# Supplementary Material: <br> Plug-and-Play Conversational Models 

Andrea Madotto ${ }^{1 *}$, Etsuko Ishii ${ }^{* 1}$, Zhaojiang Lin* ${ }^{* 1}$, Sumanth Dathathri ${ }^{*} \dagger$ Pascale Fung ${ }^{1}$<br>${ }^{1}$ The Hong Kong University of Science and Technology<br>\{amadotto, eishii,zlinao, pascale\}@ust.hk<br>sdathath@gmail.com

## A Hyperparamters

In Table 1, we report the full set of hyperparameters used in the experiments section. DialoGPT [5] medium has 345 M parameters, 24 layers and $d_{\text {model }}=1024$. For adapter we use bottleneck size $m=100$, resulting in additional 5.175 M parameters (1.5\%).

| Model | Attributes | Hyperparameters |
| :--- | :--- | :--- |
| $P P L M$ | negative, question, Business, <br> Sports, Sci/Tech | $\alpha=0.02, p=75, \gamma=1.0, \lambda_{K L}=0.01$ |
| PPLM | positive | $\alpha=0.02, p=25, \gamma=1.0, \lambda_{K L}=0.01$ |
| ADAPTER | negative, positive, question, <br> Business, Sports, Sci/Tech | lr $=6.25 \mathrm{e}-4$, batch_size $=32$, epoch $=5, \lambda_{K L}=0.5$ |

Table 1: The full set of hyperparameters used in the experiments. Here, $\lambda_{K L}$ denotes the weight of Kull-back-Leibler loss constraint for language model. All the experiments have been run on several NVIDIA 1080Ti.

## B Automatic Evaluation Datasets Statistics and Performance

In Table 2, we summarize the dataset statistics and performance of the trained scorer.

| Task | Style | \#C | Samples |  | F1-Score |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | Train | Test | Train | Test |
| $A M A Z O N 5[3]$ | Sentiment | 5 | 3 M | 650 K | 59.13 | 59.11 |
| $A G N E W S(R)[4]$ | Topic | 4 | 7600 | 120000 | 92.47 | 90.21 |

Table 2: Automatic evaluation datasets statistics and performance.

## C Additional Details of Human Evaluation

We collect human annotations for both Humanness and Attribute Consistency via crowd-sourcing platform provided by Appen Limited ${ }^{1}$. The template for human evaluation is shown in Figure 1. To get consistent observation, we use the same 30 randomly selected prefix of the dialogues across the comparisons. Each annotator is asked to judge either humanness or styleness of 5 different dialogues. In this way, we collect in total 4200 human annotation.

[^0]
Read the conversation below
Read the conversation below
UserA. The Disney lands are all different! There's also Disney Sea, which is completely unique!
UserA. The Disney lands are all different! There's also Disney Sea, which is completely unique!
User B: oh neat. I hven't heard about that robot figting show. where is that?
User B: oh neat. I hven't heard about that robot figting show. where is that?
Option 1:I don't really remember what part of town it was in. It was pretty cool though -l"m sure you can find it if you google "glant robot fighting show tokyo" haha
Option 1:I don't really remember what part of town it was in. It was pretty cool though -l"m sure you can find it if you google "glant robot fighting show tokyo" haha
Option 2: It's on Netflix. You should checkito out!
Option 2: It's on Netflix. You should checkito out!
Click here to open the original post for adititional information.
Click here to open the original post for adititional information.
Which response sounds more human? (required)
Which response sounds more human? (required)
O.Optio 1
O.Optio 1
O Both
O Both
(a) Template for judge humanness. Across all the style, the instructions kept to be same.

(b) Template for judge style. Depending on the target style, the instructions and the questions asked are modified.

Figure 1: Human evaluation template for judge humanness and style respectively.

## D Additional Details on Results

In this section, we present additional details on the experimental results, such as the attribute-specific vocabulary, and breakdown of human evaluation results per model comparison. Moreover, from Section D. 3 to Section D.8, we report breakdown results of both automatic and human evaluation per style as well as additional examples of generated responses.

## D. 1 Attribute-Specific Vocabulary

In Table 3 and Table 4, top 10 frequent attribute-specific words of adapters trained with PPLM are listed. We extract attribute-specific words from 200 dialogues per attribute by taking words that appear more than 5 times in some attribute yet never appear in the other attributes. As can be seen clearly in Table 3 and Table 4, adapters trained with PPLM are optimized to restrict the vocabulary for each style and topic. Note that the words list of the style question is not clear since it tends to ask a general question such as "What do you mean?", "How about you?", or "How much does it cost?"

| Topic | Top 10 frequent style unique words |
| :--- | :--- |
| Negative | horrible, terrible, garbage, bored, waste, lazy, loss, worst, anymore, toilet |
| Positive | amazing, excited, beautiful, awesome, happy, nice, glad, wonderful, story, fantastic |
| Question | cost, yours, u, ago, charge, hobbies, lived, ocean, N/A, N/A |

Table 3: Top 10 frequent style unique words appeared responses of $A D$ in 200 dialogues. In style question, only 8 style unique words are found.

| Topic | Top 10 frequent topic unique words |
| :--- | :--- |
| Business | oil, bank, money, gas, store, investment, insurance, grocery, station, car |
| Sports | football, hockey, soccer, basketball, baseball, fan, player, league, rugby, sport |
| Sci/Tech | computer, internet, web, software, science, android, space, programming, studying, moon |

Table 4: Top 10 frequent topic unique words appeared responses of $A D$ in 200 dialogues.

## D. 2 Human Evaluation Breakdown Per Model Comparison

In Table 5, we summarize win-tie-loss rates per comparisons on human evaluation. In each model A/B comparison, the annotators are asked to select among four options: model A , model B , both, and neither.

|  | Humanness |  |  |  | Attribute Consistensy |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | win | tie | loss | none | win | tie | loss | none |
| $D G$ vs. $H M$ | 14.2 | 64.0 | 16.1 | 5.69 | 23.8 | 8.10 | 9.05 | 59.0 |
| $W D$ vs. $H M$ | 15.2 | 62.9 | 17.6 | 4.29 | 29.0 | 4.76 | 6.19 | 60.0 |
| $P P$ vs. $H M$ | 15.2 | 61.9 | 17.6 | 5.24 | 43.3 | 9.05 | 7.14 | 40.5 |
| $A D$ vs. $H M$ | 12.4 | 70.5 | 14.8 | 2.38 | 68.1 | 9.52 | 2.38 | 20.0 |
| $W D$ vs. $D G$ | 13.7 | 66.4 | 12.3 | 7.58 | 18.1 | 16.2 | 11.9 | 53.8 |
| $P P$ vs. $D G$ | 11.4 | 63.3 | 14.3 | 11.0 | 37.1 | 16.7 | 7.14 | 39.0 |
| $A D$ vs. $D G$ | 7.14 | 75.7 | 14.3 | 2.86 | 60.0 | 16.7 | 4.29 | 19.0 |
| $P P$ vs. $W D$ | 16.7 | 55.7 | 17.1 | 10.5 | 31.4 | 18.6 | 11.9 | 38.1 |
| $A D$ vs. $W D$ | 12.4 | 78.1 | 8.10 | 1.43 | 53.8 | 21.0 | 3.81 | 21.4 |
| $A D$ vs. $P P$ | 9.52 | 77.1 | 9.52 | 3.81 | 38.6 | 40.0 | 5.71 | 15.7 |

Table 5: Win-tie-loss rates (\%) per comparison. For example, in the Attribute Consistency table, $D G$ wins $23.8 \%$, tie $8.10 \%$, loses $9.05 \%$ of the time respectively versus $H M$, and $59.0 \%$ of the time neither of them is chosen. Note that total may not become $100 \%$ due to rounding off.

## D. 3 Negative

| Model | Discr. | Ppl. | Dist. | Score | VADER | Emoji |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| HUMAN | - | 49.29 | 0.32/0.75/0.83 | 10.95 | 12.44 | $(27){ }^{(25)}(2 \cdot(21) \because(19) \ominus$ |
| $D G P T$ | 8.96 | 36.63 | 0.21/0.63/0.79 | 19.40 | 29.85 | $(33) \because(27) \odot(27) \because(19) \because$ |
| $D G P T+W D$ | 9.95 | 44.75 | 0.24/0.71/0.83 | 28.86 | 34.33 | $(32) \odot(24) \because(23) \ominus(21) \odot \cdot$ |
| PPLM | 43.78 | 45.95 | 0.26/0.69/0.83 | 39.80 | 51.74 | $(30) \ldots(27) \because(22) \because(21)=$ |
| ADAPTER | 90.55 | 39.82 | 0.18/0.60/0.80 | 73.13 | 80.10 |  |

Table 6: Automatic evaluation results for the style negative. VADER [2] is the accuracy of the external rulebased sentiment analysis tool ${ }^{2}$. Emoji predicted by pretrained DeepMoji ${ }^{3}$ [1] are listed in the rightest column with frequency. Emoji predictions indicate that the original outputs of DialoGPT [5] tend to be positive, but $P P L M$ and $A D A P T E R$ succeed in generating negative responses.


Figure 2: Human evaluation results in terms of winning rate for humanness and negativeness.

| 1. | HUMAN 1 <br> HUMAN 2$\quad$How are you? How is your day? <br> Good. Don't have much to do today, feels good. How are you? |
| :--- | :--- |
| $D G$ | I had a pretty solid day. Just watched the movie, ate pizza. |
| $W D$ | I'm good too. My job doesn t allow me to be at home today. |
| $P P$ | I'm pretty bad : I haven't had any work done yet so I guess that's why I don't feel like I've accomplished anything. |
| $A D$ | I had a really awful horrible day yesterday. It's been awful all week. So, so, so, so horrible. |
| 2. | HUMAN 1 Any plans for the weekend? |
| $D G$ | HUMAN 2 my friends are gonna visit me this weekend. we might go hiking! |
| $W D$ | That sounds awesome. Hope they get the chance to go hiking! |
| $P P$ | Sounds fun, I may have to do that. |
| $A D$ | That sounds like a really good idea, I am so sorry that you are stuck in the slum of trash and chump driven plebs. |

Table 7: Examples of generated responses for the style negative. Note that $D G$ tends to generate a positive response. In the second example, we can see that $A D$ tries to change the sentiment by adding negative words (e.g., slum, trash) in the end.

[^1]
## D. 4 Positive

| Model | Discr. | Ppl. | Dist. | Score | VADER | Emoji |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| HUMAN | - | 49.29 | 0.32/0.75/0.83 | 45.27 | 61.19 | (27) $(25) \because(21) \because(19) \Theta$ |
| $D G P T$ | 24.38 | 39.38 | 0.23/0.63/0.77 | 65.67 | 84.08 | (54) $(37) \odot(28) \ominus(26) \ominus$ |
| $D G P T+W D$ | 24.38 | 48.37 | 0.24/0.70/0.82 | 58.21 | 84.08 | (60) |
| PPLM | 48.26 | 39.68 | 0.21/0.62/0.78 | 75.12 | 86.07 | (49) |
| ADAPTER | 99.50 | 35.47 | 0.16/0.52/0.72 | 93.03 | 96.52 | $(68) \ominus(62)$ |

Table 8: Automatic evaluation results for the style positive. VADER [2] is the accuracy of the external rulebased sentiment analysis tool. Emoji predicted by DeepMoji are listed in the rightest column with frequency. Although emoji predictions for that the original outputs of DialoGPT [5] are already positive, emojis for PPLM and $A D A P T E R$ are even more positive.


Figure 3: Human evaluation results in terms of winning rate for humanness and positiveness.

| 1. | HUMAN 1 <br> HUMAN 2 | What kinds of adventures?? Any ones that I would be able to join you on? <br> Hmmm I really want to try bull riding. Do you have any interest in that? |
| :--- | :--- | :--- |
|  | No but my friend does. We could go for a ride! Hah! |  |
| $W D$ | No, bull riding just isn't what I'm looking for in the right way... |  |
| $P P$ | I do, I do. I've been wanting to learn it for awhile now, I've seen some awesome things |  |
| $A D$ | Yesss! I am excited for summer to get out there. I'm thinking of going on a few of them. |  |
| 2. | HUMAN 1 <br> HUMAN 2 | So what's your outlook? Something blue? |
| $D G$ | That was beautiful is a tranquil colour. It's a good metaphor. Do you have good advice for positivity? |  |
| $W D$ | You seem positive to us? What's your favorite thing? |  |
| $P P$ | What is the best color? And how does it affect you? You can make it a better place if you don't mind <br> sharing your experiences! |  |
| $A D$ | I think it's an excellent quote, and a great motivation to live a fulfilling life! It has been my mantra <br> since I started my job. |  |

Table 9: Examples of generated responses for the style positive. As shown in the first example, $P P$ and $A D$ can generate positive responses even if the default output of $D G$ is not that positive.

## D. 5 Question

| Model | Disc. | Ppl. | Dist. |
| ---: | :---: | :---: | :---: |
| HUMAN | - | 49.29 | $0.32 / 0.75 / 0.83$ |
| $D G P T$ | 49.75 | 41.32 | $0.23 / 0.66 / 0.77$ |
| $D G P T+W D$ | 70.15 | 60.01 | $0.25 / 0.76 / 0.85$ |
| PPLM | 84.08 | 45.53 | $0.26 / 0.66 / 0.73$ |
| ADAPTER | 99.00 | 38.40 | $0.17 / 0.48 / 0.63$ |

Table 10: Automatic evaluation results on the style question.


Figure 4: Human evaluation results in terms of winning rate for humanness and question.
\(\left.$$
\begin{array}{ll}\hline 1 . & \begin{array}{l}\text { HUMAN 1 } \\
\text { HUMAN 2 }\end{array}
$$ <br>

\hline D G \& I heard there are rainforests in southeast Asia where you can zipline from tree to tree.\end{array}\right] .\)| $W D$ | You mean zilean forest? Or am I wrong? <br> fello from jungle |
| :--- | :--- |
| $P P$ | But is that rain? Is it rain? |
| $A D$ | But is it legal? Is there any trouble? |
| 2. | HUMAN 1 Hard to keep a count. Maybe 25. |
| $D G$ | HUMAN 2 $\quad$ I don't remember, but I think it was Halloween. |
| $W D$ | I dunno.. I'm pretty sure I got it in February? Maybe March? |
| $P P$ | You know. The one where the candy corn melted? |
| $A D$ | You know. The one you're talking about. The one where they have candy and stuff... right? |

Table 11: Examples of generated responses for the style question. $A D$ can pose a question that is on topic like in the first example, as well as add a general and simple question such as "right?", "you?" after answering questions that appear in prefix.

## D. 6 Business

| Model | Disc. | Ppl. | Dist. | Score |
| ---: | :---: | :---: | :---: | :---: |
| $H U M A N$ | - | 49.29 | $0.32 / 0.75 / 0.83$ | 2.99 |
| $D G P T$ | 43.78 | 41.42 | $0.23 / 0.67 / 0.79$ | 17.41 |
| $D G P T+W D$ | 46.77 | 55.36 | $0.27 / 0.77 / 0.86$ | 19.40 |
| $P P L M$ | 75.12 | 46.41 | $0.25 / 0.68 / 0.82$ | 47.26 |
| $A D A P T E R$ | 91.54 | 41.40 | $0.18 / 0.65 / 0.85$ | 68.66 |

Table 12: Automatic evaluation results on the topic Business.


Figure 5: Human evaluation results in terms of winning rate for humanness and business.

| 1. HUMAN 1 | I don't know if there is any medication for acrophobia. I want to take plenty of it <br> if I really have to do it. |
| :--- | :--- |
| HUMAN 2 | If there isn't one, you should invent it, and then make millions |

HUMAN 2 If there isn't one, you should invent it, and then make millions

| $D G$ | I can't afford that. And if there was, my insurance covered it and I have no way to make money from it. |
| :--- | :--- |
| $W D$ | I will probably just give a million upvotes and call them a successful business. I would be a lot more <br> effective if my money was invested instead |
| $P P$ | And if you are unable to make a million dollars on the stock market, that is a good way. <br> I am making millions. |
| $A D$ | That's not how capitalism works. That's a recipe for disaster. I'd rather be rich than poor. |
| 2. | HUMAN 1 Hey, what's up? <br> HUMAN 2 Just chillin'. how are you? |
| $D G$ | I'm doing well, how do you like your coffee? |
| $W D$ | I'm good, what about you? : O |
| $P P$ | I am good! What about you? |
| $A D$ | I'm good! What about you? What is your favorite brand of hair conditioner? |

Table 13: Examples of generated responses for the Business topic. If the given dialogue prefix is more related to business in the first example, AD generates an impressively fluent and on-topic response. Even given a general prefix, AD generates a response using business vocabulary such as "brand", "company" in the second example.

## D. 7 Sports

| Model | Disc. | Ppl. | Dist. | Score |
| ---: | :---: | :---: | :---: | :---: |
| $H U M A N$ | - | 49.29 | $0.32 / 0.75 / 0.83$ | 5.97 |
| $D G P T$ | 82.09 | 38.00 | $0.22 / 0.64 / 0.78$ | 27.86 |
| $D G P T+W D$ | 86.57 | 54.98 | $0.23 / 0.74 / 0.85$ | 36.82 |
| $P P L M$ | 95.52 | 42.81 | $0.23 / 0.66 / 0.80$ | 59.20 |
| ADAPTER | 98.51 | 40.36 | $0.16 / 0.61 / 0.82$ | 83.08 |

Table 14: Automatic evaluation results on topic Sports.


Figure 6: Human evaluation results in terms of winning rate for humanness and sports.

| 1. | HUMAN 1 <br> HUMAN 2 |
| :--- | :--- |
|  | Hey there! What's up??? |
| $W D$ | I'm been pretty alright, much, how you doin? |

Table 15: Generated responses for Sports topic. In both examples, Sports vocabulary such as "soccer" or "racing" can be observed in $A D$.

## D. 8 Science \& Technology

| Model | Disc. | Ppl. | Dist. | Score |
| :---: | :---: | :---: | :---: | :---: |
| $H U M A N$ | - | 49.29 | $0.32 / 0.75 / 0.83$ | 91.04 |
| $D G P T$ | 99.00 | 39.25 | $0.23 / 0.65 / 0.78$ | 91.04 |
| $D G P T+W D$ | 99.50 | 58.60 | $0.25 / 0.77 / 0.87$ | 91.04 |
| PPLM | 100.00 | 44.43 | $0.26 / 0.72 / 0.82$ | 93.03 |
| ADAPTER | 100.00 | 41.11 | $0.19 / 0.66 / 0.84$ | 99.00 |

Table 16: Automatic evaluation results on the topic Science \& Technology.


Figure 7: Human evaluation results in terms of winning rate for humanness and Sci/Tech.

| 1. | HUMAN 1 <br> HUMAN 2$\quad$ Oh hello. Long time no talk. How's the day going for yuo? |
| :--- | :--- |

Table 17: Generated responses for $\mathrm{Sci} /$ Tech topic. Given a general starter, $D G$ generates a response containing techy words by default as in the first example. Since the provided context is less natural for the $\mathrm{Sci} / \mathrm{Tech}$ attribute in the second example, $A D$ generates a response that gives off an erratic impression.

## References

[1] Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In Conference on Empirical Methods in Natural Language Processing (EMNLP), 2017.
[2] C.J. Hutto and Eric Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. 012015.
[3] Julian McAuley and Jure Leskovec. Hidden factors and hidden topics: understanding rating dimensions with review text. In Proceedings of the 7th ACM conference on Recommender systems, pages 165-172, 2013.
[4] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In Advances in neural information processing systems, pages 649-657, 2015.
[5] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation. arXiv preprint arXiv:1911.00536, 2019.


[^0]:    * Equal Contribution
    $\dagger$ Work done primarily at the Caltech.
    ${ }^{1}$ https://client.appen.com/

[^1]:    ${ }^{2}$ https://github.com/cjhutto/vaderSentiment
    ${ }^{3}$ https://github.com/huggingface/torchMoji

