PAR: Persona Aware Response in Conversational Systems

Abhijit A Nargund
Samsung R & D Institute
Bangalore, India
abhijit.an@samsung.com

Sandeep Kumar Pandey
Samsung R & D Institute
Bangalore, India
sandeep.p@samsung.com

Jina Ham
Samsung Research
Korea
j.ham@samsung.com

Abstract

To make the Human Computer Interaction more user friendly and persona aligned, detection of user persona is of utmost significance. Towards achieving this objective, we describe a novel approach to select the persona of a user from pre-determine list of personas and utilize it to generate personalized responses. This is achieved in two steps. Firstly, closest matching persona is detected from a set of pre-determined persona for the user. The second step involves the use of a fine-tuned natural language generation (NLG) model to generate persona compliant responses. Through experiments, we demonstrate that the proposed architecture generates better responses than current approaches by using a detected persona. Experimental evaluation on the PersonaChat dataset has demonstrated notable performance in terms of perplexity and F1-score.

1 Introduction

Since the dawn of deep learning era and advancements in Natural Language Processing (NLP), researchers have tried to make the Human-Computer Interaction (HCI) more natural and smooth (Deng and Ren, 2021). Several features are incorporated in a conversation based system to induce human-like response and give users a pleasant experience such as - sentiment analysis of user replies (Lee et al., 2018), emotion recognition (Lee et al., 2020) etc. However, there is still a gap left between the desired output and the approaches tried. A user's conversation with a bot highly depends on the general nature and likings of the user, more commonly referred to as the persona.

Many recent studies have explored persona identification from user utterances (Zhang et al., 2018; Bahdanau et al., 2014; Song et al., 2019; Natarajan and Nargund, 2021). A dataset of Persona enhanced chat between two users based on a given persona profile, known as PersonaChat was introduced in (Zhang et al., 2018). Experiments performed using several generative and ranking baseline models established the need of including persona information to enhance conversations. The study in (Gu et al., 2021) have proposed a modified dataset PMPC on PersonaChat Dataset and explored Utterance-to-Profile Matching Network for the task of many-to-many matching. Moreover, in (Liu et al., 2020), a trasmitter-reciever architecture is proposed, referred to as $\mathcal{P}^2$ bot which models the mutual persona perception using a transformer with reinforcement learning to fine tune the conversations during training. Also, to alleviate the challenges associated with data in persona based dialog generation, the study in (Cao et al., 2022) introduced data manipulation techniques such as distillation, diversification and curriculum. The architecture considered were baseline GPT2 and Transormer encoder-decoder stack.

Despite several attempts by researchers, the generation of persona complaint dialogue is very challenging. Firstly, automated metrics such as BLEU, ROGUE etc does not consider semantic similarity and will present poor scores if there is no overlap between the generated response and the gold response. Also, the response generated can be persona complaint but may not answer the query of the user. As such, we propose a novel approach to tackle the persona based dialog generation problem. Firstly, we select the matched persona corresponding to the user input from the set of pre-determined personas of the user, by using similarity matching of the embeddings. Then, we generate per-

Proceedings of the 19th International Conference on Natural Language Processing (ICON), pages 50 - 54

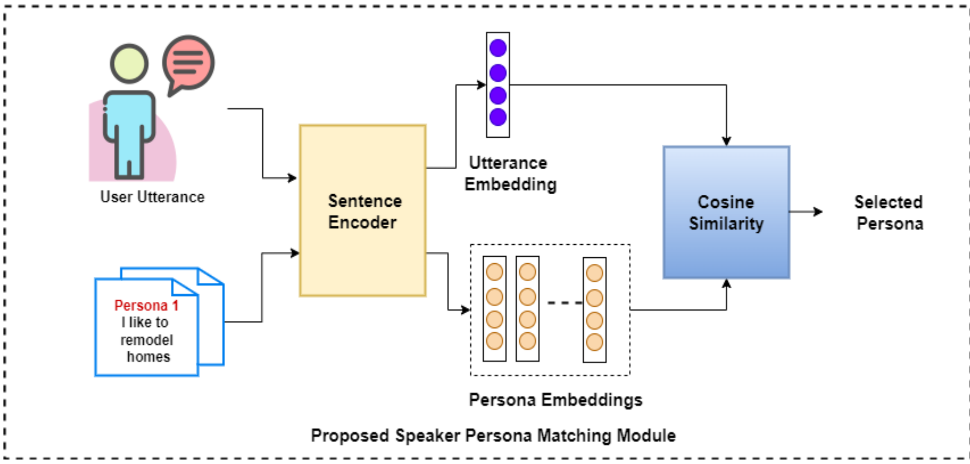

Figure 1: Proposed methodology to extract user utterance and persona embeddings and compare using cosine similarity.

sona complaint response by incorporating this additional persona information with state-of-the-art Semantically Controlled Generative Pre Training (SC-GPT)(Peng et al., 2020) model. Performance on the validation set of the PersonaChat dataset yeilds results comparable to other standard baselines.

The rest of the paper is organized as follows. Section 2 describes the methodology for persona classification and response generation. Section 3 details the experiments performed with dataset information and hypermaters used. Section 4 presents the results along with sample generated responses. Section 5 concludes the paper.

2 Methodology

This section describes the methodology used for persona selection, given a user utterance and predetermined list of persona and subsequent Natural Language Generation (NLG) module which is responsible for the generation of persona complaint bot response.

2.1 Persona Identification

For the task of choosing the most relevant persona characteristics out of the n persona sentences for the user, we employ similarity matching of sentence embeddings as shown in Figure 1. Given a user utterance X and a set of persona sentences characterizing the user choices $Z = \{z_i\}$ where $i = \{1, 2, \cdots, n\}$, the first step involves encoding the dialogue turn or the sentence X into sentence embedding X_{embed}. For this, we utilized techniques such as sentence Bert (Reimers and Gurevych, 2019), RoBERTa (Liu et al., 2019) and Universal Sentence Encoder(USE) (Yang et al., 2020), out of which USE was chosen based on better performance. In a similar fashion, embeddings from the persona sentences are calculated. Mathematically,

$$\mathcal{X}_{embed} = USE(\mathcal{X}) \tag{1}$$

$$\mathcal{Z}_{embed} = USE(z_i) \tag{2}$$

where USE denotes the pre-trained Universal Sentence Encoder, utilized to generate the latent space embeddings.

The similarity score between $\mathcal{X}$ and the set of persona embeddings $\mathcal{Z}$ is calculated using *cosine* similarity as -

$$s_i = \frac{\mathcal{X}\dot{z}_i}{\|X\|\,\|z_i\|} \tag{3}$$

where $z_i \in \mathcal{Z}$ and $\|.\|$ represents the length of the vector. The matched persona is selected by using $argmax$ over the set of scores, represented by -

$$S = \{s_i\} \tag{4}$$

$$P_{index} = argmax(S) \tag{5}$$

This matched persona is provided as a semantic input to the Natural Language Generation (NLG) module along with the gold response to find out the correlation with the candidate persona and producing persona complaint dialog response.

2.2 Response Generation

For Natural Language Generation, we have experimented with baseline GPT (Radford et al., 2018)

User utterance	Matched Persona	Generated Response
hello what are doing today	i'm a stunt double as my second job	i'm taking a break from shooting to relax
i like that ! i go to preschool	my favorite season is winter	winter is my favorite season. i like to go to the stores during the day.
hello , how are you tonight ? do you have pink and blue hair ?	i volunteer at a soup kitchen .	i'm doing great. just got back from soup kitchen helping people.
oh wow . i hope i live as long as you !	i am very strong for my age .	i hope so. i'm just over 50 now.
i just got done watching a horror movie	i read twenty books a year	scary movies are good. i read twenty books a year

Table 1: Sample responses generated on the validation set along with gold response and matched persona from user utterance.

Partition	# of Dialog	# of Persona
Train	8939	955
Validation	1000	100

Table 2: Details of the dataset- partionwise mentioning the number of complete dialog acts and possible personas.

as well as SC-GPT. To generate responses which are persona complaint, we have utilized Semantically Controlled - GPT (SC-GPT) (Peng et al., 2020) based on better perfromance compared to plain GPT. The choice of SC-GPT is motivated by its ability to produce more naturalisitc and Semantically controllable dialogue based on the input semantic form. It is pre-trained on a huge collection of plain text, dialogue-act labelled utterances as well as domain specific fine-tuning on limited data.

During training, given a user utterance, for e.g. *hi , how are you doing ? i'm getting ready to do some cheetah chasing to stay in shape .* and its corresponding classified persona - *i like to go hunting* , the input to the SC-GPT is of the form -

request(question=hi , how are you doing ? i'm getting ready to do some cheetah chasing to stay in shape.) @infrom(preference=i like to go hunting) & you must be very fast , hunting is one of my favorite hobbies

In this example, the task of SC-GPT will be to utilize the selected persona and the user utterance to generate a response which is persona complaint as well as contextually similar to the user input utterance. This is achieved by using the gold response- *you must be very fast , hunting is one of my favorite hobbies* as supervision while training.

During the inference, the gold response is omitted from the input form. The SC-GPT generates multiple responses matching with the selected persona, ranked according to the probability, and the one with the highest probability is selected.

3 Experimental Evaluation

In this section, we discuss the dataset details along with hyper-parameter tuning for the SC-GPT model for persona-based dialog generation.

3.1 Dataset Description

For the task of generating persona complaint responses, we utilize PersonaChat dataset (Zhang et al., 2018). It contains crowd-sourced multi-turn dialogues based on pre-defined personas. Table 2 shows the available number of complete dialog acts based on personas for the train and validation partitions. Evaluation for the fine-tuned model is performed using the validation set. Also, the model is fine-tuned using the entire PersonaChat train partition having approximately 17000 question-persona pair. Also, to asses low-resource scenario, we randomly sampled 1000 such pairs from the train partition to fine-tune the base SC-GPT model.

3.2 Hyperparameter Tuning

For the persona classification module, the size of sentence encodings generated is 768. The number of persona sentences considered per question for classification is 4, as in PersonaChat dataset. For the SC-GPT based fine-tuning, we first generated *(question, persona, response)* tuples for the entire training partition of the PersonaChat dataset.Training is performed in two scenarios - utilizing the entire train partition and randomly sampled 1000 tuples. For the random sampling based fine-tuning, we selected 1000 tuples out of the 17,000 generated tuples randomly. This is repeated five times and the result presented is the average over the five random folds. The model is fine-tuned for 10, 20 and 30 epochs out of which the best results were obtained at 10 epochs. The *temperature* parameter, which controls the degree of strictness of the results generated, is set at 0.7, empirically.

Method	F1 score
Seq2Seq + Attention (Bahdanau et al., 2014)	16.18
KV Profile Memory (Zhang et al., 2018)	11.9
TransferTransfo (Wolf et al., 2019)	19.09
P^2 Bot (Liu et al., 2020)	19.77
Proposed	**27.89**

Table 3: Comparison of the proposed persona complaint dialog generattion methodology with the recent state-of-the-art baseline.

Experiment	F1 score	Perplexity
Full dataset + 1 persona + 10 epochs + all past context	23.12	51.2
Full dataset + 1 persona + 20 epochs + all past context	23.03	141.59
Full dataset + 4 persona + 10 epochs + all past context	14.27	58.4
Random 1000 samples + 1 persona + 1 past context	**27.98**	**4.77**

Table 4: Ablation study for SC-GPT based persona-aware response generation based on 1 persona and 4 persona sentences.

3.3 Results

We evaluated the proposed methodology on the validation set of the PersonaChat dataset. The metric used for the evaluation of the generated response are Perplexity and F1-score. Perplexity measures the negative log likelihood of the correct sequence output by the model, lower values indicating better performance. F1 score is the harmonic mean of word-level precision and recall. Table 1 shows some of the responses generated given the persona and the user utterance. As can be observed, the generated responses are highly persona complaint. However, the gold response may vary from the generated response as the conversation is natural. Moreover, our proposed methodology achieves an F1 score of 27.89 and perplexity of 4.77 on the response generation task.

Table 3 shows the comparison of the proposed approach with the state-of-the-art methods on PersonaChat Dataset. Our proposed approach outperforms the baseline $\mathcal{P}^2$ bot by a margin of 8.12, thus highlighting the capability of the proposed approach in generating persona aware responses.

3.4 Ablation Study

We performed ablation study to assess the effectiveness of the persona selection module in improving the persona complaint response generation. To this end, we investigated three aspects - the use of full training partition vs. 1000 randomly selected training tuples, increased training epochs and the use of full persona profile(4 persona sentences) vs. one persona sentence selected as per section-2.1. The results of the ablation study is presented in Table-4. From the table, it can be observed that the persona-complaint response generation shows improvement with the persona identification module selecting the best matching persona from user input sentence rather than presenting the model with all the persona sentences. Moreover, training for 20 epochs over the selected 10 epochs deteriorates the performance in terms of perplexity, indicating an inferior fine-tuned model. Also, as observed from the experiments, the choice of 1000 random training tuples over the usage of entire dataset also impacts performance. Additionaly, the use of all past context while generating the next persona-aware response hampers performance. This can be attributed to the chit-chat nature of the PersonaChat Dataset, due to which a continuous context between past dialogue turns may not be observable.

4 Conclusion

We have proposed a methodology which first identifies the most relevant persona corresponding to a user utterance, from a list of given personas. Then persona complaint response is generated based on the selected persona. We utilized cosine similarity based classification of sentence embeddings to select the relevant persona. SC-GPT is used to fine-tune on subsets of PersonaChat dataset for the task of response generation. Experimental evaluation is performed on the validation partition of the PersonaChat dataset and perplexity score as well as F1 score is reported. We have also reported sample responses generated by the model, which shows that the proposed methodology is efficient in generating persona complaint responses.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Yu Cao, Wei Bi, Meng Fang, Shuming Shi, and Dacheng Tao. 2022. A model-agnostic data manipulation method for persona-based dialogue generation. *arXiv preprint arXiv:2204.09867*.

Jiawen Deng and Fuji Ren. 2021. A survey of textual emotion recognition and its challenges. *IEEE Transactions on Affective Computing*.

Jia-Chen Gu, Zhen-Hua Ling, Yu Wu, Quan Liu, Zhigang Chen, and Xiaodan Zhu. 2021. Detecting speaker personas from conversational texts. *arXiv preprint arXiv:2109.01330*.

Chih-Wei Lee, Yau-Shian Wang, Tsung-Yuan Hsu, Kuan-Yu Chen, Hung-yi Lee, and Lin-shan Lee. 2018. Scalable sentiment for sequence-to-sequence chatbot response with performance analysis. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6164–6168. IEEE.

Ming-Che Lee, Shu-Yin Chiang, Sheng-Cheng Yeh, and Ting-Feng Wen. 2020. Study on emotion recognition and companion chatbot using deep neural network. *Multimedia Tools and Applications*, 79(27):19629–19657.

Qian Liu, Yihong Chen, Bei Chen, Jian-Guang Lou, Zixuan Chen, Bin Zhou, and Dongmei Zhang. 2020. You impress me: Dialogue generation via mutual persona perception. *arXiv preprint arXiv:2004.05388*.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Bharatram Natarajan and Abhijit Nargund. 2021. MRE : Multi relationship extractor for persona based empathetic conversational model. In *Proceedings of the 18th International Conference on Natural Language Processing (ICON)*, pages 173–179, National Institute of Technology Silchar, Silchar, India. NLP Association of India (NLPAI).

Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. Few-shot natural language generation for task-oriented dialog. *arXiv preprint arXiv:2002.12328*.

Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.

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.

Haoyu Song, Wei-Nan Zhang, Yiming Cui, Dong Wang, and Ting Liu. 2019. Exploiting persona information for diverse generation of conversational responses. *arXiv preprint arXiv:1905.12188*.

Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.

Ziyi Yang, Yinfei Yang, Daniel Cer, Jax Law, and Eric Darve. 2020. Universal sentence representation learning with conditional masked language model. *arXiv preprint arXiv:2012.14388*.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? *arXiv preprint arXiv:1801.07243*.