



A Unified Approach to Emotion Detection and Task-Oriented Dialogue Modeling

Armand Stricker, Patrick Paroubek

► To cite this version:

Armand Stricker, Patrick Paroubek. A Unified Approach to Emotion Detection and Task-Oriented Dialogue Modeling. IWSDS, Mar 2024, Sapporo (Japon), Japan. hal-04415809

HAL Id: hal-04415809

<https://hal.science/hal-04415809>

Submitted on 24 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Unified Approach to Emotion Detection and Task-Oriented Dialogue Modeling

Armand Stricker, Patrick Paroubek

Abstract In current text-based task-oriented dialogue (TOD) systems, user emotion detection (ED) is often overlooked or is typically treated as a separate and independent task, requiring additional training. In contrast, our work demonstrates that seamlessly unifying ED and TOD modeling brings about mutual benefits, and is therefore an alternative to be considered. Our method consists in augmenting SimpleToD, an end-to-end TOD system, by extending belief state tracking to include ED, relying on a single language model. We evaluate our approach using GPT-2 and Llama-2 on the EmoWOZ benchmark, a version of MultiWOZ annotated with emotions. Our results reveal a general increase in performance for ED and task results. Our findings also indicate that user emotions provide useful contextual conditioning for system responses, and can be leveraged to further refine responses in terms of empathy.

1 Introduction

Emotional user utterances frequently occur during interactions with dialogue systems [19]. In open-domain dialogues, users directly share personal and emotional experiences [10]. In task-oriented dialogues (TODs), emotions are closely related to task progression. They usually become apparent as the task unfolds, and are contingent on the user’s expectations being met [5]. Detecting these emotions explicitly in either scenario is beneficial for several reasons: it can help with reviewing chat logs after the exchange concludes or with adjusting responses during the exchange to better align with the user’s emotional state [18]. In TODs, these more empathetic responses have demonstrated their effectiveness in compensating for system errors [12], creating the impression of a more capable system.

Existing text-based approaches to ED in TODs either require a dedicated, specifically trained component [5, 6] or assume *implicit* ED. This assumption is due to systems being commonly trained to replicate human expert responses [25, 7] which inherently convey empathy, as needed, and therefore implicit ED. In contrast, we propose an approach that eliminates the need for training an additional component while *explicitly* modeling user emotions, applicable in settings where emotion annotations are available. Our method involves predicting user emotions within a sequence of task-related components, treating ED as an extension

Armand Stricker, Patrick Paroubek
Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, 91400, Orsay, France e-mail:
firstname.lastname@lisn.upsaclay.fr

of belief state tracking. Our approach is end-to-end and relies on a single language model, optimized in a unified fashion. Specifically, we build upon SimpleTOD [8], a popular end-to-end TOD approach trained on GPT-2 [17] (Sec. 2). Recognizing the success of large language models (LLMs) [1], we also fine-tune a state-of-the-art LLM, Llama-2 [22], using LoRA [9], a parameter-efficient fine-tuning technique (Sec. 2).

Moreover, we show that when considering Llama-2, ED predictions can be readily used to more explicitly condition and refine system responses. Capitalizing on the fact that LoRA weights are optimized *on top* of Llama-2’s fixed parameters, it is possible to dynamically disable the LoRA weights after generating TOD predictions to harness the LLM’s in-context learning ability. We refine responses without further training, using a chain-of-thought prompting approach [23] (Sec. 2).

We evaluate our method on the text-based EmoWOZ benchmark [5], the only benchmark to provide emotion annotations on MultiWOZ [2]. The authors propose several competitive ED-only baselines, with ContextBERT [5] being the most successful and taking the entire dialogue history as input.

Results demonstrate that our unified approach generally leads to improved ED outcomes compared to ContextBERT, while also yielding comparable to superior task-oriented results compared with SimpleToD. Human evaluation of outputs produced by Llama-2 indicates that both refined responses and those generated by our modified pipeline are preferred by annotators when compared with SimpleToD. More specifically, refined responses display more empathy and engagingness. Our findings encourage a more direct integration of ED into the task-oriented framework, moving beyond its treatment as an independent task. Our experimental code can be found on GitHub¹.

2 Method and Experimental Setup

The EmoWOZ Dataset

This is the only text-based dataset offering emotion annotations for a large and popular corpus of 10,000 TODs. It focuses on 7 user emotions: **neutral**, **excited**, **fearful**, **satisfied**, **dissatisfied**, **apologetic**, and **abusive**. These follow the OCC model (Ortony, Clore and Collins) [15] and are based on 3 possible emotion elicitors which include the **operator** (or system), the **user**, and **events**. This emotion schema is highly specific to TODs and makes it possible to distinguish between emotions caused by the system and those provoked by external factors, independent of the system’s performance.

- Positive emotions caused by the **operator** relate to the system successfully completing the task and are classified as *satisfied*, *liking*, *appreciative*. Negative emotions caused by the operator (e.g. proposing the wrong type of restaurant) are labeled as *dissatisfied*, *disliking* or as *abusive* in some cases.
- A single emotion is attributed to the **user**, and is labeled *apologetic*. It is employed to denote negative emotions that arise when the user’s own actions or queries lead to confusion for the operator.
- Emotions elicited by **events** are tagged with the label *fearful*, *sad*, *disappointed* when negative (e.g. the user’s favorite restaurant is unavailable), and with the label *excited*, *happy*, *anticipating* otherwise (e.g. the user eagerly anticipates an upcoming trip).
- The default label when no emotion is expressed is *neutral*.

Examples of user utterances and their associated emotion label can be found in table 5 and Appendix 4. Dataset statistics are shown in Table 1. Certain emotions are notably rare and make this a challenging benchmark for ED in TODs.

¹ <https://github.com/armandstrickernlp/Emo-TOD>

	Neut.	Fear.	Diss.	Apol.	Abus.	Exci.	Sati.
proportion	71.9%	0.5%	1.3%	1.2%	0.1%	1.2%	23.8%
count	51,426	381	914	838	44	860	17,061

Table 1: (EmoWOZ label statistics. Emotions are **Neutral**; **Fearful**, sad, disappointed; **Dissatisfied**, disliking; **Apologetic**; **Abusive**; **Excited**, happy, anticipating; **Satisfied**, liking, appreciative.

Method

Our approach builds upon SimpleToD [8], a popular end-to-end strategy that employs a single language model, similar to other recent end-to-end TOD systems [7, 26]. In training, the model is exposed to instances where Context (all preceding turns $C_t = [U_0, S_0, \dots, U_t]$), Belief State (B_t), Dialogue Acts (A_t) and System Response (S_t) are concatenated into a single text sequence. To ensure database adaptability, responses are additionally delexicalized. During inference, only the Context C_t is passed as input, and the model is expected to generate the rest of the sequence. Moreover, the predicted belief state is used to query the database to retrieve values with which to replace delexicalized placeholders. This pipeline can be readily augmented, and previous studies have in fact done so with snippets of knowledge [3] and chitchat [21].

In our work, we enhance SimpleToD by **extending the belief state**, introducing an additional ED component into the pipeline. We cast both ED and TOD as tasks to be learned in a unified way with a single language modeling objective, maximizing next token probability given a sequence of input tokens: $L = \sum_i \log P(t_i | t_{<i})$.

We propose two variants of our approach. The **EMO** variant is trained on sequences which take the form $[C_t, B_t, E_t, A_t, S_t]$, where E_t is the emotion label associated with the user utterance U_t . The **PREV** variant is trained on sequences made up of $[C_t^+, B_t, E_t, A_t, S_t]$ with $C_t^+ = [(U_0, E_0), S_0, \dots, U_t]$. The notable difference between both variants is that C_t^+ concatenates previous user utterances (excluding U_t) with their corresponding emotion labels. This is meant to help the model pick up on potential emotional progression patterns in the dialogues.

During inference, both variants are expected to generate each component given only an input Context. In the case of PREV, the generated emotion label is inserted back into the Context to condition the next turn’s predictions. We show an overview of the approach as well as an example of the input format in Figure 1.




Fig. 1: (left) We expand the belief state to add emotion detection to the standard pipeline. During inference, only the dialogue Context is passed to SimpleToD and each component is generated in an autoregressive manner. (right) A turn from dialogue pmul0450, formatted for our method. Previous user emotions, highlighted in gray, are added to the Context for the PREV variant. Each component is delineated by special tokens, which helps the model in its predictions, and delexicalized placeholders in the response are in brackets.

The models used for our experiments include GPT-2² [17], typically used with SimpleToD, and the more recent, state-of-the-art Llama-2-chat³ [22]. This is an open-source LLM trained on data from publicly accessible sources which can undergo further fine-tuning or be prompted directly with instructions and/or few-shot examples. To adapt models to our task, we fully fine-tune GPT-2 and employ LoRA fine-tuning for Llama-2, preserving a significant portion of the pre-trained weights. This approach enhances regularization, generally leading to improved generalization compared to a fully fine-tuned alternative. Training details can be found in Appendix 2.

When **refining responses**, we focus on the PREV model’s predictions. Once these have been generated, we can disable the LoRA weights and, with no additional training, prompt the model using a few-shot, chain-of-thought [23] approach (REFINE). This consists in creating demonstrations which include a reasoning step to prime the LLM’s output. Each few-shot demonstration consists in **a)** the dialogue Context, **b)** the emotion to be predicted, **c)** the response to be generated, **d)** a chain-of-thought reasoning step suggesting which behavior to display. Finally, it also contains **e)** an emotion-aware snippet that can be *prepended* to the previously generated response. Compared with refining the entire response, prompting the model to produce a prependable snippet mitigates task-related errors due to substantial modifications made by the LLM. We incorporate specific instructions into the prompt and construct one exemplar per emotion, picking examples from the training set. The demonstration snippets we design are meant to convey remorse, reassurance, appeasement, or enthusiasm, depending on the user’s emotion. Appendix 3 shows an illustrative example. When generating snippets, we filter those which are overly similar to the generated response to avoid redundancy, based on Levenshtein distance. Empirically, we find that filtering out snippets with a similarity ratio $\geq 50\%$ is effective.

Automatic Metrics

For ED, we adhere to the EmoWOZ metrics, presenting **F1 scores** for each emotion alongside the **macro** and **weighted** averages excluding the *neutral* class. For end-to-end TOD, we employ the standard MultiWOZ metrics. The **joint goal accuracy** (JGA) reflects the proportion of turns where the predicted user constraints exactly match the reference ones. The **inform rate** evaluates the system’s capacity to provide the right type of entities from the database, given the user’s constraints, and the **success rate** assesses how effectively the system delivers requested attributes like phone numbers or booking references. For a more comprehensive understanding of these metrics, we direct readers to the MultiWOZ paper [2]. For response quality, we use the **BLEU** [16] score. For response diversity, we analyze the **Conditional Bigram Entropy** (CBE) and the count of **unique trigrams** to gauge the richness of vocabulary and phrasing [14].

3 Results and Discussion

With regards to **Emotion Detection** (Table 2), jointly learning ED and TOD modeling yields improvements over ContextBert, EmoWOZ’s most effective ED-only model, with the exception of one result (*fearful*, PREV-gpt).

² <https://huggingface.co/gpt2>

³ <https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>

Model	Neut.	Fear.	Diss.	Apol.	Abus.	Exci.	Sat.	Macro	Weigh.
ContextBERT [5]	95.1	35.7	36.4	70.3	19.4	34.1	90.0	47.7	83.8
EMO-gpt	95.5	38.7	37.3	71.1	27.4	40.8	90.7	51.0	84.8
PREV-gpt	95.6*	21.5	40.5*	73.4*	27.9	41.9	91.1*	49.4	85.3
EMO-llama	95.4	51.7*	34.5	71.0	21.3	39.8	90.5	51.5	84.6
PREV-llama	95.6*	55.2	31.4	74.2	36.7*	44.0*	91.1*	55.4*	85.2*

Table 2: Mean F1-scores (5 seeds) for individual emotion labels, as well as macro and weighted averages (excluding *neutral*). In each column, best values are in bold and * indicates statistical significance ($p < 0.05$, paired t-test) between best and second best values. If no statistical significance is observed, we compare second and third best values.

For EMO-gpt, a model comparable in size to the baseline, improvements can be seen in macro and weighted averages. This indicates that learning to predict task-related components, such as the belief state and dialogue acts, provides useful information for ED. This observation aligns with the idea that emotions in TODs are closely linked to task progression, making task-related information beneficial. As for the PREV-gpt variant, the addition of predicted emotions to the Context introduces noise for *fearful* but proves beneficial for *dissatisfied*. In cases where the weighted average matters most, this PREV variant is preferable.

For Llama-2, the PREV variant shows a clearer benefit. This model performs best overall, except for the *dissatisfied* emotion. By examining user utterances with this particular emotion, we find that they appear highly contextual and generally lack explicit signs of dissatisfaction, a common trait when identifying emotions from text only [20]. The regularization from LoRA in this case is a drawback compared to full fine-tuning, as annotations for this emotion are dataset-specific. However, for less common emotions characterized by more explicit expressions, such as *fearful* and *abusive*, the use of an LLM with LoRA proves advantageous. Nonetheless, opportunities for further improvements remain, particularly in the recognition of *abusive* and *dissatisfied*, two crucial emotional states.

With regards to **TOD metrics** (Table 3), we compare our variants to a SimpleTOD baseline (SIMPLE) trained with each respective language model. For both GPT-2 and Llama-2, the incorporation of ED into the pipeline yields improved or comparable results. Notably, the inclusion of emotions does not adversely affect any metric with statistical significance, while both success and inform rates see improvements for the PREV

Model	Inform	Success	JGA	CBE	Unique tri.	BLEU
SIMPLE-gpt	81.98	75.72	65.06	1.64*	2336.8	22.29
EMO-gpt	82.12	76.18	64.71	1.61	2315.8	22.4
PREV-gpt	83.56*	78.04*	65.21	1.59	2291.6	22.55*
SIMPLE-llama	78.5	70.46	64.2	1.95	4054.8	22.78
EMO-llama	78.36	70.28	64.0	1.96	4151.2	22.91
PREV-llama	83.32*	75.14*	63.08	1.96	4309.0*	22.35

Table 3: Mean MultiWOZ metrics (5 seeds) and lexical diversity metrics. Best values are in bold and are computed for each respective language model. * indicates statistical significance ($p < 0.05$, paired t-test) between best and second best values.

Model	#1	#2	#3	Mean Rank	κ
SIMPLE	35.56%	40.56%	23.89%	1.88	0.57
PREV	40.0%	51.67%	8.33%	1.68	0.43
REFINE	70.0%	22.22%	7.78%	1.38	0.41

Table 4: Mean human rankings of responses from three Llama-2-based variants, following an emotional user utterance. The first three columns represent the distribution of rankings per model in %. Rank 1 is best, so a lower rank implies preferred responses. We compute Fleiss’s κ to characterize agreement on the ranks attributed to each model.

variant. This indicates that employing this approach helps to more closely replicate reference dialogue acts and responses, which largely determine these improved metrics.

On the other hand, the EMO variant does not offer as much signal, yielding results similar to the baseline in each case. Compared to GPT-2, Llama-2 responses exhibit greater diversity, as indicated by the CBE score and the unique trigram count, while simultaneously aligning similarly to the gold responses, as reflected by BLEU scores. However, inform and success rates are slightly lower than those of PREV-gpt.

Overall, these results validate the viability of unifying ED and TOD modeling. They suggest that incorporating user emotions introduces minimal noise and, in fact, contributes meaningful signal.

Human Evaluation

To assess the impact of conditioning responses on explicit emotions, we conduct an in-house human evaluation using delexicalized Llama-2-based responses (Table 4). Our comparison involves baseline responses (SIMPLE), responses simply conditioned on user emotions through joint training (PREV), and these same PREV responses more explicitly conditioned via prompting (REFINE). We adopt a methodology inspired by [13] which involves side-by-side relative ranking, randomly selecting examples from each emotional category (*neutral* excluded). We create a collection of 60 evaluation examples, annotated by three distinct NLP practitioners.

For each example, participants are provided with the Context, the reference emotion label for the last user utterance and three responses to rank from best to worst. We allow the same rating to be applied to multiple responses if annotators perceive them as highly similar. Annotators are explicitly instructed to rate responses based on how well they account for the user’s emotion and on how well they assist the user in completing their booking. Our annotation interface is shown in Appendix 1.

Our findings reveal a preference for REFINE and PREV responses, compared to baseline SIMPLE responses. REFINE especially demonstrates a more consistent understanding of the user’s emotional state and drives engagement by making the user feel heard (Table 5). However, despite REFINE achieving a mean rank closer to 1, we identify instances where its responses are less coherent than those from the PREV variant. This discrepancy arises when the additional segments are redundant or confusing, which is sometimes caused by inaccuracies in emotion prediction. PREV proves more resilient to such errors, although it generally displays less emotion awareness. Annotators also at times observe minimal variation between responses, suggesting that while our enhanced variants stay on track with respect to the task, there is room for improvement in showcasing empathy.

<i>Dialogue History</i>	<i>Responses</i>
1. (...) <i>User</i> : Can you tell me the address to the police station? I am in trouble and I need some help asap. (<i>fearful</i>)	SIMPLE : The address is [address] and the phone number is [phone]. PREV : The police station is located at [address]. REFINE : Oh no, I'm so sorry to hear that! The police station is located at [address].
2. (...) <i>User</i> : It's our anniversary and my wife loves british food. I want it to be special, so maybe something expensive. (<i>excited</i>)	SIMPLE : I have [choice] options for you. [name] , [name] , and [name] . Do any of these sound good? PREV : I have [name] in the [area] that serves [food] food and is in the [pricerange] price range. REFINE : Of course! I'd be delighted to help you find the perfect place for your anniversary! I have [name] in the [area] that serves [food] food and is in the [pricerange] price range.
3. (...) <i>User</i> : You didn't answer my question. I need the postcode for The Cambridge Belfry and I need to know if they have free parking or not. (<i>abusive</i>)	SIMPLE : The postcode for [name] is [postcode] and they do have free parking. PREV : I'm sorry about that. The postcode for [name] is [postcode] . They do have free parking. Is there anything else I can help you with? REFINE : I apologize, I did not see the question. I'm sorry about that. The postcode for [name] is [postcode] . They do have free parking. Is there anything else I can help you with?
4. (...) <i>System</i> : Is there anything else you need help with today? <i>User</i> : That shall be all. Thanks! (<i>satisfied</i>)	SIMPLE : Thank you for using our service. PREV : Thank you for using our service. REFINE : Yes you did! You're very welcome! Thank you for using our service.

Table 5: Examples of delexicalized system responses following emotional user turns (dialogues pmul3748, pmul4894, pmul4186 and sng1016). Examples 1. and 2. demonstrate how REFINE can successfully add emotion-awareness to responses. Example 3. and 4. respectively show how the added snippets can however be slightly redundant or confusing at times. Example 3. also shows PREV providing adequate emotion awareness compared with SIMPLE.

4 Related work

Work relating to customer support dialogues focuses on the intersection of ED and TOD as these exchanges often stem from a poor user experience, making ED key. For instance, [4] annotate emotions in a small corpus of 100 spoken call-center dialogues. [27] also provide valuable insight, examining emotion intensity and frustration in customer support dialogues. However, the datasets utilized offer limited task-oriented annotations and lack explicit ties to a database, making it challenging to experiment with typical multi-domain TOD system baselines. [24] take an interest in generating emotion-aware responses. In contrast to our work, they assume the user's emotion is given and focus on a dataset of general social media post/response pairs, very different from the TODs to be found in MultiWOZ. [6] enhance emotion recognition models for chitchat and apply them to the complete EmoWOZ dataset. The full dataset contains an extra 1000 human-machine dialogues, not annotated with task information. Due to our specific approach, we focus solely on experimenting with the MultiWOZ portion of EmoWOZ and do not compare ED results with this study. Other research explores the use of emotionally-grounded user simulators for training dialogue policy models through reinforcement learning [11]. In this setting, however, the system is not expected to predict the associated emotion label, in contrast to our endeavor.

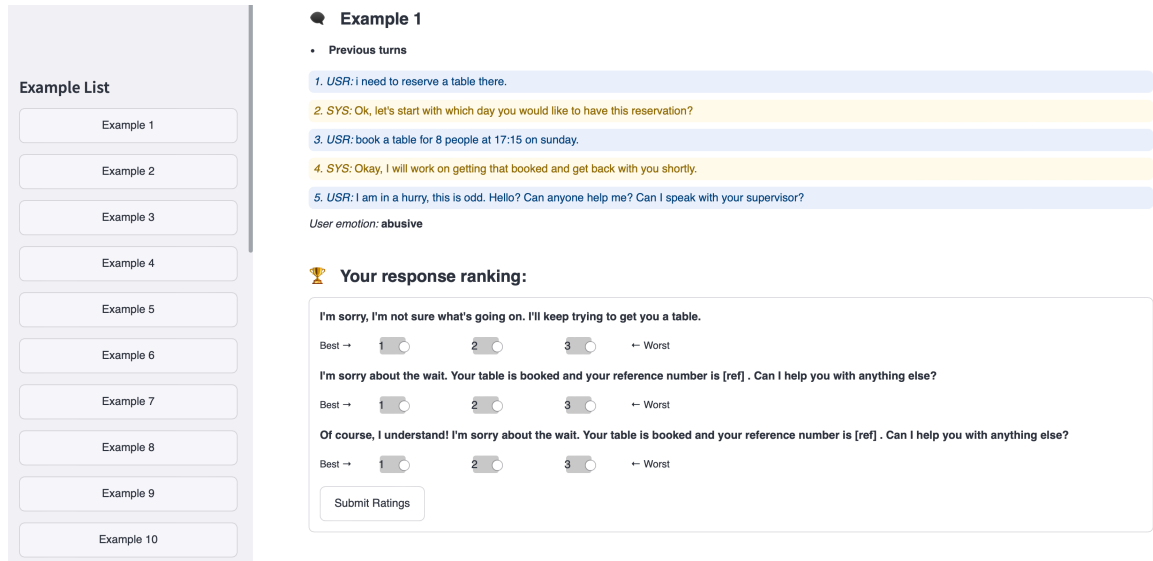
5 Conclusion

Our findings demonstrate that unifying ED and TOD modeling is overall beneficial for both tasks, producing improved results across ED and task-oriented metrics. Additionally, our experiments reveal that explicitly

predicting emotion labels provides useful grounding for system responses, resulting in a heightened display of empathy especially when refining responses. In a nutshell, our findings underscore the advantages and ease of integrating ED into TOD modeling, compared with treating it as an independent task. We believe our contribution will stimulate TOD system developments which extend beyond the sole objective of task completion, fostering more user-focused innovations.

Appendix 1: Human Evaluation Interface

We adapt the graphical user interface from [13] to fit our particular needs (Figure 2)



The interface is divided into two main sections. On the left is the 'Example List' with buttons for 'Example 1' through 'Example 10'. On the right, 'Example 1' is expanded, showing a conversation history with five turns. Turn 1 (USR) is 'i need to reserve a table there.' Turn 2 (SYS) is 'Ok, let's start with which day you would like to have this reservation?' Turn 3 (USR) is 'book a table for 8 people at 17:15 on sunday.' Turn 4 (SYS) is 'Okay, I will work on getting that booked and get back with you shortly.' Turn 5 (USR) is 'I am in a hurry, this is odd. Hello? Can anyone help me? Can I speak with your supervisor?'. Below the history, the 'User emotion' is listed as 'abusive'. The 'Your response ranking' section contains three response options, each with a 'Best' to 'Worst' slider (1 to 3). The responses are: 'I'm sorry, I'm not sure what's going on. I'll keep trying to get you a table.', 'I'm sorry about the wait. Your table is booked and your reference number is [ref] . Can I help you with anything else?', and 'Of course, I understand! I'm sorry about the wait. Your table is booked and your reference number is [ref] . Can I help you with anything else?'. A 'Submit Ratings' button is at the bottom.

Example List

- Example 1
- Example 2
- Example 3
- Example 4
- Example 5
- Example 6
- Example 7
- Example 8
- Example 9
- Example 10

Example 1

- Previous turns

1. USR: i need to reserve a table there.
2. SYS: Ok, let's start with which day you would like to have this reservation?
3. USR: book a table for 8 people at 17:15 on sunday.
4. SYS: Okay, I will work on getting that booked and get back with you shortly.
5. USR: I am in a hurry, this is odd. Hello? Can anyone help me? Can I speak with your supervisor?

User emotion: abusive

Your response ranking:

I'm sorry, I'm not sure what's going on. I'll keep trying to get you a table.

Best → 1 2 3 ← Worst

I'm sorry about the wait. Your table is booked and your reference number is [ref] . Can I help you with anything else?

Best → 1 2 3 ← Worst

Of course, I understand! I'm sorry about the wait. Your table is booked and your reference number is [ref] . Can I help you with anything else?

Best → 1 2 3 ← Worst

Submit Ratings

Fig. 2: Our graphical user interface used for human evaluations.

Appendix 2: Training and Implementation Details

We train models with the huggingface⁴ and pytorch frameworks. For each model, we perform grid search over learning rates ranging from 1e-5 to 2e-4 with an effective batch size of 32 (we use batches of 16 and 2 gradient accumulation steps). We use the validation set to pick the most effective learning rates. When training GPT-2, we use a learning rate of 8e-5 and train the model for up to 10 epochs with early stopping. When performing LoRA fine-tuning, we choose the following LoRA configuration: a rank of size 32, given the complexity of our task, a scaling α of 32, and target modules which include the query and value projection matrices, found in the self-attention module of each transformer block. We use a learning rate of 5e-5 for SIMPLE and EMO and 4e-5 for PREV, training for up to 3 epochs with early stopping. Training and inference is done on a single 80Gb A100 GPU using 5 random seeds for each experiment. We generate outputs using greedy decoding and prevent the models from repeating n-grams of size 10, to avoid repetitive loops that may arise when using greedy search. For prompting, we use sampling with a temperature of 0.9.

Appendix 3: REFINE Prompt Template

We use a few-shot chain-of-thought template to prompt Llama-2. Exemplars follow the template below, which include a *Thought* component highlighting the behavior to exhibit in the generated snippet. We concatenate one exemplar per emotion, excluding *neutral*.

Instruction:

In the provided dialogues, you'll find conversation contexts involving a user and an information system that assists with various tasks. The user's emotion is indicated. If the original response does not address the user's emotion, your task is to write a brief snippet that can be added BEFORE the original response. The snippet should convey remorse, reassurance, appeasement, or enthusiasm, depending on the user's emotion. Refrain from repeating previous snippets or the original response.

Context:

```
<|user|> While in Cambridge I need a restaurant that is moderately
    priced and serves Spanish food.
<|system|> La Tasca would be a good choice for you. Would you like
    the address or phone number?
<|user|> Yes I would like the address and phone number, thanks for
    your help!
<|system|> The address for La Tasca is 14 -16 Bridge Street and the
    phone number is 01223464630. Would you like for me to book you
    a table?
<|user|> I just want to confirm. This place has free wifi as well as
    being 3 star, correct?
<|system|> Were you also trying to book a hotel today? There are 5
    hotels in cambridge that are 3 stars and offer free wifi
<|user|> Actually, I'd like to confirm that La Tasca is a 3-star
    restaurant.
<|system|> Thank you. Have a good day!
<|user|> Wait please you haven't answered my question.
```

⁴ <https://huggingface.co/>

User is feeling:
dissatisfied, disliking

Original Response:
<|system|> [name] is the restaurant and is not assigned a star rating

Thought:
Displaying remorse would be appropriate.

Add before the original response:
Oh I'm sorry!

(...)

Appendix 4: Emotion Label Examples from EmoWOZ

User utterances and their associated emotion labels from EmoWOZ [5] can be found in table 6. We direct readers to the paper for more details.

<i>Emotion</i>	<i>Example</i>
Neutral	<i>System:</i> What sort of food would you like it to be? <i>User:</i> You choose. Book me a table for 6 people at 12:00 on Thursday.
Fearful, sad, disappointed	(Explicit expression) <i>System:</i> I have booked you clare hall <i>User:</i> What? that's not what i wanted. (Implicit expression) <i>User:</i> I need a taxi from the hotel to the museum after 23:45. <i>System:</i> Do you want the hotel reservations to begin on Monday? <i>User:</i> We're talking about a taxi now. <i>System:</i> You would love broughton house gallery. <i>User:</i> Taxi.
Apologetic	<i>System:</i> I'm afraid that's all we've got. <i>User:</i> I'm sorry I may have confused you. I would want to go from Petersborough to Cambridge on Saturday and arrive by 13:00, so anything that will get me there before that.
Abusive	<i>System:</i> (talking about something irrelevant) <i>User:</i> *****, would ya stupid thing.
Excited, happy, anticipating	<i>System:</i> There are 3 restaurants on the west side that serve British food. Do you prefer moderate or expensive priced? <i>User:</i> Expensive, please, we're splurging for my mother's birthday
Satisfied, liking, appreciative	<i>System:</i> Any time. I was happy to help. <i>User:</i> You are doing a great job! That's all I needed.

Table 6: Examples for each emotion label from EmoWOZ [5].

References

- [1] Tom B. Brown et al. *Language Models are Few-Shot Learners*. 2020. arXiv: 2005 . 14165 [cs.CL].
- [2] Paweł Budzianowski et al. “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*. Brussels, Belgium: Association for Computational Linguistics, Oct. 2018, pp. 5016–5026. DOI: 10.18653/v1/D18-1547.
- [3] Zhiyu Chen et al. “KETOD: Knowledge-Enriched Task-Oriented Dialogue”. In: *Findings of the Association for Computational Linguistics: NAACL 2022*. Seattle, United States: Association for Computational Linguistics, July 2022, pp. 2581–2593. DOI: 10.18653/v1/2022.findings-naacl.197.
- [4] L. Devillers, L. Lamel, and I. Vasilescu. “Emotion detection in task-oriented spoken dialogues”. In: *2003 International Conference on Multimedia and Expo. ICME '03. Proceedings (Cat. No.03TH8698)*. Vol. 3. 2003, pp. III–549. DOI: 10.1109/ICME.2003.1221370.
- [5] Shutong Feng et al. “EmoWOZ: A Large-Scale Corpus and Labelling Scheme for Emotion Recognition in Task-Oriented Dialogue Systems”. In: *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, June 2022, pp. 4096–4113.
- [6] Shutong Feng et al. “From Chatter to Matter: Addressing Critical Steps of Emotion Recognition Learning in Task-oriented Dialogue”. In: *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*. Ed. by David Schlangen et al. Prague, Czechia: Association for Computational Linguistics, Sept. 2023, pp. 85–103.
- [7] Donghoon Ham et al. “End-to-End Neural Pipeline for Goal-Oriented Dialogue Systems using GPT-2”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, July 2020, pp. 583–592. DOI: 10.18653/v1/2020.acl-main.54.
- [8] Ehsan Hosseini-Asl et al. “A Simple Language Model for Task-Oriented Dialogue”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle et al. Vol. 33. Curran Associates, Inc., 2020, pp. 20179–20191.
- [9] Edward J Hu et al. “Lora: Low-rank adaptation of large language models”. In: *arXiv preprint arXiv:2106.09685* (2021).
- [10] Yanran Li et al. “DailyDialog: A Manually Labelled Multi-turn Dialogue Dataset”. In: *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Taipei, Taiwan: Asian Federation of Natural Language Processing, Nov. 2017, pp. 986–995.
- [11] Hsien-Chin Lin et al. “EmoUS: Simulating User Emotions in Task-Oriented Dialogues”. In: *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2023, pp. 2526–2531.
- [12] Syaheerah Lutfi et al. “I Feel You: The Design and Evaluation of a Domotic Affect-Sensitive Spoken Conversational Agent”. In: *Sensors* 13.8 (Aug. 2013), pp. 10519–10538. ISSN: 1424-8220. DOI: 10.3390/s130810519.
- [13] Tomáš Nekvinda and Ondřej Dušek. “AARGH! End-to-end Retrieval-Generation for Task-Oriented Dialog”. In: *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Edinburgh, UK: Association for Computational Linguistics, Sept. 2022, pp. 283–297.
- [14] Tomáš Nekvinda and Ondřej Dušek. “Shades of BLEU, Flavours of Success: The Case of Multi-WOZ”. In: *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Met-*

- rics (GEM 2021)*. Online: Association for Computational Linguistics, Aug. 2021, pp. 34–46. DOI: 10.18653/v1/2021.gem-1.4.
- [15] Andrew Ortony, Gerald L. Clore, and Allan Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988. DOI: 10.1017/CBO9780511571299.
 - [16] Kishore Papineni et al. “Bleu: a Method for Automatic Evaluation of Machine Translation”. In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics, July 2002, pp. 311–318. DOI: 10.3115/1073083.1073135.
 - [17] Alec Radford et al. “Language models are unsupervised multitask learners”. In: *OpenAI blog* 1.8 (2019), p. 9.
 - [18] Hannah Rashkin et al. “Towards Empathetic Open-domain Conversation Models: A New Benchmark and Dataset”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 5370–5381. DOI: 10.18653/v1/P19-1534.
 - [19] Byron Reeves and Clifford Nass. “The media equation: How people treat computers, television, and new media like real people”. In: *Cambridge, UK* 10.10 (1996).
 - [20] Siddhant Bikram Shah, Shubham Garg, and Aikaterini Bourazeri. “Emotion Recognition in Speech by Multimodal Analysis of Audio and Text”. In: *2023 13th International Conference on Cloud Computing, Data Science & Engineering (Confluence)*. 2023, pp. 257–263. DOI: 10.1109/Confluence56041.2023.10048872.
 - [21] Kai Sun et al. “Adding Chit-Chat to Enhance Task-Oriented Dialogues”. In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, June 2021, pp. 1570–1583. DOI: 10.18653/v1/2021.naacl-main.124.
 - [22] Hugo Touvron et al. *Llama 2: Open Foundation and Fine-Tuned Chat Models*. 2023. arXiv: 2307.09288 [cs.CL].
 - [23] Jason Wei et al. “Chain-of-thought prompting elicits reasoning in large language models”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 24824–24837.
 - [24] Wei Wei et al. “Target-guided emotion-aware chat machine”. In: *ACM Transactions on Information Systems (TOIS)* 39.4 (2021), pp. 1–24.
 - [25] Anbang Xu et al. “A new chatbot for customer service on social media”. In: *Proceedings of the 2017 CHI conference on human factors in computing systems*. 2017, pp. 3506–3510.
 - [26] Yunyi Yang, Yunhao Li, and Xiaojun Quan. “UBAR: Towards fully end-to-end task-oriented dialog system with GPT-2”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 16. 2021, pp. 14230–14238.
 - [27] Janis Zuters and Viktorija Leonova. “Adaptive Vocabulary Construction for Frustration Intensity Modelling in Customer Support Dialog Texts”. In: *International Journal of Computer Science & Information Technology (IJCSIT) Vol 12* (2020).