dialogue grounded in documents**. In *Proceedings of the Second DialDoc Workshop on Document-grounded Dialogue and Conversational Question Answering*, pages 187–196, Dublin, Ireland. Association for Computational Linguistics.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.

Jason D Williams, Antoine Raux, and Matthew Henderson. 2016. The dialog state tracking challenge series: A review. *Dialogue & Discourse*, 7(3).

Te-Lin Wu, Alex Spangher, Pegah Alipoormolabashi, Marjorie Freedman, Ralph Weischedel, and Nanyun Peng. 2022. [Understanding multimodal procedural knowledge by sequencing multimodal instructional manuals](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4525–4542, Dublin, Ireland. Association for Computational Linguistics.

Li Zhang, Qing Lyu, and Chris Callison-Burch. 2020. [Reasoning about goals, steps, and temporal ordering with WikiHow](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4630–4639, Online. Association for Computational Linguistics.

Jeffrey Zhao, Raghav Gupta, Yuan Cao, Dian Yu, Mingqiu Wang, Harrison Lee, Abhinav Rastogi, Izhak Shafran, and Yonghui Wu. 2022. Description-driven task-oriented dialog modeling. *arXiv preprint arXiv:2201.08904*.

A Experiments on CookDial dataset

A.1 Incorporation of Instruction State and User Intent information

In this section, we explore the performance of models using Instruction State and Intent Information on the CookDial dataset. We fine-tuned the models adopting the same approaches as in §5 and using the same set of automatic metrics to evaluate. Instead of extracting silver labels, We use the gold Instruction State (i.e., "tracker_completed_step" in CookDial) and User Intent information from the CookDial dataset. The performance of all models is demonstrated in Table 8.

From the table, we can see that the model performance on CookDial is much higher than in our dataset, this is due to the more straightforward instruction scenarios and higher lexical similarity between the grounded recipe and system utterances in CookDial (as discussed in §2.2). In addition, similar behaviors of the models fine-tuned on CookDial and ChattyChef can also be observed here - incorporating the Instruction State information helps to improve BLEU and BLEURT scores, while incorporating the intent information helps the models generate more diverse responses.

A.2 Transfer learning from CookDial to ChattyChef

To explore whether the information learned from CookDial is helpful for our dataset, we keep fine-tuning the model from Table 8 with the corresponding settings on our dataset. As shown in Table 9, transfer learning does not show a clear improvement in terms of BLEU and BLEURT, it even hurts the performance of the model in many cases. The big difference between the two datasets (as discussed in §2.2) may be the reason why transfer learning is ineffective here. However, transfer learning also has the merit of making the system outputs more diverse, especially for GPT-J+cut and GPT-J+ctr.

Model	BLEU	BLEURT	Length	Diversity
GPT-J	34.7	65.0	11.5	13.0 / 44.3
GPT-J+int	30.9	62.5	10.4	13.8 / 45.3
GPT-J+cut	35.6	65.2	11.7	12.8 / 43.9
GPT-J+ctr	35.3	65.4	11.3	13.3 / 44.6
GPT-J+ctr+int	32.0	63.2	10.6	13.7 / 44.5

Table 8: Performance on the test set of CookDial dataset.

Model	BLEU	BLEURT	Length	Diversity	
GPT-J	3.7 (-0.4)	45.3 (+0.6)	9.7 (-1.4)	11.0 (+1.1)	39.9 (+2.0)
GPT-J+int	4.3 (+0.4)	45.5 (+0.5)	10.4 0.4	11.2 (+0.8)	40.8 (+2.3)
GPT-J+cut	4.0 (-0.3)	45.2 (0.0)	9.9 (-1.0)	11.8 (+1.9)	44.6 (+5.9)
GPT-J+ctr	4.1 (-0.6)	45.5 (-0.4)	9.6 (-2.1)	11.6 (+2.3)	42.4 (+5.8)
GPT-J+ctr+int	4.2 (0.0)	45.7 (+0.6)	9.8 (-0.5)	11.5 (+0.7)	41.9 (+2.6)

Table 9: Transfer learning performance of models on the test set of ChattyChef. The number below in each cell indicates the change to the model fine-tuned with the same setting on ChattyChef only.

B Implementation Details

For all experiments, we train models across 4 A40 GPUs (48GB each). The total GPU hours for training a GPT-J model in the Response Generation task is about 5.3 hours, and the total GPU hours for training the T5 models in the User Intent Detection task is about 4 hours.

B.1 User Intent Detection

The details about user intents and their description are reported in Table 10. Examples of input and out prompts to the T5 model are demonstrated in Table 12.

For all experiments which use T5 model, we set the maximum sequence length to 1028 and the number of training epochs to 30, we stop the training process if the perplexity of the model on the validation set does not improve after 5 epochs. We use AdamW to optimize and consider the initial learning rate $\in \{1e-5, 5e-5, 1e-4\}$. For the in-context setting, the maximum length of the input sequence is set to 1984.

For all the models, we employ beam search with a beam size of 5 for decoding. We select the model checkpoint that produces the lowest perplexity on the validation set and then apply the selected one to the test set.

B.2 Instruction State Tracking

The Instruction State Tracking Algorithm is described in Algorithm 1). In order to produce the system utterance - recipe alignment, similarity scores between the most recent system response to all the recipe steps are computed (Line 3 - 7 in Algorithm 1). After that, by comparing the similarity

Intents	Descriptions
greeting	greeting
req_temperature	ask about the cooking temperature
thank	thank
req_instruction	ask for instructions
confirm	confirm the current stage
req_repeat	ask to repeat the last information
negate	negate
req_amount	ask about the amount information
req_ingredient	ask about the ingredients
req_is_recipe_finished	ask whether the recipe is finished
req_tool	ask about the cooking tool
req_duration	ask about the cooking duration
affirm	affirm
goodbye	goodbye
req_substitute	ask for tool or ingredient substitutions
req_confirmation	ask for verification
req_description	ask for the description
req_explanation	ask to explain the reason or explain in more detail
other	other intent

Table 10: Descriptions of user intents in ChattyChef

Algorithm 1: Instruction State Tracking

Input: Recipe $\mathcal{R} = \{R_1, R_2, \dots, R_{n_r}\}$
Most recent system utterance: U_t^{sys}
Previous instruction state: T_{t-1} Threshold parameters $0 < \alpha_1 \leq \alpha_2 < 1$
Scoring function f

Output: Instruction State T_t

- 1 Initialize $score[i] = 0 \forall i = 1, 2, \dots, n_r$
- 2 Initialize $current_state = T_{t-1}$
/* compute similarity scores between the most recent system utterance and all recipe steps */
- 3 **for** $i = 1, 2, \dots, n_r$ **do**
- 4 $micro_steps \leftarrow sentence_tokenize(R_i)$
- 5 $score[i] = \max_{r \in micro_steps} f(U_t^{sys}, r)$
- 6 $best_state \leftarrow \arg \max(score)$
- 7 $max_score \leftarrow score[best_state]$
- 8 **if** $(best_state == current_state + 1$ and $max_score > \alpha_1)$ or $(max_score > \alpha_2)$ **then**
- 9 $current_state \leftarrow best_state$
- 10 $T_t \leftarrow current_state$

score to thresholds (i.e., α_1 and α_2), the algorithm decides whether the current system utterance is aligned to a new state (i.e., a new recipe step – line 9 in Algorithm 1) or is aligned with the previous state.

For the Sentence Embedding approach, we use `sentence-transformers/paraphrase-MiniLM-L6-v2` (Reimers and Gurevych, 2019) to compute the sentence embeddings of system responses and recipe steps. We use NLTK (Bird et al., 2009) to perform the word and sentence tokenization. About the threshold, we set α_1, α_2

equal to 0.2 and 0.3, respectively, for the Word Matching approach. For the sentence embedding approach, we set α_1 equal to 0.5 and α_2 equal to 0.6. The thresholds are chosen based on the accuracy of the validation set.

B.3 Response Generation

An example of the input to the GPT-J model is illustrated in Table 13.

To fine-tune all the GPT-J models, We set the maximum sequence length to 1280 and the number of training epochs to 3. We use AdamW to optimize and set the initial learning equal to $1e-5$, except for transfer learning experiments with CookDial, in which we use the learning rate of $5e-6$. We employ beam search with a beam size of 5 for decoding. We select the model checkpoint that produces the lowest perplexity on the validation set and then apply the selected one to the test set.

C Human Evaluation

In this section, we discuss the way to aggregate annotations from annotators. For the error categorization experiment, the final decision of an example is reached if all three annotators have the same annotation. When only two annotators have the same annotation, a fourth one will join and decide whether to agree with the majority. In all other situations, a discussion between annotators is held, and the final decision is based on majority voting. For the correctness rating experiment, the rating of each example is the average score of the three annotators. In cases when only two annotators rate an example as completely correct or a rating of 5, but the third one detects an error and marks the example as incorrect (i.e., belongs to one of the four error types), if the final decision from the error categorization is also incorrect, the correctness of this example is the rating of the third annotator. The same rule applies to the opposite situation, i.e., only two annotators rate an example as completely incorrect, and the third one thinks it is correct.

	Paired	Self	M-User
Avg time per turn (min)	2.35	1.77	1.60
Avg cost per turn (\$)	0.72	0.42	0.35
#Turns per recipe step	2.04	2.30	2.13
#Dialogues	86	160	21

Table 11: Statistics of the data collection methods.

Input: 0:book a table at a restaurant 1:book a hotel to stay in 2:search for police station 3:search for places to wine and dine 4:search for a medical facility or a doctor 5:search for a bus 6:search for a hotel to stay in 7:search for trains that take you places 8:search for places to see for leisure 9:book taxis to travel between places 10:book train tickets [user] I am looking for a hotel called the alpha-milton guest house. [system] Sure! I've located the guesthouse, it is located in the north area. Would you like me to book you a room? [user] No thank you but I do need the address please?

Output: [intents] 6

Input: 0:negate 1:confirm the current stage 2:ask to repeat the last information 3:ask about cooking duration 4:ask for verification 5:thank 6:ask to explain the reason or explain in more detail 7:ask about the cooking temperature 8:affirm 9:greeting 10:ask for the description 11:ask about the amount information 12:goodbye 13:ask whether the recipe is finished 14:ask for instructions 15:ask about the ingredients 16:ask about the cooking tool 17:ask for tool or ingredient substitutions 18:other intent [user] Yes, what do I need? [system] Russet potatoes, or other high starch potatoes [user] What is the first step? [system] Wash and peel the potatoes, use cold water when washing [user] What do I use to peel the potatoes?

Output: [intents] 16

Table 12: Examples of the input and output of the User Intent Detection model. The top example is from MultiWOZ 2.2 dataset, and the bottom one is from Our cooking instruction dataset). **Red:** Indexed intent descriptions. **Blue:** Conversation history.

[system] Would you like to learn how to make hash browns? [user] Yes, what do I need? [system] Russet potatoes, or other high starch potatoes [user] What is the first step? [system] Wash and peel the potatoes, use cold water when washing [user] What do I use to peel the potatoes? <|Knowledge|> - Peel the potatoes. Wash the potatoes well in cold water, then peel using a small knife or a vegetable peeler. Russet potatoes, or other potatoes with a high starch content, work best for hash browns. - Shred the potatoes. Line a bowl with a clean dishtowel, then shred the potatoes directly into the towel-lined bowl, using a cheese grater. - Squeeze out the moisture. You must squeeze out as much moisture as possible from the shredded potatoes. This is the most important step in achieving crispy (rather than mushy) hash browns. To do this, gather the corners of the dishtowel containing the shredded potatoes and twist the neck until you form a tight package. Continue twisting the cloth and squishing the potato in your fist until you've squeezed as much liquid as you can from the potato. Alternatively, you can try squeezing the moisture from the potatoes using a potato ricer. You do not need to force the potatoes through the ricer, simply use it to press out the moisture. - Heat the skillet. Heat a large skillet pan (preferably cast iron) over a medium-high heat. Add the butter to the pan and allow to melt. Once the butter has melted, add the dry, shredded potatoes to the pan and toss to coat with butter. Season with salt and pepper. - Cook the hash browns. Once the potato has been coated with butter, flatten it using a spatula to maximize contact with the hot pan. It should be no more than 1/2 an inch thick. Cook for 3-4 minutes on the first side, flip, then cook for 2-3 minutes on the other side. The hash brown potatoes are ready when each side is crisp and golden brown. - Serve. Slide the hash brown from the pan, or lift using a large spatula. Cut it into halves or quarters, if necessary. Serve on its own, with hot sauce or ketchup, or alongside bacon and eggs for a top notch breakfast. [user] want to: ask about the cooking tool. => [system] you can use a vegetable peeler, or a small knife

Table 13: An example of the prompt to the Response Generation model. **Blue:** Conversation history. **Brown:** Grounded recipe. **Green:** Intent description prompt. **Red:** Output of the model.

D Dataset construction

D.1 Collection Strategies

Even though employing two workers using different interfaces has advantages, we see that pairing workers on a task is inefficient. In particular, for each conversation, one worker would need to wait for a long time until the partner joined the task. Moreover, in some cases, some workers are uncooperative. For example, during the chat, one worker may spend too much time sending his messages or even quit the task, which will have a bad impact on his partner (Choi et al., 2018; Reddy et al., 2019). As a result, we study three different collection strategies as follows.

Paired conversations (Paired): Two workers are required for each conversation; one acts as the

agent, and the other act as the user. After two workers got paired, they would be assigned to the same cooking task; however, only the agent has access to the recipe, and all the user knows about this task is the title (i.e., what to cook).

Self-chat (Self): In this mode, only one worker is needed for each conversation. The worker will play both roles (i.e., Agent and User).

Model-User (M-User): One worker is assigned for each conversation in this mode. However, unlike Self-chat, when the worker plays the User role, he would be provided with candidate responses from a model, and he could either pick one and edit it or enter his own words.

All conversations were collected through the ParlAI API (Miller et al., 2017). The participants for this paper were recruited using Prolific

(www.prolific.co). We restricted the task to workers whose first language is English and with more than 100 submissions with at least a 99% approval rate. Crowd workers were not informed of the specific details of the dataset. However, they consented to have their responses used for this purpose through the Prolific Participation Agreement. The statistic about each collection method is reported in Table 11.

D.2 Collection Interfaces

See Figure 7 for the screenshot of our crowdsourcing interface.

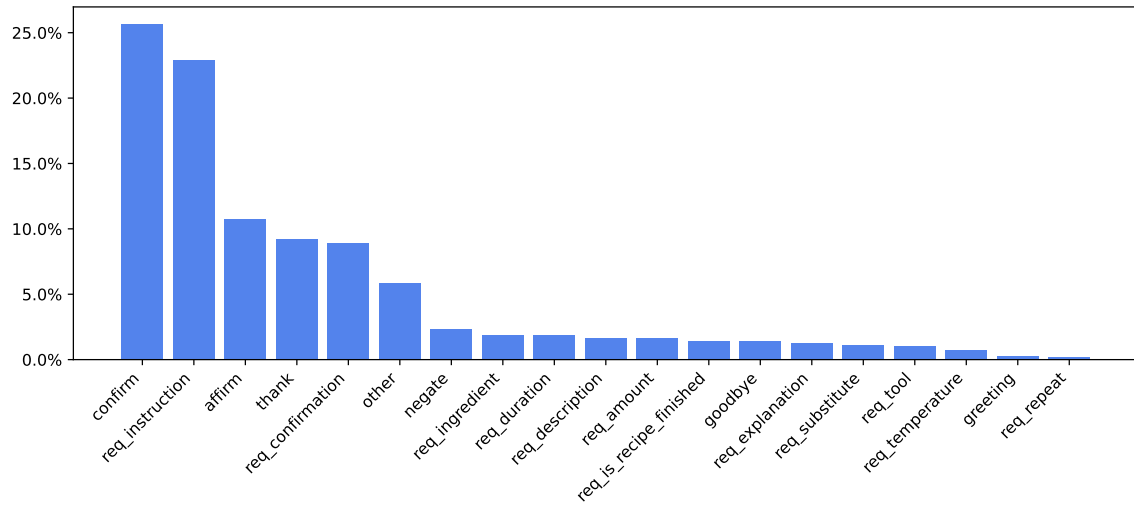



Figure 6: Statistic of user intents over 91 conversations.


How to Dress a Salad

Recipe


- Pour the dressing into the bottom of a mixing bowl.** Swirl the bowl so the dressing lines the bottom and sides. A lighter dressing mixes easily and this way you know how much dressing you're using right away. Start with a small amount of dressing. It's easier to add more, but difficult to remove. This works best with lighter dressings, like Italian or vinaigrettes.



- Put your salad greens in the bowl before any toppings.** Layer your ingredients so the heavier, sturdier ingredients are on the bottom of the bowl. The greens will absorb most of the dressing's flavor and won't drown out your other toppings. Leave out crunchy toppings that could get soggy, like croutons or bacon bits.



- Stir the salad gently with your hands.** Wash your hands before handling your salad. Carefully rotate the salad between your fingers so your dressing coats the salad evenly. You want your greens to have a light coating rather than being drenched. Tongs will bruise and brown your greens, making them less crisp. Wear gloves if you want to avoid getting your hands messy.



Agent: Hi! I'm going to teach you to dress a salad today. Are you ready to get started?

User: Cool! What's the first step?

Agent: First you need prepare a bowl for the dressing.

User: Sure, I have that bowl, what's next?

You are playing the role of **Agent**. It's YOUR TURN to chat and send utterance NOW.

Identify the type of the last utterance by your partner, and select your next action (click ? for examples):

- Which of these best describes your partner's last utterance:

 ask to move on ask question answer question other
- Select your next action:

 teach new step ask question answer question other
- Add the evidence of your utterance.

(1) Highlight SHORT text span(s) in the left recipe as evidence of your next utterance. (2) Then click "Add" to put evidence(s). Or click "Clean" to delete all. ?

Add

Clean

Hide

No Evidence/Answer from recipe
- Then enter your chatting responses:

Send

Do not refresh your page. You've earned \$0 bonus. Click ? to learn how to earn more bonus.

Figure 7: A screenshot of our crowdsourcing interface.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section Limitations
- A2. Did you discuss any potential risks of your work?
Not applicable. Left blank.
- A3. Do the abstract and introduction summarize the paper's main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 2, 3, 4, 5

- B1. Did you cite the creators of artifacts you used?
Section 2, 3, 5, Appendix B
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Section Ethical Considerations
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 2, Appendix C
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 2, 3, 4

C Did you run computational experiments?

Section 3, 4, 5, Appendix B

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix B

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- ✓ C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Appendix B
 - ✓ C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 3, 4, 5
 - ✓ C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Appendix B
- D** ✓ **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 2, Appendix C
- ✓ D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Section 2, Appendix C
 - ✓ D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Appendix C
 - ✓ D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Section 2, Appendix C
 - ✓ D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Section Ethical Considerations
 - ✓ D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Appendix C