

# P5: Plug-and-Play Persona Prompting for Personalized Response Selection

Joosung Lee<sup>1</sup>, Minsik Oh, Donghun Lee<sup>2</sup>

<sup>1</sup>NAVER

<sup>2</sup>Kakao Brain

rung.joo@navercorp.com

## Abstract

The use of persona-grounded retrieval-based chatbots is crucial for personalized conversations, but there are several challenges that need to be addressed. 1) In general, collecting persona-grounded corpus is very expensive. 2) The chatbot system does not always respond in consideration of persona at real applications. To address these challenges, we propose a plug-and-play persona prompting method. Our system can function as a standard open-domain chatbot if persona information is not available. We demonstrate that this approach performs well in the zero-shot setting, which reduces the dependence on persona-ground training data. This makes it easier to expand the system to other languages without the need to build a persona-grounded corpus. Additionally, our model can be fine-tuned for even better performance. In our experiments, the zero-shot model improved the standard model by 7.71 and 1.04 points in the original persona and revised persona, respectively. The fine-tuned model improved the previous state-of-the-art system by 1.95 and 3.39 points in the original persona and revised persona, respectively. To the best of our knowledge, this is the first attempt to solve the problem of personalized response selection using prompt sequences. Our code is available on github <sup>1</sup>.

## 1 Introduction

Designing a system that naturally communicates with humans is of great interest to researchers and is widely applied to services such as Apple Siri and Amazon Alexa. One of the critical algorithms of these services is multi-turn response selection, which selects the most appropriate response among many response candidates. Selecting a personalized response with a customized chatbot is necessary for a more human-like conversational system.

<sup>1</sup><https://github.com/rungjoo/plug-and-play-prompt-persona>

Indeed, Zhang et al. (2018a) shows that dialog context alone is insufficient for response selection.

Zhang et al. (2018a) released PERSONA-CHAT where the speakers have each persona. The persona is expressed in multiple sentences, and they get to know each other through conversation. PERSONA-CHAT can be used for research on personalized response generation and selection. However, the following challenges exist in developing personalized response selection for a real application. 1) Building conversations based on Persona is very expensive. PERSONA-CHAT is data from the research environment in English, and persona-grounded corpus in other languages is difficult to access. That is, there is a challenge to build data for real applications. 2) In general domains, the persona may not need to be reflected. However, since previous approaches (Gu et al., 2019; Hua et al., 2020; Zhu et al., 2021; Gu et al., 2020b, 2021; Xu et al., 2022; Das et al., 2022) are trained in combination with personas, the model can select a response only given a persona. Therefore, previous approaches always have the disadvantage of reflecting personal information. For example, when the persona is related to a favorite food, it is not helpful knowledge when answering the other topics (i.e., weather). Ideally, a chatbot system needs the ability to consider persona as an option while maintaining standard response selection capabilities.

We propose P5 (Plug-and-Play Persona Prompting for Personalized Response Selection) to solve the above challenges. First, we assume that there is no expensive persona-based corpus. Therefore, we can train only standard response selection models that do not consider persona. Then, we show that the standard response selection model combined with persona prompting allows response selection to reflect persona, which is a zero-shot inference strategy. Persona prompting improves the performance of standard response selection in persona-based conversations. Also,

the model uses persona prompting as optional information because it is a plug-and-play method. If no persona is given to the model, the model acts as a standard response selection model. So we can optionally combine model and persona. Persona sentences to be used for prompting are selected by measuring the similarity to the response. We use a pre-trained model as our similarity model. Only top- $k$  persona sentences are used in order of highest similarity score. In addition, we introduce a zero-shot baseline SoP (Similarity of Persona) based on the similarity score.

To our best knowledge, previous studies only provide fine-tuned models. For comparison in these same experimental settings, we show the experimental results for fine-tuned P5. Our method further improves the performance of the fine-tuned strategy as well as the zero-shot strategy. Fine-tuned P5 achieves state-of-the-art, which proves that persona prompting is effective in learning the relationship between persona and context. We evaluate our methods on PERSONA-CHAT (Zhang et al., 2018a) and Focus (Jang et al., 2022). PERSONA-CHAT provides 19 negative responses and 1 positive response for personalized response selection. Focus is given only one positive response as response candidates. Therefore, we build the data by sampling 19 negative candidates.

## 2 Related Work

### 2.1 Standard Response Selection

In dialog systems, retrieval-based response selection is an important module. Earlier retrieval-based methods (Hu et al., 2014; Wang et al., 2015) attempted to select a response based on a single-turn context. However, since multi-turn response selection is a scenario for a more realistic service, recent studies (Gu et al., 2020a; Whang et al., 2021; Han et al., 2021) focus on selecting a response based on multi-turn context. These multi-turn response selection models leverage pre-trained language models to reflect context. It also improves performance by training a language model to understand conversations through a post-training strategy or multi-task learning. These studies are generally conducted on ubuntu (Lowe et al., 2015), douban (Wu et al., 2017), and e-commerce (Zhang et al., 2018b) corpus. Since these datasets are not given a persona, we refer to relevant studies as standard response selection models.

### 2.2 Personalized Response Selection

Standard response selection models suffer from a coherent personality by being trained on general conversations. Zhang et al. (2018a) releases PERSONA-CHAT dataset, which is a dialogue corpus built on persona. Jang et al. (2022) releases Focus dataset, which is a dialogue corpus built on persona and knowledge.

Recently, many studies introduce fine-tuned models in PERSONA-CHAT:

Hua et al. (2020) proposes an approach that detects only context-related persona. **RSM-DCK** (Response Selection Model that can Detect the relevant parts of the Context and Knowledge collection) introduces context selectors and knowledge selectors, which are soft-selection of persona through attention weights. Gu et al. (2020b) also performs soft-selection of persona, and iteratively referring not only between context and response representations but also between knowledge and response representations to collect deep matching features for scoring response candidates (**FIRE**: Filtering before Iteratively REferring). Zhu et al. (2021) introduces hard-selection of context-related persona, and shows that recent utterances in context are more important for response selection (**CSN**: Content Selection Network). Gu et al. (2021) shows that partner-persona as well as self-persona are important for response selection. **BERT-CRA** (BERT with Context-Response-Aware Persona) also achieves high performance by combining persona with BERT for context-aware-persona. Xu et al. (2022) suggests **COSPLAY** (COnccept Set guided Personalized dialogue generation Across both partY personas), which considers both speakers as "team". COSPLAY utilizes both self-persona and partner-persona, and proposes a Concept Set framework with a suite of knowledge-enhanced operations to process them such as set algebras, set expansion, and set distance. Das et al. (2022) first learns emotion and intent classifiers respectively with external data. Then, **BERT-EmA** (Emotion Aware Fusion) and **BERT-P-EnA** (Entailment Aware Fusion) are learned by adding the predicted emotion and intent of the utterance as input of BERT. These methods are based on BERT-CRA.

## 3 Task Definition

In the training phase, the dialogue dataset  $D_{train} = \{c, r, y\}$  is given where context  $c =$

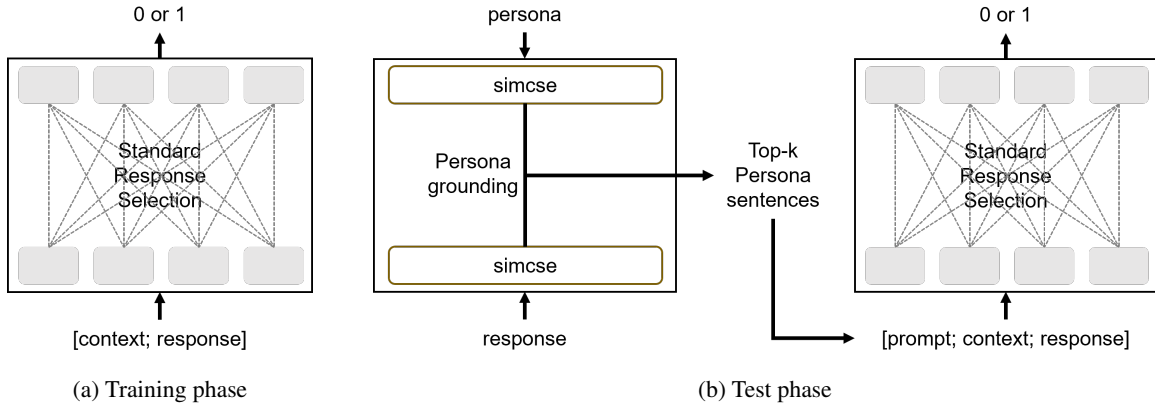


Figure 1: The overview architecture of our proposed P5 model

$\{u_1, u_2, \dots, u_n\}$ ,  $u_i$  is the  $i$ th utterance,  $r$  is a response candidate, and  $y \in \{0, 1\}$  is a label.  $y = 1$  indicates a positive response, the opposite is  $y = 0$ . In the test phase, the dialogue dataset  $D_{test} = \{c, r, p, y\}$  is given, where persona  $p = \{p_1, p_2, \dots, p_m\}$  are  $m$  persona sentences. Our goal is to learn a matching model without persona from  $D_{train}$  and select an appropriate response by reflecting persona from  $D_{test}$ .

## 4 Approach

Figure 1 shows the proposed approach. In the training phase, the model is trained to perform a multi-turn response selection task with  $D_{train}$  similar to Gu et al. (2020a); Han et al. (2021), which is called standard response selection. Since our goal is not to improve the performance of standard response selection, we do not use any particular strategy (i.e., post-training). The test phase consists of two steps. The first step is persona grounding corresponding to the response. The second step is calculating the matching score of  $(c, r, p)$  with the persona-prompted standard response selection model. Our approach improves the performance by combining the persona prompting with the standard response selection not trained with persona, which can also be utilized as a fine-tuned method.

### 4.1 Standard Response Selection (SRS)

The standard response selection model is trained with  $D_{train}$  without a persona. The standard response selection model follows pre-trained language models such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), which are Transformer encoders (Vaswani et al., 2017). We use RoBERTa as the backbone and train with binary classification for multi-turn response selection. The

input format ( $x_{st}$ ) is as follows:

$$x_{st} = \text{concat}([SEP], u_1, [SEP], u_2, \dots, [CLS], [SEP], r) \quad (1)$$

where  $u_i$  is  $i$ th utterance, [SEP] is a special token to distinguish each utterance, and [CLS] is prepended before the response. Most studies using pre-trained language models for response selection tasks prepend [CLS] to the front of the entire input sequence. However, since the distance between [CLS] and the response changes dynamically, an additional special token must be used to inform which response to classify. We intuitively change the position of [CLS] to represent the input as Equation 1.

The response score using the standard response selection model is calculated as follows:

$$s_{st} = W(\text{PLM}(x_{st})) \quad (2)$$

where PLM is a pre-trained language model used in the standard response selection model, and  $W$  is a matrix that projects the output vector of [CLS] into a two-dimensional vector.  $s_{st}$  is the score vector corresponding to  $x_{st}$ .

The standard response selection model is trained to minimize cross-entropy loss:

$$L = \frac{1}{N} \sum_{j=1}^N \text{CE}(s_{st}^j, y^j) \quad (3)$$

where  $j$  means the  $j$ th training sample, and  $N$  is the number of training data.

### 4.2 Persona Grounding

Zhao et al. (2019); Gu et al. (2019); Hua et al. (2020); Gu et al. (2021) use all given persona

the number of (#)	PERSON-CHAT			Focus	
	Train	Val	Test	Train	Val
# conversations	8,939	1,000	968	12,484	1,000
# turns	65,719	7,801	7,512	70,332	5,639

Table 1: Statistics of the two datasets

sentences as an input regardless of context and response. Instead, these approaches perform soft-selection of a persona by assigning attention weights between persona sentences and context embeddings. In other words, small weights are assigned to less relevant persona sentences, affecting response selection less.

Gu et al. (2020b); Zhu et al. (2021) mention that assigning a low weight to a less relevant persona is possible, but the cumulative weight of an irrelevant persona can be significant. Therefore, persona sentences used through the threshold are hard-selected based on the attention weight, which is an essential key for the model. However, since the hard-selection method extracts feature vectors through attention between persona and context, persona sentences are essential inputs for the trained model. That is, the previous frameworks require at least one persona sentence, which is different from selectively combining persona sentences like in our framework. Also, the hard-selection method requires training data, and our persona grounding is sufficient without training data.

We introduce an approach to select only the top- $k$  persona sentences simply and efficiently before combining the standard response selection model and persona in the test phase. A personalized response contains relevant personal information. Therefore, we find the persona sentence related to the response through the similarity between the response and the persona sentence.

$$e_r = \text{SimModel}([CLS], r) \quad (4)$$

$$e_{p_i} = \text{SimModel}([CLS], p_i) \quad (5)$$

$$s_{rp_i} = \frac{e_r \cdot e_{p_i}}{\|e_r\| \cdot \|e_{p_i}\|} \quad (6)$$

where  $r$  is the response,  $p_i$  is the  $i$ th persona sentence,  $e_r$  and  $e_{p_i}$  are the output vectors of [CLS] passed through the similarity model, and  $s_{rp_i}$  is the similarity score between the two vectors.

A persona sentence with a high similarity score ( $s_{rp_i}$ ) is considered to help select the corresponding response. Therefore, we combine only the top- $k$  persona sentences with high similarity scores with

the standard response selection model. In our default setting,  $k = 2$ , but it can be set dynamically. We used unsupervised simcse (Gao et al., 2021b) and supervised bert-nli (Reimers and Gurevych, 2019) as models for calculating similarity, and the difference between the two is introduced in Section 5.5.

### 4.3 Persona Prompting

Recently, GPT-3 (Brown et al., 2020) model leverages the natural-language prompt to improve few-shot performance very effectively. Gao et al. (2021a) achieves high performance by prompt-based fine-tuning of small language models on a small number of training data, which is a more practical scenario. Han et al. (2022) introduces a dialog prompt that is created using several utterances of fictional characters, in which a pre-trained language model that is not trained in character styles generates attractive responses that mimic the characters.

Inspired by these studies, we propose a persona prompting for personalized response selection. The prompt sequence asks and answers the speaker’s persona, simply composed of a prompt question and persona sentences. So the input format ( $x_p$ ) is as follows:

$$x_p = \text{concat}([SEP], p_q, [SEP], p_1, p_2, \dots, [SEP], u_1, [SEP], u_2, \dots, [CLS], [SEP], r) \quad (7)$$

where  $p_q$  is a prompt question, which defaults to "what is your personality?" sentences are used and  $p_i$  ( $i \in \{1, \dots, k\}$ ) are grounded persona sentences that result from Section 4.2. Other prompt questions are described in Section 5.6.

In Equation 2, the input is changed to  $x_p$ , and the response score is calculated as:

$$s_p = W(\text{PLM}(x_p)) \quad (8)$$

where  $s_p$  is the score vector of the response considering persona and context. The response selection model does not learn about persona fusion but naturally recognizes the prompt sequence as part of the context.

### 4.4 Baseline: Similarity of Persona (SoP)

In the previous approaches, only fine-tuned approaches have been studied under the assumption

that persona is given. Therefore, we introduce a baseline for a simple zero-shot setting, which utilizes a similarity score used to find a persona related to a response. That is, the final response score is the weighted sum of the response score of the standard response selection model and the similarity scores.

$$s_f = s_{st} + \alpha F(s_{rp_1}, \dots, s_{rp_m}) \quad (9)$$

where  $s_{rp_i}$  is obtained from Equation 6 as a similarity score between the  $i$ th persona sentence and the response, and  $s_{st}$  is obtained from Equation 2 as a score using the standard response selection model, and  $s_f$  is the final score.  $F$  is the function to aggregate  $s_{rp_i}$  and  $\alpha$  is the weight. We tried top- $k$  average function for  $F$ , but it was most effective to use top-1  $s_{rp_i}$ . Therefore, the max function is used for  $F$  in the experiment.

## 5 Experiments

### 5.1 Datasets

We experiment on two benchmark datasets. Table 1 shows the statistics of the datasets, and Table 6, 7 in Appendix A are examples.

**PERSONA-CHAT** The first dataset is PERSONA-CHAT, where each speaker is described with multiple persona sentences. PERSONA-CHAT is a dataset mainly used in previous studies, and 1 positive response and 19 negative response candidates corresponding to the context. Response selection is performed for every turn of dialogue. PERSONA-CHAT provides two versions of persona, original and revised. The revised persona is data that makes the task more difficult by rephrasing the original persona.

**Focus** The second dataset is Focus, a dialogue created using persona and knowledge. Focus was created from the motivation that more appropriate utterances can be generated by considering persona and knowledge together. However, for personalized response selection, we only use persona. Since only a positive response is given to the context in Focus, 19 negative response candidates are sampled and formatted according to the response selection task. The sampling strategy follows two steps. (1) Context sampling from the speaker’s previous utterances. In this case, utterances using the same persona sentence are sampled first. (2) Random sampling from the corpus. In (1), 2 candidates are sampled, and in (2), 17 candidates are sampled. Therefore, models can achieve high performance by considering both the appropriateness

of response and persona in Focus. Also, Focus has a label for persona grounding when constructing a positive response, so it is used to measure the performance of our proposed persona grounding. There can be multiple persona sentences labeled as "True".

### 5.2 Evaluation Metric

We use the evaluation metric used in previous works (Gu et al., 2021, 2020b) for a fair comparison. Each model checks whether the candidate with the highest ranking score is a positive response, denoted by  $R@1$ . Specifically, both PERSONA-CHAT and Focus are  $R_{20}@1$  because 1 positive response and 19 negative responses are given as candidates.

### 5.3 Training Setup

We use a pre-trained model from the huggingface library <sup>2</sup>. For the standard response selection model, we use AdamW as the optimizer. The learning rate is an initial value of 1e-6, and `get_linear_schedule_with_warmup` provided by the huggingface library is used for the learning rate scheduler. The maximum value of 10 is used for the gradient clipping. The training epoch is 10, the model is evaluated on the validation data for each epoch, and the best model is selected. In prompt-based fine-tuning, the training epoch is 5, and the rest are the same as the standard response selection. All experiments are performed on one A100 GPU, and the results are for a single turn because there is little variation between each run.

### 5.4 Results

Table 2 shows the evaluation results of the previous and the proposed method for two benchmarks. Previous methods are described in Section 2.2. SRS is a standard response selection model that does not consider persona. SoP is a baseline using the similarity score of persona introduced in Section 4.4. P5 is our proposed method using persona prompting.

SRS does not consider persona but achieves a performance of 72.4 in PERSONA-CHAT, which is an unsatisfactory performance in the original persona, but a good performance in the revised persona. We believe that many previous models (RSM-DCK, FIRE, CSN-word, COSPLAY) were

<sup>2</sup><https://huggingface.co/docs/transformers/index>

Method	Model	Original Persona	Revised Persona	Focus
Fine-tuning	RSM-DCK (Hua et al., 2020)	79.65	71.85	-
	FIRE (Gu et al., 2020b)	81.6	74.8	
	CSN-word (Zhu et al., 2021)	78.1	70.1	
	BERT-CRA (Gu et al., 2021)	84.3	79.4	
	COSPLAY (Xu et al., 2022)	85.5	74.4	
	BERT-EmA (Das et al., 2022)	84.6	79.8	
	BERT-P-EnA (Das et al., 2022)	85.3	80.5	
	<b>P5</b>	<b>87.45</b>	<b>82.79</b>	
Zero-shot	SRS	72.4	72.4	91.65
	SoP	76.41	73.16	94.27
	<b>P5</b>	<b>80.11</b>	<b>73.44</b>	<b>97.85</b>

Table 2: Evaluation results on the test sets of PERSONA-CHAT and validation sets of Focus. Performance is measured as  $R@1$ . Bold text indicates the best performance in each part. In PERSONA-CHAT and Focus, RoBERTa-base are used as PLM.

effective at fusing the original persona, but did not fuse the revised persona well.

SoP is a simple baseline we propose, which improves the performance of SRS. In Equation 9,  $\alpha$  is 0.5 and 0.05 in the original persona (or Focus) and revised persona, respectively, and an appropriate value for  $\alpha$  was selected through an experiment. Since the original persona has many examples that directly overlap the words of the response, the similarity score is more effective. However, since a revised persona is a rephrased sentence, simply scoring similarity with persona sentences is less effective.

P5 identifies the speaker’s persona from the prompt in the form of simple dialogue and uses it for response ranking. P5 achieves the best performance in both zero-shot and fine-tuning. Zero-shot P5 improves the performance of the SRS through persona prompting by 7.71 points in the original persona and 1.04 points in the revised persona. The zero-shot inference strategy is more effective for the original persona (or Focus) than the revised persona, which is considered difficult for the SRS to understand the revised persona as a prompt context. When a persona is given as training data, the fine-tuned P5 achieves 87.45 and 82.79 performance in the original persona and revised persona, respectively, which is a remarkable performance improvement compared to previous models. Also, our proposed prompt is a plug-and-play module and has the advantage of being turned on and off according to the real application.

We further experimented with the zero-shot setting in Focus to verify the effectiveness of our model. Focus is structured similarly to the original persona in PERSONA-CHAT, and the experimen-




Figure 2: Performance of similarity model on persona grounding in Focus

tal results also show the same aspect as the original persona. Since zero-shot P5 has already achieved satisfactory performance in Focus, fine-tuned P5 has not been tested.

## 5.5 Effects of Persona Grounding

With a persona format similar to the original persona in PERSONA-CHAT, Focus provides a persona grounding label. Figure 2 shows the performance of persona grounding using simcse and bert-nli models in Focus, where both simcse and bert-nli models are pre-trained models. The evaluation metric is  $R@k$ , where the value is considered 100 if the number of "True" candidates is 0. As  $k$  increases,  $R@k$  improves, which increases the probability that persona sentences related to the response are reflected. However, as  $k$  increases, irrelevant persona sentences are also entered as input. Therefore, an appropriate value of  $k$  is required.

Model	prompt question	similarity model	order of persona	Original Persona	Revised Persona
Zero-shot P5	<i>what is your personality?</i>	<b>simcse</b>	<b>ascending</b>	<b>80.11</b>	<b>73.44</b>
	<i>tell me your personality.</i>	simcse	ascending	79.73	73.34
	<i>tell me more about yourself.</i>	simcse	ascending	79.81	73.32
	random utterance	simcse	ascending	78.89	72.63
	empty string	simcse	ascending	79.71	73.31
	<i>what is your personality?</i>	nli	ascending	79.86	73.4
	<i>what is your personality?</i>	simcse	descending	79.7	73.34

Table 3: Performance with variants of persona prompting

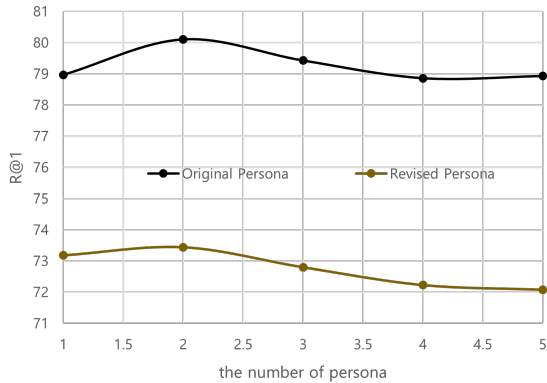


Figure 3: Performance of zero-shot P5 change according to the number of persona sentences used

In Table 2, P5 is the evaluation results using only 2 persona sentences (top-2). Since PERSONA-CHAT does not have a label for persona grounding, we cannot confirm the ground truth persona sentences reflected in the response. Figure 3 is zero-shot P5 performance change according to the number of persona sentences used in PERSONA-CHAT. The best performance is achieved when two persona sentences are used in both the original persona and the revised persona. More persona sentences increase the persona grounding performance but confuse the personality to be reflected by the model and degrade the performance. Simcse has a higher rate of finding used persona sentences than the bert-nli model, and similar results can be expected in PERSONA-CHAT. Comparing the 1st and 6th rows in Table 3, it can be seen that simcse is more effective than bert-nli in PERSONA-CHAT.

## 5.6 Structure of Persona Prompting

Table 3 shows zero-shot P5 performance for the variant of persona prompting. We changed the prompt question "what is your personality?" to the rephrase "tell me your personality." and "tell me more about yourself.". The performance difference

according to the prompt question is not large, and it is not easy to find the optimal prompt question in the discrete space.

The random utterance (4th row) indicates that the prompt question was randomly sampled from the training utterances. The empty string (5th row) indicates that there is no prompt question, which means that the prompt sequence consists only of persona sentences. These two methods make it difficult to know whether the persona sentences represent the speaker's personality. Performance is slightly lower than when the prompt question is "what is your personality?", but it doesn't show a huge difference. That is, persona sentences are more important to performance than prompt questions.

We also experiment with two methods for the input sequence of a grounded persona. The first is the ascending method, from lowest to highest similarity score, which means that the position of the most similar persona and response is close (1st row). The second is the descending method, from highest to lowest similarity score, which means that the position of the most similar persona and response is far (7th row). Experimental results show that the ascending method achieves higher performance. Therefore, the closer the distance between the relevant persona and the response, the more effective it is for the model.

## 5.7 Other Standard Response Selection Model

Table 4 shows the results for another backbone in PERSONA-CHAT. We experimented by changing the backbone of the SRS from RoBERTa-base to RoBERTa-large. The performance of the SRS-large is better than that of the SRS-base. Zero-shot P5-large improves the performance by 7.17 and 1.65 points, respectively, in the original persona and revised persona compared to SRS. Regardless of the performance of the SRS, the P5 effectively combines persona to improve performance. Table 8 in Appendix B show experiments on more diverse

Method	Model (large)	Original Persona	Revised Persona
Fine-tuning	P5	90.54	86.62
Zero-shot	SRS	77.96	77.96
	P5	85.13	79.61

Table 4: Experimental results for a large model.

backbones.

Fine-tuned P5-base achieves the best performance, but fine-tuned P5-large achieves SoTA with a larger margin. In addition, zero-shot P5-large achieves competitive performance with previous fine-tuning approaches. That is, with better SRS, we observe that even the zero-shot approach can achieve remarkable performance.

## 5.8 Ablation Study

We perform several ablation studies when testing to know the importance of each part of the framework. Without persona grounding (G), P5 considers the personality to be all persona sentences, and the order is randomized to form the prompt sequence. This is similar to the experiment in Figure 3 where the number of persona = 5. Without prompt question (Q), P5 uses only persona sentences as the prompt sequence, which is the same as the 5th row in Table 3. Without a prompt sequence (P), P5 does not consider persona as context, which is equivalent to standard response selection. -D indicates that SRS does not have access to PERSONA-CHAT dialogue, so it is trained with external data. We use dailydialog (Li et al., 2017) as training data and 10 negative candidate responses are randomly sampled.

In the zero-shot setting, ablation studies are performed on persona grounding, prompt question, prompt sequence (prompt question+persona sentence), and dialogue corpus. All components clearly show differences in the original persona. The absence of persona grounding and prompt questions reduces performance, but these are considered minor components. However, persona sentences are an important component of performance, and using them as prompt sequences is our major contribution. We also assumed that the model is inaccessible to dialogue as well as persona sentences from PERSONA-CHAT. So SRS is trained as dailydialog. Zero-shot P5 (-D&P) (i.e. SRS w/ dailydialog), without using persona sentences in the test, achieves a performance of 44.5. Zero-shot P5 (-D) utilizing persona sentences achieves the performance of 59.94 and 50.23 in the original

Method	Model	Original Persona	Revised Persona
Zero-shot	P5	80.11	73.44
	-G	78.93	72.08
	-Q	79.71	73.31
	-P	72.4	72.4
	-D	59.94	50.23
	-D&P	44.5	44.5
Fine-tuning	P5	87.45	82.79
	-G	87.13	81.55
	-P	68.18	70.38

Table 5: Ablation study for each module of the framework. G stands for persona grounding, Q stands for prompt question, P stands for prompt sequence, and D stands for PERSONA-CHAT dialogue.

persona and revised persona, respectively, which is a much larger performance improvement than shown in Table 2. With the same conclusion as in Section 5.7, the proposed prompting proves to lead to a large performance improvement regardless of SRS.

In the fine-tuning setting, ablation studies are performed on persona grounding, prompt sequences when testing. Fine-tuning P5 (-G) achieves 87.13 and 81.55 performance on the original and revised, respectively, showing that the performance difference due to persona grounding is smaller than the zero-shot method. In addition, it has the advantage of not requiring additional computation for persona grounding. This is because the model gains the ability to attend to the appropriate persona when selecting a response through learning from the persona corpus. Therefore, our prompting method operates effectively in a fair comparison with other frameworks. Fine-tuning P5 (-P) achieves a performance of 68.18 and 70.38 in the original and revised versions, respectively, which is worse than zero-shot P5 (-P). Therefore, we find that fine-tuning P5 exhibits a strong dependence on the persona sentence when selecting responses. These limitations will have similar limitations to fine-tuned models as in previous studies.

## 6 Conclusion

In this paper, we present a method called P5, which functions as a plug-and-play system that only incorporates persona when desired. Our approach involves identifying related persona sentences through their similarity to a given response, and then adding these sentences as a prompt to the input. This allows the standard response selection model to better match context and response



by taking into account the persona. We evaluate our method on two benchmark datasets using both fine-tuning and zero-shot settings. Fine-tuned P5 outperforms previous studies by a significant margin. Zero-shot P5 also effectively improves performance when compared to standard response selection models. Even the zero-shot P5-large shows performance that is comparable to previous fine-tuning approaches.

P5 is only evaluated using persona-based corpus, however, in real-world applications, persona information is not always available. Therefore, it is important that the standard response selection model can be combined with persona in a dynamic manner. One way to achieve this is by only incorporating persona sentences that have a similarity score above a certain threshold. We plan to investigate other options for reflecting persona in future studies.

## Limitations

P5 is only evaluated using persona-based corpus, however, in real-world applications, persona information is not always available. Therefore, it is important that the standard response selection model can be combined with persona in a dynamic manner. One way to achieve this is by only incorporating persona sentences that have a similarity score above a certain threshold. We plan to investigate other options for reflecting persona in future studies.

The importance of a standard response selection model outweighs the use of persona sentences in personalized response selection. In Table 5, the P5 (-D) performance improves with persona prompting, however, it is still lower than that of P5 (-P). The low performance of the standard response selection model (P5 (-D&P)) is the reason for this. To improve zero-shot P5 performance, it is crucial to improve the standard response selection performance. Therefore, we will conduct further research on enhancing the performance of zero-shot standard response selection models that do not utilize PERSONA-CHAT.

## References

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris 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](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Souvik Das, Sougata Saha, and Rohini K. Srihari. 2022. [Using multi-encoder fusion strategies to improve personalized response selection](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 532–541, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Tianyu Gao, Adam Fisch, and Danqi Chen. 2021a. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.

Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021b. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Jia-Chen Gu, Tianda Li, Quan Liu, Zhen-Hua Ling, Zhiming Su, Si Wei, and Xiaodan Zhu. 2020a. [Speaker-aware bert for multi-turn response selection in retrieval-based chatbots](#). In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM '20*, page 2041–2044, New York, NY, USA. Association for Computing Machinery.

Jia-Chen Gu, Zhen-Hua Ling, Xiaodan Zhu, and Quan Liu. 2019. [Dually interactive matching network for personalized response selection in retrieval-based chatbots](#). 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 1845–1854, Hong Kong, China. Association for Computational Linguistics.

- Jia-Chen Gu, Zhenhua Ling, Quan Liu, Zhigang Chen, and Xiaodan Zhu. 2020b. [Filtering before iteratively referring for knowledge-grounded response selection in retrieval-based chatbots](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1412–1422, Online. Association for Computational Linguistics.
- Jia-Chen Gu, Hui Liu, Zhen-Hua Ling, Quan Liu, Zhigang Chen, and Xiaodan Zhu. 2021. [Partner matters! an empirical study on fusing personas for personalized response selection in retrieval-based chatbots](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 565–574, New York, NY, USA. Association for Computing Machinery.
- Janghoon Han, Taesuk Hong, Byoungjae Kim, Youngjoong Ko, and Jungyun Seo. 2021. [Fine-grained post-training for improving retrieval-based dialogue systems](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1549–1558, Online. Association for Computational Linguistics.
- Seungju Han, Beomsu Kim, Jin Yong Yoo, Seokjun Seo, Sangbum Kim, Enkhbayar Erdenee, and Buru Chang. 2022. [Meet your favorite character: Open-domain chatbot mimicking fictional characters with only a few utterances](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5114–5132, Seattle, United States. Association for Computational Linguistics.
- Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, page 2042–2050, Cambridge, MA, USA. MIT Press.
- Kai Hua, Zhiyuan Feng, Chongyang Tao, Rui Yan, and Lu Zhang. 2020. [Learning to detect relevant contexts and knowledge for response selection in retrieval-based dialogue systems](#). *CIKM '20*, page 525–534, New York, NY, USA. Association for Computing Machinery.
- Yoonna Jang, Jungwoo Lim, Yuna Hur, Dongsuk Oh, Suhyune Son, Yeonsoo Lee, Dong-Hoon Shin, Seungryong Kim, and Heuiseok Lim. 2022. [Call for customized conversation: Customized conversation grounding persona and knowledge](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 10803–10812. AAAI Press.
- Zi-Hang Jiang, Weihao Yu, Daquan Zhou, Yunpeng Chen, Jiashi Feng, and Shuicheng Yan. 2020. [Convbert: Improving bert with span-based dynamic convolution](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 12837–12848. Curran Associates, Inc.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- 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.
- 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](#). *CoRR*, abs/1907.11692.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. [The Ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems](#). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294, Prague, Czech Republic. Association for Computational Linguistics.
- 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 and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Mingxuan Wang, Zhengdong Lu, Hang Li, and Qun Liu. 2015. [Syntax-based deep matching of short texts](#). In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, page 1354–1361. AAAI Press.
- Taesun Whang, Dongyub Lee, Dongsuk Oh, Chanhee Lee, Kijong Han, Donghun Lee, and Saebyeok Lee. 2021. [Do response selection models really know what's next? utterance manipulation strategies for multi-turn response selection](#). In *AAAI*.

Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. [Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 496–505, Vancouver, Canada. Association for Computational Linguistics.

Chen Xu, Piji Li, Wei Wang, Haoran Yang, Siyun Wang, and Chuangbai Xiao. 2022. [Cosplay: Concept set guided personalized dialogue generation across both party personas](#). SIGIR '22, page 201–211, New York, NY, USA. Association for Computing Machinery.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018a. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Zhuosheng Zhang, Jiangtong Li, Pengfei Zhu, Hai Zhao, and Gongshen Liu. 2018b. [Modeling multi-turn conversation with deep utterance aggregation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3740–3752, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Xueliang Zhao, Chongyang Tao, Wei Wu, Can Xu, Dongyan Zhao, and Rui Yan. 2019. [A document-grounded matching network for response selection in retrieval-based chatbots](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5443–5449. International Joint Conferences on Artificial Intelligence Organization.

Yutao Zhu, Jian-Yun Nie, Kun Zhou, Pan Du, and Zhicheng Dou. 2021. [Content selection network for document-grounded retrieval-based chatbots](#). Berlin, Heidelberg. Springer-Verlag.

## A Dataset Example

PERSONA-CHAT has {8,939, 1,000, 968} conversations in {train, validation, test} and Focus has {12,484, 1,000} conversations in {train, validation}. Table 6, 7 show a example of PERSONA-CHAT and Focus, respectively.

## B Additional Standard Response Selection Model

We prove that the proposed prompt sequence is effective through more extensive experiments with various backbones. In addition to RoBERTa, we experiment with BERT (Devlin et al., 2019), ALBERT (Lan et al., 2020), and ConvBERT (Jiang

et al., 2020) as backbones. **ALBERT** (A Lite BERT) is a language model for natural language processing that was designed to be more efficient than its predecessor, BERT. ALBERT reduces the amount of computation through matrix decomposing through factorized embedding parameterization and uses Sentence Order Prediction (SOP) as loss. **ConvBERT** is a language model for natural language processing that combines the power of convolutional neural networks with transformer architecture. It utilizes a hierarchical convolutional structure to capture local dependencies and long-range dependencies, resulting in improved computational efficiency and better performance on various language tasks. The experimental settings are the same for all backbones.

SRS is a response selection model trained without persona sentences. P5 is a model in which SRS receives both the prompt sequence and context as inputs and selects a response when testing. "with DD" means that the performance was measured based on the SRS learned by dailydialog, which is the same as -D and -D&P in Table 5. Depending on the backbone model, the effect of improving performance is different, but in the original persona, the model's ability improves by a large margin. In the revised persona, P5 performance improves less than in the original persona. The large backbone did not differ from the base backbone. SRS w/ DD trained with dailydialog has lower performance, but the effect of the prompt sequence is significant.

We demonstrate that prompt sequences improve the performance of SRS on all backbones. We believe that pre-trained language models help to perform this zero-shot inference because they are trained from a large corpus and thus have the ability to understand context. The prompt sequence helps to achieve higher performance regardless of the performance of SRS, which is a very simple but powerful method. However, the effect is relatively small in the revised persona, so further research will be done in this area.

	Pesona 1	Persona 2
Original	i love to meet new people. my favorite sport is ultimate frisbee. autumn is my favorite season. my parents are living in bora bora. i have a turtle named timothy.	i work as an accountant. i live in ohio. i am a single mom of two boys. i drive a honda civic. i like to go hiking in my spare time.
Revised	i like getting friends. i love to run around and get out my energy. i love watching the leaves change colors. my family lives on a island. reptiles make good pets.	they call me a bean counter. i am from the north. i am raising sons all on my own. i own a small car. i enjoy nature walks.
Dialogue	<p>person 1: hi , i am kera and i am a social butterfly</p> <p>person 2: hi . i am more the mousy type . numbers are my world at my day job . you ?</p> <p>person 1: i work for a tech firm , i am a tom girl</p> <p>person 2: i am just an ohio mom with two amazing sons . not married though .</p> <p>person 1: cool . i have no kids just my pet turtle timothy</p> <p>person 2: great pet name . i do not have any pets unless you count my car , sally .</p>	

Table 6: An example from PERSONA-CHAT dataset

	Persona 2	Persona Grouding
Original	I like to go to Church.	FALSE
	I am Roman Catholic.	FALSE
	I wish to go to El Paso.	TRUE
	I would like to go to Texas.	FALSE
	I love the United States.	FALSE
Dialogue	person 1: Wow, this is amazing! What is this?	
	person 2: It is a Church, something which you like.	
	person 1: What is the name of this place?	
	person 2: The name of this place is The Roman Catholic Diocese of El Paso, remember you are Roman Catholic.	
	<p>person 1: Where is this place?</p> <p>person 2: It is located in El Paso, a city which you wish to go.</p>	

Table 7: An example from Focus dataset. The label of persona grounding is for the last utterance.

Backbone	Method	Model	Original Persona	Revised Persona
RoBERTa-base	Zero-shot	SRS	72.4	72.4
		P5	80.11	73.44
		SRS w/ DD	44.5	44.5
		P5 w/ DD	59.94	50.23
RoBERTa-large		SRS	77.96	77.96
		P5	85.13	79.61
BERT-base		SRS	69.61	69.61
		P5	76.66	69.74
ALBERTv2-base	SRS	67.37	67.37	
	P5	76.58	68.89	
ConvBERT-base	SRS	69.68	69.68	
	P5	78.49	71.49	

Table 8: Experimental results from all the backbones we experimented with.