

Response-conditioned Turn-taking Prediction

Bing'er Jiang and Erik Ekstedt and Gabriel Skantze

Division of Speech, Music and Hearing, KTH Royal Institute of Technology
{binger, erikekst, skantze}@kth.se

Abstract

Previous approaches to turn-taking and response generation in conversational systems have treated it as a two-stage process: First, the end of a turn is detected (based on conversation history), then the system generates an appropriate response. Humans, however, do not take the turn just because it is likely, but also consider whether what they want to say fits the position. In this paper, we present a model (an extension of TurnGPT) that conditions the end-of-turn prediction on both conversation history and what the next speaker wants to say. We find that our model consistently outperforms the baseline model on a variety of metrics. The improvement is most prominent in two scenarios where turn predictions can be ambiguous solely from the conversation history: 1) when the current utterance contains a statement followed by a question; 2) when the end of the current utterance semantically matches the response. Treating the turn-prediction and response-ranking as a one-stage process, our findings suggest that our model can be used as an incremental response ranker, which can be applied in various settings.

1 Introduction

A fundamental component of spoken dialogue system (SDS) is turn-taking, i.e., the decision of when to take turns at appropriate places, without causing long response delays or interrupting the user. In other words, the system must be able to correctly identify when the user is yielding the turn, and it is appropriate to make a response, and when the user is simply making a mid-utterance pause (Skantze, 2021). Traditionally, this has been done using a simple silence threshold. However, silence is not a very good indicator of turn-shifts and more modern approaches instead use various cues known to be important in human-human turn-taking, such as lexico-syntactic cues, prosody, or gaze (Gravano and Hirschberg, 2011; Ishii et al., 2016; Lala et al., 2019; Ekstedt and Skantze, 2022).

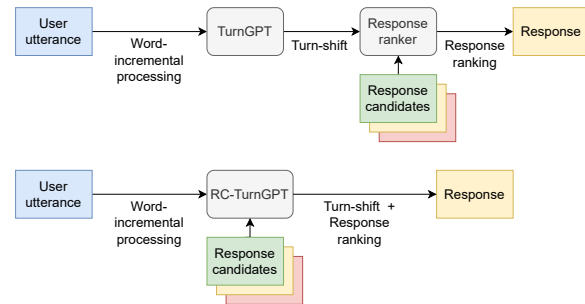


Figure 1: Response-conditioned turn-taking prediction (bottom) compared to traditional turn-taking prediction (top), in a dialogue system using a response ranker.

Ekstedt and Skantze (2020) proposed TurnGPT, a transformer-based language model that incrementally processes words in the user’s utterance and predicts the probability of a turn-shift after each word. This is similar to the notion of syntactic or pragmatic completion points that have been identified in conversation analysis (Ford and Thompson, 1996). In their analysis of TurnGPT, Ekstedt and Skantze (2020) found that the 20% of the model’s attention is directed towards utterances earlier than the current one, indicating that it is sensitive to pragmatic aspects of dialogue.

While such models are indeed a step forward, there is still an important component missing that we will address in this paper. When humans make a decision to take the turn, it is not just based on whether there are enough turn-yielding cues in the interlocutor’s utterance. Sacks et al. (1974) use the notion of transition-relevant places, or TRP, for places where a transition could potentially take place (but does not have to). Thus, many places for turn-shifts are highly optional. To partly address this problem, Ishii et al. (2022) annotated the willingness of the next speaker to take the turn, and built a model that could predict this willingness based on multimodal cues.

Whether a turn-shift takes place or not also depends on the intention of the next speaker, and what

they want to say. For dialogue systems, this means that the system should not automatically take the turn once the transition-probability passes a certain threshold, and only then decide what it should respond. Instead, the system should take the potential response into account when deciding whether it is appropriate to take the turn or not.

We call this **response-conditioned turn-taking prediction**, illustrated in Figure 1. We present a model called **RC-TurnGPT**, which is an extension of TurnGPT. Note that the current study does *not* intend to address *how* and *when* the next speaker comes up with what they would like to say. This depends of course on the exact implementation of the dialogue system, which could for example be response-ranking (Gao et al., 2020) or an intent-based planning approach (FAIR et al., 2022). In Figure 1, we have assumed that a response ranker is used. If so, a traditional system would first use a model like TurnGPT to decide when to take the turn, and then ask the response ranker which response would fit best. In such a setting, it might be the case that none of the candidates would be a good fit from the system’s perspective, but the system would produce a response anyway. In such a setting, RC-TurnGPT could instead be used to incrementally rank or score potential responses to see whether they fit well from a turn-taking perspective, or pass taking the turn if none of them has a high enough utility.

In this paper, we take a first step towards such an approach, and investigate to what extent and under what scenarios such response-conditioning would help to predict turn-shifts. Similar to TurnGPT, we do not model acoustic information, as our focus is to investigate how the semantic and pragmatic aspects of the dialogue affect turn-shift prediction. Instead, we use written dialogues as a stand-in for audio for incremental end-of-turn prediction. We leave the incorporation of acoustic information (cf. Ekstedt and Skantze 2022) for future work.

2 Methods

TurnGPT is a unidirectional transformer-based language model (LM) optimized through cross-entropy to predict the next token in a sequence. It is a pre-trained GPT-2 (base) model (Radford et al., 2019), finetuned on *unpunctuated* dialogue corpora, with a special turn-shift token (TS) that delimits consecutive turns. RC-TurnGPT is an extension of this model, by also conditioning the prediction on

the response.

While the RC-TurnGPT model is architecturally equivalent to TurnGPT, it differs in the training objective through a simple data transformation. This transformation permutes the ordering of turns in a similar approach as the FIM pre-training objective of Bavarian et al. (2022). We consider turn-based dialogue sequences to consist of three parts: the context/history (**H**), the current utterance (**CU**) and the next response (**R**). The task is to correctly predict the location of the turn-shift token in the current utterance, CU_i , given the history, H_i , and the next response, R_i , over all samples i in the dataset, D . The samples $i \in D_I$ are extracted by applying a turn-based sliding window approach with a step size of 1 and a window size of 3 turns.

However, instead of the uniform left-to-right next token prediction task of regular LMs, the RC-TurnGPT model train on ordered sequences of $\{R, H, CU\}$, masking the loss over R and H to solely learn over the CU turns. This enables the model to use information of both H and R while keeping the original left-to-right next token prediction setup.

Finally, the TurnGPT model utilized three special tokens in addition to the original GPT-2 vocabulary, the aforementioned TS token and two speaker tokens. The speaker tokens are similar to positional embeddings and are added to the word embeddings to encode the speaker identity over each word. Because of the permuted ordering of the RC-TurnGPT setup we also include a fourth special response-token that are added to the words of the response to distinguish them from the actual context. Both the base model and the datasets were implemented using Huggingface (Wolf et al., 2020; Lhoest et al., 2021).

2.1 Data

We train RC-TurnGPT and the baseline TurnGPT on two types of data sets based on Ekstedt and Skantze (2020): **Assistant** and **Written Social**. The former constitutes of three task-oriented dialogue corpora: Taskmaster (Byrne et al., 2019), MetaLWOZ (Lee et al., 2019), and MultiWoz (Zang et al., 2020). The latter includes two corpora constructed by human-human written dialogues: CuriosityDialogs (Rodriguez et al., 2020) and DailyDialog (Li et al., 2017). All datasets are written dialogues with clearly defined turns. The resulting full dataset contains 106,830 dialogues for training, 9,362 for validation, and

7,897 for test, with an average number of turns being 13.69.

2.2 Evaluation

To evaluate the models, we propose five turn-level based metrics that measures the turn-shift performance in various ways. The models are considered to make a turn-shift prediction when the probability exceeds a certain threshold optimized for performance over the validation split, for each model independently.

First, we define turn-level accuracy (**TL-Acc**) to be the percentage of turns where the turn-shift probability exceeds the threshold at, and *only* at, the ground-truth end of turn. Second, the no response rate (**NRR**) is the percentage of turns where the threshold is never exceeded and the model fails to make a response. The third metric is defined to measure the barge-in rate (**BR**), the percentage of turns where the models would make a turn-shift prediction before the actual turn-shift.

We also investigate instances where the two models make different turn-taking decisions to see how well the response would fit, using perplexity as a measure. We use the TurnGPT model to calculate the average perplexity over the response (**R-PPL**).

Lastly, we define the ordinal spike rate (**OSR**) to be the percentage of turns where the probability is the greatest at the end of the turn. This metric does not consider a threshold but simply measures how many times the highest probability is located at the correct turn-shift location.

3 Results

3.1 Aggregate results

Table 1 shows that RC-TurnGPT performs better in all evaluations metrics, although the improvement is not large overall. While 55.77% turn-level accuracy may not seem very high, it should be noted that even predictions different from ground-truth turn-shift can also be valid in everyday conversations, especially in long utterances where several completion points are likely. While the threshold-based binary metric is low, the probability-based OSR is much higher, indicating that the model is indeed able to detect end of turn reflected by assigning the highest probability. Furthermore, the perplexity of the response also decreases, showing that when one or both of the two models make a mistake, the response fits better with the context for the turn-shifts RC-TurnGPT takes.

Metric	Turn-GPT	RC-TurnGPT
TL-Acc ↑	53.93%	55.77%
NRR ↓	20.90%	19.23%
BR ↓	25.17%	24.75%
R-PPL ↓	1.923	1.918
OSR ↑	88.57%	89.17%

Table 1: The turn-level accuracy (TL-Acc), no response rate (NRR), barge-in rate (BR), response perplexity (R-PPL) and the ordinal spike rate (OSR) performance for TurnGPT and RC-TurnGPT. Best performance are bold.

3.2 Model analysis

In order to better understand *when* conditioning on the response helps turn-shift prediction and when it does not, we proceed to analyse cases where only RC-TurnGPT makes the correct prediction, and where both models are successful.

We extract all turns in the test set where TurnGPT makes a pre-mature turn-shift prediction but RC-TurnGPT correctly predicts the end of the turn. We sort the turns by the difference in probability assigned by the two models at the TurnGPT-predicted turn-shift. We then investigate the difference between the top and bottom 1000 cases. By comparing these two subsets, we can better understand when conditioning on the response makes the biggest difference. We identified two scenarios which we hypothesized would be important: 1) statement to question; 2) semantic matching.

Statement to question refers to cases where the current utterance consists of at least one statement and ends with a question. As there are more than one natural completion point, TurnGPT will be greedy while RC-TurnGPT will take the response into consideration and choose a later completion point as turn shift. Consider the following dialogue in Figure 2 (Current Utterance plotted, Response in caption):

Figure 2 shows that without conditioning on the response, TurnGPT spikes at an early completion point interrupting the current speaker. However, as the response clearly corresponds to an answer to a request, RC-TurnGPT waits until the speaker finishes their request.

In order to quantify this effect, we use punctuations to calculate how often TurnGPT makes a mistake by missing a question. We use the top/bottom subsets and ask GPT3¹ (Brown et al.,

¹Model version: “text-curie-001”

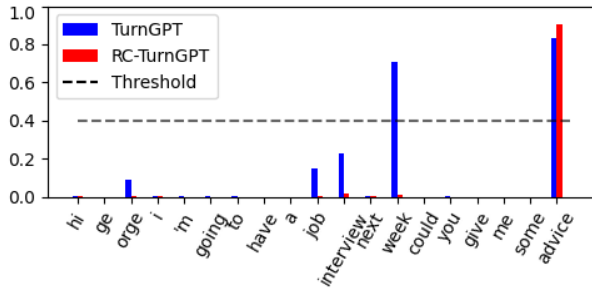


Figure 2: Different turn-taking predictions: TurnGPT predicts the turn-shift at the end of a statement; RC-TurnGPT predicts the end of a question. Response: sure first of all it's very important for you not to be late

2020) to insert punctuation over the ground truth turns (advice in this example) and the incomplete TurnGPT predicted turns (week in this example). We then calculate the ratio of cases where the former ends with a question mark while the latter does not. The top cases contain 36.3% statements to questions and the bottom 11.7%. The higher ratio in the top cases indicates that the RC-TurnGPT model recognizes this pattern and uses the response conditioning to wait for the appropriate moment to take the turn.

Semantic matching refers to cases where the response semantically corresponds to the specification made in the later parts of the current utterance. Consider the dialogue in Figure 3:

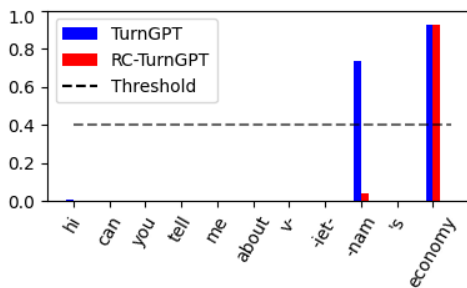


Figure 3: Different turn-taking predictions: RC-TurnGPT's prediction allows closer semantic matching between current utterance and response. Response: sure vietnam achieved an 8% gdp growth between 1990 and 1997

As the response clearly addresses the topic of economy, Figure 3 shows that RC-TurnGPT would spike only after economy is specified, whereas TurnGPT has two spikes at both places and would predict the turn shift after v-iet-nam. It is important to note that while the response has no lexical

overlap, the model still manages to find the semantic correlation.

In order to investigate whether RC-TurnGPT consistently recognizes such pattern, we use Sentence-Bert (Reimers and Gurevych, 2019) to measure the Semantic Textual Similarity between the Response and the last part of the actual turns missed by TurnGPT (here, 's economy). The average cosine distance for the top and bottom subsets are 0.293 and 0.209 respectively. This indicates that where RC-TurnGPT outperforms TurnGPT, it does consider the semantic content of the response and delays predicting a turn-shift until the relevant semantic information has been stated.

Non-ambiguous turn-completions. In addition, there are also a large number of cases where the current utterance has a fairly simple structure and hence it is not ambiguous where to take the turn. In those cases, conditioning on the next response obviously makes a very small difference. As illustrated in Figure 4, given that there is only one completion point, both models predict the turn shift correctly. This also explains why there are no drastic improvements for RC-TurnGPT when looking at aggregate results on the whole test set, as most of the task-oriented dialogues contain such simple utterances, which TurnGPT can perform well on.

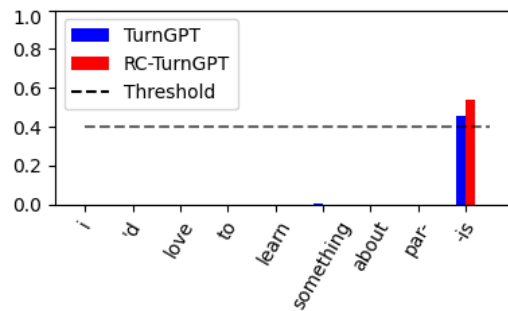


Figure 4: Similar turn-taking predictions for a simple utterance. Response: it is the capital of france

4 Discussion and conclusion

In this study, we examined how turn-taking prediction can be improved when conditioned on the response. We found that the response conditioning is particularly helpful under two circumstances, mainly by preventing greedy turn-taking at earlier completion point: 1) when the current utterance contains statements followed by questions; 2) when the end of the current utterance semantically

matches the response. However, for simple utterances with fewer completion points, TurnGPT is already capable of predicting the correct turn shift, and there is no additional help from conditioning on the response.

We should again stress that this paper does not address the question of how and when the system comes up with a potential response. However, our analysis shows that it is indeed possible to find a more suitable transition-point, when conditioning on the response. As we have suggested, the decision *what* to say and *when* to say it should be considered as a joint decision rather than a two-step process. We acknowledge the fact that this would be problematic if one assume a system using a response generator such as GPT (Brown et al., 2020), as such models generate responses conditioned on a turn-shift already being decided.

However, the RC-TurnGPT model could be used as an *incremental response ranker*, which does not only consider different responses at each step, but which can also decide *not* to respond and wait for more input. For instance, it can be applied in an interview setting where the model (interviewer) asks questions (ranking from a list of interview questions) and take the turn at appropriate places. For future work, it would also be interesting to involve the *utility* of the candidate responses (from the system’s perspective). In the interview scenario, this could for example mean that the system can find moments where certain important questions can be asked, and which also fit well from a turn-taking perspective.

Limitations

As mentioned above, the current study is limited to the question of whether (and when) conditioning turn-taking prediction on the response improves the performance. It does not yet show how the model could be incorporated in a spoken dialogue system. Moreover, this study focuses only on written conversations without incorporating spoken dialogues. Thus, the interpretations can be limited to dialogues that are relatively ‘formal’ without hesitations, repetitions, etc. Note also that we only analyse lexical cues to turn-taking (just like with TurnGPT), and leave out other modalities for future work.

Ethics Statement

The current study does not involve any human subjects and we do not foresee any ethical con-

sequences.

Acknowledgements

This work was supported by Riksbankens Jubileumsfond (RJ), through the project *Understanding predictive models of turn-taking in spoken interaction* (P20-0484), as well as the Swedish Research Council, through the project *Prediction and Coordination for Conversational AI* (2020-03812).

References

- Mohammad Bavarian, Heewoo Jun, Nikolas Tezak, John Schulman, Christine McLeavey, Jerry Tworek, and Mark Chen. 2022. [Efficient training of language models to fill in the middle](#).
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. 2019. [Taskmaster-1: Toward a realistic and diverse dialog dataset](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4516–4525, Hong Kong, China. Association for Computational Linguistics.
- Erik Ekstedt and Gabriel Skantze. 2020. [TurnGPT: a transformer-based language model for predicting turn-taking in spoken dialog](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2981–2990, Online. Association for Computational Linguistics.
- Erik Ekstedt and Gabriel Skantze. 2022. How Much Does Prosody Help Turn-taking? Investigations using Voice Activity Projection Models. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 541–551, Edinburgh, UK. Association for Computational Linguistics.
- FAIR, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, Athul Paul Jacob, Mojtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander H.

- Miller, Sasha Mitts, Adithya Renduchintala, Stephen Roller, Dirk Rowe, Weiyan Shi, Joe Spisak, Alexander Wei, David Wu, Hugh Zhang, and Markus Zijlstra. 2022. [Human-level play in the game of *Diplomacy* by combining language models with strategic reasoning](#). *Science*, 378(6624):1067–1074.
- C Ford and S Thompson. 1996. Interactional units in conversation: syntactic, intonational, and pragmatic resources for the management of turns. In E Ochs, E Schegloff, and A Thompson, editors, *Interaction and grammar*, Studies in interactional sociolinguistics 13, chapter 3, pages 134–184. Cambridge University Press, Cambridge.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and Bill Dolan. 2020. [Dialogue response ranking training with large-scale human feedback data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 386–395, Online. Association for Computational Linguistics.
- Agustin Gravano and Julia. Hirschberg. 2011. Turn-taking cues in task-oriented dialogue. *Computer Speech & Language*, 25(3):601–634.
- Ryo Ishii, Kazuhiro Otsuka, Shiro Kumano, and Junji Yamato. 2016. [Prediction of who will be the next speaker and when using gaze behavior in multiparty meetings](#). *ACM Transactions on Interactive Intelligent Systems*.
- Ryo Ishii, Xutong Ren, Michal Muszynski, and Louis-Philippe Morency. 2022. [Trimodal prediction of speaking and listening willingness to help improve turn-changing modeling](#). *Frontiers in Psychology*, 13:774547.
- Divesh Lala, Koji Inoue, and Tatsuya Kawahara. 2019. [Smooth turn-taking by a robot using an online continuous model to generate turn-taking cues](#). In *2019 International Conference on Multimodal Interaction, ICMI '19*, page 226–234, New York, NY, USA. Association for Computing Machinery.
- Sungjin Lee, Hannes Schulz, Adam Atkinson, Jianfeng Gao, Kaheer Suleman, Layla El Asri, Mahmoud Adada, Minlie Huang, Shikhar Sharma, Wendy Tay, and Xiujun Li. 2019. [Multi-domain task-completion dialog challenge](#). In *Dialog System Technology Challenges 8*.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Pedro Rodriguez, Paul Crook, Seungwhan Moon, and Zhiguang Wang. 2020. Information seeking in the spirit of learning: a dataset for conversational curiosity. In *Empirical Methods in Natural Language Processing*.
- H Sacks, Emanuel Schegloff, and G Jefferson. 1974. [A simplest systematics for the organization of turn-taking for conversation](#). *Language*, 50:696–735.
- Gabriel Skantze. 2021. [Turn-taking in Conversational Systems and Human-Robot Interaction : A Review](#). *Computer Speech & Language*, 67.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. Multiwoz 2.2: A dialogue dataset with additional annotation corrections and state tracking baselines. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI, ACL 2020*, pages 109–117.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
5. Limitations
- A2. Did you discuss any potential risks of your work?
Our work is on turn-shift prediction and does not have potential risks.
- A3. Do the abstract and introduction summarize the paper's main claims?
Section 1.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
No response.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
No response.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
No response.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
No response.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
No response.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
No response.

C Did you run computational experiments?

2. methods; 3. results

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Our model is fine-tuned based on GPT2, a widely-used model, so there is no need to further specify those information.

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

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

2. methods

C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

3. results.

C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

2. methods.

D **Did you use human annotators (e.g., crowdworkers) or research with human participants?**

Left blank.

D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.