

# Dually Interactive Matching Network for Personalized Response Selection in Retrieval-Based Chatbots

Jia-Chen Gu<sup>1</sup>, Zhen-Hua Ling<sup>1</sup>, Xiaodan Zhu<sup>2</sup>, Quan Liu<sup>1,3</sup>

<sup>1</sup>National Engineering Laboratory for Speech and Language Information Processing, University of Science and Technology of China, Hefei, China

<sup>2</sup>ECE, Queen's University, Kingston, Canada

<sup>3</sup>State Key Laboratory of Cognitive Intelligence, iFLYTEK Research, Hefei, China  
gujcc@mail.ustc.edu.cn, zhling@ustc.edu.cn,  
xiaodan.zhu@queensu.ca, quanliu@ustc.edu.cn

## Abstract

This paper proposes a dually interactive matching network (DIM) for presenting the personalities of dialogue agents in retrieval-based chatbots. This model develops from the interactive matching network (IMN) which models the matching degree between a context composed of multiple utterances and a response candidate. Compared with previous *persona* fusion approaches which enhance the representation of a context by calculating its similarity with a given *persona*, the DIM model adopts a dual matching architecture, which performs interactive matching between responses and contexts and between responses and *personas* respectively for ranking response candidates. Experimental results on PERSONA-CHAT dataset show that the DIM model outperforms its baseline model, i.e., IMN with *persona* fusion, by a margin of 14.5% and outperforms the current state-of-the-art model by a margin of 27.7% in terms of top-1 accuracy **hits@1**.

## 1 Introduction

Building a conversation system with intelligence is challenging. Response selection, which aims to select a potential response from a set of candidates given the context of a conversation, is an important technique to build retrieval-based chatbots (Zhou et al., 2018). Many previous studies on single-turn (Wang et al., 2013) or multi-turn response selection (Lowe et al., 2015; Zhou et al., 2018; Gu et al., 2019) rank response candidates according to their semantic relevance with the given context.

With the emergence and popular use of personal assistants such as Apple Siri, Google Now and Microsoft Cortana, the techniques of making personalized dialogues has attracted much research attention in recent years (Li et al., 2016; Zhang et al., 2018; Mazaré et al., 2018). Zhang et al. (2018) constructed a PERSONA-CHAT dataset

for building personalized dialogue agents, where each *persona* was represented as multiple sentences of profile description. An example dialogue conditioned on given profiles from this dataset is given in Table 1 for illustration.

A *persona* fusion method for personalized response selection was also proposed by Zhang et al. (2018). In this method, given a context and a *persona* composed of several profile sentences, the similarities between the context representation and all profile sentences are computed first using attention to get the *persona* representation. Then, the *persona* representation is applied to enhance the context representation by a simple concatenation or addition operation. Finally, the enhanced context representation is used to rank response candidates. This method has two main deficiencies. First, the context is treated as a whole for calculating its attention towards profile sentences. However, each context is composed of multiple utterances and these utterances may play different roles when matching different profile sentences. Second, the interactions between the *persona* and each response candidate are ignored when deriving the *persona* representation.

In this paper, the interactive matching network (IMN) (Gu et al., 2019) is adopted as the fundamental architecture to build our baseline and improved models for personalized response selection. The baseline model follows the *persona* fusion method proposed by Zhang et al. (2018) and two improved models are then proposed. First, an IMN-based *persona* fusion model with fine-grained context-*persona* interaction is designed. In this model, each utterance in a context, instead of the whole context, is used to calculate its similarity with each profile sentence in a *persona*. Second, a dually interactive matching network (DIM) is proposed by formulating the task of personalized response selection as a dual matching problem,

Persona 1		Persona 2	
Original	I just bought a brand new house. I like to dance at the club. I run a dog obedience school. I have a big sweet tooth. I like taking and posting selkies.	Original	I love to meet new people. I have a turtle named timothy. My favorite sport is ultimate frisbee. My parents are living in bora bora. Autumn is my favorite season.
Revised	I have purchased a home. Just go dancing at the nightclub, it is fun! I really enjoy animals. I enjoy chocolate. I pose for pictures and put them online.	Revised	I like getting friends. Reptiles make good pets. I love to run around and get out my energy. My family lives on a island. I love watching the leaves change colors.

#### Dialogue

Person 1: Hello, how are you doing tonight?  
Person 2: I am well an loving this interaction how are you?  
Person 1: I am great. I just got back from the club.  
Person 2: This is my favorite time of the year season wise.  
Person 1: I would rather eat chocolate cake during this season.  
Person 2: What club did you go to? Me an timothy watched tv.  
Person 1: I went to club chino. What show are you watching?  
Person 2: LOL oh okay kind of random.  
Person 1: I love those shows. I am really craving cake.  
Person 2: Why does that matter any? I went outdoors to play frisbee.  
Person 1: It matters because I have a sweet tooth.  
Person 2: So? LOL I want to meet my family at home in bora.  
Person 1: My family lives in alaska. It is freezing down there.  
Person 2: I bet it is oh I could not.

Table 1: An example dialogue from the PERSONA-CHAT dataset.

i.e., finding a response that can properly match the given context and persona simultaneously. The DIM model calculates the interactions between the context and the response, and between the persona and the response in parallel, in order to derive the final matching feature for response selection.

We test our proposed methods on the PERSONA-CHAT dataset (Zhang et al., 2018). Results show that the IMN-based utterance-level persona fusion model and the DIM model can obtain a top-1 accuracy **hits@1** improvement of 2.4% and 14.5%, respectively, over the baseline model, i.e., the IMN-based context-level persona fusion model. Finally, our proposed DIM model outperforms the current state-of-the-art model by a margin of 27.7% in terms of top-1 accuracy **hits@1** on the PERSONA-CHAT dataset.

In summary, the contributions of this paper are three-fold. (1) An IMN-based fine-grained persona fusion model is designed in order to consider the utterance-level interactions between contexts and personas. (2) A dually interactive matching network (DIM) is proposed by formulating the task of personalized response selection as a dual matching problem, aiming to find a response that can properly match the given context and persona simultaneously. (3) Experimental results

on the PERSONA-CHAT dataset demonstrate that our proposed models outperform the baseline and state-of-the-art models by large margins on the accuracy of response selection.

## 2 Related Work

### 2.1 Response Selection

Response selection is an important problem in building retrieval-based chatbots. Existing work on response selection can be categorized into single-turn (Wang et al., 2013) and multi-turn dialogues (Lowe et al., 2015; Zhou et al., 2018; Gu et al., 2019). Early studies have been more on single-turn dialogues, considering only the last utterance of a context for response matching. More recently, the research focus has been shifted to multi-turn conversations, a more practical setup for real applications. Wu et al. (2017) proposed the sequential matching network (SMN) which first matched the response with each context utterance and then accumulated the matching information by a recurrent neural network (RNN). Zhou et al. (2018) proposed the deep attention matching network (DAM) to construct representations at different granularities with stacked self-attention. Gu et al. (2019) proposed the interactive matching network (IMN)

to enhance the representations of the context and response at both the word-level and sentence-level, and to perform the bidirectional and global interactions between the context and response in order to derive the matching feature vector.

## 2.2 Persona for Chatbots

Chit-chat models suffer from a lack of a consistent personality as they are typically trained over many dialogues, each with different speakers, and a lack of explicit long-term memory as they are typically trained to produce an utterance given only a very recent dialogue history. Li et al. (2016) proposed a persona-based neural conversation model to capture individual characteristics such as background information and speaking style. Miller et al. (2016) proposed the key-value memory network, where the keys were dialogue histories, i.e., contexts, and the values were next dialogue utterances. Zhang et al. (2018) proposed the profile memory network by considering the dialogue history as input and then performing attention over the persona to be combined with the dialogue history. Mazaré et al. (2018) proposed the fine-tuned persona-chat (FT-PC) model which first pretrained a model using a large-scale corpus with external knowledge and then fine-tuned it on the PERSONA-CHAT dataset.

In general, all these methods adopted a context-level persona fusion strategy, which first obtained the embedding vector of a context and then computed the similarities between the whole context and each profile sentence to acquire the persona representation. However, such persona fusion is relatively too coarse. The utterance-level representations of contexts are not leveraged. The interactions between the persona and each response candidate are also ignored when deriving the persona representation.

## 3 Task Definition

Given a dialogue dataset  $\mathcal{D}$  with personas, an example of the dataset can be represented as  $(c, p, r, y)$ . Specifically,  $c = \{u_1, u_2, \dots, u_{n_c}\}$  represents a context with  $\{u_m\}_{m=1}^{n_c}$  as its utterances and  $n_c$  as the utterance number.  $p = \{p_1, p_2, \dots, p_{n_p}\}$  represents a persona with  $\{p_n\}_{n=1}^{n_p}$  as its profile sentences and  $n_p$  as the profile number.  $r$  represents a response candidate.  $y \in \{0, 1\}$  denotes a label.  $y = 1$  indicates that  $r$  is a proper response for  $(c, p)$ ; otherwise,  $y = 0$ .

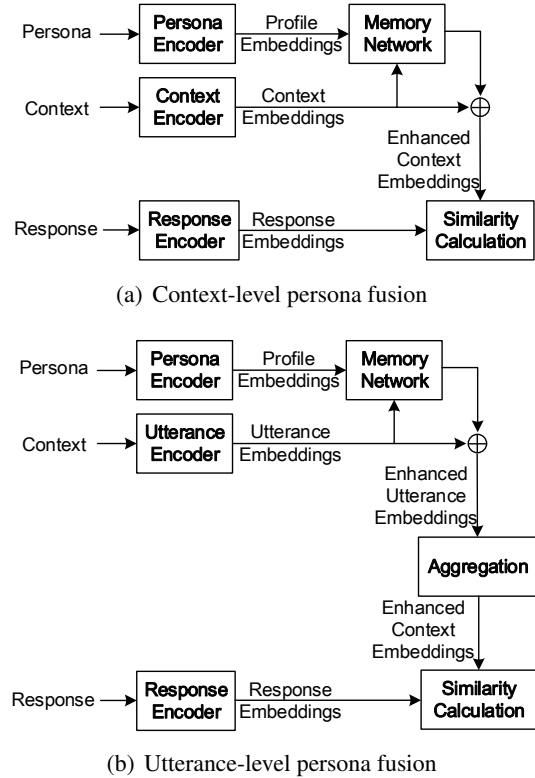


Figure 1: Comparison of the model architectures for (a) context-level persona fusion and (b) utterance-level persona fusion.

Our goal is to learn a matching model  $g(c, p, r)$  from  $\mathcal{D}$ . For any context-persona-response triple  $(c, p, r)$ ,  $g(c, p, r)$  measures the matching degree between  $(c, p)$  and  $r$ . A softmax output layer over all response candidates is adopted in this model. The model parameters are trained by minimizing a multi-class cross-entropy loss function on  $\mathcal{D}$ .

## 4 IMN-Based Persona Fusion

The model architecture used by previous methods with persona fusion (Zhang et al., 2018; Mazaré et al., 2018) is shown in Figure 1(a). It first obtains the context representation and then computes the similarities between the whole context and each profile sentence in a persona. Attention weights are calculated for all profile sentences to obtain the persona representation. Finally, the persona representation is combined with the context representation through concatenation or addition operations.

Formally, the representations of the whole context which is the concatenation of utterances, the context utterances, and the profile sentences are denoted as  $\mathbf{c}$ ,  $\{\mathbf{u}_m\}_{m=1}^{n_c}$  and  $\{\mathbf{p}_n\}_{n=1}^{n_p}$  respectively,

where  $\mathbf{c}$ ,  $\mathbf{u}_m$  and  $\mathbf{p}_n \in \mathbb{R}^d$ . In previous context-level persona fusion methods, the enhanced context representation  $\mathbf{c}^+$  fused with persona information is calculated as

$$\mathbf{c}^+ = \mathbf{c} + \sum_n \text{Softmax}(\mathbf{c} \cdot \mathbf{p}_n) \mathbf{p}_n. \quad (1)$$

Then, the similarity between  $\mathbf{c}^+$  and the response representation are computed to get the matching degree of  $(c, p, r)$ .

In this paper, we build our baseline model based on IMN (Gu et al., 2019). After the context and response embeddings are obtained in the IMN model, the context-level persona fusion architecture shown in Figure 1(a) is applied to integrate persona information. All model parameters are estimated in an end-to-end manner. This baseline model is denoted as  $\text{IMN}_{ctx}$  in this paper.

Considering each context is composed of multiple utterances and these utterances may play different roles when matching different profile sentences, we propose to improve the baseline model by fusing the persona information at a fine-grained utterance-level as shown in Figure 1(b). This model is denoted as  $\text{IMN}_{utr}$  in this paper. First, the similarities between each context utterance and each profile sentence are computed and the enhanced representation  $\mathbf{u}_m^+$  of each context utterance is calculated as

$$\mathbf{u}_m^+ = \mathbf{u}_m + \sum_n \text{Softmax}(\mathbf{u}_m \cdot \mathbf{p}_n) \mathbf{p}_n. \quad (2)$$

Then, these enhanced utterance representations are aggregated into the enhanced context representation as

$$\mathbf{c}^+ = \text{Aggregation}(\{\mathbf{u}_m^+\}_{m=1}^{n_c}), \quad (3)$$

where either RNN or attention-based aggregation (Gu et al., 2019) can be employed.

## 5 Dually Interactive Matching Network

### 5.1 Model Overview

Previous studies on personalized response selection treat personas as supplementary information to enhance context representations by attention-based interaction. In this paper, we formulate the task of personalized response selection as a dual matching problem. The selected response is expected to properly match the given context and persona respectively. Here, personas are considered as equally important counterparts to contexts

for ranking response candidates. The interactive matching between the context and response, and that between the persona and response constitute the dually interactive matching network (DIM).

The DIM model is composed of five layers. Figure 2 shows an overview of the architecture. Details about each layer are provided in the following subsections.

### 5.2 Word Representation Layer

We follow the setting used in IMN (Gu et al., 2019), which constructs word representations by combining general pre-trained word embeddings, those estimated on the task-specific training set, as well as character-level embeddings, in order to deal with the out-of-vocabulary issue.

Formally, embeddings of the  $m$ -th utterance in a context, the  $n$ -th profile sentence in a persona and a response candidate are denoted as  $\mathbf{U}_m = \{\mathbf{u}_{m,i}\}_{i=1}^{l_{u_m}}$ ,  $\mathbf{P}_n = \{\mathbf{p}_{n,j}\}_{j=1}^{l_{p_n}}$  and  $\mathbf{R} = \{\mathbf{r}_k\}_{k=1}^{l_r}$  respectively, where  $l_{u_m}$ ,  $l_{p_n}$  and  $l_r$  are the numbers of words in  $\mathbf{U}_m$ ,  $\mathbf{P}_n$  and  $\mathbf{R}$  respectively. Each  $\mathbf{u}_{m,i}$ ,  $\mathbf{p}_{n,j}$  or  $\mathbf{r}_k$  is an embedding vector of  $d$ -dimensions.

### 5.3 Sentence Encoding Layer

The context utterances, profile sentences and response candidate are encoded by bidirectional long short-term memories (BiLSTMs) (Hochreiter and Schmidhuber, 1997). We denote the calculations as follows,

$$\bar{\mathbf{u}}_{m,i} = \text{BiLSTM}(\mathbf{U}_m, i), i \in \{1, \dots, l_{u_m}\}, \quad (4)$$

$$\bar{\mathbf{p}}_{n,j} = \text{BiLSTM}(\mathbf{P}_n, j), j \in \{1, \dots, l_{p_n}\}, \quad (5)$$

$$\bar{\mathbf{r}}_k = \text{BiLSTM}(\mathbf{R}, k), k \in \{1, \dots, l_r\}, \quad (6)$$

where  $\bar{\mathbf{U}}_m = \{\bar{\mathbf{u}}_{m,i}\}_{i=1}^{l_{u_m}}$ ,  $\bar{\mathbf{P}}_n = \{\bar{\mathbf{p}}_{n,j}\}_{j=1}^{l_{p_n}}$  and  $\bar{\mathbf{R}} = \{\bar{\mathbf{r}}_k\}_{k=1}^{l_r}$ . The parameters of these three BiLSTMs are shared in our implementation.

### 5.4 Matching Layer

The interactions between the context and the response and those between the persona and the response can provide useful matching information for deciding the matching degree between them. Here, the DIM model adopts the same strategy as in the IMN model (Gu et al., 2019) which considers the global and bidirectional interactions between two sequences.

Take the context-response matching as an example. First, the context representation  $\bar{\mathbf{C}} = \{\bar{\mathbf{c}}_i\}_{i=1}^{l_c}$  with  $l_c = \sum_{m=1}^{n_c} l_{u_m}$  is formed by




Figure 2: An overview of our proposed DIM model.

concatenating the set of utterance representations  $\{\bar{\mathbf{U}}_m\}_{m=1}^{n_c}$ . Then, a soft alignment is performed by computing the attention weight between each tuple  $\{\tilde{\mathbf{c}}_i, \tilde{\mathbf{r}}_k\}$  as

$$e_{ik} = (\tilde{\mathbf{c}}_i)^\top \cdot \tilde{\mathbf{r}}_k. \quad (7)$$

After that, local inference is determined by the attention weights computed above to obtain the local relevance between a context and a response bidirectionally. For a word in the context, its relevant representation carried by the response is identified and composed using  $e_{ik}$  as

$$\tilde{\mathbf{c}}_i = \sum_{k=1}^{l_r} \frac{\exp(e_{ik})}{\sum_{l=1}^{l_r} \exp(e_{il})} \tilde{\mathbf{r}}_k, i \in \{1, \dots, l_c\}, \quad (8)$$

where the contents in  $\{\tilde{\mathbf{r}}_k\}_{k=1}^{l_r}$  that are relevant to  $\tilde{\mathbf{c}}_i$  are selected to form  $\tilde{\mathbf{c}}_i$ . Then, we define  $\tilde{\mathbf{C}} = [\tilde{\mathbf{c}}_1, \dots, \tilde{\mathbf{c}}_{l_c}]$ . The same calculation is performed for each word in the response to form its relevant representation carried by the context as

$$\tilde{\mathbf{r}}_k = \sum_{i=1}^{l_c} \frac{\exp(e_{ik})}{\sum_{l=1}^{l_c} \exp(e_{lk})} \tilde{\mathbf{c}}_i, k \in \{1, \dots, l_r\}, \quad (9)$$

and we define  $\tilde{\mathbf{R}} = [\tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_{l_r}]$ . To further enhance the collected information, the differences and element-wise products between  $\{\tilde{\mathbf{C}}, \tilde{\mathbf{C}}\}$  and between  $\{\tilde{\mathbf{R}}, \tilde{\mathbf{R}}\}$  are computed, and are then concatenated with the original vectors to obtain the enhanced representations as follows,

$$\hat{\mathbf{C}} = [\tilde{\mathbf{C}}; \tilde{\mathbf{C}}; \tilde{\mathbf{C}} - \tilde{\mathbf{C}}; \tilde{\mathbf{C}} \odot \tilde{\mathbf{C}}], \quad (10)$$

$$\hat{\mathbf{R}} = [\tilde{\mathbf{R}}; \tilde{\mathbf{R}}; \tilde{\mathbf{R}} - \tilde{\mathbf{R}}; \tilde{\mathbf{R}} \odot \tilde{\mathbf{R}}]. \quad (11)$$

So far we have collected the relevant information between the context and response. The

enhanced context representation  $\hat{\mathbf{C}}$  is further converted back to matching matrices of separated utterances as  $\{\hat{\mathbf{U}}_m\}_{m=1}^{n_c}$ .

The persona-response matching is conducted identically to the context-response matching introduced above, where the representations of profile sentences  $\{\hat{\mathbf{P}}_n\}_{n=1}^{n_p}$  are used, instead of the representations of context utterances  $\{\bar{\mathbf{U}}_m\}_{m=1}^{n_c}$ . The results of persona-response matching are denoted as  $\{\hat{\mathbf{P}}_n\}_{n=1}^{n_p}$  and  $\hat{\mathbf{R}}^*$ .

## 5.5 Aggregation Layer

The aggregation layer converts the matching matrices of context utterances, profile sentences and response into a final matching feature vector.

First, each matching matrix  $\hat{\mathbf{U}}_m, \hat{\mathbf{R}}, \hat{\mathbf{P}}_n$  and  $\hat{\mathbf{R}}^*$  are processed by BiLSTMs as

$$\hat{\mathbf{u}}_{m,i}^{utr} = \text{BiLSTM}(\hat{\mathbf{U}}_m, i), i \in \{1, \dots, l_{u_m}\}, \quad (12)$$

$$\hat{\mathbf{r}}_k^{utr} = \text{BiLSTM}(\hat{\mathbf{R}}, k), k \in \{1, \dots, l_r\}, \quad (13)$$

$$\hat{\mathbf{p}}_{n,j}^{utr} = \text{BiLSTM}(\hat{\mathbf{P}}_n, j), j \in \{1, \dots, l_{p_n}\}, \quad (14)$$

$$\hat{\mathbf{r}}_k^{utr*} = \text{BiLSTM}(\hat{\mathbf{R}}^*, k), k \in \{1, \dots, l_r\}, \quad (15)$$

where the four BiLSTMs share the same parameters in our implementation. Then, the aggregated embeddings are calculated by max pooling and last-hidden-state pooling operations as

$$\hat{\mathbf{u}}_m^{agr} = [\hat{\mathbf{u}}_{m,max}^{utr}; \hat{\mathbf{u}}_{m,l_{u_m}}^{utr}], m \in \{1, \dots, n_c\}, \quad (16)$$

$$\hat{\mathbf{r}}^{agr} = [\hat{\mathbf{r}}_{max}^{utr}; \hat{\mathbf{r}}_{l_r}^{utr}], \quad (17)$$

$$\hat{\mathbf{p}}_n^{agr} = [\hat{\mathbf{p}}_{n,max}^{utr}; \hat{\mathbf{p}}_{n,l_{p_n}}^{utr}], n \in \{1, \dots, n_p\}, \quad (18)$$

$$\hat{\mathbf{r}}^{agr*} = [\hat{\mathbf{r}}_{max}^{utr*}; \hat{\mathbf{r}}_{l_r}^{utr*}]. \quad (19)$$

Next, the sequences of  $\hat{\mathbf{u}}_m^{agr}$  and  $\hat{\mathbf{p}}_n^{agr}$  are further aggregated to get the embedding vectors for the context and the persona respectively.

**Context aggregation** As the utterances in a context are chronologically ordered, the utterance embeddings  $\hat{\mathbf{U}}^{agr} = \{\hat{\mathbf{u}}_m^{agr}\}_{m=1}^{n_c}$  are sent into another BiLSTM following the chronological order of utterances in the context. Combined max pooling and last-hidden-state pooling operations are then performed to obtain the context embeddings as

$$\hat{\mathbf{u}}_m^{ctx} = \text{BiLSTM}(\hat{\mathbf{U}}^{agr}, m), m \in \{1, \dots, n_c\}, \quad (20)$$

$$\hat{\mathbf{u}}^{agr} = [\hat{\mathbf{u}}_{max}^{ctx}; \hat{\mathbf{u}}_{n_c}^{ctx}]. \quad (21)$$

**Persona aggregation** As the profile sentences in a persona are independent to each other, an attention-based aggregation is designed to derive the persona embeddings as follows,

$$\alpha_n = \text{ReLU}(\mathbf{w}^\top \cdot \hat{\mathbf{p}}_n^{agr} + b), \quad (22)$$

$$\hat{\mathbf{p}}^{agr} = \sum_{n=1}^{n_p} \frac{e^{\alpha_n}}{\sum_{k=1}^{n_p} e^{\alpha_k}} \hat{\mathbf{p}}_n^{agr}, \quad (23)$$

where  $\mathbf{w}$  and  $b$  are parameters need to be estimated during training.

Last, the final matching feature vector is the concatenation of context, persona and response embeddings as

$$\mathbf{m} = [\hat{\mathbf{c}}^{agr}; \hat{\mathbf{r}}^{agr}; \hat{\mathbf{p}}^{agr}; \hat{\mathbf{r}}^{agr*}], \quad (24)$$

where the first two features describe the context-response matching, and the last two describe the persona-response matching.

## 5.6 Prediction Layer

The final matching feature vector is then sent into a multi-layer perceptron (MLP) classifier with softmax output. Here, the MLP is designed to predict whether a  $(c, p, r)$  triple match appropriately based on the derived matching feature vector. Finally, the MLP returns a probability to denote the matching degree.

## 6 Experiments

### 6.1 Dataset

We tested our proposed methods on the PERSONA-CHAT dataset (Zhang et al., 2018) which contains multi-turn dialogues conditioned on personas. The dataset consists of 8939 complete dialogues for training, 1000 for validation, and 968 for testing. Response selection is performed at every turn of a complete dialogue, which results in 65719 dialogues for

training, 7801 for validation, and 7512 for testing in total. Positive responses are true responses from humans and negative ones are randomly sampled. The ratio between positive and negative responses is 1:19 in the training, validation, and testing sets. There are 955 possible personas for training, 100 for validation, and 100 for testing, each consisting of 3 to 5 profile sentences. To make this task more challenging, a version of revised persona descriptions are also provided by rephrasing, generalizing, or specializing the original ones. Since the personas of both speakers in a dialogue are available, the response selection task can be conditioned on the speaker’s persona (“self persona”) or the dialogue partner’s persona (“their persona”) respectively.

### 6.2 Evaluation Metrics

We used the same evaluation metrics as in the previous work (Zhang et al., 2018). Each model aimed to select the best-matched response from available candidates for the given context  $c$  and persona  $p$ . We calculated the recall of the true positive replies, denoted as **hits@1**. In addition, the mean reciprocal rank (**MRR**) (Voorhees, 1999) metric was also adopted to take the rank of the correct response over all candidates into consideration.

### 6.3 Training Details

For building the IMN,  $\text{IMN}_{ctx}$ ,  $\text{IMN}_{utr}$  and DIM models, the Adam method (Kingma and Ba, 2015) was employed for optimization with a batch size of 16. The initial learning rate was 0.001 and was exponentially decayed by 0.96 every 5000 steps. Dropout (Srivastava et al., 2014) with a rate of 0.2 was applied to the word embeddings and all hidden layers. A word representation is a concatenation of a 300-dimensional GloVe embedding (Pennington et al., 2014), a 100-dimensional embedding estimated on the training set using the Word2Vec algorithm (Mikolov et al., 2013), and 150-dimensional character-level embeddings with window sizes  $\{3, 4, 5\}$ , each consisting of 50 filters. The word embeddings were not updated during training. All hidden states of the LSTM have 200 dimensions. The MLP at the prediction layer have 256 hidden units with ReLU (Nair and Hinton, 2010) activation. The maximum number of characters in a word, that of words in a context utterance, of utterances in a context, and of words in a response were set to be 18, 20, 15, and

	<b>hits@1</b>	<b>MRR</b>
IR baseline	21.4	-
Starspace	31.8	-
Profile	31.8	-
KV Profile	34.9	-
IMN	63.8	75.8

Table 2: Evaluation results of the IMN model and previous methods on PERSONA-CHAT dataset without using personas. All the results except ours are copied from Zhang et al. (2018).

20, respectively. We padded with zeros if the number of utterances in a context was less than 15; otherwise, we kept the last 15 utterances. For the  $IMN_{ctx}$ ,  $IMN_{utr}$  and the DIM models, the maximum number of words in a profile sentence and that of profile sentences in a persona were set to be 15 and 5, respectively. Similarly, we padded with zeros if the number of profile sentences in a persona was less than 5. The development set was used to select the best model for testing.

All code was implemented in the TensorFlow framework (Abadi et al., 2016) and is published to help replicate our results <sup>1</sup>.

#### 6.4 Experimental Results

Table 2 presents the evaluation results of our reproduced IMN model (Gu et al., 2019) and previous methods on PERSONA-CHAT dataset without using personas. It can be seen that the IMN model outperformed other models on this dataset by a margin larger than 28.9% in terms of **hits@1**. As introduced above, our proposed models for personalized response selection were all built on IMN.

Table 3 presents the evaluation results of our proposed and previous methods on PERSONA-CHAT under various persona configurations. The t-test shows that the differences between our proposed models, i.e.,  $IMN_{utt}$  and DIM, and the baseline model, i.e.  $IMN_{ctx}$ , were both statistically significant with  $p$ -value  $< 0.01$ . We can see that the fine-grained persona fusion at the utterance level rendered a **hits@1** improvement of 2.4% and an **MRR** improvement of 1.9% by comparing  $IMN_{ctx}$  and  $IMN_{utr}$  conditioned on original self personas. The DIM model outperformed its baseline  $IMN_{ctx}$  by a margin of 14.5% in terms of **hits@1** and 10.5% in terms of **MRR**. Compared

with the FT-PC model (Mazaré et al., 2018) which was first pretrained using a large-scale corpus and then fine-tuned on the PERSONA-CHAT dataset, the DIM model outperformed it by a margin of 10.0% in terms of **hits@1** conditioned on revised self personas. Another advantage of DIM is that it was trained in an end-to-end mode without pre-training and using any external knowledge. Lastly, the DIM model outperforms previous models by margins larger than 27.7% in terms of **hits@1** conditioned on original self personas.

**Improvement of Using Personas** Examining the numbers which indicate the gains or losses after adding persona conditions in Table 3, we can see that the context-level persona fusion improves the performance of previous models significantly when original self personas are used. However, the gain achieved by the  $IMN_{ctx}$  model is limited. One possible reason is that the IMN model performs attention-based interactions between the context and the response in order to get their matching feature for response selection. Thus, the context embeddings shown in Fig. 1(a) contained the information from both the context and the response, which may be inappropriate for the following context-level persona fusion shown in Eq. (1). The improvement achieved by the DIM model is much higher because it adopts a dual matching framework to address this issue.

**Original vs. Revised** Comparing with using original personas, it is more difficult for the models conditioned on the revised personas to extract useful persona information, as shown by the limited improvement achieved by the previous models shown in Table 3. One possible reason is that there are fewer shared words between the response and the persona revised by rephrasing, generalizing, or specializing, which increases the difficulty of understanding the persona and its relationships with the response. For example, it is easier for models to judge the matching degree between the original profile “Autumn is my favorite season.” and the response “This is my favorite time of the year season wise.” than between the revised profile “I love watching the leaves change colors.” and the response. On the contrary, our proposed DIM model still obtains a **hits@1** improvement of 6.9% and an **MRR** improvement of 5.4% when conditioned on the revised self personas, which can be attributed to the direct and interactive

<sup>1</sup><https://github.com/JasonForJoy/DIM>

	Self Persona				Their Persona			
	Original		Revised		Original		Revised	
	hits@1	MRR	hits@1	MRR	hits@1	MRR	hits@1	MRR
IR baseline	41.0 (+19.6)	-	20.7 (-0.7)	-	18.1 (-3.3)	-	18.1 (-3.3)	-
Starspace	48.1 (+16.3)	-	32.2 (+0.4)	-	24.5 (-7.3)	-	26.1 (-5.7)	-
Profile	47.3 (+15.5)	-	35.4 (+3.6)	-	28.3 (-3.5)	-	29.4 (-2.4)	-
KV Profile	51.1 (+16.2)	-	35.1 (+0.2)	-	29.1 (-5.8)	-	28.9 (-6.0)	-
FT-PC	-	-	60.7 (-)	-	-	-	-	-
IMN <sub>ctx</sub>	64.3 (+0.5)	76.2 (+0.4)	63.8 (+0.0)	75.8 (+0.0)	63.7 (-0.1)	75.8 (+0.0)	63.5 (-0.3)	75.7 (-0.1)
IMN <sub>utr</sub>	66.7 (+2.9)	78.1 (+2.3)	64.0 (+0.2)	76.0 (+0.2)	63.9 (+0.1)	75.9 (+0.1)	63.7 (-0.1)	75.7 (-0.1)
DIM	<b>78.8 (+15.0)</b>	<b>86.7 (+10.9)</b>	<b>70.7 (+6.9)</b>	<b>81.2 (+5.4)</b>	<b>64.0 (+0.2)</b>	<b>76.1 (+0.3)</b>	<b>63.9 (+0.1)</b>	<b>76.0 (+0.2)</b>

Table 3: Performance of the proposed and previous methods on the PERSONA-CHAT under various persona configurations. The meanings of ‘‘Self Persona’’, ‘‘Their Persona’’, ‘‘Original’’, and ‘‘revised’’ can be found in Section 6.1. All results except ours are copied from Zhang et al. (2018); Mazaré et al. (2018). Numbers in parentheses indicate the gains or losses after adding the persona conditions.

persona-response matching used in this model.

**Self vs. Their** As shown in Table 3, no significant gains can be obtained when the models are conditioned on the personas of dialogue partners. Note that there are no significant performance losses with our proposed methods, while the losses of previous models are 2.4% to 7.3% in terms of hits@1.

## 7 Analysis

### 7.1 Ablations

	hits@1	MRR
DIM	78.8	86.7
- persona	63.8	75.8
- context	48.8	60.9

Table 4: Ablation tests of removing either persona-response matching or context-response matching in the DIM model conditioned on original self personas.

To demonstrate the importance of the dual matching framework followed by our proposed DIM model, ablation tests were performed using the original self personas and the results are shown in Table 4. We can see that both the persona-response matching and the context-response matching contribute to the performance of the DIM model. It is reasonable that the context-persona matching is more important because contexts provide the fundamental semantic descriptions for response selection. On the other hand, the single persona-response matching can also achieve a hits@1 of 48.8% and an MRR of 60.9%, which shows the usefulness of utilizing persona information to select the best-matched response.

### 7.2 Interactive Matching in DIM

In order to investigate the effectiveness of the interaction matching between the context and the response and that between the persona and the response in the DIM model, a case study was conducted by visualizing the response-to-context and response-to-persona attention weights used in Eq. (9). The results are shown in Fig. 3. We can see that some important words such as ‘‘dogs’’ in the response selected its relevant words such as ‘‘animals’’ in the context to derive the context-response matching features. Some important profile texts such as ‘‘I love animals and have two dogs.’’ also obtained large attention weights for getting the persona-response matching features. This experimental result confirms our formulation of the task of personalized response selection as a dual matching problem.


### 7.3 Transfer Test

Train \ Test	Original	Revised
	Original	78.8
Revised	77.6	70.7


Table 5: hits@1 results of transfer tests on the DIM model.

Transfer tests were conducted by training and evaluating the DIM model using mismatched types of personas. The results are reported in Table 5. It shows that the DIM model achieved a better performance when testing on the same type of personas as training. Meanwhile, the model trained on the revised personas and tested on the original personas achieved less loss than the ones trained





(a) Response-to-context attention weights



(b) Response-to-persona attention weights

Figure 3: Visualizations of (a) response-to-context or (b) response-to-persona attention weights at the matching layer for a test sample. The darker units correspond to larger values.

on the original personas and tested on the revised personas, which shows that the revised personas can provide a better generalization ability to the DIM model than the original ones.

## 8 Conclusions

In this paper, we formulate the task of personalized response selection as a dual matching problem to search for a response that can properly match the given context and persona simultaneously. A new model named Dually Interactive Matching Network (DIM) is proposed, which performs the interaction matching between the context and response as well as between the persona and response in parallel, in order to derive the final matching features for personalized response selection. Experimental results show that DIM improves over the IMN models with context-level or utterance-level persona fusion, outperforming previous methods and achieving a new state-of-the-art performance on the PERSONA-CHAT dataset. In the future, we will explore models to make better use of dialogue partners' persona for response selection.

## Acknowledgments

We thank the anonymous reviewers for their valuable comments.

## References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek Gordon Murray, Benoit Steiner, Paul A. Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. [Tensorflow: A system for large-scale machine learning](#). In *12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016, Savannah, GA, USA, November 2-4, 2016.*, pages 265–283.
- Jia-Chen Gu, Zhen-Hua Ling, and Quan Liu. 2019. Interactive matching network for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management. ACM*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P. Spithourakis, Jianfeng Gao, and William B. Dolan. 2016. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- 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 SIGDIAL 2015 Conference, The 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 2-4 September 2015, Prague, Czech Republic*, pages 285–294.
- Pierre-Emmanuel Mazaré, Samuel Humeau, Martin Raison, and Antoine Bordes. 2018. [Training millions of personalized dialogue agents](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2775–2779.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. [Distributed](#)

- representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 3111–3119.
- Alexander H. Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. [Key-value memory networks for directly reading documents](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1400–1409.
- Vinod Nair and Geoffrey E. Hinton. 2010. [Rectified linear units improve restricted boltzmann machines](#). In *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 807–814.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar; A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543.
- Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: a simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15(1):1929–1958.
- Ellen M. Voorhees. 1999. [The TREC-8 question answering track report](#). In *Proceedings of The Eighth Text REtrieval Conference, TREC 1999, Gaithersburg, Maryland, USA, November 17-19, 1999*.
- Hao Wang, Zhengdong Lu, Hang Li, and Enhong Chen. 2013. [A dataset for research on short-text conversations](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 935–945.
- 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, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 496–505.
- 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?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2204–2213.
- Xiangyang Zhou, Lu Li, Daxiang Dong, Yi Liu, Ying Chen, Wayne Xin Zhao, Dianhai Yu, and Hua Wu. 2018. [Multi-turn response selection for chatbots with deep attention matching network](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1118–1127.