

Commonsense-Aware Prompting for Controllable Empathetic Dialogue Generation

Yiren Liu ¹, Halil Kilicoglu ²

¹ Informatics, University of Illinois Urbana-Champaign

² School of Information Sciences, University of Illinois Urbana-Champaign
yiren12@illinois.edu, halil@illinois.edu

Abstract

Improving the emotional awareness of pre-trained language models is an emerging important problem for dialogue generation tasks. Although prior studies have introduced methods to improve empathetic dialogue generation, few have discussed how to incorporate commonsense knowledge into pre-trained language models for controllable dialogue generation. In this study, we propose a novel framework that improves empathetic dialogue generation using pre-trained language models by 1) incorporating commonsense knowledge through prompt verbalization, and 2) controlling dialogue generation using a strategy-driven future discriminator. We conducted experiments to reveal that both the incorporation of social commonsense knowledge and enforcement of control over generation help to improve generation performance. Finally, we discuss the implications of our study for future research.

Introduction

Empathetic dialogue generation has been an important task found to be beneficial for the applications of conversational agents in many domains, such as healthcare (Kim, Kaplowitz, and Johnston 2004; Nadarzynski et al. 2019; Athota et al. 2020) and mental health consulting (Lee et al. 2020). Recent research about pre-trained language models (PLMs, e.g., T5 and GPT-2) has advanced task performance related to dialogue generation. However, such language models still struggle to comprehend commonsense knowledge and emotions within dialogue contexts, as shown in Table. 1.

Research (Li et al. 2022) has pointed out that the ability to understand users’ emotions is crucial for dialogue systems to perform empathetic conversations. For humans, the perception of emotions is based on understanding and inference of commonsense knowledge that is often implicit during conversational interactions (Brown et al. 2020). However, more studies are needed to understand how commonsense knowledge can be incorporated into pre-trained language models during dialogue generation tasks. A promising avenue in this

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.
Accepted to Workshop on Knowledge Augmented Methods for Natural Language Processing, in conjunction with AAAI 2023.

Speaker:	<i>It was 100% their fault but they hit the water barrels and survived. They had no injuries but they almost ran me off the road.</i>
Listener:	<i>Did you suffer any injuries?</i>
Speaker:	<i>No I was not hit. It turned out they were drunk. I felt guilty but realized it was his fault.</i>
Listener (GPT-J 6B):	Why was it your fault?
Prompt:	The speaker was almost run over by a group with a vehicle. The speaker reacts with guilt.

Listener (GPT-J 6B):	Are you still angry at them?

Table 1: Example of how prompting commonsense knowledge can improve the empathy of dialogue generation.

respect is prompt-based methods, which enable direct injection of external knowledge by augmenting model input. In this work, we conduct experiments to provide further insight into whether and how commonsense knowledge injection using prompt-based methods can benefit the language model’s emotional awareness.

Additionally, Plug-and-Play methods (Dathathri et al. 2020) are also discovered recently as a promising direction for text generation. Yang and Klein (2021) introduced the method of using future discriminators (FUDGE), which enables conditioning of text generation over a given attribute based on pre-trained language models. The method has been evaluated over several different text generation tasks, including poetry generation, topic-based generation and formality transfer (Yang and Klein 2021). However, few have discussed the possibility of using plug-and-play methods to control generation in a dialogue context.

Two main contributions of this study include:

- introducing a prompt-based method to incorporate social commonsense knowledge into pre-trained language model for dialogue generation;
- proposing a strategy-controlled dialogue generation method that can be used on language models with or without finetuning.

Related Work

Using Commonsense Knowledge for Empathetic Dialogue Generation

Recent works have explored the ability of pre-trained language models in dialogue generations (Zhang et al. 2020b). Prior works have shown that these models tend to have difficulty in comprehending emotional expressions from users (Li et al. 2017). Sabour, Zheng, and Huang (2021) considered both commonsense knowledge (cognitive) and emotional factors (affective) when encoding knowledge for empathetic dialogue generation using a seq2seq model. This work is built upon the cognitive and affective aspects of the theory of empathy from (Davis 1983). Later work (Li et al. 2022) also explored a graph-based approach for encoding commonsense knowledge in combination with emotional intensity based on VAD (Valence, Arousal and Dominance) vectors. However, these works did not use PLMs but traditional encoder-decoder models that need to be trained from scratch, thus limiting their comprehension ability to the training dataset used.

Controllable Dialogue Generation

Prior studies have also discussed using external signals to condition dialogue generation. Xu, Wu, and Wu (2018) proposed the idea of controlling dialogue generation using dialogue acts, which is found to significantly improve the quality of generated dialogues. Wang et al. (2020) has also explored the possibility of using topics to guide dialogue generation. Further work by Wu et al. (2021) focused on using external knowledge from Wikipedia to avoid fact hallucination in dialogue generation. Yang and Klein (2021) proposed the method of FUDGE, an attribute future discriminator, that enables strong plug-and-play control over text generation. In this work, we propose to use plug-and-play methods in order to control empathetic dialogue generation, which enables controllable generation with or without finetuning the PLM.

Preliminaries

Problem Formulation

Empathetic Response Generation We formulate the task of empathetic dialogue response generation as follows: Given the dialogue history $C = [X_1, X_2, \dots, X_{T-1}]$, where T is the total number of dialogue utterances and X_i is the i -th utterance text sequence, generate the next dialogue response Y . Each utterance $X_i = [x_1, x_2, \dots, x_m]$ and the target utterance $Y = X_T = [y_1, y_2, \dots, y_n]$, where m and n denote the total number of word tokens of each utterance.

Dialogue Strategy Prediction The generation of empathetic responses can benefit from the guidance of proper dialogue strategies. In this work, we refer to the conversational skills introduced in Hill (2009) needed in delivering emotional support as dialogue strategies (Xu, Wu, and Wu 2018), e.g. *providing suggestions* and *self-disclosure*. We denote the total dialogue strategy set as $S = \{s_i | 1 \leq i \leq T\}$, where T is the total number of strategies available. With the

dialogue history C , we need to predict the next correct dialogue strategy $s_Y \in S$ as a conditional signal for generating the correct response Y .

COMET Commonsense

We utilize the COMET-ATOMIC₂₀²⁰ (Hwang et al. 2021) commonsense reasoning language model to augment dialogue history input with social-related commonsense knowledge. The commonsense knowledge is represented as relation-entailment tuples. For example, $([xReact, depressed])$ indicates that the subject in the dialogue (PersonX) has the reaction of feeling depressed. The COMET-ATOMIC₂₀²⁰ dataset contains 23 relation types that are categorized into three categorical types including social commonsense relations, physical relations, and event-centric relations. In this study, we use dialogue history as the context, and used the COMET-BART model pre-trained on the COMET-ATOMIC₂₀²⁰ dataset to generate the commonsense entailment.

Approach

As shown in Figure. 1, Our proposed method consists of three major components: 1) commonsense prompting, 2) dialogue strategy predictor and 3) FUDGE-controlled response generation.

Commonsense Prompt Verbalizer

We first use the BART model pre-trained on the COMET-ATOMIC₂₀²⁰ dataset to obtain implicit commonsense knowledge that can be deduced from the dialogue context. We decode the commonsense entailment using 10 social relations from COMET-ATOMIC₂₀²⁰, including $[oEffect]$, $[oReact]$, $[oWant]$, $[xAttr]$, $[xEffect]$, $[xIntent]$, $[xNeed]$, $[xReact]$, $[xReason]$ and $[xWant]$. Relations starting with x refers to the main subject of the conversation, and o refers to others besides the subject. The commonsense knowledge is represented in the form of relation-entailment tuples $(r_{i,j}, e_{i,j})$ for each dialogue history $X_i \in C$, where $1 \leq j \leq 10$ refers to the j -th social relation. We generate entailment with all 10 relations for each dialogue history.

To better help the language model to comprehend the commonsense relations obtained from COMET-BART, we verbalize the commonsense tuples with natural language templates using the templates proposed by Hosseini, Broniatowski, and Diab (2022). For example, we convert the tuple $([xReact], depressed)$ into a sentence “As a result, PersonX feels depressed.”. We verbalize each tuple obtained from entailment generation, and append it to dialogue history as the input text.

Dialogue Strategy Predictor

To provide stronger control over the response generation, we first need to identify the correct strategy to be deployed in the next utterance given an existing dialogue history. We approach the problem as a text classification task. In addition to using the LM for generating the strategy with constrained

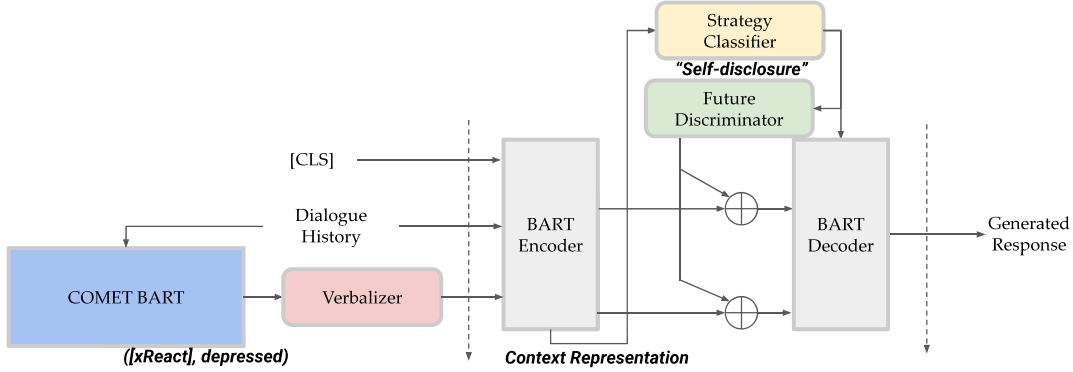


Figure 1: The Overall Framework of the Proposed Model.

decoding, we propose two methods to predict the next utterance strategies. First, we use the complete dialogue histories as the input text sequences and ground truth strategies as the labels to train an end-to-end text classification model. This method is agnostic to the generation LM and can ideally be swapped with any off-the-shelf text classification models available. The other approach is to adopt the encoder of the generation LM and use the encoder representation of the dialogue history text to train a classification module jointly with a generation objective. In this case, the training loss of the generation model for each utterance can be written as a combination of classification and generation loss as follows:

$$\ell_{strategy} = -\log p(s|Encoder([x_1, x_2, \dots, x_{t-1}])) \quad (1)$$

$$\ell_{LM} = -\sum_{t=1}^n \log p(x_t|x_1, x_2, \dots, x_{t-1}, s) \quad (2)$$

$$\ell = \ell_{LM} + \alpha \ell_{strategy} \quad (3)$$

where x_i is the i -th token within an utterance, α is the weight used to control the emphasis of strategy prediction over text generation. Note that strategy s remains the same for each target utterance. In this study, we experimented using a simple dense layer module as the strategy predictor, in which case we extract the encoder hidden representation at the [CLS] position as the context vector.

Future Discriminator

We then apply additional mechanisms to control the response generation with the predicted strategy. We adopt the method of FUDGE (Yang and Klein 2021) to enforce the strategy control. We introduce a future discriminator specifically for identifying dialogue strategies.

The general goal of the future discriminator is to classify the potential strategy category of a given dialogue utterance at each step of the token sequence. More specifically, for an utterance text sequence $X_i = [x_1, x_2, \dots, x_n]$, we train a text classifier to predict the ground truth strategy type of $\{[x_1, x_2, \dots, x_t] | 1 \leq t \leq n\}$ at every token step. Namely, we optimize for the log-likelihood sum for sequences from the first word to each word in each utterance with the loss:

$$\ell_F = -\sum_{t=1}^n \log p(s|x_1, x_2, \dots, x_t) \quad (4)$$

where n is the length of the target utterance sequence. While the objective is determined, the model used for classification can be any model separately trainable for text classification tasks. In our experiments, we follow Yang and Klein (2021)’s approach and discuss the performance of controlled generation using a light-weight LSTM text classifier, which requires a comparatively low computational cost to train and has been found by prior work (Schuurmans and Frasincar 2020) to have strong performance over dialogue intent classification task.

Experimental Settings

We experiment with our proposed method using the empathetic dialogue dataset ESConv (Liu et al. 2021). The ESConv dataset contains a total of 1,053 multi-turn dialogues with 31,410 utterances. The dataset consists of dialogue history between pairs of help-seekers and emotional support providers, including the cause and annotated dialogue strategies.

We compare the results with BART (Lewis et al. 2019) text generation model without any control. We use the BART-large model pre-trained with the objective of text denoising using corrupted input. In addition to BART, we also experimented with the COMET-BART model introduced by Hwang et al. (2021). Since COMET-BART is a text generation model pre-trained for commonsense reasoning, by further fine-tuning the model for dialogue generation we want to examine whether COMET-BART has a stronger capability in understanding commonsense knowledge input we provided. We use automatic evaluation metrics including BLEU (Papineni et al. 2002), ROUGE-L (Lin 2004), and BERTScore (Zhang et al. 2020a) to measure the dialogue response generation performance.

Results and Discussion

Main Results

We investigate the extent to which the introduction of commonsense knowledge input and controlled generation can improve the generation of dialogue response. Our results are shown in Table. 2. Detailed comparisons of generated examples are also provided in the Technical Appendix.

Incorporating commonsense knowledge helps the generation. We inform the language models with commonsense

Model	BLEU-1	BLEU-2	BLEU-3	BLEU-4	ROUGE-L	BERTScore	Strategy Accuracy
BART	17.79	5.74	2.43	1.18	14.15	90.77	30.97%
COMET-BART	16.39	6.73	3.47	2.11	16.23	90.97	31.94%
With COMET input (verbalized with templates)							
BART	16.31	5.90	2.71	1.45	15.31	90.86	30.46%
COMET-BART	18.53	7.80	4.12	2.54	17.79	91.21	30.40%
With strategy classifier (linear classifier with classification + LM loss)							
BART	20.12	8.19	4.16	2.47	15.67	90.97	16.29%
COMET-BART	9.28	3.90	2.13	1.35	13.56	90.36	16.29%
BART (+ FUDGE)	17.53	5.83	2.51	1.24	14.29	91.00	16.29%
COMET-BART (+ FUDGE)	6.60	2.11	1.05	0.59	11.82	90.72	16.29%
With strategy classifier (BERT classifier, trained separately from the LM)							
BART (+ FUDGE)	14.00	5.53	2.96	1.86	14.86	90.09	18.46%
COMET-BART (+ FUDGE)	4.74	1.71	0.86	0.52	11.36	90.08	18.46%

Table 2: The experimental results. The **Strategy Accuracy** column refers to the accuracy of predicted dialogue strategies. For the linear classifier, we train the classifier using a joint loss with LM, while the strategy classifier is trained separately.

	Dialogue strategy	Generated response
BART	[Providing Suggestions]	I understand that. I have been in a similar situation before and it's hard to know what to do.
+ FUDGE oracle	[Self-disclosure]	I had a friend die in a car accident and I was able to talk to him about how to deal with it.
COMET-BART	[Affirmation and Reassurance]	I can understand how you feel.
+ FUDGE oracle	[Self-disclosure]	I had a friend die in a car accident and I was there with him, and I was able to talk to him.

Table 3: Examples of generated Responses w/o FUDGE control.

Model	B-2	B-4	R-L	BERTScore
BART + oracle	5.02	1.32	18.34	90.59
COMET-BART + oracle	3.83	0.75	12.04	90.34
BART + oracle + FUDGE	7.46	2.01	21.18	91.11
COMET-BART + oracle + FUDGE	5.76	1.78	20.82	90.66

Table 4: Generation results using oracle strategies.

knowledge by concatenating verbalized COMET-generated reasoning knowledge with the dialogue history. By prompting the BART models with verbalized COMET commonsense knowledge, the results of models **with COMET input** in Table. 2 show consistently improved results in both BART and COMET-BART models. COMET-BART is observed to perform better than vanilla BART regardless of whether the commonsense knowledge is explicitly prompted. This might imply that 1) commonsense knowledge learned by the COMET-BART during pretraining already benefits dialogue response generation; 2) additionally, prompting commonsense knowledge further improves generation performance, but COMET-BART comprehends the knowledge better than the vanilla BART model since it has already been pre-trained on commonsense reasoning task.

FUDGE strategy control improves the generation. We compare the generation results from baseline models with models controlled with FUDGE. We first provide ground truth oracle strategies to condition the generation for comparison. As shown in the results of **BART/COMET-BART + oracle + FUDGE** in Table. 4, using FUDGE as a control for dialogue generation is able to improve the performance compared to simply using the language models. This implies that FUDGE enforces stronger control over generation given the correct strategy, which can also be observed from the examples shown in Table. 3. However, this means generation would largely rely on the strategy predictor. This

corresponds to the results listed in Table. 2. Using a separate strategy classifier does not fulfill the goal of correctly predicting the strategy, thus leading to worse generation results. Future work needs to be done to improve the performance of dialogue strategy prediction.

Jointly training with generation and strategy prediction loss improves generation performance on vanilla BART. Another interesting observation is that when finetuning the language model with a joint loss of generation and strategy classification, the performance of vanilla BART improved when not using FUDGE control (as shown in the **BART with strategy classifier** results in Table. 2). The strategy prediction accuracy dropped while generation performance improved. This might be because separating the training objective allows models to have stronger adaptability to generate strategy-specific responses. However, this improvement was not observed in the COMET-BART model.

Conclusion

In this work, we introduced a new language model-based framework for controllable empathetic dialogue response generation. We use experiment results on a public dataset to show that: 1) Our method of incorporating commonsense knowledge through verbalized prompts can improve the dialogue generation quality, 2) By separating the dialogue strategy prediction with the generation, we can use future discriminator to enforce stronger control over generated dialogue responses. Our method can be used on off-the-shelf generative language models with or without fine-tuning. Future work can further explore how dialogue strategy prediction can be improved to benefit dialogue generation, and introduce better prompting methods in addition to using templates to verbalize commonsense knowledge.

References

- Athota, L.; Shukla, V. K.; Pandey, N.; and Rana, A. 2020. Chatbot for Healthcare System Using Artificial Intelligence. In *2020 8th International Conference on Reliability, Information Technologies and Optimization (Trends and Future Directions) (ICRITO)*, 619–622.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, 1877–1901. Curran Associates, Inc.
- Dathathri, S.; Madotto, A.; Lan, J.; Hung, J.; Frank, E.; Molino, P.; Yosinski, J.; and Liu, R. 2020. Plug and play language models: A simple approach to controlled text generation. In *International conference on learning representations*.
- Davis, M. H. 1983. Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology*, 44(1): 113–126. Place: US Publisher: American Psychological Association.
- Hill, C. E. 2009. *Helping skills: Facilitating, exploration, insight, and action, 3rd ed.* Helping skills: Facilitating, exploration, insight, and action, 3rd ed. Washington, DC, US: American Psychological Association. ISBN 978-1-4338-0451-9. Pages: xix, 430.
- Hosseini, P.; Broniatowski, D. A.; and Diab, M. 2022. Knowledge-Augmented Language Models for Cause-Effect Relation Classification. In *Proceedings of the First Workshop on Commonsense Representation and Reasoning (CSRR 2022)*, 43–48. Dublin, Ireland: Association for Computational Linguistics.
- Hwang, J. D.; Bhagavatula, C.; Bras, R. L.; Da, J.; Sakaguchi, K.; Bosselut, A.; and Choi, Y. 2021. (Comet-) Atomic 2020: On Symbolic and Neural Commonsense Knowledge Graphs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7): 6384–6392. Text copyright: Copyright (c) 2021 Association for the Advancement of Artificial Intelligence.
- Kim, S. S.; Kaplowitz, S.; and Johnston, M. V. 2004. The Effects of Physician Empathy on Patient Satisfaction and Compliance. *Evaluation & the Health Professions*, 27(3): 237–251. Publisher: SAGE Publications Inc.
- Lee, Y.-C.; Yamashita, N.; Huang, Y.; and Fu, W. 2020. "I Hear You, I Feel You": Encouraging Deep Self-disclosure through a Chatbot. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, 1–12. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-6708-0.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2019. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. ArXiv:1910.13461 [cs, stat].
- Li, J.; Monroe, W.; Shi, T.; Jean, S.; Ritter, A.; and Jurafsky, D. 2017. Adversarial Learning for Neural Dialogue Generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2157–2169. Copenhagen, Denmark: Association for Computational Linguistics.
- Li, Q.; Li, P.; Ren, Z.; Ren, P.; and Chen, Z. 2022. Knowledge bridging for empathetic dialogue generation. Publication Title: AAAI.
- Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, 74–81. Barcelona, Spain: Association for Computational Linguistics.
- Liu, S.; Zheng, C.; Demasi, O.; Sabour, S.; Li, Y.; Yu, Z.; Jiang, Y.; and Huang, M. 2021. Towards emotional support dialog systems. In *ACL/IJCNLP (1)*, 3469–3483. Text date: 1609459200000 text.crossref: conf/acl/2021-1.
- Nadarzynski, T.; Miles, O.; Cowie, A.; and Ridge, D. 2019. Acceptability of artificial intelligence (AI)-led chatbot services in healthcare: A mixed-methods study. *DIGITAL HEALTH*, 5: 2055207619871808. Publisher: SAGE Publications Ltd.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics.
- Sabour, S.; Zheng, C.; and Huang, M. 2021. CEM: Commonsense-aware Empathetic Response Generation. ArXiv: 2109.05739 [cs] Number: arXiv:2109.05739.
- Schuermans, J.; and Frasincar, F. 2020. Intent Classification for Dialogue Utterances. *IEEE Intelligent Systems*, 35(1): 82–88. Conference Name: IEEE Intelligent Systems.
- Wang, Y.; Si, P.; Lei, Z.; and Yang, Y. 2020. Topic Enhanced Controllable CVAE for Dialogue Generation (Student Abstract). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(10): 13955–13956. Number: 10.
- Wu, Z.; Galley, M.; Brockett, C.; Zhang, Y.; Gao, X.; Quirk, C.; Koncel-Kedziorski, R.; Gao, J.; Hajishirzi, H.; Ostendorf, M.; and Dolan, B. 2021. A Controllable Model of Grounded Response Generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16): 14085–14093. Number: 16.
- Xu, C.; Wu, W.; and Wu, Y. 2018. Towards Explainable and Controllable Open Domain Dialogue Generation with Dialogue Acts. ArXiv:1807.07255 [cs].
- Yang, K.; and Klein, D. 2021. FUDGE: Controlled Text Generation With Future Discriminators. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3511–3535. Online: Association for Computational Linguistics.

Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2020a. BERTScore: Evaluating Text Generation with BERT. ArXiv:1904.09675 [cs].

Zhang, Y.; Sun, S.; Galley, M.; Chen, Y.-C.; Brockett, C.; Gao, X.; Gao, J.; Liu, J.; and Dolan, B. 2020b. DIALOGPT: Large-Scale Generative Pre-training for Conversational Response Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 270–278. Online: Association for Computational Linguistics.

Acknowledgments

This research used the Delta advanced computing and data resource which is supported by the National Science Foundation (award OAC 2005572) and the State of Illinois. Delta is a joint effort of the University of Illinois Urbana-Champaign and its National Center for Supercomputing Applications.

Context	Here! Last Christmas I bought the condo of my dreams. This April 2020 I lost my job due Covid 19. My unemployment is ending and my savings are depleted. I've recived an eviction notice and don't know what to do. Help!
Seeker	Have you ever been evicted?
Helper	[Others] Where are you that you are able to be evicted?
Helper	[Providing Suggestions] You should not be able to be evicted from your home during this time in pretty much every state. Have you tried contacting anyone for help other than your landlord?
Seeker	I'm still in my condo but I no longer have a job, umemployment has run out and my savings are all gone.
Helper	[Providing Suggestions] Well the plus side is they cannot kick you out right now. Have you talk to your landlord about working out some payment arrangement. Have you tried applying for any job? Now is not the time for pride to get in the way, i would accept anything
Seeker	When you say work out a plan what do you mean? I have no job. I have been applying for jobs but I haven't received a callback or an offer of employment. Have you ever been in this situation?
Seeker	Also, I'm calling a helpline so Pride is not my issue.
Helper	[Others] I have but the referendum on evictions with covid makes this situation different. I am sorry if that pride comment came off as crass. I was just trying to say desperate times come with desperate measures
Ground truth	
Helper	[Providing Suggestions] They could work out a way you can pay extra monthly when you get back on your feet.
Baseline models	
BART	[Providing Suggestions] This is not your issue, i would say that you should be proud of yourself for being able to get out of your home. I am sorry if that came off as crass.
COMET-BART	[Affirmation and Reassurance] I have been in this situation. I was not able to get out of it due to covid. I was able to get out of it because I had a plan.
With COMET output (verbalized with templates)	
BART	[Providing Suggestions] i would say you should try to find a job and see if you can get a raise. I would also suggest talking to your landlord about working out some payment arrangement.
COMET-BART	[Self-disclosure] I have been in your shoes.
BART (+ FUDGE)	[Providing Suggestions] Have you tried contacting a local government to see if they can help you out?
COMET-BART (+ FUDGE)	[Self-disclosure] my I have been in this situation too. I was in a similar situation and i was able to get out of my home and move to a new one.
With strategy classifier (Linear CLF with classification + LM loss)	
BART	[Providing Suggestions] I am sorry to hear that. I have been in this situation before and it is not my issue. I would not be able to help you. I am sorry if that pride comment came off as crass.
COMET-BART	[Providing Suggestions] I am sorry. I was just trying to say desperate times come with desperate measures
BART (+ FUDGE)	[Providing Suggestions] I could tell you are feeling a bit down about the situation. I would suggest you try to find a job that pays more than you can afford.
COMET-BART (+ FUDGE)	[Providing Suggestions] Have you tried to talk to your landlord about a payment plan?
With strategy classifier (BERT classifier, trained separately from the LM)	
BART (+ FUDGE)	[Question] I do not think you should feel ashamed of your pride. What you are doing is the right thing.
COMET-BART (+ FUDGE)	[Question] Do you have any friends or family that you can talk to about this?
strategy oracle	
BART (+ FUDGE)	[Providing Suggestions] This is not your issue, i would just say that you should be prepared for the worst. I have been in this situation before and it was a lot worse than i thought.
COMET-BART (+ FUDGE)	[Providing Suggestions] Have you tried to talk to your landlord about a payment plan?

Table 5: An Example for comparison of different generation methods.