

CoPrUS: Consistency Preserving Utterance Synthesis towards more realistic benchmark dialogues

Sebastian Steindl
Ostbayerische Technische
Hochschule Amberg-Weiden
Germany
s.steindl@oth-aw.de

Ulrich Schäfer
Ostbayerische Technische
Hochschule Amberg-Weiden
Germany
u.schaefer@oth-aw.de

Bernd Ludwig
University Regensburg
Germany
bernd.ludwig@ur.de

Abstract

Large-scale Wizard-Of-Oz dialogue datasets have enabled the training of deep learning-based dialogue systems. While they are successful as benchmark datasets, they lack certain types of utterances, which would make them more realistic. In this work, we investigate the creation of synthetic communication errors in an automatic pipeline. Based on linguistic theory, we propose and follow a simple error taxonomy. We focus on three types of miscommunications that could happen in real-world dialogues but are underrepresented in the benchmark dataset: misunderstandings, non-understandings and vaguely related questions. Our two-step approach uses a state-of-the-art Large Language Model (LLM) to first create the error and secondly the repairing utterance. We perform Language Model-based evaluation to ensure the quality of the generated utterances. We apply the method to the MultiWOZ dataset and evaluate it both qualitatively and empirically as well as with human judges. Our results indicate that current LLMs can aid in adding post-hoc miscommunications to benchmark datasets as a form of data augmentation. We publish the resulting dataset¹, in which nearly 1900 dialogues have been modified, as CoPrUS-MultiWOZ to facilitate future work on dialogue systems.

1 Introduction

Task-Oriented Dialogue (TOD) systems are a special case of chatbots that help the user achieve certain tasks, such as booking a hotel, through conversation. This requires the TOD system to understand the user’s utterance, decide on a policy, and utter a response to the user (He et al., 2022).

With the great advances in model architectures, these systems nowadays utilize deep learning and transformer models instead of being rule-based (Su

et al., 2022; Bang et al., 2023; Zhao et al., 2022). This, however, requires large amounts of labeled training data. The gold standard for collecting these dialogues is the Wizard-Of-Oz (WOZ) technique (Kelley, 1984), in which two humans simulate a human-machine interaction. One widespread benchmark dataset that was created with Wizard-Of-Oz is MultiWOZ (Budzianowski et al., 2018).

Even though the dialogues were created by two humans, we find them *sterile* and unrealistic. More precisely, they are very close to what would be considered the idealistic *happy path*, where everything works as expected without sidetracking. Previous work has also identified multiple caveats within the MultiWOZ dataset, which we will discuss in section 3. Nonetheless, to make dialogue systems generalize well into real-world usage, it is of high priority that the training data is realistic, since the system will only be able to adequately answer utterances that are close to the training distribution.

The current state-of-the-art Large Language Models (LLMs) are known for their strong capabilities in Natural Language Generation (NLG), instruction-following and in-context learning (Ouyang et al., 2022; Sanh et al., 2022). In this study, we thus investigate whether LLMs can aid in making task-oriented dialogue datasets more realistic by adding synthetic miscommunications. We see this as a form of data augmentation that is not focused on the quantity of the data, but on its quality in the sense of one of its properties, i.e., lack of miscommunications.

We specifically focus on three types of miscommunication turns and their repairing, that are lacking in the benchmark: non-understanding, misunderstanding, and vaguely related questions. These should not change the goal or consistency of the dialogue, since their only motivation is clarification. The inclusion of these types of turns can make the dialogues more realistic and improve the ease of use for the dialogue system. This could allow for

¹https://github.com/sebastian-steindl/CoPrUS_data

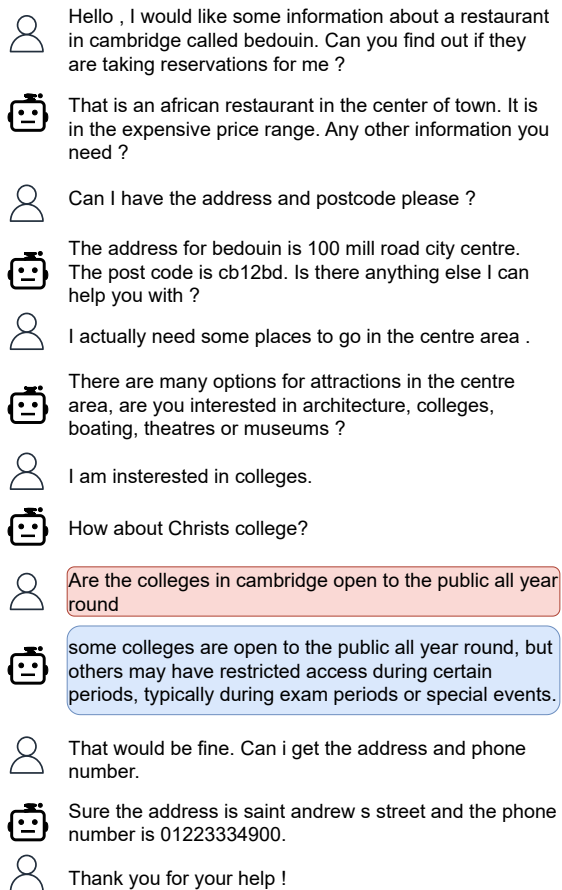


Figure 1: Example MultiWOZ dialogue after application of CoPrUS. Highlighted red is the miscommunication utterance, in this case, a vaguely related question, and in blue the repairing attempt by the system.

a higher level of technology acceptance (Mitzner et al., 2010) and reduce frustration in all user age groups (Van Der Goot and Pilgrim, 2020).

We thus propose consistency-preserving utterance synthesis (CoPrUS) to make benchmark dialogues more realistic post-hoc. We evaluate our method empirically and with a qualitative analysis. Moreover, we publish the dataset that results from applying CoPrUS to the MultiWOZ 2.1 (Eric et al., 2020) dataset in the unified format of Zhu et al. (2022).

2 Background and Related Work

In the following, we will first give an overview of problems with the used TOD benchmark dataset and TOD systems. Then, we will go into the differences between human-human and human-machine conversations and finally talk about taxonomies for conversational errors.

2.1 Task-Oriented Dialogue Systems

Initially, TOD systems addressed each subtask (Natural Language Understanding, Policy Planning, NLG) with a specific (rule-based) model (Young et al., 2013). Yet the publication of larger datasets, such as MultiWOZ (Budzianowski et al., 2018), in combination with general progress in deep learning, led to the usage of (end-to-end) deep learning approaches, e.g., as proposed by Lin et al. (2020); Peng et al. (2021); He et al. (2022); Sun et al. (2023); Bang et al. (2023). Lately, LLMs have also been investigated as zero-shot TOD systems (Ulmer et al., 2024; Li et al., 2024; Chung et al., 2023).

With this change towards learning the dialogue system from conversational data instead of using rule-based approaches, it becomes even more important to have dialogue flows beyond the happy path. If miscommunications are inherent in the training data, the model can learn to act accordingly and initialize repairing.

2.2 Human-Human and Human-Machine Dialogue Differences

Skantze (2005) shows that human-human conversations behave differently than human-machine dialogues. They attribute this to the fact that the human is aware of talking to a machine and that the machine has limited language capabilities. Furthermore, Nass and Moon (2000) argue that the way a human talks in a human-machine dialogue is also dependent on how human-like the machine acts. Still, research has assumed that humans apply social rules from human-human conversation to human-computer dialogues as well (Nass and Moon, 2000). While traditionally, a human could easily identify if he was talking to a human or a machine, e.g., due to Automatic Speech Recognition (ASR) errors or unnatural Text-To-Speech (TTS) (Nass and Gong, 1999), this has become much more difficult with recent technical advances.

Additionally, if we focus on strictly written dialogues, ASR and TTS cannot give away that it is a human-machine interaction. Given the remarkable NLG and Natural Language Understanding (NLU) capabilities of LLMs, we believe that even today, humans will not always be able to tell if they are talking to a machine. The implication of this is that in the near future, users will expect the ability to have a human-machine conversation in the same way that they have human-human dialogues.

2.3 Taxonomy for Conversational Errors

In the past, researchers have proposed various taxonomies for conversational errors. An error, in that sense, is a deviation from how a conversation should be. This quality of conversation has been studied by Grice (1975), who derives from these the maxims of conversation. It follows that an error in a conversation is a deviation from those maxims.

Based on this, Paek (2003) derives a general taxonomy for conversations with four levels, in which each level has to be achieved to reach the next. The taxonomy follows the assumption that conversational partners need to establish the mutual belief that they have been understood by one another. Paek (2003) discern between four hierarchical levels (in ascending order): *channel*, *signal*, *intention* and *conversation*. Errors at the lower levels will lead to uncertainty and *repairing* attempts. Ideally, a conversation would always be at the highest level, and no error would occur. Alas, due to the complexity of conversation, this is unrealistic.

Research has already adapted these ideas to dialogue systems and thus derived taxonomies of errors. Mostly, they focus on errors by the dialogue system, to be able to explicitly avoid them in the future, e.g., by special rules. Higashinaka et al. (2015) propose a first taxonomy for chat-oriented dialogue systems that the authors refined in later work (Higashinaka et al., 2019, 2021). For task-oriented dialogue systems, such taxonomies have been proposed earlier. Dybkjær et al. (1996) present a taxonomy for spoken WOZ dialogues leaning onto Grice (1975). Möller et al. (2007) also define an error taxonomy for spoken dialogue systems. They distinguish five categories of error: goal-level, task-level, command-level, concept-level, and recognition-level. Moreover, when chatbots should imitate a specific person, new types of errors need to be addressed, such as self-recognition or other-recognition errors (Mitsuda et al., 2022).

In summary, many studies have tried to create taxonomies, and much can be traced back to the effort in communication research by Grice (1975). We build on these understandings to focus on three error types: misunderstandings, non-understandings and vaguely related questions.

3 Unrealistic Wizard-Of-Oz Conversations

To the best of our knowledge, no research has been done on the details of the general differences between real conversations and WOZ conversations. However, we believe that the fact that they exist, at least in the MultiWOZ (Budzianowski et al., 2018) dataset, is evident and can be both deducted from the setting and seen in the resulting dialogues.

One can argue that for most workflows that can be treated as a TOD, there exist one or multiple *happy paths*. These are idealistic paths through the workflow, where no exception or error happens within the system. Commonly, it is not enough to ensure that a system is successful on such a happy path but also in its variations and deviations.

The MultiWOZ (Budzianowski et al., 2018) benchmark has, thanks to its size, enabled much research into deep learning-based dialogue modeling, becoming one of the most prevalent datasets in this field and successfully showing the potential of learning conversations from large amounts of data. However, previous work has started to look more closely into certain properties of the dataset that will be important for real-world usage and generalization. For example, Qian et al. (2021) found the MultiWOZ dialogues to be biased towards certain entities, e.g., with Cambridge making up 50% of the train destinations. Furthermore, the dialogues lack *turnback* utterances, in which users change their minds about slot values later in the conversation (Kim et al., 2022). Similarly, Yang et al. (2022) note that the dataset lacks negative feedback for when the dialogue system makes a mistake. Additionally, Yu et al. (2022) argue that in real life, TODs still contain utterances that would be regarded as chit-chat, but these are lacking in the TOD benchmark. What's more, Kim et al. (2023) identify a user familiarity bias in this benchmark dataset, arising because the crowd-workers were given a specific user goal and are assumed to know the dialogue system and its capabilities well enough. For users with a less defined goal, this can be problematic.

To these remarks, we want to add that within the MultiWOZ conversations, nearly no miscommunication is happening. On a side note, this is even more unrealistic for this specific dataset since every dialogue was created by not two but multiple crowd-workers.

We assume that these properties could be a result

Error Type	Explanation
MU	Misunderstanding: The user misunderstood a system utterance but continued the conversation because he was not aware of it. This awareness comes only later, leading the user to ask for clarification.
NU	Non-understanding: The user misunderstood a system utterance and did not continue the conversation because he became aware of the lack of understanding immediately. Thus, he directly asks for clarification of the last system utterance.
VQ	Vaguely related Question: The user asks the system a question that is part of the domain but not directly related to the dialogue goal or intended workflow. The system needs to answer and then continue with the dialogue.

Table 1: Proposed error taxonomy for CoPrUS.

of the collection setting, in which crowd-workers are paid to have a conversation. Since they know that they are participating in such a study, and they are additionally incentivized to deliver high-quality dialogues. These might be so good that they become unrealistic in that they only represent the happy path. We see this effect as akin to the social-desirability bias, which is well-known in social science research (Grimm, 2010). Thus, because of the somewhat clinical setting of being paid to have the WOZ conversation, we argue that users will not behave as they would in a real conversation but more in a way that they deem desirable. This is to be seen in parallel to the user familiarity bias (Kim et al., 2023).

4 Method

Our method utilizes an LLM and a two-step inference to add post-hoc miscommunications and repair utterances to task-oriented dialogues in an effort to make them more natural and realistic. We showcase our method on the MultiWOZ 2.1 (Eric et al., 2020) dataset. The following will introduce the taxonomy of miscommunications that we will assume for our work. Secondly, it will demonstrate the creation of miscommunication utterances with a two-step prompt and an automatic, LLM-based quality assurance.

4.1 CoPrUS Miscommunications

The general Gricean Cooperative Principle (GCP) states that a participant in a conversation should make his “conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange” (Grice, 1975).

Furthermore, when two speakers hold a conversation, they need to ensure that they are heard and

understood (Paek, 2003). Paek (2003) characterizes this mutual understanding by four levels: The lower levels *channel* and *signal* can be seen as achieved by the definition of the chatbot setting, since, e.g., not hearing someone well enough cannot happen. However, the higher levels *intention*, where semantics get interpreted, and *conversation*, where actions happen, are not as straight-forward, and we thus focus on these.

Based on the GCP, we define three types of miscommunications in this work: misunderstanding (MU), non-understanding (NU), and vaguely related question (VQ). Since we focus on textual chat data, the understanding on the channel and signal level is given by design. The three error types studied are therefore residing on the higher levels of intention and conversation.

The last type, VQ, might be considered a rather soft communication error, in the sense that most would not consider it as an error per se. However, as in previous work (cf. Sec. 2), we understand errors as deviations from a maxim. For example, if a participant deviates from the maxim of avoiding ambiguity (Grice, 1975), which falls under the category of manner and follows from the GCP, this can lead to a MU or NU. If, on the other hand, a speaker poses a VQ, he does not adhere to the maxim of making his contribution appropriate to the immediate needs of the current situation. This maxim is part of the relation category of the GCP.

Therefore, while VQs can be seen as a soft type of error, they are nonetheless relevant to investigate as part of our taxonomy. From a different perspective, the VQ error type might be regarded as a chit-chat utterance, which would align with the work of Yu et al. (2022). One could also study the similar case of completely unrelated questions, which would be a harder error. Yet, since this would be

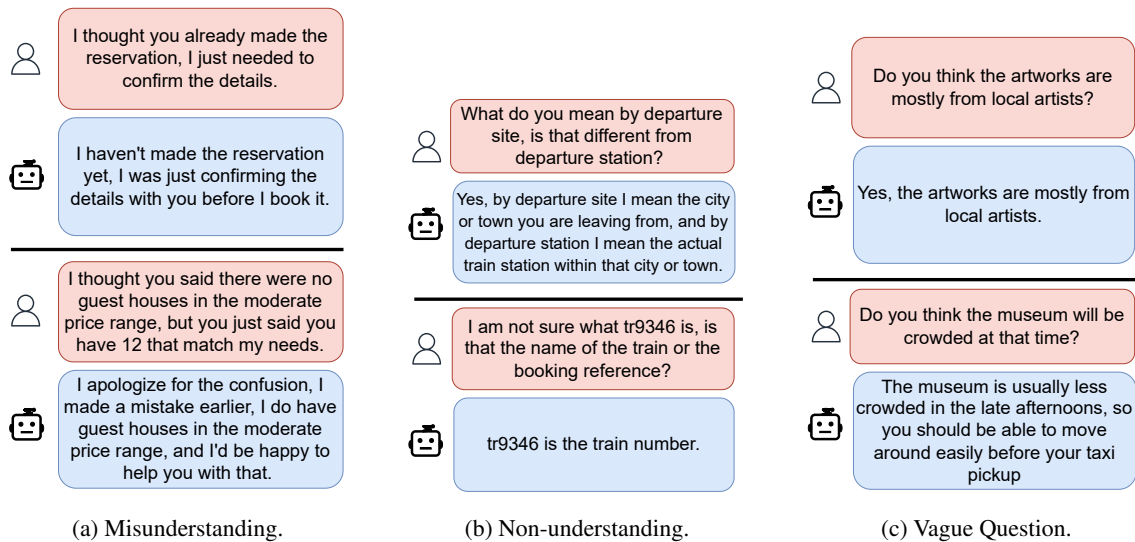


Figure 2: Two examples of miscommunication and repairing utterances for each type.

part of the out-of-scope definition of a chatbot and thus be handled differently, we only study VQs.

Furthermore, the distinction between misunderstandings and non-understandings is made by the speaker’s awareness (Weigand, 1999; Hirst et al., 1994). That is, if a listener is aware that he did not understand the other speaker, be it acoustically or semantically, the miscommunication is classified as a non-understanding. In many cases, he will also be aware of the cause of the non-understanding, i.e., was it acoustically, or does he not know the meaning of a certain word that the other has used. Since the listener is aware of this directly after the utterance, he can initialize the repairing immediately. For example, he might ask his dialogue partner to repeat or clarify the last utterance.

On the contrary, a misunderstanding is not noticed immediately, but the listener thinks he understood the speaker correctly (Hirst et al., 1994). It is only later in the conversation, that the listener gains awareness of the misunderstanding due to the further dialogue adding context or not fitting the understanding the listener had.

CoPrUS handles the generation of different types of communication errors by having a specific prompt for each one of them.

4.2 Repairing utterances

During a successful conversation, a miscommunication should lead to an attempt to resolve this trouble. This is known as repairing (Schegloff et al., 1977). One can differentiate between self-initiated repairing, e.g., when the speaker rephrases what he

said to make it clearer, or other-initiated repairing, which, for example, could be that a native speaker repairs an unintelligible utterance of a non-native speaker (Schegloff et al., 1977; Schegloff, 2000). In the studied scenario, it is sensible to only focus on other-initiated repairs, where the user asks the dialogue system for clarification. We thus differentiate between a miscommunication utterance, which should always be uttered by the user, and a repair utterance, uttered by the dialogue system.

There are multiple types of repairing that have been identified. These include, for example, repetition, word replacement, paraphrasing, and ambiguity reduction (Gonzalez-Lloret, 2005; Brédart, 1991; Rieger, 2003). This work focuses mostly on clarification requests as a repairing initiation (Gonzalez-Lloret, 2005), which is well-suited for booking conversations.

Once the miscommunication has happened or the user has become aware of it, the repairing utterance has to come into effect. The repairing utterance attempts to rectify the user’s interpretation. This error handling thus needs to be different for the three types of miscommunications that we study. To begin with, simply repeating the utterance, as would be a relevant strategy for spoken dialogue systems where errors on the channel and signal level can be handled this way (Skantze, 2005; Giesemann, 2006), is insufficient for chat data. Since the user can permanently see the previous utterances, the error should not have happened on the first two levels and thus needs a different repairing utterance.

If the error type was VQ, then the chatbot should

either answer the user’s question or state that it cannot answer it. Additionally, it could redirect the user to the dialogue goal. On the other hand, if the error type was NU or MU, the dialogue system needs to clarify the situation. This could potentially include rephrasing the previous content, explaining certain words, or giving additional context.

4.3 Affected Dialogues

To create a realistic scenario, we cannot assume that every dialogue should be affected by such a communication error. Möller et al. (2007) found for their data analysis that, depending on the specific dialogue situation, between 28% and 32% of the conversations contained a communication error. Since this work is done on spoken dialogues, which opens up the possibility of more errors than written dialogues, one needs to assume a reduced error frequency. Möller et al. (2007) report an error rate of 19% without communication errors that are the result of ASR errors. Based on the error rates reported in these studies, we opt for a slightly lower frequency of 18% for the affected dialogues, since the rather simple booking scenario using everyday language allows for less knowledge-based NUs or MUs, which should reduce the amount of affected dialogues, but on the other hand we also consider VQs as errors, which will increase the number.

Since no prior work studies the same setting, we have to decide on the distribution of the different types of miscommunications. Based on the combined prior error taxonomies (cf. Sec. 2.3) and the specifics of the dataset, we posit that NU should be the most common type of error in such a chat-based booking scenario, and that MU and VQ should be equal in their frequency. We thus use the following probabilities to sample an error type: $p_{MU} = 0.2$, $p_{VQ} = 0.2$, and $p_{NU} = 0.6$.

4.4 Creation of Miscommunication Utterances

In order to add post-hoc miscommunications to the dataset, we utilize LLaMA-3.1-70B-Instruct, which is the newest iteration of the open-source LLM from the LLaMA model family first introduced by Dubey et al. (2024). In preliminary experiments, we also investigated alternatives such as the LLaMA-2-Instruct (Touvron et al., 2023) models, Mixtral-8x7B-Instruct-v0.1 (Jiang et al., 2024) and Phi-3-mini-128k-instruct (Abdin et al., 2024). However, we found that LLaMA-3.1-70B created the most satisfactory utterances. The others showed

undesirable behavior, such as beginning every utterance with “Ah, i see!”, adhering too closely to the example to the one provided in the prompt, or generating another miscommunication instead of a repair utterance. We run the model on a DGX A100 with eight 80 GB GPUs.

4.5 Automatic Quality Assurance

Having a human-in-the-loop would naturally be an adequate approach to ensure the quality of the generated utterances. However, the required effort would be very high, and one could then use humans to generate the miscommunications and repairs directly.

To keep in line with the goal of having an automatic pipeline, we therefore evaluate the model’s responses with another LLM. Language Model-based evaluation has seen increased interest as an approach to mimic human judges (Chan et al., 2023; Liu et al., 2023; Zheng et al., 2023). Prometheus 2 is an open-source LLM specialized in evaluating the output of other LLMs, that shows high alignment with human judges (Kim et al., 2024).

We frame the evaluation as a direct assessment task, meaning that Prometheus 2 will assign each candidate utterance a score on a scale from 1 to 5. This is done based on the prompt to the LLaMA 3.1 model, the LLaMA output, and a description of the scoring rubrics. We generate up to ten candidates for each utterance, accepting the first one that scored a 4 or 5. If no candidate was accepted within ten tries, we use the highest rated candidate. The full prompt is shown in Fig. 7 the appendix.

To evaluate the validity of this automatic approach, we perform a small scale human evaluation, described in Sec. 7.

4.6 Two-Step Prompting

CoPrUS uses a two-step prompting approach, where we first sample and generate the miscommunication and then, in a second step, the repairing utterance. All prompts are shown in full in the appendix.

For the first step, we created three similar but individual prompts, one for each MU, NU, and VQ. Each prompt contains a description of the task, the type of miscommunication to generate, and an example. Furthermore, it contains the previous utterances as context.

The context includes all previous turns up to a maximum of 5. Limiting the context is a design

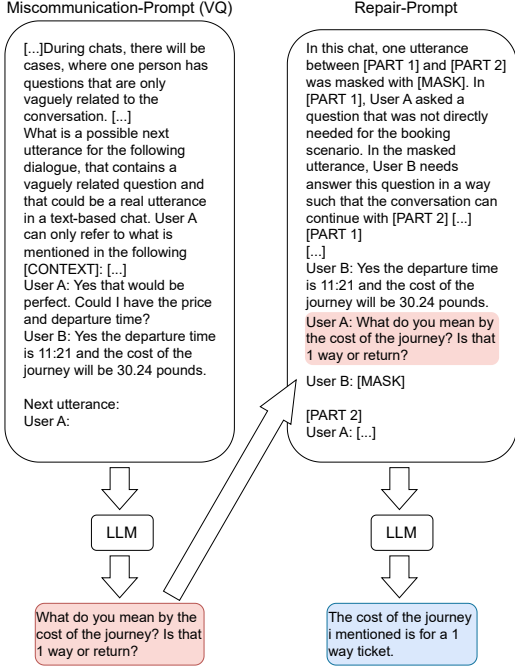


Figure 3: Prompting procedure with shortened prompts. Full prompts are shown in the appendix.

choice we made since both generation steps within CoPrUS should focus on the latest context and not, e.g., the beginning of the conversation. We found the context to greatly improve the coherence and relevance of the generated utterances. The prompt for MU also specifies that the generated utterance should not be in reference to the last system turn, but an earlier one (otherwise it would be a NU).

We sample 18% of the dialogues, and within each of those one random system turn index with the condition that it cannot be the first turn. This is done to ensure that some context can be given in the prompt for CoPrUS.

Let $D = \{u_0, s_1, u_1, s_2, u_2, \dots, s_n, u_n\}$ be a MultiWOZ dialogue of length n , where the user always sends the first and last message as per the definition of the dataset. We then sample an index i , which will determine the entry point for CoPrUS. We then use all previous utterances up to a maximum of 5 turns, excluding the last user utterance, as the first context with $j = \max(i - 5, 0)$ as

$$c = \bigcup_j^i \{s_j, u_j\} \setminus u_i. \quad (1)$$

In a first step, we use this context c and a sampled error type e to prompt the LLM $g(c, e)$ to generate the miscommunication user utterance: $\tilde{u}_i = g(c, e)$.

We append this result and the next real utterance u_{i+1} to c to create the context for the second step (cf. Fig. 3), the generation of the repairing system utterance: $\tilde{s}_{i+1} = g((c_1 \cup \tilde{u}_i), e)$.

We insert \tilde{u}_i and \tilde{s}_{i+1} at index i to create the final CoPrUS dialogue as:

$$\tilde{D} = \{u_0, \dots, u_{i-1}, s_i, \tilde{u}_i, \tilde{s}_{i+1}, u_i, s_i, \dots, s_n, u_n\}. \quad (2)$$

Within the second-step prompt, we split the context into two parts. The first part is the previous context plus the synthetic user utterance, and the second part is the next ground-truth user utterance. They are separated by a pseudo-masked system utterance. In this way, we frame the task as replacing the [MASK] in a suitable way.

5 CoPrUS-MultiWOZ

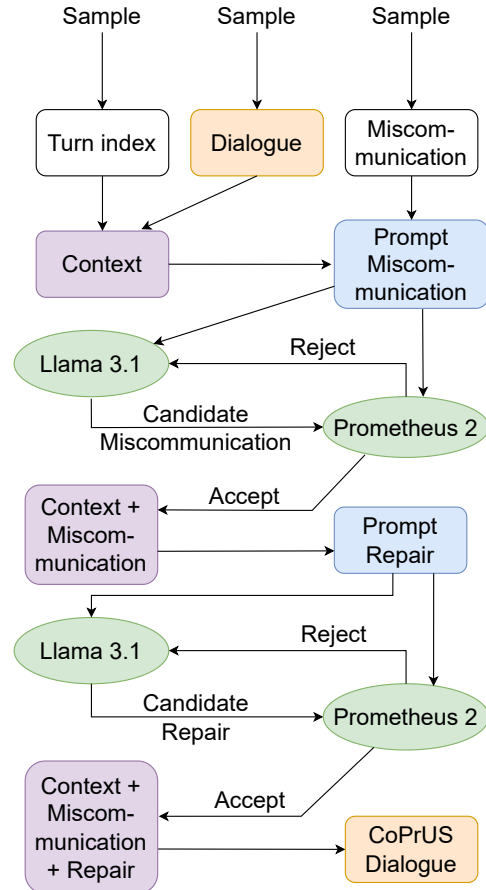


Figure 4: Overview of the CoPrUS workflow.

We apply the CoPrUS method to the MultiWOZ 2.1 (Eric et al., 2020) dataset to create CoPrUS-MultiWOZ. MultiWOZ 2.1 is an updated version

Model	Task	F1		BLEU		JGA	
		CoPrUS	MWZ	CoPrUS	MWZ	CoPrUS	MWZ
jointBERT	NLU	89.67	89.59	-	-	-	-
t5-small	NLU	83.51	83.45	-	-	-	-
SCGPT	NLG	-	-	23.48	24.63	-	-
t5-small	NLG	-	-	29.54	29.50	-	-
SetSUMBT	DST	92.18	92.27	-	-	50.83	51.12
t5-small	DST	91.81	91.76	-	-	53.04	52.67

Table 2: Empirical evaluation on the CoPrUS-MultiWOZ and original MultiWOZ dataset. We compare the NLG, NLU and DST task and report the F1, BLEU and JGA metrics. For DST, F1 refers to the slot F1 score and for NLU to the dialogue act F1 score.

of the original MultiWOZ (Budzianowski et al., 2018) dataset that fixes annotation errors. It contains roughly ten thousand annotated dialogues and a fixed train, develop, and test split.

Since the CoPrUS utterances should not change the user’s goal, we do not change the existing annotation but propagate the annotation of the real utterances with the utterances themselves.

We randomly sample 18% of each split of the original dataset and apply CoPrUS individually. Thus, it affects nearly 1900 dialogues in total. We adhere to the unified format of (Zhu et al., 2022) to facilitate further usage. We also provide the type of miscommunication for future work.

6 Qualitative Analysis

Given the strong performances of current LLMs, it is not surprising that the model follows the instructions and the language quality of the generated utterances is good enough to not give away that they are synthetically generated, as can be seen from the examples in Fig. 2. During our research, we found that problematic utterances were mostly illogical dialogue flows. For MU and NU, this primarily concerns the miscommunication utterance in the first step of our approach, where the user’s question is not consistent with the user’s understanding that can be inferred from the dialogue history. For VQ, we saw that problems were rather due to the repair utterance, which did not fit well with the following (gold-standard) utterance.

With the application of the automatic quality assurance through the LLM-based evaluation, we were able to mitigate these problems. Nevertheless, it has to be expected that these synthetic utterances are not perfect substitutes for human-generated gold-standard data. In this sense, our qualitative analysis shows that some dialogues remain where

the generated utterances seem illogical.

Metric	Miscom.	Repair	Total
EM	0.22	0.34	0.28
Difference	1.58	1.68	1.61
FP	0.22	0.14	0.18
FN	0.30	0.36	0.33
Human	3.46	4.20	3.83
LLM	3.24	3.12	3.18

Table 3: Result of the evaluation of human alignment with the LLM-judge. We show the EM, average score difference, FP and FN rate as well as the average scores from human and LLM judges for the miscommunications and repairs separately as well as in total.

7 Human Evaluation

To estimate if the automatic, LLM-based quality assurance is effective, we perform a small-scale human evaluation to measure its alignment to human perspective. For this, the human judges were presented with nearly the same instructions as the LLM-judge, apart from some changes in the form of presentation. These judges rated 100 examples, comprising both miscommunication and repair candidates and both candidates that were accepted and rejected by the LLM in equal amounts. The judges were recruited non-author volunteers with sufficient language proficiency. Each dialogue is used at most twice for each inference step (miscommunication or repair), one rejected and one accepted candidate.

We report the Exact Match (EM, human and LLM gave the same score), the average score difference, the percentage of False Positives (FP, model accepts, human rejects) and False Negatives (FN, model rejects, human accepts), as well as the average rating of humans and LLM. As before, we

define the threshold to accept a candidate as a rating of at least 4 out of 5.

The results in Tab. 3 show that the alignment is not perfect but sufficiently good. Especially, the low FP rate is encouraging, since these are accepted candidates that humans would have rejected. We can see that the average difference for repair candidates is higher, even though the EM is better. This is due to the fact that the second-most common score given by the LLM was 1 (the most common was 5), while the human judges gave no score of 1 at all. This is also evidenced by the high average human score for repair candidates. Interestingly, the human judges gave higher scores on average than the LLM, speaking for the high quality of the generated candidates.

8 Empiric Evaluation

Even though the synthetic data is not of gold-standard quality, it can be beneficial for downstream TOD systems. Therefore, it is crucial to evaluate if the addition of CoPrUS utterances harms the performance of a TOD system, especially with regard to *model collapse* (Shumailov et al., 2024). To this end, we evaluate the three classical tasks NLG, NLU and Dialogue State Tracking (DST). For each task, we train both a t5-small (Raffel et al., 2020) model and another, specific architecture. For NLG this is SCGPT (Peng et al., 2020), for NLU jointBERT (Zhu et al., 2020) and for DST Set-SUMBT (van Niekerk et al., 2021). Every model is trained with ConvLab-3 (Zhu et al., 2022) and the hyperparameters defined there. We evaluate the models with the slot F1 Score and Joint Goal Accuracy (JGA) in the DST task, BLEU in the NLG task, and dialogue act F1 score for the NLU task. Our evaluation is thus not focused on the model’s error recovery ability. The results of this evaluation are shown in Tab. 2. They show that the CoPrUS utterances did not impede the learning for any architecture or task. The results across the different metrics are minimal, with CoPrUS-trained models being marginally better in most cases.

9 Conclusion

This paper introduces CoPrUS, a method to add miscommunication turns to dialogues in order to make them more realistic by allowing them to deviate from the happy path. To this end, we introduce a simple error taxonomy that is based on previous communication research. We utilize a two-step

prompting approach and a state-of-the-art LLM with another LLM performing quality assurance to generate these utterances.

Our evaluation shows that the generated utterances have a good quality. Moreover, the LLM mostly adheres to the task and generates sensible utterances, both for miscommunications and repairing. This data can help future research into error recovery for dialogue systems. However, as was made clear by the qualitative analysis, one cannot expect data of gold-standard quality. For example, the generated MU or NU statements can be illogical in some instances. Utilizing a LLM for quality assurance has mitigated this issue.

Our investigation into the alignment of human and LLM ratings show that utilizing a LLM can be a valid approach for quality assurance in our scenario.

10 Limitations

This work focuses on only three types of miscommunications. This could be extended in future work. Furthermore, we had to posit a distribution for the types of miscommunications. An empiric study on such a taxonomy in real-world, natural TODs could shed light on such a distribution.

While our results show that it is generally feasible to adapt the dialogues by adding miscommunications, our evaluation shows that some problems remain. Mostly, the dialogue flow can be disrupted from a logical point of view. This goes to show that LLMs still have limited implicit reasoning capabilities. The error recovery abilities of end-to-end TOD systems are currently understudied. Our evaluation does not measure or try to improve this ability but rather lays the foundation for such research in future work.

References

Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Qin Cai, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Yen-Chun Chen, Yi-Ling Chen, Parul Chopra, Xiyang Dai, Allie Del Giorno, Gustavo de Rosa, Matthew Dixon, Ronen Eldan, Victor Fragoso, Dan Iter, Mei Gao, Min Gao, Jianfeng Gao, Amit Garg, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Jamie

- Huynh, Mojan Javaheripi, Xin Jin, Piero Kauffmann, Nikos Karampatziakis, Dongwoo Kim, Mahoud Khademi, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Ce Liu, Mengchen Liu, Weishung Liu, Eric Lin, Zeqi Lin, Chong Luo, Piyush Madan, Matt Mazzola, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacrose, Shital Shah, Ning Shang, Hiteshi Sharma, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Xin Wang, Lijuan Wang, Chunyu Wang, Yu Wang, Rachel Ward, Guanhua Wang, Philipp Witte, Haiping Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Sonali Yadav, Fan Yang, Jianwei Yang, Ziyi Yang, Yifan Yang, Donghan Yu, Lu Yuan, Chengruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. 2024. [Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone](#). *Preprint*, arxiv:2404.14219.
- Namo Bang, Jeehyun Lee, and Myoung-Wan Koo. 2023. [Task-optimized adapters for an end-to-end task-oriented dialogue system](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7355–7369, Toronto, Canada. Association for Computational Linguistics.
- Serge Brédart. 1991. Word interruption in self-repairing. *Journal of psycholinguistic research*, 20:123–138.
- 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.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. [Chateval: Towards better llm-based evaluators through multi-agent debate](#). *Preprint*, arXiv:2308.07201.
- Willy Chung, Samuel Cahyawijaya, Bryan Wilie, Holy Lovenia, and Pascale Fung. 2023. [InstructTODS: Large language models for end-to-end task-oriented dialogue systems](#). In *Proceedings of the Second Workshop on Natural Language Interfaces*, pages 1–21, Bali, Indonesia. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Laila Dybkjær, Niels Ole Bernsen, and Hans Dybkjær. 1996. Grice incorporated: Cooperativity in spoken dialogue. In *COLING 1996 Volume 1: The 16th International Conference on Computational Linguistics*.
- 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 Twelfth Language Resources and Evaluation Conference*, pages 422–428, Marseille, France. European Language Resources Association.
- Petra Gieselmann. 2006. Comparing error-handling strategies in human-human and human-robot dialogues. In *Proc. 8th Conf. Nat. Language Process.(KONVENS)*. Konstanz, Germany, pages 24–31.
- Marta Gonzalez-Lloret. 2005. Reconstructing ns/nns communication. *Book?? The Consequences of Mobility: Linguistic and Sociocultural Contact Zones?*, pages 1–27.
- H. P. Grice. 1975. Logic and Conversation. *Syntax and Semantics*, 3:43–58.
- Pamela Grimm. 2010. Social desirability bias. *Wiley international encyclopedia of marketing*.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, and Luo Si. 2022. [Galaxy: A generative pre-trained model for task-oriented dialog with semi-supervised learning and explicit policy injection](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10749–10757.
- Ryuichiro Higashinaka, Masahiro Araki, Hiroshi Tsukahara, and Masahiro Mizukami. 2019. [Improving Taxonomy of Errors in Chat-Oriented Dialogue Systems](#). In Luis Fernando D’Haro, Rafael E. Banchs, and Haizhou Li, editors, *9th International Workshop on Spoken Dialogue System Technology*, volume 579, pages 331–343. Springer Singapore, Singapore.
- Ryuichiro Higashinaka, Masahiro Araki, Hiroshi Tsukahara, and Masahiro Mizukami. 2021. [Integrated taxonomy of errors in chat-oriented dialogue systems](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 89–98, Singapore and Online. Association for Computational Linguistics.
- Ryuichiro Higashinaka, Kotaro Funakoshi, Masahiro Araki, Hiroshi Tsukahara, Yuka Kobayashi, and Masahiro Mizukami. 2015. Towards taxonomy of errors in chat-oriented dialogue systems. In *Proceedings of the 16th annual meeting of the special interest group on discourse and dialogue*, pages 87–95.
- Graeme Hirst, Susan McRoy, Peter Heeman, Philip Edmonds, and Diane Horton. 1994. [Repairing conversational misunderstandings and non-understandings](#). *Speech Communication*, 15(3):213–229.

- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L elio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2024. *Mixtral of Experts*. Preprint, arxiv:2401.04088.
- J. F. Kelley. 1984. *An iterative design methodology for user-friendly natural language office information applications*. *ACM Transactions on Information Systems*, 2(1):26–41.
- Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham Neubig, Moontae Lee, Kyungjae Lee, and Minjoon Seo. 2024. *Prometheus 2: An open source language model specialized in evaluating other language models*. Preprint, arXiv:2405.01535.
- Takyoung Kim, Yukyung Lee, Hoonsang Yoon, Pilsung Kang, Junseong Bang, and Misuk Kim. 2022. *Oh My Mistake!: Toward Realistic Dialogue State Tracking including Turnback Utterances*. In *Proceedings of the Towards Semi-Supervised and Reinforced Task-Oriented Dialog Systems (SereTOD)*, pages 1–12, Abu Dhabi, Beijing (Hybrid). Association for Computational Linguistics.
- Takyoung Kim, Jamin Shin, Young-Ho Kim, Sanghwan Bae, and Sungdong Kim. 2023. *Revealing User Familiarity Bias in Task-Oriented Dialogue via Interactive Evaluation*. Preprint, arxiv:2305.13857.
- Zekun Li, Zhiyu Chen, Mike Ross, Patrick Huber, Seungwhan Moon, Zhaojiang Lin, Xin Dong, Adithya Sagar, Xifeng Yan, and Paul Crook. 2024. *Large language models as zero-shot dialogue state tracker through function calling*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8688–8704, Bangkok, Thailand. Association for Computational Linguistics.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020. *MinTL: Minimalist transfer learning for task-oriented dialogue systems*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3391–3405, Online. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. *G-eval: NLG evaluation using gpt-4 with better human alignment*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Koh Mitsuda, Ryuichiro Higashinaka, Tingxuan Li, and Sen Yoshida. 2022. *Investigating person-specific errors in chat-oriented dialogue systems*. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 464–469, Dublin, Ireland. Association for Computational Linguistics.
- Tracy L. Mitzner, Julie B. Boron, Cara Bailey Fausset, Anne E. Adams, Neil Charness, Sara J. Czaja, Katinka Dijkstra, Arthur D. Fisk, Wendy A. Rogers, and Joseph Sharit. 2010. *Older adults talk technology: Technology usage and attitudes*. *Computers in human behavior*, 26(6):1710–1721.
- Sebastian M oller, Klaus-Peter Engelbrecht, and Antti Oulasvirta. 2007. *Analysis of communication failures for spoken dialogue systems*. In *INTER-SPEECH*, pages 134–137.
- Clifford Nass and Li Gong. 1999. *Maximized modality or constrained consistency?* In *AVSP’99-International Conference on Auditory-Visual Speech Processing*.
- Clifford Nass and Youngme Moon. 2000. *Machines and mindlessness: Social responses to computers*. *Journal of social issues*, 56(1):81–103.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. *Training language models to follow instructions with human feedback*. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Tim Paek. 2003. *Toward a taxonomy of communication errors*. In *ISCA Tutorial and Research Workshop on Error Handling in Spoken Dialogue Systems*.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2021. *Soloist: Building task bots at scale with transfer learning and machine teaching*. *Transactions of the Association for Computational Linguistics*, 9:807–824.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. *Few-shot natural language generation for task-oriented dialog*. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 172–182, Online. Association for Computational Linguistics.
- Kun Qian, Ahmad Beirami, Zhouhan Lin, Ankita De, Alborz Geramifard, Zhou Yu, and Chinnadhurai Sankar. 2021. *Annotation inconsistency and entity bias in MultiWOZ*. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 326–337. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi

- Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Caroline L Rieger. 2003. Repetitions as self-repair strategies in english and german conversations. *Journal of pragmatics*, 35(1):47–69.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Emanuel A Schegloff. 2000. When others’ initiate repair. *Applied linguistics*, 21(2):205–243.
- Emanuel A Schegloff, Gail Jefferson, and Harvey Sacks. 1977. The preference for self-correction in the organization of repair in conversation. *Language*, 53(2):361–382.
- Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. 2024. [AI models collapse when trained on recursively generated data](#). *Nature*, 631(8022):755–759.
- Gabriel Skantze. 2005. Exploring human error recovery strategies: Implications for spoken dialogue systems. *Speech Communication*, 45(3):325–341.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2022. [Multi-task pre-training for plug-and-play task-oriented dialogue system](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4661–4676, Dublin, Ireland. Association for Computational Linguistics.
- Haipeng Sun, Junwei Bao, Youzheng Wu, and Xiaodong He. 2023. [Mars: Modeling Context & State Representations with Contrastive Learning for End-to-End Task-Oriented Dialog](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 11139–11160, Toronto, Canada. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#). *Preprint*, arxiv:2307.09288.
- Dennis Ulmer, Elman Mansimov, Kaixiang Lin, Lijia Sun, Xibin Gao, and Yi Zhang. 2024. [Bootstrapping LLM-based task-oriented dialogue agents via self-talk](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9500–9522, Bangkok, Thailand. Association for Computational Linguistics.
- Margot J. Van Der Goot and Tyler Pilgrim. 2020. [Exploring Age Differences in Motivations for and Acceptance of Chatbot Communication in a Customer Service Context](#). In Asbjørn Følstad, Theo Araujo, Symeon Papadopoulos, Effie Lai-Chong Law, Ole-Christoffer Granmo, Ewa Luger, and Petter Bae Brandtzaeg, editors, *Chatbot Research and Design*, volume 11970, pages 173–186. Springer International Publishing, Cham.
- Carel van Niekerk, Andrey Malinin, Christian Geisshauer, Michael Heck, Hsien-chin Lin, Nurul Lubis, Shutong Feng, and Milica Gasic. 2021. [Uncertainty measures in neural belief tracking and the effects on dialogue policy performance](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7901–7914, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Edda Weigand. 1999. [Misunderstanding: The standard case](#). *Journal of Pragmatics*, 31(6):763–785.
- Puhai Yang, Heyan Huang, Wei Wei, and Xian-Ling Mao. 2022. [Toward Real-life Dialogue State Tracking Involving Negative Feedback Utterances](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2222–2232, Washington DC USA. ACM.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D. Williams. 2013. [POMDP-Based statistical spoken dialog systems: A review](#). *Proceedings of the IEEE*, 101(5):1160–1179.
- Changhong Yu, Chunhong Zhang, and Qi Sun. 2022. [A Chit-Chats Enhanced Task-Oriented Dialogue Corpora for Fuse-Motive Conversation Systems](#). *Preprint*, arxiv:2205.05886.

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](#). *Preprint*, arxiv:2201.08904.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623. Curran Associates, Inc.

Qi Zhu, Christian Geishauser, Hsien chin Lin, Carel van Niekerk, Baolin Peng, Zheng Zhang, Michael Heck, Nurul Lubis, Dazhen Wan, Xiaochen Zhu, Jianfeng Gao, Milica Gašić, and Minlie Huang. 2022. [Convlab-3: A flexible dialogue system toolkit based on a unified data format](#). *arXiv preprint arXiv:2211.17148*.

Qi Zhu, Zheng Zhang, Yan Fang, Xiang Li, Ryuichi Takanobu, Jinchao Li, Baolin Peng, Jianfeng Gao, Xiaoyan Zhu, and Minlie Huang. 2020. [ConvLab-2: An open-source toolkit for building, evaluating, and diagnosing dialogue systems](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 142–149, Online. Association for Computational Linguistics.

A Appendix

You have a deep understanding of how humans communicate. During conversations, there will be cases, where one person's understanding doesn't align with what the other person meant.

For example see the last turn of User A in the [EXAMPLE]. Here, User A thought he had to pay by 15th of August, but User B was only talking about the delivery and the payment is due earlier:

[EXAMPLE]:

User A: Yes, please order the new game for me.
User B: Okay, 15th of August works for you?
User A: Sure, go ahead!
User B: Great, I will have the game delivered to you. And payment is due in three days.
User A: Sorry, I am confused. I assumed I have to pay on 15th of August as well?

Now, create an utterance that is different from the above [EXAMPLE]. Do NOT talk about the game ordering.

This dialogue is about a booking scenario. User A is the client and User B is the booking assistant. Do not use any emojis. Answer as short as possible without explanations. Your response is only the fitting utterance, nothing else.

User A cannot change his responses from before, for example he cannot change the type of hotel he is looking for or cancel a reservation. User A is only looking for clarification for what was said in the second-to-last utterance by User B. It is important, that User A is not surprised by something he himself said or initiated earlier. For example, if User A asked for the reference number, the created utterance should not ask, what he needs the reference number for.

What is a possible next utterance for User A in the following dialogue, that shows that User A had a different understanding from the second-to-last User B utterance which led to a misunderstanding by User A. User A can only refer to what is mentioned in the following [CONTEXT]:

[CONTEXT]
{context}

Next utterance:

You have a deep understanding of how humans communicate. During conversations, there will be cases, where one person doesn't comprehend what the other person meant. Not because he couldn't hear them, but because he cannot understand the meaning of the utterance.

For example see the last turn of User A in the [EXAMPLE]. Here, User A does not understand why User B mentions a certain date:

[EXAMPLE]:

User A: Yes, please order the new game for me.
User B: Okay, 15th of August works for you?
User A: I'm not sure what you are talking about. Would this date work for what? Are you going to order on 15th of August?

Now, create an utterance that is different from the above [EXAMPLE]. Do NOT talk about the game ordering.

This dialogue is about a booking scenario. User A is the client and User B is the booking assistant. Do not use any emojis. Answer as short as possible without explanations. Your response is only the fitting utterance, nothing else.

It is important, that User A is not surprised by something he himself said or initiated earlier. For example, if User A asked for the reference number, the created utterance should not ask, what he needs the reference number for.

What is a possible next utterance for User A in the following dialogue, that shows that User A did not comprehend what User B said in the last utterance? User A can only refer to what is mentioned in the following [CONTEXT], especially he should react to the last utterance by User B and ask for clarification:

[CONTEXT]
{context}

Next utterance:

(a) Misunderstanding.

(b) Non-understanding.

You have a deep understanding of how humans communicate and can predict following utterances in textual dialogues. During chats, there will be cases, where one person has questions that are only vaguely related to the conversation. They fit the domain but don't bring the user closer to his goal.

For example, see the last turn of User A here where he asks "Have you ever played the game". This utterance is not directly related to the process of ordering the game, but it is still related to the domain of games:

User A: Yes, please order the new game for me.
User B: Okay, 15th of August works for you?
User A: Have you ever played the game?
User B: No, I have not yet played this game, but it sounds really interesting! Is delivery by the 15th of August okay for you?

Now, create an utterance that is different from the above example. Do NOT talk about the game ordering.

This dialogue is about a booking scenario. User A is the client and User B is the booking assistant. Do not use any emojis. Answer as short as possible without explanations. Your response is only the fitting utterance, nothing else.

What is a possible next utterance for the following dialogue, that contains a vaguely related question and that could be a real utterance in a text-based chat. User A can only refer to what is mentioned in the following [CONTEXT]:

[CONTEXT]
{context}

Next utterance:

(c) Vague Question.

Figure 5: The full prompts used for the miscommunication generation.

You have a deep understanding of how humans communicate and can predict utterances that fit into textual chats.

In this chat, one utterance between [PART 1] and [PART 2] was masked with [MASK]. In [PART 1] {typeOfErrorExplanation}

Respond with exactly one utterance that can replace the mask seamlessly. It should reference the last utterance from User A in [PART 1]. It should also make sense with what is being said in [PART 2], however it absolutely cannot use information that is only given in [PART 2], because [PART 2] in the conversation happens AFTER the [MASK].

User B is a booking assistant who is helping User A. User B should therefore never ask User A for information that he cannot have. He should only ask User A for more information on what he wants. In most cases, User B will not ask a question, but respond to User A's question.

User B cannot say in his utterance, that he will look up something, e.g. a restaurant, and tell User A to wait. Instead, User B should directly propose the result of looking up, e.g., give a restaurant name. Moreover, User B cannot change his responses from before, for example he cannot change the type of hotel he is offering. User B can only offer clarification for what he said before.

Do not use any emojis. Answer as short as possible without explanations. Your response is only the fitting utterance, nothing else.

What utterance can replace the [MASK] here so that User B is helping User A and responding to his question?

```
[PART 1]
(contextA)
User B: [MASK]
```

```
[PART 2]
(contextB)
```

The utterance to replace the [MASK] is:

Figure 6: The full prompt used for the repair generation.

```

###Task Description:
    An instruction (might include an Input inside
it), a response to evaluate and a score rubric
representing a evaluation criteria are given.
    1. Write a detailed feedback that assess the
quality of the response strictly based on the
given score rubric, not evaluating in general.
    2. After writing a feedback, write a score
that is an integer between 1 and 5. You should
refer to the score rubric.
    3. The output format should look as follows:
"Feedback: (write a feedback for criteria)
[RESULT] (an integer number between 1 and 5)"
    4. Please do not generate any other opening,
closing, and explanations.

###The instruction to evaluate:
{instruction}

###Response to evaluate:
{response}

###Score Rubrics:
"criteria": "Did the model provide a realistic
continuation for the dialogue that follows the
instructions, where User A did not comprehend what
User B said in his last turn? This non-
understanding should not be auditive, but because
the User A did not grasp the meaning or sense of
what User B said.",
"score1_description": "The model neglects to
follow the instructions completely. The utterance
is not relevant or not realistic.",
"score2_description": "The model acknowledges
most of the instructions, but not sufficiently.
For example, the generated utterance might be
illogical, if User A did not understand why he
needs something that he requested himself
earlier.",
"score3_description": "The model follows the
instructions sufficiently. The utterance is
fitting the context, is logical, and shows a non-
understanding. However, the utterance is not
perfectly realistic.",
"score4_description": "The model follows the
instructions very closely. The utterance fits well
into the context and is realistic. It correctly
shows a non-understanding.",
"score5_description": "The model excels in
following the instructions perfectly and creates
an utterance that perfectly fits the context and
is indistinguishable from a real non-understanding
that can happen in a conversation between two
humans."

###Feedback: ""

```

Figure 7: The prompt used for Prometheus 2. {instruction} is replaced by the prompt used for the Llama model, and {response} by its output. We show the score rubrics for a NU type. The rubrics for the other types and the repairing are analog to this.